

EIGHTH EDITION

**OPTIONS, FUTURES,
AND OTHER DERIVATIVES**

EIGHTH EDITION

OPTIONS, FUTURES, AND OTHER DERIVATIVES

John C. Hull

*Maple Financial Group Professor of Derivatives and Risk Management
Joseph L. Rotman School of Management
University of Toronto*

Prentice Hall

Boston Columbus Indianapolis New York San Francisco Upper Saddle River
Amsterdam Cape Town Dubai London Madrid Milan Munich Paris Montreal Toronto
Delhi Mexico City São Paulo Sydney Hong Kong Seoul Singapore Taipei Tokyo

Editor in Chief: Donna Battista
Acquisitions Editor: Tessa O'Brien
Editorial Assistant: Elissa Senra-Sargent
Marketing Director: Patrice Jones
Marketing Coordinator: Ian Gold
Managing Editor: Nancy Fenton
Production Project Manager: Alison Eusden
Senior Manufacturing Buyer: Carol Melville
Senior Media Manufacturing Buyer: Ginny Michaud
Media Project Manager: Nicole Sackin
Cover Image: Norebbo/Shutterstock
Composition: The Geometric Press
Printer/Binder: Edwards Brothers
Cover Printer: Lehigh/Phoenix
Text Font: 10/12 pt Monotype Times

Credits and acknowledgments borrowed from other sources and reproduced, with permission, in this textbook appear on appropriate page within text.

Copyright © 2012, 2009, 2006, 2003, 2000 by Pearson Education, Inc., publishing as Prentice Hall. All rights reserved. Printed in the United States of America. This publication is protected by Copyright and permission should be obtained from the publisher prior to any prohibited reproduction, storage in a retrieval system, or transmission in any form or by any means, electronic, mechanical, photocopying, recording, or likewise. To obtain permission(s) to use material from this work, please submit a written request to Pearson Education, Inc., Rights and Contracts Department, 501 Boylston Street, Suite 900, Boston, MA 02116, fax your request to 617-671-3447, or e-mail at <http://www.pearsoned.com/legal/permission.htm>.

Many of the designations by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book and the publisher was aware of a trademark claim, the designations have been printed in initial caps or all caps.

Library of Congress Cataloging-in-Publication Data

Hull, John, 1946–

Options, futures, and other derivatives / John C. Hull.—8th ed.
p. cm.

Includes bibliographical references and index.

ISBN 978-0-13-216494-8 (hbk.)

1. Futures. 2. Stock options. 3. Derivative securities. I. Title.

HG6024.A3H85 2012

332.64'52—dc22

2010047005

Prentice Hall
is an imprint of



10 9 8 7 6 5 4 3 2 1
ISBN-13: 978-0-13-216494-8
ISBN-10: 0-13-216494-9

To Michelle

CONTENTS IN BRIEF

| | |
|---|--------------|
| List of Business Snapshots | <i>xvii</i> |
| List of Technical Notes | <i>xviii</i> |
| Preface | <i>xix</i> |
| 1. Introduction | 1 |
| 2. Mechanics of futures markets | 22 |
| 3. Hedging strategies using futures | 47 |
| 4. Interest rates | 75 |
| 5. Determination of forward and futures prices | 101 |
| 6. Interest rate futures | 129 |
| 7. Swaps | 148 |
| 8. Securitization and the credit crisis of 2007 | 180 |
| 9. Mechanics of options markets | 194 |
| 10. Properties of stock options | 214 |
| 11. Trading strategies involving options | 234 |
| 12. Binomial trees | 253 |
| 13. Wiener processes and Itô's lemma | 280 |
| 14. The Black–Scholes–Merton model | 299 |
| 15. Employee stock options | 332 |
| 16. Options on stock indices and currencies | 345 |
| 17. Futures options | 361 |
| 18. The Greek letters | 377 |
| 19. Volatility smiles | 409 |
| 20. Basic numerical procedures | 427 |
| 21. Value at risk | 471 |
| 22. Estimating volatilities and correlations | 498 |
| 23. Credit risk | 521 |
| 24. Credit derivatives | 547 |
| 25. Exotic options | 574 |
| 26. More on models and numerical procedures | 599 |
| 27. Martingales and measures | 630 |
| 28. Interest rate derivatives: The standard market models | 648 |
| 29. Convexity, timing, and quanto adjustments | 668 |
| 30. Interest rate derivatives: Models of the short rate | 682 |
| 31. Interest rate derivatives: HJM and LMM | 715 |
| 32. Swaps Revisited | 733 |
| 33. Energy and commodity derivatives | 748 |
| 34. Real options | 765 |
| 35. Derivatives mishaps and what we can learn from them | 779 |
| Glossary of terms | 791 |
| DerivaGem software | 812 |
| Major exchanges trading futures and options | 817 |
| Tables for $N(x)$ | 818 |
| Author index | 820 |
| Subject index | 824 |

Contents

| | | |
|-------------------|--|--------------|
| | List of Business Snapshots..... | xvii |
| | List of Technical Notes..... | xviii |
| | Preface | xix |
| Chapter 1. | Introduction | 1 |
| | 1.1 Exchange-traded markets | 2 |
| | 1.2 Over-the-counter markets | 3 |
| | 1.3 Forward contracts..... | 5 |
| | 1.4 Futures contracts | 7 |
| | 1.5 Options | 7 |
| | 1.6 Types of traders..... | 9 |
| | 1.7 Hedgers..... | 10 |
| | 1.8 Speculators | 13 |
| | 1.9 Arbitrageurs..... | 15 |
| | 1.10 Dangers | 16 |
| | Summary..... | 16 |
| | Further reading | 18 |
| | Practice questions..... | 18 |
| | Further questions..... | 20 |
| Chapter 2. | Mechanics of futures markets..... | 22 |
| | 2.1 Background | 22 |
| | 2.2 Specification of a futures contract..... | 24 |
| | 2.3 Convergence of futures price to spot price | 26 |
| | 2.4 The operation of margins..... | 27 |
| | 2.5 OTC markets | 30 |
| | 2.6 Market quotes..... | 33 |
| | 2.7 Delivery | 36 |
| | 2.8 Types of traders and types of orders..... | 37 |
| | 2.9 Regulation | 38 |
| | 2.10 Accounting and tax..... | 39 |
| | 2.11 Forward vs. futures contracts..... | 41 |
| | Summary..... | 42 |
| | Further reading | 43 |
| | Practice questions..... | 43 |
| | Further questions..... | 45 |
| Chapter 3. | Hedging strategies using futures..... | 47 |
| | 3.1 Basic principles..... | 47 |
| | 3.2 Arguments for and against hedging | 49 |
| | 3.3 Basis risk..... | 52 |
| | 3.4 Cross hedging | 56 |

| | | |
|-------------------|---|------------|
| 3.5 | Stock index futures..... | 60 |
| 3.6 | Stack and roll..... | 65 |
| | Summary..... | 67 |
| | Further reading..... | 68 |
| | Practice questions..... | 69 |
| | Further questions..... | 71 |
| | Appendix: Capital asset pricing model..... | 73 |
| Chapter 4. | Interest rates..... | 75 |
| 4.1 | Types of rates..... | 75 |
| 4.2 | Measuring interest rates..... | 77 |
| 4.3 | Zero rates..... | 80 |
| 4.4 | Bond pricing..... | 80 |
| 4.5 | Determining Treasury zero rates..... | 82 |
| 4.6 | Forward rates..... | 84 |
| 4.7 | Forward rate agreements..... | 86 |
| 4.8 | Duration..... | 89 |
| 4.9 | Convexity..... | 92 |
| 4.10 | Theories of the term structure of interest rates..... | 93 |
| | Summary..... | 96 |
| | Further reading..... | 97 |
| | Practice questions..... | 97 |
| | Further questions..... | 99 |
| Chapter 5. | Determination of forward and futures prices..... | 101 |
| 5.1 | Investment assets vs. consumption assets..... | 101 |
| 5.2 | Short selling..... | 102 |
| 5.3 | Assumptions and notation..... | 103 |
| 5.4 | Forward price for an investment asset..... | 104 |
| 5.5 | Known income..... | 107 |
| 5.6 | Known yield..... | 109 |
| 5.7 | Valuing forward contracts..... | 109 |
| 5.8 | Are forward prices and futures prices equal?..... | 111 |
| 5.9 | Futures prices of stock indices..... | 112 |
| 5.10 | Forward and futures contracts on currencies..... | 114 |
| 5.11 | Futures on commodities..... | 117 |
| 5.12 | The cost of carry..... | 120 |
| 5.13 | Delivery options..... | 121 |
| 5.14 | Futures prices and the expected future spot price..... | 121 |
| | Summary..... | 123 |
| | Further reading..... | 125 |
| | Practice questions..... | 125 |
| | Further questions..... | 127 |
| Chapter 6. | Interest rate futures..... | 129 |
| 6.1 | Day count and quotation conventions..... | 129 |
| 6.2 | Treasury bond futures..... | 132 |
| 6.3 | Eurodollar futures..... | 137 |
| 6.4 | Duration-based hedging strategies using futures..... | 142 |
| 6.5 | Hedging portfolios of assets and liabilities..... | 143 |
| | Summary..... | 144 |
| | Further reading..... | 145 |
| | Practice questions..... | 145 |
| | Further questions..... | 147 |

| | |
|--|------------|
| Chapter 7. Swaps | 148 |
| 7.1 Mechanics of interest rate swaps | 148 |
| 7.2 Day count issues..... | 154 |
| 7.3 Confirmations | 155 |
| 7.4 The comparative-advantage argument | 156 |
| 7.5 The nature of swap rates..... | 158 |
| 7.6 Determining the LIBOR/swap zero rates | 159 |
| 7.7 Valuation of interest rate swaps..... | 160 |
| 7.8 Overnight indexed swaps..... | 164 |
| 7.9 Currency swaps | 165 |
| 7.10 Valuation of currency swaps | 168 |
| 7.11 Credit risk | 171 |
| 7.12 Other types of swaps | 173 |
| Summary..... | 175 |
| Further reading | 176 |
| Practice questions..... | 176 |
| Further questions..... | 178 |
| Chapter 8. Securitization and the Credit Crisis of 2007 | 180 |
| 8.1 Securitization | 180 |
| 8.2 The US housing market..... | 184 |
| 8.3 What went wrong?..... | 188 |
| 8.4 The aftermath | 190 |
| Summary..... | 191 |
| Further reading | 192 |
| Practice questions..... | 193 |
| Further questions..... | 193 |
| Chapter 9. Mechanics of options markets | 194 |
| 9.1 Types of options..... | 194 |
| 9.2 Option positions | 196 |
| 9.3 Underlying assets..... | 198 |
| 9.4 Specification of stock options | 199 |
| 9.5 Trading | 203 |
| 9.6 Commissions..... | 204 |
| 9.7 Margins | 205 |
| 9.8 The options clearing corporation..... | 206 |
| 9.9 Regulation | 207 |
| 9.10 Taxation..... | 207 |
| 9.11 Warrants, employee stock options, and convertibles | 209 |
| 9.12 Over-the-counter markets..... | 210 |
| Summary..... | 210 |
| Further reading | 211 |
| Practice questions..... | 211 |
| Further questions..... | 213 |
| Chapter 10. Properties of stock options | 214 |
| 10.1 Factors affecting option prices | 214 |
| 10.2 Assumptions and notation | 218 |
| 10.3 Upper and lower bounds for option prices..... | 218 |
| 10.4 Put–call parity | 221 |
| 10.5 Calls on a non-dividend-paying stock | 225 |
| 10.6 Puts on a non-dividend-paying stock | 226 |
| 10.7 Effect of dividends | 229 |

| | | |
|--|--|------------|
| | Summary | 230 |
| | Further reading..... | 231 |
| | Practice questions..... | 231 |
| | Further questions | 233 |
| Chapter 11. Trading strategies involving options..... | | 234 |
| 11.1 | Principal-protected notes | 234 |
| 11.2 | Trading an option and the underlying asset | 236 |
| 11.3 | Spreads..... | 238 |
| 11.4 | Combinations | 246 |
| 11.5 | Other payoffs..... | 249 |
| | Summary | 249 |
| | Further reading..... | 250 |
| | Practice questions..... | 250 |
| | Further questions | 252 |
| Chapter 12. Binomial trees | | 253 |
| 12.1 | A one-step binomial model and a no-arbitrage argument | 253 |
| 12.2 | Risk-neutral valuation..... | 257 |
| 12.3 | Two-step binomial trees | 259 |
| 12.4 | A put example..... | 262 |
| 12.5 | American options..... | 263 |
| 12.6 | Delta..... | 264 |
| 12.7 | Matching volatility with u and d | 265 |
| 12.8 | The binomial tree formulas..... | 267 |
| 12.9 | Increasing the number of steps..... | 268 |
| 12.10 | Using DerivaGem | 269 |
| 12.11 | Options on other assets..... | 269 |
| | Summary | 272 |
| | Further reading..... | 273 |
| | Practice questions..... | 274 |
| | Further questions | 275 |
| | Appendix: Derivation of the Black–Scholes–Merton option-pricing formula from a binomial tree..... | 276 |
| Chapter 13. Wiener processes and Itô’s lemma | | 280 |
| 13.1 | The Markov property | 280 |
| 13.2 | Continuous-time stochastic processes..... | 281 |
| 13.3 | The process for a stock price | 286 |
| 13.4 | The parameters..... | 289 |
| 13.5 | Correlated processes | 290 |
| 13.6 | Itô’s lemma | 291 |
| 13.7 | The lognormal property | 292 |
| | Summary | 293 |
| | Further reading..... | 294 |
| | Practice questions..... | 294 |
| | Further questions | 295 |
| | Appendix: Derivation of Itô’s lemma | 297 |
| Chapter 14. The Black–Scholes–Merton model..... | | 299 |
| 14.1 | Lognormal property of stock prices | 300 |
| 14.2 | The distribution of the rate of return | 301 |
| 14.3 | The expected return..... | 302 |
| 14.4 | Volatility..... | 303 |
| 14.5 | The idea underlying the Black–Scholes–Merton differential equation..... | 307 |

| | | |
|--------------------|---|------------|
| 14.6 | Derivation of the Black–Scholes–Merton differential equation | 309 |
| 14.7 | Risk-neutral valuation | 311 |
| 14.8 | Black–Scholes–Merton pricing formulas | 313 |
| 14.9 | Cumulative normal distribution function | 315 |
| 14.10 | Warrants and employee stock options | 316 |
| 14.11 | Implied volatilities..... | 318 |
| 14.12 | Dividends | 320 |
| | Summary..... | 323 |
| | Further reading | 324 |
| | Practice questions..... | 325 |
| | Further questions..... | 328 |
| | Appendix: Proof of Black–Scholes–Merton formula using risk-neutral valuation | 329 |
| Chapter 15. | Employee stock options..... | 332 |
| 15.1 | Contractual arrangements..... | 332 |
| 15.2 | Do options align the interests of shareholders and managers?..... | 334 |
| 15.3 | Accounting issues | 335 |
| 15.4 | Valuation..... | 336 |
| 15.5 | Backdating scandals | 341 |
| | Summary..... | 342 |
| | Further reading | 343 |
| | Practice questions..... | 343 |
| | Further questions..... | 344 |
| Chapter 16. | Options on stock indices and currencies | 345 |
| 16.1 | Options on stock indices | 345 |
| 16.2 | Currency options | 347 |
| 16.3 | Options on stocks paying known dividend yields | 350 |
| 16.4 | Valuation of European stock index options..... | 352 |
| 16.5 | Valuation of European currency options..... | 355 |
| 16.6 | American options | 356 |
| | Summary..... | 357 |
| | Further reading | 357 |
| | Practice questions..... | 358 |
| | Further questions..... | 360 |
| Chapter 17. | Futures options..... | 361 |
| 17.1 | Nature of futures options | 361 |
| 17.2 | Reasons for the popularity of futures options | 364 |
| 17.3 | European spot and futures options | 364 |
| 17.4 | Put–call parity | 365 |
| 17.5 | Bounds for futures options..... | 366 |
| 17.6 | Valuation of futures options using binomial trees | 367 |
| 17.7 | Drift of a futures prices in a risk-neutral world | 369 |
| 17.8 | Black’s model for valuing futures options..... | 370 |
| 17.9 | American futures options vs. American spot options..... | 372 |
| 17.10 | Futures-style options | 372 |
| | Summary..... | 373 |
| | Further reading | 374 |
| | Practice questions..... | 374 |
| | Further questions..... | 376 |
| Chapter 18. | The Greek letters..... | 377 |
| 18.1 | Illustration | 377 |

| | | |
|--------------------|---|------------|
| 21.4 | Linear model | 481 |
| 21.5 | Quadratic model | 486 |
| 21.6 | Monte Carlo simulation..... | 488 |
| 21.7 | Comparison of approaches..... | 489 |
| 21.8 | Stress testing and back testing..... | 490 |
| 21.9 | Principal components analysis..... | 490 |
| | Summary | 494 |
| | Further reading | 494 |
| | Practice questions..... | 495 |
| | Further questions..... | 497 |
| Chapter 22. | Estimating volatilities and correlations | 498 |
| 22.1 | Estimating volatility | 498 |
| 22.2 | The exponentially weighted moving average model..... | 500 |
| 22.3 | The GARCH (1,1) model | 502 |
| 22.4 | Choosing between the models..... | 503 |
| 22.5 | Maximum likelihood..... | 504 |
| 22.6 | Using GARCH (1,1) to forecast future volatility | 509 |
| 22.7 | Correlations | 512 |
| 22.8 | Application of EWMA to four-index example..... | 515 |
| | Summary | 517 |
| | Further reading | 517 |
| | Practice questions..... | 518 |
| | Further questions..... | 519 |
| Chapter 23. | Credit risk..... | 521 |
| 23.1 | Credit ratings | 521 |
| 23.2 | Historical default probabilities | 522 |
| 23.3 | Recovery rates | 523 |
| 23.4 | Estimating default probabilities from bond prices | 524 |
| 23.5 | Comparison of default probability estimates..... | 526 |
| 23.6 | Using equity prices to estimate default probabilities | 530 |
| 23.7 | Credit risk in derivatives transactions..... | 531 |
| 23.8 | Credit risk mitigation | 534 |
| 23.9 | Default correlation..... | 536 |
| 23.10 | Credit VaR | 540 |
| | Summary | 542 |
| | Further reading | 543 |
| | Practice questions..... | 543 |
| | Further questions..... | 545 |
| Chapter 24. | Credit derivatives | 547 |
| 24.1 | Credit default swaps..... | 548 |
| 24.2 | Valuation of credit default swaps | 551 |
| 24.3 | Credit indices | 555 |
| 24.4 | The use of fixed coupons | 556 |
| 24.5 | CDS forwards and options..... | 557 |
| 24.6 | Basket credit default swaps..... | 558 |
| 24.7 | Total return swaps | 558 |
| 24.8 | Collateralized debt obligations | 559 |
| 24.9 | Role of correlation in a basket CDS and CDO | 561 |
| 24.10 | Valuation of a synthetic CDO..... | 562 |
| 24.11 | Alternatives to the standard market model..... | 569 |
| | Summary..... | 570 |

| | | |
|---|--|------------|
| | Further reading..... | 571 |
| | Practice questions..... | 571 |
| | Further questions..... | 573 |
| Chapter 25. Exotic options | | 574 |
| 25.1 | Packages..... | 574 |
| 25.2 | Nonstandard American options..... | 575 |
| 25.3 | Gap options..... | 575 |
| 25.4 | Forward start options..... | 576 |
| 25.5 | Cliquet options..... | 577 |
| 25.6 | Compound options..... | 577 |
| 25.7 | Chooser options..... | 578 |
| 25.8 | Barrier options..... | 579 |
| 25.9 | Binary options..... | 581 |
| 25.10 | Lookback options..... | 582 |
| 25.11 | Shout options..... | 584 |
| 25.12 | Asian options..... | 584 |
| 25.13 | Options to exchange one asset for another..... | 586 |
| 25.14 | Options involving several assets..... | 587 |
| 25.15 | Volatility and variance swaps..... | 588 |
| 25.16 | Static options replication..... | 591 |
| | Summary..... | 593 |
| | Further reading..... | 594 |
| | Practice questions..... | 594 |
| | Further questions..... | 596 |
| Chapter 26. More on models and numerical procedures..... | | 599 |
| 26.1 | Alternatives to Black–Scholes–Merton..... | 600 |
| 26.2 | Stochastic volatility models..... | 605 |
| 26.3 | The IVF model..... | 607 |
| 26.4 | Convertible bonds..... | 608 |
| 26.5 | Path-dependent derivatives..... | 611 |
| 26.6 | Barrier options..... | 615 |
| 26.7 | Options on two correlated assets..... | 618 |
| 26.8 | Monte Carlo simulation and American options..... | 621 |
| | Summary..... | 625 |
| | Further reading..... | 626 |
| | Practice questions..... | 627 |
| | Further questions..... | 628 |
| Chapter 27. Martingales and measures..... | | 630 |
| 27.1 | The market price of risk..... | 631 |
| 27.2 | Several state variables..... | 634 |
| 27.3 | Martingales..... | 635 |
| 27.4 | Alternative choices for the numeraire..... | 636 |
| 27.5 | Extension to several factors..... | 640 |
| 27.6 | Black’s model revisited..... | 641 |
| 27.7 | Option to exchange one asset for another..... | 642 |
| 27.8 | Change of numeraire..... | 643 |
| | Summary..... | 644 |
| | Further reading..... | 645 |
| | Practice questions..... | 645 |
| | Further questions..... | 647 |

| | |
|--|------------|
| Chapter 28. Interest rate derivatives: The standard market models | 648 |
| 28.1 Bond options | 648 |
| 28.2 Interest rate caps and floors..... | 653 |
| 28.3 European swap options..... | 659 |
| 28.4 Generalizations..... | 663 |
| 28.5 Hedging interest rate derivatives | 663 |
| Summary..... | 664 |
| Further reading | 665 |
| Practice questions..... | 665 |
| Further questions..... | 666 |
| Chapter 29. Convexity, timing, and quanto adjustments | 668 |
| 29.1 Convexity adjustments | 668 |
| 29.2 Timing adjustments..... | 672 |
| 29.3 Quantos | 674 |
| Summary..... | 677 |
| Further reading | 677 |
| Practice questions..... | 678 |
| Further questions..... | 679 |
| Appendix: Proof of the convexity adjustment formula | 681 |
| Chapter 30. Interest rate derivatives: models of the short rate | 682 |
| 30.1 Background | 682 |
| 30.2 Equilibrium models..... | 683 |
| 30.3 No-arbitrage models..... | 689 |
| 30.4 Options on bonds | 694 |
| 30.5 Volatility structures | 695 |
| 30.6 Interest rate trees | 696 |
| 30.7 A general tree-building procedure..... | 698 |
| 30.8 Calibration..... | 707 |
| 30.9 Hedging using a one-factor model..... | 709 |
| Summary..... | 710 |
| Further reading | 710 |
| Practice questions..... | 711 |
| Further questions..... | 713 |
| Chapter 31. Interest rate derivatives: HJM and LMM | 715 |
| 31.1 The Heath, Jarrow, and Morton model..... | 715 |
| 31.2 The LIBOR market model | 718 |
| 31.3 Agency mortgage-backed securities | 728 |
| Summary..... | 730 |
| Further reading | 731 |
| Practice questions..... | 731 |
| Further questions..... | 732 |
| Chapter 32. Swaps Revisited | 733 |
| 32.1 Variations on the vanilla deal | 733 |
| 32.2 Compounding swaps | 735 |
| 32.3 Currency swaps | 736 |
| 32.4 More complex swaps | 737 |
| 32.5 Equity swaps..... | 740 |
| 32.6 Swaps with embedded options | 742 |
| 32.7 Other swaps..... | 744 |
| Summary..... | 745 |
| Further reading | 746 |

| | | |
|--|--|------------|
| | Practice questions | 746 |
| | Further questions | 747 |
| Chapter 33. Energy and commodity derivatives..... | | 748 |
| 33.1 | Agricultural commodities | 748 |
| 33.2 | Metals | 749 |
| 33.3 | Energy products..... | 750 |
| 33.4 | Modeling commodity prices..... | 752 |
| 33.5 | Weather derivatives..... | 758 |
| 33.6 | Insurance derivatives..... | 759 |
| 33.7 | Pricing weather and insurance derivatives..... | 760 |
| 33.8 | How an energy producer can hedge risks | 761 |
| | Summary | 762 |
| | Further reading..... | 762 |
| | Practice questions | 763 |
| | Further question | 764 |
| Chapter 34. Real options | | 765 |
| 34.1 | Capital investment appraisal | 765 |
| 34.2 | Extension of the risk-neutral valuation framework | 766 |
| 34.3 | Estimating the market price of risk | 768 |
| 34.4 | Application to the valuation of a business | 769 |
| 34.5 | Evaluating options in an investment opportunity | 769 |
| | Summary | 776 |
| | Further reading..... | 776 |
| | Practice questions | 777 |
| | Further questions | 777 |
| Chapter 35. Derivatives mishaps and what we can learn from them | | 779 |
| 35.1 | Lessons for all users of derivatives | 779 |
| 35.2 | Lessons for financial institutions..... | 783 |
| 35.3 | Lessons for nonfinancial corporations | 788 |
| | Summary | 790 |
| | Further reading..... | 790 |
| | Glossary of terms..... | 791 |
| | DerivaGem software | 812 |
| | Major exchanges trading futures and options..... | 817 |
| | Table for $N(x)$ when $x \leq 0$..... | 818 |
| | Table for $N(x)$ when $x \geq 0$..... | 819 |
| | Author index | 820 |
| | Subject index..... | 824 |

BUSINESS SNAPSHOTS

| | | |
|------|---|-----|
| 1.1 | The Lehman Bankruptcy | 3 |
| 1.2 | Hedge Funds | 11 |
| 1.2 | SocGen's Big Loss in 2008 | 17 |
| 2.1 | The Unanticipated Delivery of a Futures Contract | 23 |
| 2.2 | Long-Term Capital Management's Big Loss | 31 |
| 2.3 | Systemic Risk | 33 |
| 3.1 | Hedging By Gold Mining Companies | 52 |
| 3.2 | Metallgesellschaft: Hedging Gone Awry | 67 |
| 4.1 | What is the Risk-Free Rate? | 77 |
| 4.2 | Orange County's Yield Curve Plays | 87 |
| 4.3 | Liquidity and the 2007–2009 Financial Crisis | 96 |
| 5.1 | Kidder Peabody's Embarrassing Mistake | 106 |
| 5.2 | A Systems Error? | 111 |
| 5.3 | The CME Nikkei 225 Futures Contract | 113 |
| 5.4 | Index Arbitrage in October 1987 | 114 |
| 6.1 | Day Counts Can Be Deceptive | 130 |
| 6.2 | The Wild Card Play | 135 |
| 6.3 | Asset–Liability Management by Banks | 144 |
| 7.1 | Extract from Hypothetical Swap Confirmation | 155 |
| 7.2 | The Hammersmith and Fulham Story | 173 |
| 8.1 | Basel I, Basel II, and Basel III | 191 |
| 9.1 | Gucci Group's Large Dividend | 202 |
| 9.2 | Tax Planning Using Options | 209 |
| 10.1 | Put–Call Parity and Capital Structure | 224 |
| 11.1 | Losing Money with Box Spreads | 244 |
| 11.2 | How to Make Money from Trading Straddles | 248 |
| 14.1 | Mutual Fund Returns Can be Misleading | 304 |
| 14.2 | What Causes Volatility? | 307 |
| 14.3 | Warrants, Employee Stock Options, and Dilution | 317 |
| 16.1 | Can We Guarantee that Stocks Will Beat Bonds in the Long Run? | 354 |
| 18.1 | Dynamic Hedging in Practice | 396 |
| 18.2 | Was Portfolio Insurance to Blame for the Crash of 1987? | 403 |
| 19.1 | Making Money from Foreign Currency Options | 413 |
| 19.2 | Crashophobia | 416 |
| 20.1 | Calculating Pi with Monte Carlo Simulation | 446 |
| 20.2 | Checking Black–Scholes–Merton in Excel | 449 |
| 21.1 | How Bank Regulators Use VaR | 472 |
| 23.1 | Downgrade Triggers and Enron's Bankruptcy | 537 |
| 24.1 | Who Bears the Credit Risk? | 548 |
| 24.2 | The CDS Market | 550 |
| 24.3 | Is the CDS Market a Fair Game? | 556 |
| 25.1 | Is Delta Hedging Easier or More Difficult for Exotics? | 592 |
| 28.1 | Put–Call Parity for Caps and Floors | 655 |
| 28.2 | Swaptions and Bond Options | 660 |
| 29.1 | Siegel's Paradox | 676 |
| 31.1 | IOs and POs | 730 |
| 32.1 | Hypothetical Confirmation for Nonstandard Swap | 734 |
| 32.2 | Hypothetical Confirmation for Compounding Swap | 735 |
| 32.3 | Hypothetical Confirmation for Equity Swap | 741 |
| 32.4 | Procter and Gamble's Bizarre Deal | 745 |
| 34.1 | Valuing Amazon.com | 770 |
| 35.1 | Big Losses by Financial Institutions | 780 |
| 35.2 | Big Losses by Nonfinancial Organizations | 781 |

TECHNICAL NOTES

Available on the Author's Website
www.rotman.utoronto.ca/~hull/TechnicalNotes

1. Convexity Adjustments to Eurodollar Futures
2. Properties of the Lognormal Distribution
3. Warrant Valuation When Value of Equity plus Warrants Is Lognormal
4. Exact Procedure for Valuing American Calls on Stocks Paying a Single Dividend
5. Calculation of the Cumulative Probability in a Bivariate Normal Distribution
6. Differential Equation for Price of a Derivative on a Stock Paying a Known Dividend Yield
7. Differential Equation for Price of a Derivative on a Futures Price
8. Analytic Approximation for Valuing American Options
9. Generalized Tree-Building Procedure
10. The Cornish–Fisher Expansion to Estimate VaR
11. Manipulation of Credit Transition Matrices
12. Calculation of Cumulative Noncentral Chi-Square Distribution
13. Efficient Procedure for Valuing American-Style Lookback Options
14. The Hull–White Two-Factor Model
15. Valuing Options on Coupon-Bearing Bonds in a One-Factor Interest Rate Model
16. Construction of an Interest Rate Tree with Nonconstant Time Steps and Nonconstant Parameters
17. The Process for the Short Rate in an HJM Term Structure Model
18. Valuation of a Compounding Swap
19. Valuation of an Equity Swap
20. Changing the Market Price of Risk for Variables That Are Not the Prices of Traded Securities
21. Hermite Polynomials and Their Use for Integration
22. Valuation of a Variance Swap
23. The Black, Derman, Toy Model
24. Proof that Forward and Futures Prices are Equal When Interest Rates Are Constant
25. A Cash-Flow Mapping Procedure
26. A Binomial Measure of Credit Correlation
27. Calculation of Moments for Valuing Asian Options
28. Calculation of Moments for Valuing Basket Options
29. Proof of Extensions to Itô's Lemma
30. The Return of a Security Dependent on Multiple Sources of Uncertainty

Preface

It is sometimes hard for me to believe that the first edition of this book was only 330 pages and 13 chapters long! The book has grown and been adapted to keep up with the fast pace of change in derivatives markets.

Like earlier editions, the book serves several markets. It is appropriate for graduate courses in business, economics, financial mathematics, and financial engineering. It can be used on advanced undergraduate courses when students have good quantitative skills. Also, many practitioners who are involved in derivatives markets find the book useful. I am delighted that half the purchasers of the book are analysts, traders, and other professionals in derivatives and risk management.

One of the key decisions that must be made by an author who is writing in the area of derivatives concerns the use of mathematics. If the level of mathematical sophistication is too high, the material is likely to be inaccessible to many students and practitioners. If it is too low, some important issues will inevitably be treated in a rather superficial way. I have tried to be particularly careful about the way I use both mathematics and notation in the book. Nonessential mathematical material has been either eliminated or included in end-of-chapter appendices and in the technical notes on my website. Concepts that are likely to be new to many readers have been explained carefully, and many numerical examples have been included.

Options, Futures, and Other Derivatives can be used for a first course in derivatives or for a more advanced course. There are many different ways it can be used in the classroom. Instructors teaching a first course in derivatives are likely to want to spend most classroom time on the first half of the book. Instructors teaching a more advanced course will find that many different combinations of chapters in the second half of the book can be used. I find that the material in Chapter 35 works well at the end of either an introductory or an advanced course.

What's New?

Material has been updated and improved throughout the book. The changes in the eighth edition include the following:

1. There is a new chapter (Chapter 8) devoted to securitization and the credit crisis. The events in financial markets since the seventh edition was published make these topics particularly relevant.
2. There is more discussion (Chapter 33) of the way commodity prices are modeled and how commodity derivatives are valued. Energy derivatives and other commodity derivatives have become progressively more important in recent years.

3. The chapter on hedging using futures (Chapter 3) has been simplified and an appendix explaining the capital asset pricing model has been included. This was suggested by a number of instructors.
4. Material on central clearing, liquidity risk, and overnight indexed swaps has been included. Following the credit crisis, these are features of derivatives markets that all students need to understand.
5. An appendix to Chapter 12 shows that the Black–Scholes–Merton formula can be derived as the limiting case of a binomial tree. Some instructors like to introduce the Black–Scholes–Merton result this way.
6. The material on value at risk is developed using an example involving real data taken from the credit crisis. Spreadsheets for the example are on my website. This change makes the material more interesting for readers and allows richer assignment questions to be used by instructors.
7. New material has been added on topics such as principal-protected notes, gap options, cliquet options, and jump processes, reflecting their importance in derivatives markets.
8. More material has been added on applications of the Vasicek and CIR models. This material provides a way in which readers can improve their understanding of key concepts. It is particularly important for actuarial students and fund managers.
9. There are a number of enhancements to the DerivaGem software. The software now covers credit derivatives. A version of the software is provided that can be used with Open Office by Mac and Linux users. In response to many requests from users, the code is provided for the DerivaGem functions. The software is now much easier to install and a “Getting Started” section is included on page 812.
10. The Test Bank available to adopting instructors has been improved.
11. New end-of-chapter problems have been added.

Software

DerivaGem version 2.01 is included with this book. It consists of two Excel applications: the *Options Calculator* and the *Applications Builder*. The Options Calculator consists of easy-to-use software for valuing a wide range of options. The Applications Builder consists of a number of Excel functions from which users can build their own applications. A number of sample applications are included to enable students to explore the properties of options and numerical procedures more easily. The Applications Builder also allows more interesting assignments to be designed.

The latest version of the software allows credit derivatives to be valued. A version of the software’s functions that is compatible with Open Office for Mac and Linux users is now provided, and users can now access the code for the functions underlying DerivaGem.

The description of the software starting on page 812 includes a “Getting Started” section. Updates to the software can be downloaded from my website:

www.rotman.utoronto.ca/~hull.

Slides

Several hundred PowerPoint slides can be downloaded from Pearson’s Instructor Resource Center (www.pearsonhighered.com/irc) or from my website (www.rotman.utoronto.ca). Instructors who adopt the text may adapt the slides to meet their own needs.

Test Bank

The Test Bank has been improved and provides a wealth of multiple-choice and short-calculation questions that can be used by instructors for testing. It can be downloaded from the Instructor Resource Center at www.pearsonhighered.com/irc.

Solutions Manual

End-of-chapter problems are divided into two groups: “Questions and Problems” and “Further Questions”. Solutions to the Questions and Problems are in *Options, Futures, and Other Derivatives 8e: Solutions Manual* (ISBN 0-13-216496-5) which is published by Pearson and can be purchased by students.

Instructors Manual

The Instructors Manual contains solutions to all questions (both “Practice Questions” and “Further Questions”), notes on the teaching of each chapter, test bank questions, notes on course organization, and some relevant Excel worksheets. It is available for download from the Instructor Resource Center at www.pearsonhighered.com/irc.

Technical Notes

Technical Notes are used to elaborate on points made in the text. They are referred to in the text and can be downloaded from www.rotman.utoronto.ca/~hull/TechnicalNotes. By not including the Technical Notes in the book, I am able to streamline the presentation of material so that it is more student-friendly.

Acknowledgments

Many people have played a part in the development of successive editions of this book. Indeed, the list of people who have provided me with feedback on the book is now so long that it is not possible to mention everyone. I have benefited from the advice of many academics who have taught from the book and from the comments of many derivatives practitioners. I would like to thank the students on my courses at the University of Toronto who have made many suggestions on how the material can be improved. Eddie Mizzi from The Geometric Press did an excellent job editing the final manuscript and handling page composition. Emilio Barone from Luiss Guido Carli University in Rome provided many detailed comments.

Alan White, a colleague at the University of Toronto, deserves a special acknowledgement. Alan and I have been carrying out joint research and consulting in the areas of derivatives and risk management for over 25 years. During that time, we have spent many hours discussing key issues. Many of the new ideas in this book, and many of the new ways used to explain old ideas, are as much Alan’s as mine. Alan has done most of the development work on the DerivaGem software.

Special thanks are due to many people at Pearson, particularly Tessa O’Brien, Donna Battista, and Nancy Fenton for their enthusiasm, advice, and encouragement. I welcome comments on the book from readers. My e-mail address is:

hull@rotman.utoronto.ca

John Hull
Joseph L. Rotman School of Management
University of Toronto

1 CHAPTER

Introduction



In the last 30 years, derivatives have become increasingly important in finance. Futures and options are actively traded on many exchanges throughout the world. Many different types of forward contracts, swaps, options, and other derivatives are entered into by financial institutions, fund managers, and corporate treasurers in the over-the-counter market. Derivatives are added to bond issues, used in executive compensation plans, embedded in capital investment opportunities, used to transfer risks in mortgages from the original lenders to investors, and so on. We have now reached the stage where those who work in finance, and many who work outside finance, need to understand how derivatives work, how they are used, and how they are priced.

Whether you love derivatives or hate them, you cannot ignore them! The derivatives market is huge—much bigger than the stock market when measured in terms of underlying assets. The value of the assets underlying outstanding derivatives transactions is several times the world gross domestic product. As we shall see in this chapter, derivatives can be used for hedging or speculation or arbitrage. They play a key role in transferring a wide range of risks in the economy from one entity to another.

A *derivative* can be defined as a financial instrument whose value depends on (or derives from) the values of other, more basic, underlying variables. Very often the variables underlying derivatives are the prices of traded assets. A stock option, for example, is a derivative whose value is dependent on the price of a stock. However, derivatives can be dependent on almost any variable, from the price of hogs to the amount of snow falling at a certain ski resort.

Since the first edition of this book was published in 1988 there have been many developments in derivatives markets. There is now active trading in credit derivatives, electricity derivatives, weather derivatives, and insurance derivatives. Many new types of interest rate, foreign exchange, and equity derivative products have been created. There have been many new ideas in risk management and risk measurement. Capital investment appraisal now often involves the evaluation of what are known as *real options*. The book has kept up with all these developments.

Derivatives markets have come under a great deal of criticism because of their role in the credit crisis that started in 2007. Derivative products were created from portfolios of risky mortgages in the United States using a procedure known as securitization. Many of the products that were created became worthless when house prices declined. Financial institutions, and investors throughout the world, lost a huge

amount of money and the world was plunged into the worst recession it had experienced for many generations. Chapter 8, new to this edition, explains how securitization works and why such big losses occurred. As a result of the credit crisis, derivatives markets are now more heavily regulated than they used to be. For example, banks are required to keep more capital for the risks they are taking and to pay more attention to liquidity.

In this opening chapter, we take a first look at forward, futures, and options markets and provide an overview of how they are used by hedgers, speculators, and arbitrageurs. Later chapters will give more details and elaborate on many of the points made here.

1.1 EXCHANGE-TRADED MARKETS

A derivatives exchange is a market where individuals trade standardized contracts that have been defined by the exchange. Derivatives exchanges have existed for a long time. The Chicago Board of Trade (CBOT) was established in 1848 to bring farmers and merchants together. Initially its main task was to standardize the quantities and qualities of the grains that were traded. Within a few years, the first futures-type contract was developed. It was known as a *to-arrive contract*. Speculators soon became interested in the contract and found trading the contract to be an attractive alternative to trading the grain itself. A rival futures exchange, the Chicago Mercantile Exchange (CME), was established in 1919. Now futures exchanges exist all over the world. (See table at the end of the book.) CME and CBOT have merged to form the CME Group (www.cmegroup.com), which also includes the New York Mercantile Exchange.

The Chicago Board Options Exchange (CBOE, www.cboe.com) started trading call option contracts on 16 stocks in 1973. Options had traded prior to 1973, but the CBOE succeeded in creating an orderly market with well-defined contracts. Put option contracts started trading on the exchange in 1977. The CBOE now trades options on over 2,500 stocks and many different stock indices. Like futures, options have proved to be very popular contracts. Many other exchanges throughout the world now trade options. (See table at the end of the book.) The underlying assets include foreign currencies and futures contracts as well as stocks and stock indices.

Electronic Markets

Traditionally derivatives exchanges have used what is known as the *open outcry system*. This involves traders physically meeting on the floor of the exchange, shouting, and using a complicated set of hand signals to indicate the trades they would like to carry out. Exchanges are increasingly replacing the open outcry system by *electronic trading*. This involves traders entering their desired trades at a keyboard and a computer being used to match buyers and sellers. The open outcry system has its advocates, but, as time passes, it is becoming less and less used.

Electronic trading has led to a growth in algorithmic trading (also known as black-box trading, automated trading, high-frequency trading, or robo trading). This involves the use of computer programs to initiate trades, often without human intervention.

Business Snapshot 1.1 The Lehman Bankruptcy

On September 15, 2008, Lehman Brothers filed for Chapter 11 bankruptcy protection. This was the largest bankruptcy filing in US history and its ramifications were felt throughout derivatives markets. Almost until the end, it seemed as though there was a good chance that Lehman would survive. A number of companies (e.g., the Korean Development Bank, Barclays Bank in the UK, and Bank of America) expressed interest in buying it, but none of these was able to close a deal. Many people thought that Lehman was “too big to fail” and that the US government would have to bail it out if no purchaser could be found. This proved not to be the case.

How did this happen? It was a combination of high leverage, risky investments, and liquidity problems. Commercial banks that take deposits are subject to regulations on the amount of capital they must keep. Lehman was an investment bank and not subject to these regulations. By 2007, its leverage ratio had increased to 31:1, which means that a 3–4% decline in the value of its assets would wipe out its capital. Dick Fuld, Lehman’s Chairman and Chief Executive Officer, encouraged an aggressive deal-making, risk-taking culture. He is reported to have told his executives: “Every day is a battle. You have to kill the enemy.” The Chief Risk Officer at Lehman was competent, but did not have much influence and was even removed from the executive committee in 2007. The risks taken by Lehman included large positions in the instruments created from subprime mortgages, which will be described in Chapter 8. Lehman funded much of its operations with short-term debt. When there was a loss of confidence in the company, lenders refused to roll over this funding, forcing it into bankruptcy.

Lehman was very active in the over-the-counter derivatives markets. It had hundreds of thousands of transactions outstanding with about 8,000 different counterparties. Lehman’s counterparties were often required to post collateral and this collateral had in many cases been used by Lehman for various purposes. It is easy to see that sorting out who owes what to whom in this type of situation is a nightmare!

1.2 OVER-THE-COUNTER MARKETS

Not all trading of derivatives is done on exchanges. The *over-the-counter market* is an important alternative to exchanges and, measured in terms of the total volume of trading, has become much larger than the exchange-traded market. It is a telephone- and computer-linked network of dealers. Trades are done over the phone and are usually between two financial institutions or between a financial institution and one of its clients (typically a corporate treasurer or fund manager). Financial institutions often act as market makers for the more commonly traded instruments. This means that they are always prepared to quote both a bid price (a price at which they are prepared to buy) and an offer price (a price at which they are prepared to sell).

Telephone conversations in the over-the-counter market are usually taped. If there is a dispute about what was agreed, the tapes are replayed to resolve the issue. Trades in the over-the-counter market are typically much larger than trades in the exchange-traded market. A key advantage of the over-the-counter market is that the terms of a contract do not have to be those specified by an exchange. Market participants are free

to negotiate any mutually attractive deal. A disadvantage is that there is usually some credit risk in an over-the-counter trade (i.e., there is a small risk that the contract will not be honored). As we shall see in the next chapter, exchanges have organized themselves to eliminate virtually all credit risk.


Lehman Brothers was a very active trader of over-the-counter derivatives. As discussed in Business Snapshot 1.1, its bankruptcy in 2008 provided a dramatic test for the market.

Market Size

Both the over-the-counter and the exchange-traded market for derivatives are huge. Although the statistics that are collected for the two markets are not exactly comparable, it is clear that the over-the-counter market is much larger than the exchange-traded market. The Bank for International Settlements (www.bis.org) started collecting statistics on the markets in 1998. Figure 1.1 compares (a) the estimated total principal amounts underlying transactions that were outstanding in the over-the-counter markets between June 1998 and December 2009 and (b) the estimated total value of the assets underlying exchange-traded contracts during the same period. Using these measures, we see that, by December 2009, the over-the-counter market had grown to \$614.7 trillion and the exchange-traded market had grown to \$73.1 trillion.

In interpreting these numbers, we should bear in mind that the principal underlying an over-the-counter transaction is not the same as its value. An example of an over-the-counter contract is an agreement to buy 100 million US dollars with British pounds at a predetermined exchange rate in 1 year. The total principal amount underlying this transaction is \$100 million. However, the value of the contract might be only \$1 million. The Bank for International Settlements estimates the gross market value of all over-the-counter contracts outstanding in December 2009 to be about \$21.6 trillion.¹

Figure 1.1 Size of over-the-counter and exchange-traded derivatives markets.



¹ A contract that is worth \$1 million to one side and $-\$1$ million to the other side would be counted as having a gross market value of \$1 million.

1.3 FORWARD CONTRACTS

A relatively simple derivative is a *forward contract*. It is an agreement to buy or sell an asset at a certain future time for a certain price. It can be contrasted with a *spot contract*, which is an agreement to buy or sell an asset today. A forward contract is traded in the over-the-counter market—usually between two financial institutions or between a financial institution and one of its clients.

One of the parties to a forward contract assumes a *long position* and agrees to buy the underlying asset on a certain specified future date for a certain specified price. The other party assumes a *short position* and agrees to sell the asset on the same date for the same price.

Forward contracts on foreign exchange are very popular. Most large banks employ both spot and forward foreign-exchange traders. Spot traders are trading a foreign currency for almost immediate delivery. Forward traders are trading for delivery at a future time. Table 1.1 provides the quotes on the exchange rate between the British pound (GBP) and the US dollar (USD) that might be made by a large international bank on May 24, 2010. The quote is for the number of USD per GBP. The first row indicates that the bank is prepared to buy GBP (also known as sterling) in the spot market (i.e., for virtually immediate delivery) at the rate of \$1.4407 per GBP and sell sterling in the spot market at \$1.4411 per GBP. The second, third, and fourth rows indicate that the bank is prepared to buy sterling in 1, 3, and 6 months at \$1.4408, \$1.4410, and \$1.4416 per GBP, respectively, and to sell sterling in 1, 3, and 6 months at \$1.4413, \$1.4415, and \$1.4422 per GBP, respectively.

Forward contracts can be used to hedge foreign currency risk. Suppose that, on May 24, 2010, the treasurer of a US corporation knows that the corporation will pay £1 million in 6 months (i.e., on November 24, 2010) and wants to hedge against exchange rate moves. Using the quotes in Table 1.1, the treasurer can agree to buy £1 million 6 months forward at an exchange rate of 1.4422. The corporation then has a long forward contract on GBP. It has agreed that on November 24, 2010, it will buy £1 million from the bank for \$1.4422 million. The bank has a short forward contract on GBP. It has agreed that on November 24, 2010, it will sell £1 million for \$1.4422 million. Both sides have made a binding commitment.


Payoffs from Forward Contracts

Consider the position of the corporation in the trade we have just described. What are the possible outcomes? The forward contract obligates the corporation to buy £1 million

Table 1.1 Spot and forward quotes for the USD/GBP exchange rate, May 24, 2010 (GBP = British pound; USD = US dollar; quote is number of USD per GBP).

| | <i>Bid</i> | <i>Offer</i> |
|-----------------|------------|--------------|
| Spot | 1.4407 | 1.4411 |
| 1-month forward | 1.4408 | 1.4413 |
| 3-month forward | 1.4410 | 1.4415 |
| 6-month forward | 1.4416 | 1.4422 |

Figure 1.2 Payoffs from forward contracts: (a) long position, (b) short position. Delivery price = K ; price of asset at contract maturity = S_T .



for \$1,442,200. If the spot exchange rate rose to, say, 1.5000, at the end of the 6 months, the forward contract would be worth \$57,800 ($= \$1,500,000 - \$1,442,200$) to the corporation. It would enable £1 million to be purchased at an exchange rate of 1.4422 rather than 1.5000. Similarly, if the spot exchange rate fell to 1.3500 at the end of the 6 months, the forward contract would have a negative value to the corporation of \$92,200 because it would lead to the corporation paying \$92,200 more than the market price for the sterling.

In general, the payoff from a long position in a forward contract on one unit of an asset is

$$S_T - K$$

where K is the delivery price and S_T is the spot price of the asset at maturity of the contract. This is because the holder of the contract is obligated to buy an asset worth S_T for K . Similarly, the payoff from a short position in a forward contract on one unit of an asset is

$$K - S_T$$

These payoffs can be positive or negative. They are illustrated in Figure 1.2. Because it costs nothing to enter into a forward contract, the payoff from the contract is also the trader's total gain or loss from the contract.

In the example just considered, $K = 1.4422$ and the corporation has a long contract. When $S_T = 1.5000$, the payoff is \$0.0578 per £1; when $S_T = 1.3500$, it is $-\$0.0922$ per £1.

Forward Prices and Spot Prices

We shall be discussing in some detail the relationship between spot and forward prices in Chapter 5. For a quick preview of why the two are related, consider a stock that pays no dividend and is worth \$60. You can borrow or lend money for 1 year at 5%. What should the 1-year forward price of the stock be?

The answer is \$60 grossed up at 5% for 1 year, or \$63. If the forward price is more than this, say \$67, you could borrow \$60, buy one share of the stock, and sell it forward for \$67. After paying off the loan, you would net a profit of \$4 in 1 year. If the forward price is less than \$63, say \$58, an investor owning the stock as part of a portfolio would sell the stock for \$60 and enter into a forward contract to buy it back for \$58 in 1 year. The proceeds of investment would be invested at 5% to earn \$3. The investor would end up \$5 better off than if the stock were kept in the portfolio for the year.

1.4 FUTURES CONTRACTS

Like a forward contract, a futures contract is an agreement between two parties to buy or sell an asset at a certain time in the future for a certain price. Unlike forward contracts, futures contracts are normally traded on an exchange. To make trading possible, the exchange specifies certain standardized features of the contract. As the two parties to the contract do not necessarily know each other, the exchange also provides a mechanism that gives the two parties a guarantee that the contract will be honored.

The largest exchanges on which futures contracts are traded are the Chicago Board of Trade (CBOT) and the Chicago Mercantile Exchange (CME), which have now merged to form the CME Group. On these and other exchanges throughout the world, a very wide range of commodities and financial assets form the underlying assets in the various contracts. The commodities include pork bellies, live cattle, sugar, wool, lumber, copper, aluminum, gold, and tin. The financial assets include stock indices, currencies, and Treasury bonds. Futures prices are regularly reported in the financial press. Suppose that, on September 1, the December futures price of gold is quoted as \$1,080. This is the price, exclusive of commissions, at which traders can agree to buy or sell gold for December delivery. It is determined in the same way as other prices (i.e., by the laws of supply and demand). If more traders want to go long than to go short, the price goes up; if the reverse is true, then the price goes down.

Further details on issues such as margin requirements, daily settlement procedures, delivery procedures, bid–offer spreads, and the role of the exchange clearing house are given in Chapter 2.

1.5 OPTIONS

Options are traded both on exchanges and in the over-the-counter market. There are two types of option. A *call option* gives the holder the right to buy the underlying asset by a certain date for a certain price. A *put option* gives the holder the right to sell the underlying asset by a certain date for a certain price. The price in the contract is known as the *exercise price* or *strike price*; the date in the contract is known as the *expiration date* or *maturity*. *American options* can be exercised at any time up to the expiration date. *European options* can be exercised only on the expiration date itself.² Most of the options that are traded on exchanges are American. In the exchange-traded equity option market, one contract is usually an agreement to buy or sell 100 shares. European

² Note that the terms *American* and *European* do not refer to the location of the option or the exchange. Some options trading on North American exchanges are European.

options are generally easier to analyze than American options, and some of the properties of an American option are frequently deduced from those of its European counterpart.

It should be emphasized that an option gives the holder the right to do something. The holder does not have to exercise this right. This is what distinguishes options from forwards and futures, where the holder is obligated to buy or sell the underlying asset. Whereas it costs nothing to enter into a forward or futures contract, there is a cost to acquiring an option.

The largest exchange in the world for trading stock options is the Chicago Board Options Exchange (CBOE; www.cboe.com). Table 1.2 gives the bid and offer quotes for some of the call options trading on Google (ticker symbol: GOOG) on June 15, 2010. Table 1.3 does the same for put options trading on Google on that date. The quotes are taken from the CBOE website. The Google stock price at the time of the quotes was bid 497.02, offer 497.25. The bid–offer spread on an option is usually greater than that on the underlying stock and depends on the volume of trading. The option strike prices are \$460, \$480, \$500, \$520, \$540, and \$560. The maturities are July 2010, September 2010, and December 2010. The July options expire on July 17, 2010, the September options on September 18, 2010, and the December options on December 18, 2010.

The tables illustrate a number of properties of options. The price of a call option decreases as the strike price increases, while the price of a put option increases as the strike price increases. Both types of option tend to become more valuable as their time to maturity increases. These properties of options will be discussed further in Chapter 10.

Suppose an investor instructs a broker to buy one December call option contract on Google with a strike price of \$520. The broker will relay these instructions to a trader at the CBOE and the deal will be done. The (offer) price is \$32.00, as indicated in Table 1.2. This is the price for an option to buy one share. In the United States, an option contract is a contract to buy or sell 100 shares. Therefore, the investor must arrange for \$3,200 to be remitted to the exchange through the broker. The exchange will then arrange for this amount to be passed on to the party on the other side of the transaction.

In our example, the investor has obtained at a cost of \$3,200 the right to buy 100 Google shares for \$520 each. If the price of Google does not rise above \$520 by December 18, 2010, the option is not exercised and the investor loses \$3,200.³ But if

Table 1.2 Prices of call options on Google, June 15, 2010; stock price: bid \$497.07; offer \$497.25 (*Source*: CBOE).

| <i>Strike price</i> (\$) | <i>July 2010</i> | | <i>September 2010</i> | | <i>December 2010</i> | |
|-----------------------------|------------------|--------------|-----------------------|--------------|----------------------|--------------|
| | <i>Bid</i> | <i>Offer</i> | <i>Bid</i> | <i>Offer</i> | <i>Bid</i> | <i>Offer</i> |
| 460 | 43.30 | 44.00 | 51.90 | 53.90 | 63.40 | 64.80 |
| 480 | 28.60 | 29.00 | 39.70 | 40.40 | 50.80 | 52.30 |
| 500 | 17.00 | 17.40 | 28.30 | 29.30 | 40.60 | 41.30 |
| 520 | 9.00 | 9.30 | 19.10 | 19.90 | 31.40 | 32.00 |
| 540 | 4.20 | 4.40 | 12.70 | 13.00 | 23.10 | 24.00 |
| 560 | 1.75 | 2.10 | 7.40 | 8.40 | 16.80 | 17.70 |

³ The calculations here ignore commissions paid by the investor.

Table 1.3 Prices of put options on Google, June 15, 2010; stock price: bid \$497.07; offer \$497.25 (Source: CBOE).

| Strike price (\$) | July 2010 | | September 2010 | | December 2010 | |
|----------------------|-----------|-------|----------------|-------|---------------|-------|
| | Bid | Offer | Bid | Offer | Bid | Offer |
| 460 | 6.30 | 6.60 | 15.70 | 16.20 | 26.00 | 27.30 |
| 480 | 11.30 | 11.70 | 22.20 | 22.70 | 33.30 | 35.00 |
| 500 | 19.50 | 20.00 | 30.90 | 32.60 | 42.20 | 43.00 |
| 520 | 31.60 | 33.90 | 41.80 | 43.60 | 52.80 | 54.50 |
| 540 | 46.30 | 47.20 | 54.90 | 56.10 | 64.90 | 66.20 |
| 560 | 64.30 | 66.70 | 70.00 | 71.30 | 78.60 | 80.00 |

Google does well and the option is exercised when the bid price for the stock is \$600, the investor is able to buy 100 shares at \$520 and immediately sell them for \$600 for a profit of \$8,000, or \$4,800 when the initial cost of the options is taken into account.⁴

An alternative trade for the investor would be to sell one September put option contract with a strike price of \$480. This would lead to an immediate cash inflow of $100 \times 22.20 = \$2,220$. If the Google stock price stays above \$480, the option is not exercised and the investor makes a profit of this amount. However, if stock price falls and the option is exercised when the stock price is \$420, then there is a loss. The investor must buy 100 shares at \$480 when they are worth only \$420. This leads to a loss of \$6,000, or \$3,780 when the initial amount received for the option contract is taken into account.

The stock options trading on the CBOE are American. If we assume for simplicity that they are European, so that they can be exercised only at maturity, the investor's profit as a function of the final stock price for the two trades we have considered is shown in Figure 1.3.

Further details about the operation of options markets and how prices such as those in Tables 1.2 and 1.3 are determined by traders are given in later chapters. At this stage we note that there are four types of participants in options markets:

1. Buyers of calls
2. Sellers of calls
3. Buyers of puts
4. Sellers of puts.


Buyers are referred to as having *long positions*; sellers are referred to as having *short positions*. Selling an option is also known as *writing the option*.

1.6 TYPES OF TRADERS

Derivatives markets have been outstandingly successful. The main reason is that they have attracted many different types of traders and have a great deal of liquidity. When

⁴ The calculations here ignore the effect of discounting. Theoretically, the \$8,000 should be discounted from the time of exercise to June 15, 2010, when calculating the profit.

Figure 1.3 Net profit per share from (a) purchasing a contract consisting of 100 Google December call options with a strike price of \$520 and (b) selling a contract consisting of 100 Google September put options with a strike price of \$480.



an investor wants to take one side of a contract, there is usually no problem in finding someone who is prepared to take the other side.

Three broad categories of traders can be identified: hedgers, speculators, and arbitrageurs. Hedgers use derivatives to reduce the risk that they face from potential future movements in a market variable. Speculators use them to bet on the future direction of a market variable. Arbitrageurs take offsetting positions in two or more instruments to lock in a profit. As described in Business Snapshot 1.2, hedge funds have become big users of derivatives for all three purposes.

In the next few sections, we will consider the activities of each type of trader in more detail.

1.7 HEDGERS

In this section we illustrate how hedgers can reduce their risks with forward contracts and options.

Hedging Using Forward Contracts

Suppose that it is May 24, 2010, and ImportCo, a company based in the United States, knows that it will have to pay £10 million on August 24, 2010, for goods it has purchased from a British supplier. The USD–GBP exchange rate quotes made by a financial institution are shown in Table 1.1. ImportCo could hedge its foreign exchange risk by buying pounds (GBP) from the financial institution in the 3-month forward market at 1.4415. This would have the effect of fixing the price to be paid to the British exporter at \$14,415,000.

Consider next another US company, which we will refer to as ExportCo, that is exporting goods to the United Kingdom and, on May 24, 2010, knows that it will receive £30 million 3 months later. ExportCo can hedge its foreign exchange risk by selling £30 million in the 3-month forward market at an exchange rate of 1.4410. This would have the effect of locking in the US dollars to be realized for the sterling at \$43,230,000.

Note that a company might do better if it chooses not to hedge than if it chooses to hedge. Alternatively, it might do worse. Consider ImportCo. If the exchange rate is

Business Snapshot 1.2 Hedge Funds

Hedge funds have become major users of derivatives for hedging, speculation, and arbitrage. They are similar to mutual funds in that they invest funds on behalf of clients. However, they accept funds only from financially sophisticated individuals and do not publicly offer their securities. Mutual funds are subject to regulations requiring that the shares be redeemable at any time, that investment policies be disclosed, that the use of leverage be limited, that no short positions be taken, and so on. Hedge funds are relatively free of these regulations. This gives them a great deal of freedom to develop sophisticated, unconventional, and proprietary investment strategies. The fees charged by hedge fund managers are dependent on the fund's performance and are relatively high—typically 1 to 2% of the amount invested plus 20% of the profits. Hedge funds have grown in popularity, with about \$1 trillion being invested in them throughout the world. “Funds of funds” have been set up to invest in a portfolio of hedge funds.

The investment strategy followed by a hedge fund manager often involves using derivatives to set up a speculative or arbitrage position. Once the strategy has been defined, the hedge fund manager must:

1. Evaluate the risks to which the fund is exposed
2. Decide which risks are acceptable and which will be hedged
3. Devise strategies (usually involving derivatives) to hedge the unacceptable risks.

Here are some examples of the labels used for hedge funds together with the trading strategies followed:

Long/Short Equities: Purchase securities considered to be undervalued and short those considered to be overvalued in such a way that the exposure to the overall direction of the market is small.

Convertible Arbitrage: Take a long position in a convertible bond combined with an actively managed short position in the underlying equity.

Distressed Securities: Buy securities issued by companies in or close to bankruptcy.

Emerging Markets: Invest in debt and equity of companies in developing or emerging countries and in the debt of the countries themselves.

Global Macro: Carry out trades that reflect anticipated global macroeconomic trends.

Merger Arbitrage: Trade after a merger or acquisition is announced so that a profit is made if the announced deal takes place.

1.3000 on August 24 and the company has not hedged, the £10 million that it has to pay will cost \$13,000,000, which is less than \$14,415,000. On the other hand, if the exchange rate is 1.5000, the £10 million will cost \$15,000,000—and the company will wish that it had hedged! The position of ExportCo if it does not hedge is the reverse. If the exchange rate in August proves to be less than 1.4410, the company will wish that it had hedged; if the rate is greater than 1.4410, it will be pleased that it has not done so.

This example illustrates a key aspect of hedging. The purpose of hedging is to reduce risk. There is no guarantee that the outcome with hedging will be better than the outcome without hedging.

Hedging Using Options


Options can also be used for hedging. Consider an investor who in May of a particular year owns 1,000 Microsoft shares. The share price is \$28 per share. The investor is concerned about a possible share price decline in the next 2 months and wants protection. The investor could buy ten July put option contracts on Microsoft on the Chicago Board Options Exchange with a strike price of \$27.50. This would give the investor the right to sell a total of 1,000 shares for a price of \$27.50. If the quoted option price is \$1, then each option contract would cost $100 \times \$1 = \100 and the total cost of the hedging strategy would be $10 \times \$100 = \$1,000$.

The strategy costs \$1,000 but guarantees that the shares can be sold for at least \$27.50 per share during the life of the option. If the market price of Microsoft falls below \$27.50, the options will be exercised, so that \$27,500 is realized for the entire holding. When the cost of the options is taken into account, the amount realized is \$26,500. If the market price stays above \$27.50, the options are not exercised and expire worthless. However, in this case the value of the holding is always above \$27,500 (or above \$26,500 when the cost of the options is taken into account). Figure 1.4 shows the net value of the portfolio (after taking the cost of the options into account) as a function of Microsoft's stock price in 2 months. The dotted line shows the value of the portfolio assuming no hedging.

A Comparison

There is a fundamental difference between the use of forward contracts and options for hedging. Forward contracts are designed to neutralize risk by fixing the price that the hedger will pay or receive for the underlying asset. Option contracts, by contrast, provide insurance. They offer a way for investors to protect themselves against adverse price movements in the future while still allowing them to benefit from favorable price movements. Unlike forwards, options involve the payment of an up-front fee.

Figure 1.4 Value of Microsoft holding in 2 months with and without hedging.



1.8 SPECULATORS

We now move on to consider how futures and options markets can be used by speculators. Whereas hedgers want to avoid exposure to adverse movements in the price of an asset, speculators wish to take a position in the market. Either they are betting that the price of the asset will go up or they are betting that it will go down.

Speculation Using Futures

Consider a US speculator who in February thinks that the British pound will strengthen relative to the US dollar over the next 2 months and is prepared to back that hunch to the tune of £250,000. One thing the speculator can do is purchase £250,000 in the spot market in the hope that the sterling can be sold later at a higher price. (The sterling once purchased would be kept in an interest-bearing account.) Another possibility is to take a long position in four CME April futures contracts on sterling. (Each futures contract is for the purchase of £62,500.) Table 1.4 summarizes the two alternatives on the assumption that the current exchange rate is 1.4470 dollars per pound and the April futures price is 1.4410 dollars per pound. If the exchange rate turns out to be 1.5000 dollars per pound in April, the futures contract alternative enables the speculator to realize a profit of $(1.5000 - 1.4410) \times 250,000 = \$14,750$. The spot market alternative leads to 250,000 units of an asset being purchased for \$1.4470 in February and sold for \$1.5000 in April, so that a profit of $(1.5000 - 1.4470) \times 250,000 = \$13,250$ is made. If the exchange rate falls to 1.4000 dollars per pound, the futures contract gives rise to a $(1.4410 - 1.4000) \times 250,000 = \$10,250$ loss, whereas the spot market alternative gives rise to a loss of $(1.4470 - 1.4000) \times 250,000 = \$11,750$. The spot market alternative appears to give rise to slightly worse outcomes for both scenarios. But this is because the calculations do not reflect the interest that is earned or paid.

What then is the difference between the two alternatives? The first alternative of buying sterling requires an up-front investment of \$361,750 ($= 250,000 \times 1.4470$). In contrast, the second alternative requires only a small amount of cash to be deposited by the speculator in what is termed a “margin account”. (The operation of margin accounts is explained in Chapter 2.) In Table 1.4, the initial margin requirement is assumed to be \$5,000 per contract, or \$20,000 in total. The futures market allows the speculator to obtain leverage. With a relatively small initial outlay, the investor is able to take a large speculative position.

Table 1.4 Speculation using spot and futures contracts. One futures contract is on £62,500. Initial margin on four futures contracts = \$20,000.

| | <i>Possible trades</i> | |
|-------------------------------|---|---|
| | <i>Buy £250,000</i> <i>Spot price = 1.4470</i> | <i>Buy 4 futures contracts</i> <i>Futures price = 1.4410</i> |
| Investment | \$361,750 | \$20,000 |
| Profit if April spot = 1.5000 | \$13,250 | \$14,750 |
| Profit if April spot = 1.4000 | -\$11,750 | -\$10,250 |

Speculation Using Options

Options can also be used for speculation. Suppose that it is October and a speculator considers that a stock is likely to increase in value over the next 2 months. The stock price is currently \$20, and a 2-month call option with a \$22.50 strike price is currently selling for \$1. Table 1.5 illustrates two possible alternatives, assuming that the speculator is willing to invest \$2,000. One alternative is to purchase 100 shares; the other involves the purchase of 2,000 call options (i.e., 20 call option contracts). Suppose that the speculator's hunch is correct and the price of the stock rises to \$27 by December. The first alternative of buying the stock yields a profit of

$$100 \times (\$27 - \$20) = \$700$$

However, the second alternative is far more profitable. A call option on the stock with a strike price of \$22.50 gives a payoff of \$4.50, because it enables something worth \$27 to be bought for \$22.50. The total payoff from the 2,000 options that are purchased under the second alternative is

$$2,000 \times \$4.50 = \$9,000$$

Subtracting the original cost of the options yields a net profit of

$$\$9,000 - \$2,000 = \$7,000$$

The options strategy is, therefore, 10 times more profitable than directly buying the stock.

Options also give rise to a greater potential loss. Suppose the stock price falls to \$15 by December. The first alternative of buying stock yields a loss of

$$100 \times (\$20 - \$15) = \$500$$

Because the call options expire without being exercised, the options strategy would lead to a loss of \$2,000—the original amount paid for the options. Figure 1.5 shows the profit or loss from the two strategies as a function of the stock price in 2 months.

Options like futures provide a form of leverage. For a given investment, the use of options magnifies the financial consequences. Good outcomes become very good, while bad outcomes result in the whole initial investment being lost.


A Comparison

Futures and options are similar instruments for speculators in that they both provide a way in which a type of leverage can be obtained. However, there is an important difference between the two. When a speculator uses futures, the potential loss as well as

Table 1.5 Comparison of profits from two alternative strategies for using \$2,000 to speculate on a stock worth \$20 in October.

| <i>Investor's strategy</i> | <i>December stock price</i> | |
|----------------------------|-----------------------------|-------------|
| | <i>\$15</i> | <i>\$27</i> |
| Buy 100 shares | −\$500 | \$700 |
| Buy 2,000 call options | −\$2,000 | \$7,000 |

Figure 1.5 Profit or loss from two alternative strategies for speculating on a stock currently worth \$20.



the potential gain is very large. When options are used, no matter how bad things get, the speculator's loss is limited to the amount paid for the options.

1.9 ARBITRAGEURS

Arbitrageurs are a third important group of participants in futures, forward, and options markets. Arbitrage involves locking in a riskless profit by simultaneously entering into transactions in two or more markets. In later chapters we will see how arbitrage is sometimes possible when the futures price of an asset gets out of line with its spot price. We will also examine how arbitrage can be used in options markets. This section illustrates the concept of arbitrage with a very simple example.

Let us consider a stock that is traded on both the New York Stock Exchange (www.nyse.com) and the London Stock Exchange (www.stockex.co.uk). Suppose that the stock price is \$140 in New York and £100 in London at a time when the exchange rate is \$1.4300 per pound. An arbitrageur could simultaneously buy 100 shares of the stock in New York and sell them in London to obtain a risk-free profit of

$$100 \times [(\$1.43 \times 100) - \$140]$$

or \$300 in the absence of transactions costs. Transactions costs would probably eliminate the profit for a small investor. However, a large investment bank faces very low transactions costs in both the stock market and the foreign exchange market. It would find the arbitrage opportunity very attractive and would try to take as much advantage of it as possible.

Arbitrage opportunities such as the one just described cannot last for long. As arbitrageurs buy the stock in New York, the forces of supply and demand will cause

the dollar price to rise. Similarly, as they sell the stock in London, the sterling price will be driven down. Very quickly the two prices will become equivalent at the current exchange rate. Indeed, the existence of profit-hungry arbitrageurs makes it unlikely that a major disparity between the sterling price and the dollar price could ever exist in the first place. Generalizing from this example, we can say that the very existence of arbitrageurs means that in practice only very small arbitrage opportunities are observed in the prices that are quoted in most financial markets. In this book most of the arguments concerning futures prices, forward prices, and the values of option contracts will be based on the assumption that no arbitrage opportunities exist.

1.10 DANGERS

Derivatives are very versatile instruments. As we have seen, they can be used for hedging, for speculation, and for arbitrage. It is this very versatility that can cause problems. Sometimes traders who have a mandate to hedge risks or follow an arbitrage strategy become (consciously or unconsciously) speculators. The results can be disastrous. One example of this is provided by the activities of Jérôme Kerviel at Société Générale (see Business Snapshot 1.3).

To avoid the sort of problems Société Générale encountered, it is very important for both financial and nonfinancial corporations to set up controls to ensure that derivatives are being used for their intended purpose. Risk limits should be set and the activities of traders should be monitored daily to ensure that these risk limits are adhered to.

Unfortunately, even when traders follow the risk limits that have been specified, big mistakes can happen. Some of the activities of traders in the derivatives market during the period leading up to the start of the credit crisis in July 2007 proved to be much riskier than they were thought to be by the financial institutions they worked for. As will be discussed in Chapter 8, house prices in the United States had been rising fast. Most people thought that the increases would continue—or, at worst, that house prices would simply level off. Very few were prepared for the steep decline that actually happened. Furthermore, very few were prepared for the high correlation between mortgage default rates in different parts of the country. Some risk managers did express reservations about the exposures of the companies for which they worked to the US real estate market. But, when times are good (or appear to be good), there is an unfortunate tendency to ignore risk managers and this is what happened at many financial institutions during the 2006–2007 period. The key lesson from the credit crisis is that financial institutions should always be dispassionately asking “What can go wrong?”, and they should follow that up with the question “If it does go wrong, how much will we lose?”

SUMMARY

One of the exciting developments in finance over the last 30 years has been the growth of derivatives markets. In many situations, both hedgers and speculators find it more attractive to trade a derivative on an asset than to trade the asset itself. Some derivatives

Business Snapshot 1.3 SocGen's big loss in 2008

Derivatives are very versatile instruments. They can be used for hedging, speculation, and arbitrage. One of the risks faced by a company that trades derivatives is that an employee who has a mandate to hedge or to look for arbitrage opportunities may become a speculator.

Jérôme Kerviel joined Société Générale (SocGen) in 2000 to work in the compliance area. In 2005, he was promoted and became a junior trader in the bank's Delta One products team. He traded equity indices such as the German DAX index, the French CAC 40, and the Euro Stoxx 50. His job was to look for arbitrage opportunities. These might arise if a futures contract on an equity index was trading for a different price on two different exchanges. They might also arise if equity index futures prices were not consistent with the prices of the shares constituting the index. (This type of arbitrage is discussed in Chapter 5.)

Kerviel used his knowledge of the bank's procedures to speculate while giving the appearance of arbitraging. He took big positions in equity indices and created fictitious trades to make it appear that he was hedged. In reality, he had large bets on the direction in which the indices would move. The size of his unhedged position grew over time to tens of billions of euros.

In January 2008, his unauthorized trading was uncovered by SocGen. Over a three-day period, the bank unwound his position for a loss of 4.9 billion euros. This was at the time the biggest loss created by fraudulent activity in the history of finance. (Later in the year, a much bigger loss from Bernard Madoff's Ponzi scheme came to light.)

Rogue trader losses were not unknown at banks prior to 2008. For example, in the 1990s, Nick Leeson, who worked at Barings Bank, had a mandate similar to that of Jérôme Kerviel. His job was to arbitrage between Nikkei 225 futures quotes in Singapore and Osaka. Instead he found a way to make big bets on the direction of the Nikkei 225 using futures and options, losing \$1 billion and destroying the 200-year old bank in the process. In 2002, it was found that John Rusnak at Allied Irish Bank had lost \$700 million from unauthorized foreign exchange trading. The lessons from these losses are that it is important to define unambiguous risk limits for traders and then to monitor what they do very carefully to make sure that the limits are adhered to.

are traded on exchanges; others are traded by financial institutions, fund managers, and corporations in the over-the-counter market, or added to new issues of debt and equity securities. Much of this book is concerned with the valuation of derivatives. The aim is to present a unifying framework within which all derivatives—not just options or futures—can be valued.

In this chapter we have taken a first look at forward, futures, and options contracts. A forward or futures contract involves an obligation to buy or sell an asset at a certain time in the future for a certain price. There are two types of options: calls and puts. A call option gives the holder the right to buy an asset by a certain date for a certain price. A put option gives the holder the right to sell an asset by a certain date for a certain price. Forwards, futures, and options trade on a wide range of different underlying assets.

Derivatives have been very successful innovations in capital markets. Three main types of traders can be identified: hedgers, speculators, and arbitrageurs. Hedgers are in

the position where they face risk associated with the price of an asset. They use derivatives to reduce or eliminate this risk. Speculators wish to bet on future movements in the price of an asset. They use derivatives to get extra leverage. Arbitrageurs are in business to take advantage of a discrepancy between prices in two different markets. If, for example, they see the futures price of an asset getting out of line with the cash price, they will take offsetting positions in the two markets to lock in a profit.

FURTHER READING

Chancellor, E. *Devil Take the Hindmost—A History of Financial Speculation*. New York: Farrar Straus Giroux, 2000.

Merton, R. C. “Finance Theory and Future Trends: The Shift to Integration,” *Risk*, 12, 7 (July 1999): 48–51.

Miller, M. H. “Financial Innovation: Achievements and Prospects,” *Journal of Applied Corporate Finance*, 4 (Winter 1992): 4–11.

Zingales, L., “Causes and Effects of the Lehman Bankruptcy,” Testimony before Committee on Oversight and Government Reform, United States House of Representatives, October 6, 2008.

Practice Questions (Answers in Solutions Manual)

- 1.1. What is the difference between a long forward position and a short forward position?
- 1.2. Explain carefully the difference between hedging, speculation, and arbitrage.
- 1.3. What is the difference between entering into a long forward contract when the forward price is \$50 and taking a long position in a call option with a strike price of \$50?
- 1.4. Explain carefully the difference between selling a call option and buying a put option.
- 1.5. An investor enters into a short forward contract to sell 100,000 British pounds for US dollars at an exchange rate of 1.4000 US dollars per pound. How much does the investor gain or lose if the exchange rate at the end of the contract is (a) 1.3900 and (b) 1.4200?
- 1.6. A trader enters into a short cotton futures contract when the futures price is 50 cents per pound. The contract is for the delivery of 50,000 pounds. How much does the trader gain or lose if the cotton price at the end of the contract is (a) 48.20 cents per pound and (b) 51.30 cents per pound?
- 1.7. Suppose that you write a put contract with a strike price of \$40 and an expiration date in 3 months. The current stock price is \$41 and the contract is on 100 shares. What have you committed yourself to? How much could you gain or lose?
- 1.8. What is the difference between the over-the-counter market and the exchange-traded market? What are the bid and offer quotes of a market maker in the over-the-counter market?
- 1.9. You would like to speculate on a rise in the price of a certain stock. The current stock price is \$29 and a 3-month call with a strike price of \$30 costs \$2.90. You have \$5,800 to invest. Identify two alternative investment strategies, one in the stock and the other in an option on the stock. What are the potential gains and losses from each?

- 1.10. Suppose that you own 5,000 shares worth \$25 each. How can put options be used to provide you with insurance against a decline in the value of your holding over the next 4 months?
- 1.11. When first issued, a stock provides funds for a company. Is the same true of a stock option? Discuss.
- 1.12. Explain why a futures contract can be used for either speculation or hedging.
- 1.13. Suppose that a March call option to buy a share for \$50 costs \$2.50 and is held until March. Under what circumstances will the holder of the option make a profit? Under what circumstances will the option be exercised? Draw a diagram illustrating how the profit from a long position in the option depends on the stock price at maturity of the option.
- 1.14. Suppose that a June put option to sell a share for \$60 costs \$4 and is held until June. Under what circumstances will the seller of the option (i.e., the party with the short position) make a profit? Under what circumstances will the option be exercised? Draw a diagram illustrating how the profit from a short position in the option depends on the stock price at maturity of the option.
- 1.15. It is May and a trader writes a September call option with a strike price of \$20. The stock price is \$18 and the option price is \$2. Describe the trader's cash flows if the option is held until September and the stock price is \$25 at that time.
- 1.16. A trader writes a December put option with a strike price of \$30. The price of the option is \$4. Under what circumstances does the trader make a gain?
- 1.17. A company knows that it is due to receive a certain amount of a foreign currency in 4 months. What type of option contract is appropriate for hedging?
- 1.18. A US company expects to have to pay 1 million Canadian dollars in 6 months. Explain how the exchange rate risk can be hedged using (a) a forward contract and (b) an option.
- 1.19. A trader enters into a short forward contract on 100 million yen. The forward exchange rate is \$0.0080 per yen. How much does the trader gain or lose if the exchange rate at the end of the contract is (a) \$0.0074 per yen and (b) \$0.0091 per yen?
- 1.20. The Chicago Board of Trade offers a futures contract on long-term Treasury bonds. Characterize the traders likely to use this contract.
- 1.21. "Options and futures are zero-sum games." What do you think is meant by this?
- 1.22. Describe the profit from the following portfolio: a long forward contract on an asset and a long European put option on the asset with the same maturity as the forward contract and a strike price that is equal to the forward price of the asset at the time the portfolio is set up.
- 1.23. In the 1980s, Bankers Trust developed *index currency option notes* (ICONS). These are bonds in which the amount received by the holder at maturity varies with a foreign exchange rate. One example was its trade with the Long Term Credit Bank of Japan. The ICON specified that if the yen–US dollar exchange rate, S_T , is greater than 169 yen per dollar at maturity (in 1995), the holder of the bond receives \$1,000. If it is less than 169 yen per dollar, the amount received by the holder of the bond is

$$1,000 - \max\left[0, 1,000\left(\frac{169}{S_T} - 1\right)\right]$$

When the exchange rate is below 84.5, nothing is received by the holder at maturity. Show that this ICON is a combination of a regular bond and two options.

- 1.24. On July 1, 2011, a company enters into a forward contract to buy 10 million Japanese yen on January 1, 2012. On September 1, 2011, it enters into a forward contract to sell 10 million Japanese yen on January 1, 2012. Describe the payoff from this strategy.
- 1.25. Suppose that USD/sterling spot and forward exchange rates are as follows:

| | |
|-----------------|--------|
| Spot | 1.4580 |
| 90-day forward | 1.4556 |
| 180-day forward | 1.4518 |

What opportunities are open to an arbitrageur in the following situations?

- (a) A 180-day European call option to buy £1 for \$1.42 costs 2 cents.
 (b) A 90-day European put option to sell £1 for \$1.49 costs 2 cents.

Further Questions

- 1.26. Trader A enters into a forward contract to buy gold for \$1,000 an ounce in one year. Trader B buys a call option to buy gold for \$1,000 an ounce in one year. The cost of the option is \$100 an ounce. What is the difference between the positions of the traders? Show the profit per ounce as a function of the price of gold in one year for the two traders.
- 1.27. In March, a US investor instructs a broker to sell one July put option contract on a stock. The stock price is \$42 and the strike price is \$40. The option price is \$3. Explain what the investor has agreed to. Under what circumstances will the trade prove to be profitable? What are the risks?
- 1.28. A US company knows it will have to pay 3 million euros in three months. The current exchange rate is 1.4500 dollars per euro. Discuss how forward and options contracts can be used by the company to hedge its exposure.
- 1.29. A stock price is \$29. An investor buys one call option contract on the stock with a strike price of \$30 and sells a call option contract on the stock with a strike price of \$32.50. The market prices of the options are \$2.75 and \$1.50, respectively. The options have the same maturity date. Describe the investor's position.
- 1.30. The price of gold is currently \$1,000 per ounce. The forward price for delivery in 1 year is \$1,200. An arbitrageur can borrow money at 10% per annum. What should the arbitrageur do? Assume that the cost of storing gold is zero and that gold provides no income.
- 1.31. The current price of a stock is \$94, and 3-month European call options with a strike price of \$95 currently sell for \$4.70. An investor who feels that the price of the stock will increase is trying to decide between buying 100 shares and buying 2,000 call options (= 20 contracts). Both strategies involve an investment of \$9,400. What advice would you give? How high does the stock price have to rise for the option strategy to be more profitable?
- 1.32. On July 15, 2010, an investor owns 100 Google shares. As indicated in Table 1.3, the share price is about \$497 and a December put option with a strike price of \$460 costs \$27.30. The investor is comparing two alternatives to limit downside risk. The first involves buying one December put option contract with a strike price of \$460. The second involves instructing a broker to sell the 100 shares as soon as Google's price reaches \$460. Discuss the advantages and disadvantages of the two strategies.

- 1.33. A bond issued by Standard Oil some time ago worked as follows. The holder received no interest. At the bond's maturity the company promised to pay \$1,000 plus an additional amount based on the price of oil at that time. The additional amount was equal to the product of 170 and the excess (if any) of the price of a barrel of oil at maturity over \$25. The maximum additional amount paid was \$2,550 (which corresponds to a price of \$40 per barrel). Show that the bond is a combination of a regular bond, a long position in call options on oil with a strike price of \$25, and a short position in call options on oil with a strike price of \$40.
- 1.34. Suppose that in the situation of Table 1.1 a corporate treasurer said: "I will have £1 million to sell in 6 months. If the exchange rate is less than 1.41, I want you to give me 1.41. If it is greater than 1.47, I will accept 1.47. If the exchange rate is between 1.41 and 1.47, I will sell the sterling for the exchange rate." How could you use options to satisfy the treasurer?
- 1.35. Describe how foreign currency options can be used for hedging in the situation considered in Section 1.7 so that (a) ImportCo is guaranteed that its exchange rate will be less than 1.4600, and (b) ExportCo is guaranteed that its exchange rate will be at least 1.4200. Use DerivaGem to calculate the cost of setting up the hedge in each case assuming that the exchange rate volatility is 12%, interest rates in the United States are 5%, and interest rates in Britain are 5.7%. Assume that the current exchange rate is the average of the bid and offer in Table 1.1.
- 1.36. A trader buys a European call option and sells a European put option. The options have the same underlying asset, strike price, and maturity. Describe the trader's position. Under what circumstances does the price of the call equal the price of the put?



2 CHAPTER

Mechanics of Futures Markets

In Chapter 1 we explained that both futures and forward contracts are agreements to buy or sell an asset at a future time for a certain price. Futures contracts are traded on an organized exchange, and the contract terms are standardized by that exchange. By contrast, forward contracts are private agreements between two financial institutions or between a financial institution and one of its clients.

This chapter covers the details of how futures markets work. We examine issues such as the specification of contracts, the operation of margin accounts, the organization of exchanges, the regulation of markets, the way in which quotes are made, and the treatment of futures transactions for accounting and tax purposes. We compare futures contracts with forward contracts and explain the difference between the payoffs realized from them.

2.1 BACKGROUND

As we saw in Chapter 1, futures contracts are now traded actively all over the world. The Chicago Board of Trade, the Chicago Mercantile Exchange, and the New York Mercantile Exchange have merged to form the CME Group (www.cmegroup.com). Other large exchanges include NYSE Euronext (www.euronext.com), Eurex (www.eurexchange.com), BM&F BOVESPA (www.bmfbovespa.com.br), and the Tokyo International Financial Futures Exchange (www.tfx.co.jp). A table at the end of this book provides a more complete list of exchanges.

We examine how a futures contract comes into existence by considering the corn futures contract traded by the CME Group. On March 5 a trader in New York might call a broker with instructions to buy 5,000 bushels of corn for delivery in July of the same year. The broker would immediately issue instructions to a trader to buy (i.e., take a long position in) one July corn contract. (Each corn contract on CBOT is for the delivery of exactly 5,000 bushels.) At about the same time, another trader in Kansas might instruct a broker to sell 5,000 bushels of corn for July delivery. This broker would then issue instructions to sell (i.e., take a short position in) one corn contract. A price would be determined and the deal would be done. Under the traditional open outcry system, floor traders representing each party would physically meet to determine the price. With electronic trading, a computer would match the traders.

Business Snapshot 2.1 The Unanticipated Delivery of a Futures Contract

This story (which may well be apocryphal) was told to the author of this book by a senior executive of a financial institution. It concerns a new employee of the financial institution who had not previously worked in the financial sector. One of the clients of the financial institution regularly entered into a long futures contract on live cattle for hedging purposes and issued instructions to close out the position on the last day of trading. (Live cattle futures contracts are traded by the CME Group and each contract is on 40,000 pounds of cattle.) The new employee was given responsibility for handling the account.

When the time came to close out a contract the employee noted that the client was long one contract and instructed a trader at the exchange to buy (not sell) one contract. The result of this mistake was that the financial institution ended up with a long position in two live cattle futures contracts. By the time the mistake was spotted trading in the contract had ceased.

The financial institution (not the client) was responsible for the mistake. As a result, it started to look into the details of the delivery arrangements for live cattle futures contracts—something it had never done before. Under the terms of the contract, cattle could be delivered by the party with the short position to a number of different locations in the United States during the delivery month. Because it was long, the financial institution could do nothing but wait for a party with a short position to issue a *notice of intention to deliver* to the exchange and for the exchange to assign that notice to the financial institution.

It eventually received a notice from the exchange and found that it would receive live cattle at a location 2,000 miles away the following Tuesday. The new employee was sent to the location to handle things. It turned out that the location had a cattle auction every Tuesday. The party with the short position that was making delivery bought cattle at the auction and then immediately delivered them. Unfortunately the cattle could not be resold until the next cattle auction the following Tuesday. The employee was therefore faced with the problem of making arrangements for the cattle to be housed and fed for a week. This was a great start to a first job in the financial sector!

The trader in New York who agreed to buy has a *long futures position* in one contract; the trader in Kansas who agreed to sell has a *short futures position* in one contract. The price agreed to is the current *futures price* for July corn, say 300 cents per bushel. This price, like any other price, is determined by the laws of supply and demand. If, at a particular time, more traders wish to sell rather than buy July corn, the price will go down. New buyers then enter the market so that a balance between buyers and sellers is maintained. If more traders wish to buy rather than sell July corn, the price goes up. New sellers then enter the market and a balance between buyers and sellers is maintained.

Closing Out Positions

The vast majority of futures contracts do not lead to delivery. The reason is that most traders choose to close out their positions prior to the delivery period specified in the

contract. Closing out a position means entering into the opposite trade to the original one. For example, the New York investor who bought a July corn futures contract on March 5 can close out the position by selling (i.e., shorting) one July corn futures contract on, say, April 20. The Kansas investor who sold (i.e., shorted) a July contract on March 5 can close out the position by buying one July contract on, say, May 25. In each case, the investor's total gain or loss is determined by the change in the futures price between March 5 and the day when the contract is closed out.

Delivery is so unusual that traders sometimes forget how the delivery process works (see Business Snapshot 2.1). Nevertheless we will spend part of this chapter reviewing the delivery arrangements in futures contracts. This is because it is the possibility of final delivery that ties the futures price to the spot price.¹

2.2 SPECIFICATION OF A FUTURES CONTRACT

When developing a new contract, the exchange must specify in some detail the exact nature of the agreement between the two parties. In particular, it must specify the asset, the contract size (exactly how much of the asset will be delivered under one contract), where delivery will be made, and when delivery will be made.

Sometimes alternatives are specified for the grade of the asset that will be delivered or for the delivery locations. As a general rule, it is the party with the short position (the party that has agreed to sell the asset) that chooses what will happen when alternatives are specified by the exchange. When the party with the short position is ready to deliver, it files a *notice of intention to deliver* with the exchange. This notice indicates selections it has made with respect to the grade of asset that will be delivered and the delivery location.

The Asset

When the asset is a commodity, there may be quite a variation in the quality of what is available in the marketplace. When the asset is specified, it is therefore important that the exchange stipulate the grade or grades of the commodity that are acceptable. The IntercontinentalExchange (ICE) has specified the asset in its orange juice futures contract as frozen concentrates that are US Grade A with Brix value of not less than 62.5 degrees.

For some commodities a range of grades can be delivered, but the price received depends on the grade chosen. For example, in the CME Group's corn futures contract, the standard grade is "No. 2 Yellow," but substitutions are allowed with the price being adjusted in a way established by the exchange. No. 1 Yellow is deliverable for 1.5 cents per bushel more than No. 2 Yellow. No. 3 Yellow is deliverable for 1.5 cents per bushel less than No. 2 Yellow.

The financial assets in futures contracts are generally well defined and unambiguous. For example, there is no need to specify the grade of a Japanese yen. However, there are some interesting features of the Treasury bond and Treasury note futures contracts traded on the Chicago Board of Trade. The underlying asset in the Treasury bond contract is any long-term US Treasury bond that has a maturity of greater than 15 years

¹ As mentioned in Chapter 1, the spot price is the price for almost immediate delivery.

and is not callable within 15 years. In the Treasury note futures contract, the underlying asset is any long-term Treasury note with a maturity of no less than 6.5 years and no more than 10 years from the date of delivery. In both cases, the exchange has a formula for adjusting the price received according to the coupon and maturity date of the bond delivered. This is discussed in Chapter 6.

The Contract Size

The contract size specifies the amount of the asset that has to be delivered under one contract. This is an important decision for the exchange. If the contract size is too large, many investors who wish to hedge relatively small exposures or who wish to take relatively small speculative positions will be unable to use the exchange. On the other hand, if the contract size is too small, trading may be expensive as there is a cost associated with each contract traded.

The correct size for a contract clearly depends on the likely user. Whereas the value of what is delivered under a futures contract on an agricultural product might be \$10,000 to \$20,000, it is much higher for some financial futures. For example, under the Treasury bond futures contract traded by the CME Group, instruments with a face value of \$100,000 are delivered.

In some cases exchanges have introduced “mini” contracts to attract smaller investors. For example, the CME Group’s Mini Nasdaq 100 contract is on 20 times the Nasdaq 100 index, whereas the regular contract is on 100 times the index. (We will cover futures on indices more fully in Chapter 3.)

Delivery Arrangements

The place where delivery will be made must be specified by the exchange. This is particularly important for commodities that involve significant transportation costs. In the case of the ICE frozen concentrate orange juice contract, delivery is to exchange-licensed warehouses in Florida, New Jersey, or Delaware.

When alternative delivery locations are specified, the price received by the party with the short position is sometimes adjusted according to the location chosen by that party. The price tends to be higher for delivery locations that are relatively far from the main sources of the commodity.

Delivery Months

A futures contract is referred to by its delivery month. The exchange must specify the precise period during the month when delivery can be made. For many futures contracts, the delivery period is the whole month.

The delivery months vary from contract to contract and are chosen by the exchange to meet the needs of market participants. For example, corn futures traded by the CME Group have delivery months of March, May, July, September, and December. At any given time, contracts trade for the closest delivery month and a number of subsequent delivery months. The exchange specifies when trading in a particular month’s contract will begin. The exchange also specifies the last day on which trading can take place for a given contract. Trading generally ceases a few days before the last day on which delivery can be made.

Price Quotes

The exchange defines how prices will be quoted. For example, in the US, crude oil futures prices are quoted in dollars and cents. Treasury bond and Treasury note futures prices are quoted in dollars and thirty-seconds of a dollar.

Price Limits and Position Limits

For most contracts, daily price movement limits are specified by the exchange. If in a day the price moves down from the previous day's close by an amount equal to the daily price limit, the contract is said to be *limit down*. If it moves up by the limit, it is said to be *limit up*. A *limit move* is a move in either direction equal to the daily price limit. Normally, trading ceases for the day once the contract is limit up or limit down. However, in some instances the exchange has the authority to step in and change the limits.

The purpose of daily price limits is to prevent large price movements from occurring because of speculative excesses. However, limits can become an artificial barrier to trading when the price of the underlying commodity is advancing or declining rapidly. Whether price limits are, on balance, good for futures markets is controversial.

Position limits are the maximum number of contracts that a speculator may hold. The purpose of these limits is to prevent speculators from exercising undue influence on the market.

2.3 CONVERGENCE OF FUTURES PRICE TO SPOT PRICE

As the delivery period for a futures contract is approached, the futures price converges to the spot price of the underlying asset. When the delivery period is reached, the futures price equals—or is very close to—the spot price.


To see why this is so, we first suppose that the futures price is above the spot price during the delivery period. Traders then have a clear arbitrage opportunity:

1. Sell (i.e., short) a futures contract
2. Buy the asset
3. Make delivery.

These steps are certain to lead to a profit equal to the amount by which the futures price exceeds the spot price. As traders exploit this arbitrage opportunity, the futures price will fall. Suppose next that the futures price is below the spot price during the delivery period. Companies interested in acquiring the asset will find it attractive to enter into a long futures contract and then wait for delivery to be made. As they do so, the futures price will tend to rise.

The result is that the futures price is very close to the spot price during the delivery period. Figure 2.1 illustrates the convergence of the futures price to the spot price. In Figure 2.1(a) the futures price is above the spot price prior to the delivery period. In Figure 2.1(b) the futures price is below the spot price prior to the delivery period. The circumstances under which these two patterns are observed are discussed in Chapter 5.

Figure 2.1 Relationship between futures price and spot price as the delivery period is approached: (a) Futures price above spot price; (b) futures price below spot price.



2.4 THE OPERATION OF MARGINS

If two investors get in touch with each other directly and agree to trade an asset in the future for a certain price, there are obvious risks. One of the investors may regret the deal and try to back out. Alternatively, the investor simply may not have the financial resources to honor the agreement. One of the key roles of the exchange is to organize trading so that contract defaults are avoided. This is where margins come in.

Daily Settlement

To illustrate how margins work, we consider an investor who contacts his or her broker to buy two December gold futures contracts on the COMEX division of the New York Mercantile Exchange (NYMEX), which is part of the CME Group. We suppose that the current futures price is \$1,250 per ounce. Because the contract size is 100 ounces, the investor has contracted to buy a total of 200 ounces at this price. The broker will require the investor to deposit funds in a *margin account*. The amount that must be deposited at the time the contract is entered into is known as the *initial margin*. We suppose this is \$6,000 per contract, or \$12,000 in total. At the end of each trading day, the margin account is adjusted to reflect the investor's gain or loss. This practice is referred to as *daily settlement* or *marking to market*.

Suppose, for example, that by the end of the first day the futures price has dropped by \$9 from \$1,250 to \$1,241. The investor has a loss of \$1,800 ($= 200 \times \9), because the 200 ounces of December gold, which the investor contracted to buy at \$1,250, can now be sold for only \$1,241. The balance in the margin account would therefore be reduced by \$1,800 to \$10,200. Similarly, if the price of December gold rose to \$1,259 by the end of the first day, the balance in the margin account would be increased by \$1,800 to \$13,800. A trade is first settled at the close of the day on which it takes place. It is then settled at the close of trading on each subsequent day.

Note that daily settlement is not merely an arrangement between broker and client. When there is a decrease in the futures price so that the margin account of an investor with a long position is reduced by \$1,800, the investor's broker has to pay the exchange \$1,800 and the exchange passes the money on to the broker of an investor with a short position. Similarly, when there is an increase in the futures price, brokers for parties with short positions pay money to the exchange and brokers for parties with long positions receive money from the exchange. Later we will examine in more detail the mechanism by which this happens.

The investor is entitled to withdraw any balance in the margin account in excess of the initial margin. To ensure that the balance in the margin account never becomes negative a *maintenance margin*, which is somewhat lower than the initial margin, is set. If the balance in the margin account falls below the maintenance margin, the investor receives a margin call and is expected to top up the margin account to the initial margin level by the end of the next day. The extra funds deposited are known as a *variation margin*. If the investor does not provide the variation margin, the broker closes out the position. In the case of the investor considered earlier, closing out the position would involve neutralizing the existing contract by selling 200 ounces of gold for delivery in December.

Table 2.1 illustrates the operation of the margin account for one possible sequence of futures prices in the case of the investor considered earlier. The maintenance margin is assumed to be \$4,500 per contract, or \$9,000 in total. On Day 7, the balance in the margin account falls \$1,020 below the maintenance margin level. This drop triggers a

Table 2.1 Operation of margins for a long position in two gold futures contracts.

The initial margin is \$6,000 per contract, or \$12,000 in total; the maintenance margin is \$4,500 per contract, or \$9,000 in total. The contract is entered into on Day 1 at \$1,250 and closed out on Day 16 at \$1226.90.

| Day | Trade price (\$) | Settlement price (\$) | Daily gain (\$) | Cumulative gain (\$) | Margin account balance (\$) | Margin call (\$) |
|-----|------------------|-----------------------|-----------------|----------------------|-----------------------------|------------------|
| 1 | 1,250.00 | | | | 12,000 | |
| 1 | | 1,241.00 | -1,800 | -1,800 | 10,200 | |
| 2 | | 1,238.30 | -540 | -2,340 | 9,660 | |
| 3 | | 1,244.60 | 1,260 | -1,080 | 10,920 | |
| 4 | | 1,241.30 | -660 | -1,740 | 10,260 | |
| 5 | | 1,240.10 | -240 | -1,980 | 10,020 | |
| 6 | | 1,236.20 | -780 | -2,760 | 9,240 | |
| 7 | | 1,229.90 | -1,260 | -4,020 | 7,980 | 4,020 |
| 8 | | 1,230.80 | 180 | -3,840 | 12,180 | |
| 9 | | 1,225.40 | -1,080 | -4,920 | 11,100 | |
| 10 | | 1,228.10 | 540 | -4,380 | 11,640 | |
| 11 | | 1,211.00 | -3,420 | -7,800 | 8,220 | 3,780 |
| 12 | | 1,211.00 | 0 | -7,800 | 12,000 | |
| 13 | | 1,214.30 | 660 | -7,140 | 12,660 | |
| 14 | | 1,216.10 | 360 | -6,780 | 13,020 | |
| 15 | | 1,223.00 | 1,380 | -5,400 | 14,400 | |
| 16 | 1,226.90 | | 780 | -4,620 | 15,180 | |

margin call from the broker for an additional \$4,020 to bring the account balance up to the initial margin level of \$12,000. It is assumed that the investor provides this margin by the close of trading on Day 8. On Day 11, the balance in the margin account again falls below the maintenance margin level, and a margin call for \$3,780 is sent out. The investor provides this margin by the close of trading on Day 12. On Day 16, the investor decides to close out the position by selling two contracts. The futures price on that day is \$1,226.90, and the investor has a cumulative loss of \$4,620. Note that the investor has excess margin on Days 8, 13, 14, and 15. It is assumed that the excess is not withdrawn.

Further Details

Most brokers pay investors interest on the balance in a margin account. The balance in the account does not, therefore, represent a true cost, provided that the interest rate is competitive with what could be earned elsewhere. To satisfy the initial margin requirements, but not subsequent margin calls, an investor can usually deposit securities with the broker. Treasury bills are usually accepted in lieu of cash at about 90% of their face value. Shares are also sometimes accepted in lieu of cash, but at about 50% of their market value.

Whereas a forward contract is settled at the end of its life, a futures contract is, as we have seen, settled daily. At the end of each day, the investor's gain (loss) is added to (subtracted from) the margin account, bringing the value of the contract back to zero. A futures contract is in effect closed out and rewritten at a new price each day.

Minimum levels for initial and maintenance margins are set by the exchange. Individual brokers may require greater margins from their clients than those specified by the exchange. However, they cannot require lower margins than those specified by the exchange. Margin levels are determined by the variability of the price of the underlying asset. The higher this variability, the higher the margin levels. The maintenance margin is usually about 75% of the initial margin.

Margin requirements may depend on the objectives of the trader. A bona fide hedger, such as a company that produces the commodity on which the futures contract is written, is often subject to lower margin requirements than a speculator. The reason is that there is deemed to be less risk of default. Day trades and spread transactions often give rise to lower margin requirements than do hedge transactions. In a *day trade* the trader announces to the broker an intent to close out the position in the same day. In a *spread transaction* the trader simultaneously buys (i.e., takes a long position in) a contract on an asset for one maturity month and sells (i.e., takes a short position in) a contract on the same asset for another maturity month.

Note that margin requirements are the same on short futures positions as they are on long futures positions. It is just as easy to take a short futures position as it is to take a long one. The spot market does not have this symmetry. Taking a long position in the spot market involves buying the asset for immediate delivery and presents no problems. Taking a short position involves selling an asset that you do not own. This is a more complex transaction that may or may not be possible in a particular market. It is discussed further in Chapter 5.

The Clearing House and Clearing Margins

A *clearing house* acts as an intermediary in futures transactions. It guarantees the performance of the parties to each transaction. The clearing house has a number of

members, who must post funds with the clearing house. Brokers who are not members themselves must channel their business through a member. The main task of the clearing house is to keep track of all the transactions that take place during a day, so that it can calculate the net position of each of its members.

Just as an investor is required to maintain a margin account with a broker, the broker is required to maintain a margin account with a clearing house member and the clearing house member is required to maintain a margin account with the clearing house. The latter is known as a *clearing margin*. The margin accounts for clearing house members are adjusted for gains and losses at the end of each trading day in the same way as are the margin accounts of investors. However, in the case of the clearing house member, there is an original margin, but no maintenance margin. Every day the account balance for each contract must be maintained at an amount equal to the original margin times the number of contracts outstanding. Thus, depending on transactions during the day and price movements, the clearing house member may have to add funds to its margin account at the end of the day or it may find it can remove funds from the account at this time. Brokers who are not clearing house members must maintain a margin account with a clearing house member.

In determining clearing margins, the exchange clearing house calculates the number of contracts outstanding on either a gross or a net basis. When the gross basis is used, the number of contracts equals the sum of the long and short positions. When the net basis is used, these are offset against each other. Suppose a clearing house member has two clients: one with a long position in 20 contracts, the other with a short position in 15 contracts. Gross margining would calculate the clearing margin on the basis of 35 contracts; net margining would calculate the clearing margin on the basis of 5 contracts. Most exchanges currently use net margining.

Credit Risk

The whole purpose of the margining system is to ensure that funds are available to pay traders when they make a profit. Overall the system has been very successful. Traders entering into contracts at major exchanges have always had their contracts honored. Futures markets were tested on October 19, 1987, when the S&P 500 index declined by over 20% and traders with long positions in S&P 500 futures found they had negative margin balances. Traders who did not meet margin calls were closed out but still owed their brokers money. Some did not pay and as a result some brokers went bankrupt because, without their clients' money, they were unable to meet margin calls on contracts they entered into on behalf of their clients. However, the clearing houses had sufficient funds to ensure that everyone who had a short futures position on the S&P 500 got paid off.

2.5 OTC MARKETS

Credit risk has traditionally been a feature of the over-the-counter markets. There is always a chance that the party on the other side of an over-the-counter trade will default. It is interesting that, in an attempt to reduce credit risk, the over-the-counter market has adopted, or has been compelled to adopt, some of the procedures used by exchanges.

Business Snapshot 2.2 Long-Term Capital Management's Big Loss

Long-Term Capital Management (LTCM), a hedge fund formed in the mid-1990s, always collateralized its transactions. The hedge fund's investment strategy was known as convergence arbitrage. A very simple example of what it might do is the following. It would find two bonds, X and Y, issued by the same company that promised the same payoffs, with X being less liquid (i.e., less actively traded) than Y. The market always places a value on liquidity. As a result the price of X would be less than the price of Y. LTCM would buy X, short Y, and wait, expecting the prices of the two bonds to converge at some future time.

When interest rates increased, the company expected both bonds to move down in price by about the same amount, so that the collateral it paid on bond X would be about the same as the collateral it received on bond Y. Similarly, when interest rates decreased, LTCM expected both bonds to move up in price by about the same amount, so that the collateral it received on bond X would be about the same as the collateral it paid on bond Y. It therefore expected that there would be no significant outflow of funds as a result of its collateralization agreements.

In August 1998, Russia defaulted on its debt and this led to what is termed a "flight to quality" in capital markets. One result was that investors valued liquid instruments more highly than usual and the spreads between the prices of the liquid and illiquid instruments in LTCM's portfolio increased dramatically. The prices of the bonds LTCM had bought went down and the prices of those it had shorted increased. It was required to post collateral on both. The company experienced difficulties because it was highly leveraged. Positions had to be closed out and LTCM lost about \$4 billion. If the company had been less highly leveraged, it would probably have been able to survive the flight to quality and could have waited for the prices of the liquid and illiquid bonds to move back closer to each other.

Collateralization

Collateralization has been used in OTC markets for some time and is similar to the practice of posting margin in futures markets.

Consider two companies, A and B, that have entered into an OTC derivatives transaction such as a forward. A collateralization agreement applying to the transaction might involve the transaction being valued each day. If, from one day to the next, the value of the transaction to company A increases by a positive amount X (so that the value to company B decreases by X), company B is required to pay X to company A. Similarly, if the value to company B increases by a positive amount X (so that the value to company A decreases by X), company A is required to pay X to company B. The contract is not settled daily, as in the case of futures. The payments are a security deposit designed to ensure that obligations will be honored. Interest is paid on the full amount of the funds that have been deposited by one party with the other.

There are many variations on this simple arrangement and collateralization is discussed more fully in Chapter 23. Collateralization significantly reduces the credit risk in OTC contracts. As discussed in Business Snapshot 2.2, it was used by the hedge fund Long-Term Capital Management (LTCM) in the 1990s. As a result LTCM's counterparties were prepared to accept LTCM's credit risk.

The Use of Clearing Houses in OTC Markets

Since the 2007–2009 crisis, governments in the US and elsewhere have passed legislation requiring clearing houses to be used for some OTC transactions.


The way in which clearing houses work in the OTC market is as follows. An OTC transaction is negotiated between two parties, A and B, in the usual way. It is then presented to a clearing house (sometimes called a central clearing party). Assuming the clearing house accepts the transaction, it becomes the counterparty to both A and B. (This is similar to the way the clearing house for a futures exchange becomes the counterparty to the two sides of a futures trade.) The clearing house takes on the credit risk of both A and B. It manages this risk by requiring an initial margin and daily variation margins from them.

The OTC market has traditionally been a series of bilateral agreements between market participants as illustrated in Figure 2.2a. If all OTC contracts were cleared in the way that has just been described, the OTC market would move to the situation where each participant deals with one or more clearing houses, as illustrated in Figure 2.2b. In practice, because not all OTC transactions are routed through clearing houses, the market has elements of both Figure 2.2a and 2.2b.

A number of arguments have been cited for the use of clearing houses in OTC markets. Collateral will automatically have to be posted; credit risk in the financial system will (hopefully) be reduced;² and the trades taking place in the OTC market will become more transparent. A major concern of governments since the credit crisis of 2007 is *systemic risk*. This is the risk that a failure by a large financial institution will lead to failures by other large financial institutions and a collapse of the financial system. The way this can happen is described in Business Snapshot 2.3.

One of the motivations for the legislation requiring that clearing houses be used for OTC transactions is what might be termed the “AIG fiasco.” During the period

Figure 2.2 (a) The traditional way in which OTC markets have operated: a series of bilateral agreements between market participants; (b) how OTC markets would operate with a single central clearing house.



² The impact of clearing houses on credit risk depends on the number of clearing houses and the proportion of all OTC trades that are cleared through them. See D. Duffie and H. Zhu (2010), “Does a Central Clearing Counterparty Reduce Counterparty Risk?” Working Paper, Stanford University.

Business Snapshot 2.3 Systemic risk

Systemic risk is the risk that a default by one financial institution will create a “ripple effect” that leads to defaults by other financial institutions and threatens the stability of the financial system. There are huge numbers of over-the-counter transactions between banks. If Bank A fails, Bank B may take a huge loss on the transactions it has with Bank A. This in turn could lead to Bank B failing. Bank C that has many outstanding transactions with both Bank A and Bank B might then take a large loss and experience severe financial difficulties; and so on.

The financial system has survived defaults such as Drexel in 1990 and Lehman Brothers in 2008, but regulators continue to be concerned. During the market turmoil of 2007 and 2008, many large financial institutions were bailed out, rather than being allowed to fail, because governments were concerned about systemic risk.

leading up to the credit crisis, the insurance company AIG provided protection to other financial institutions against a huge volume of credit risks that were related to subprime mortgages. Since AIG had a AAA credit rating at the time the transactions were negotiated, it was not required to post collateral by its counterparties. The transactions resulted in big losses for AIG and led to an \$85 billion bailout of the company by the US government. Whether the clearing house legislation by itself will prevent companies taking risks as large as those of AIG in the future is doubtful. This is because the legislation applies only to “standardized” OTC transactions and AIG’s transactions were nonstandard. However, mandatory collateralization for nonstandard OTC contracts will go a long way toward preventing another AIG occurring in the future.

2.6 MARKET QUOTES

Futures quotes are available from exchanges and from several online sources (see, for example, futures.tradingcharts.com/marketquotes). Table 2.2 shows quotes provided by exchanges for a number of different commodities on May 26, 2010. Quotes for index, currency, and interest rate futures are given in Chapters 3, 5, and 6, respectively.

The asset underlying the futures contract, the exchange that the contract is traded on, the contract size, and how the price is quoted are all shown at the top of each section in Table 2.2. The first asset is gold, traded on COMEX (a division of the New York Mercantile Exchange, which is now part of the CME Group). The contract size is 100 ounces, and the price is quoted in dollars per ounce. The maturity month of the contract is shown in the first column.

Prices

The first three numbers in each row show the opening price, the highest price achieved in trading during the day, and the lowest price achieved in trading during the day. The opening price is representative of the prices at which contracts were trading immediately after the start of trading. For June 2010 gold, the opening price on May 26, 2010, was \$1,203.80. During the day, the price traded between \$1,201.00 and \$1,216.90.

Table 2.2 Futures quotes for a selection of CME Group contracts on commodities on May 26, 2010.

| | <i>Open</i> | <i>High</i> | <i>Low</i> | <i>Settlement</i> | <i>Change</i> | <i>Volume</i> | <i>Open interest</i> |
|---|-------------|-------------|------------|-------------------|---------------|---------------|----------------------|
| Gold 100 oz, \$ per oz | | | | | | | |
| June 2010 | 1203.80 | 1216.90 | 1201.00 | 1213.40 | 15.40 | 194,461 | 156,156 |
| July 2010 | 1205.00 | 1217.50 | 1202.00 | 1214.20 | 15.50 | 838 | 714 |
| Aug. 2010 | 1205.00 | 1218.70 | 1202.70 | 1215.30 | 15.50 | 130,676 | 240,074 |
| Oct. 2010 | 1208.30 | 1220.20 | 1205.30 | 1217.50 | 15.60 | 2,445 | 21,792 |
| Dec. 2010 | 1208.80 | 1222.90 | 1207.50 | 1219.90 | 15.60 | 7,885 | 61,497 |
| June 2011 | 1215.90 | 1228.00 | 1215.20 | 1227.80 | 15.80 | 408 | 13,461 |
| Crude oil 1,000 barrels, \$ per barrel | | | | | | | |
| July 2010 | 70.06 | 71.70 | 69.21 | 71.51 | 2.76 | 6,315 | 388,902 |
| Aug. 2010 | 71.25 | 72.77 | 70.42 | 72.54 | 2.44 | 3,746 | 115,305 |
| Dec. 2010 | 74.00 | 75.34 | 73.17 | 75.23 | 2.19 | 5,055 | 196,033 |
| Dec. 2011 | 77.01 | 78.59 | 76.51 | 78.53 | 2.00 | 4,175 | 100,674 |
| Dec. 2012 | 78.50 | 80.21 | 78.50 | 80.18 | 1.86 | 1,258 | 70,126 |
| Corn 5,000 bushels, cents per bushel | | | | | | | |
| July 2010 | 369.50 | 372.00 | 368.75 | 371.50 | 7.25 | 122,528 | 491,587 |
| Sept. 2010 | 379.50 | 381.00 | 379.00 | 381.00 | 7.75 | 24,186 | 175,798 |
| Dec. 2010 | 389.00 | 390.75 | 380.25 | 390.75 | 8.00 | 47,428 | 373,026 |
| Mar. 2011 | 400.00 | 403.25 | 400.00 | 403.25 | 7.75 | 4,581 | 55,836 |
| May 2011 | 410.50 | 411.50 | 410.50 | 411.50 | 7.25 | 830 | 8,995 |
| July 2011 | 417.50 | 419.50 | 417.50 | 419.50 | 7.50 | 3,491 | 31,939 |
| Dec. 2011 | 416.25 | 418.00 | 415.75 | 418.00 | 7.25 | 4,760 | 59,061 |
| Soybeans 5,000 bushels, cents per bushel | | | | | | | |
| July 2010 | 934.25 | 939.75 | 933.00 | 938.00 | 7.50 | 41,816 | 220,712 |
| Aug. 2010 | 922.00 | 931.50 | 922.00 | 929.50 | 8.50 | 4,881 | 15,674 |
| Sept. 2010 | 914.50 | 918.75 | 912.50 | 916.50 | 7.00 | 1,935 | 12,983 |
| Nov. 2010 | 906.00 | 912.50 | 905.00 | 910.00 | 7.00 | 18,908 | 157,826 |
| Jan. 2011 | 917.75 | 921.50 | 914.75 | 919.75 | 7.00 | 2,621 | 12,391 |
| Mar. 2011 | 926.00 | 930.00 | 925.00 | 928.50 | 8.00 | 1,406 | 5,857 |
| May 2011 | 933.50 | 935.50 | 931.00 | 933.50 | 7.50 | 942 | 5,626 |
| Wheat 5,000 bushels, cents per bushel | | | | | | | |
| July 2010 | 462.75 | 472.00 | 459.00 | 461.75 | 1.25 | 45,283 | 246,683 |
| Sept. 2010 | 480.00 | 489.00 | 476.50 | 479.00 | 1.00 | 13,941 | 90,257 |
| Dec. 2010 | 510.75 | 519.50 | 507.25 | 510.00 | 1.25 | 9,756 | 70,618 |
| Mar. 2011 | 541.50 | 548.50 | 536.75 | 539.00 | 0.75 | 2,748 | 27,879 |
| May 2011 | 557.00 | 563.50 | 552.75 | 555.50 | 1.25 | 923 | 8,199 |
| July 2011 | 574.25 | 583.00 | 571.00 | 573.75 | 0.75 | 4,938 | 34,300 |
| Live Cattle 40,000 lb, cents per lb | | | | | | | |
| June 2010 | 90.800 | 90.850 | 90.450 | 90.800 | 0.775 | 12,410 | 51,817 |
| Aug. 2010 | 89.700 | 90.050 | 89.525 | 89.925 | 0.850 | 19,341 | 144,587 |
| Oct. 2010 | 91.100 | 91.150 | 90.750 | 91.100 | 0.750 | 7,718 | 78,300 |
| Dec. 2010 | 92.100 | 92.250 | 91.875 | 92.175 | 0.800 | 3,347 | 42,102 |
| Feb. 2011 | 93.200 | 93.550 | 93.200 | 93.550 | 0.800 | 792 | 18,428 |

Settlement Price

The fourth number is the *settlement price*. This is the price used for calculating daily gains and losses and margin requirements. It is usually calculated as the price at which the contract traded immediately before the end of a day's trading session (1:30 p.m. for gold). The fifth number is the change in the settlement price from the previous day. For the June 2010 gold futures contract, the settlement price on May 26, 2010, was \$1,213.40, up \$15.40 from the previous trading day. In this case, an investor with a long position in one contract would find his or her margin account balance increased by \$1,540 (= 100 × \$15.40) on May 26, 2010. Similarly, an investor with a short position in one contract would find that the margin balance decreased by \$1,540 on this date.

The numbers in the fifth column show that, by chance, settlement prices for all the contracts considered increased between May 25 and May 26, 2010.

Trading Volume and Open Interest

The final two columns in Table 2.2 show the trading volume for the day and the open interest at the end of the previous day. The trading volume is the number of contracts traded. The open interest is the number of contracts outstanding, that is, the number of long positions or, equivalently, the number of short positions.

Trading volume can be greater than both the beginning-of-day and end-of-day open interest. (This was the case for June 2010 gold on May 26, 2010.) This indicates that many traders who entered into positions during the day closed them out before the end of the day. (Traders who do this are referred to as *day traders*.)

Patterns of Futures Prices

Futures prices can show a number of different patterns. The futures price of gold generally increases with the maturity of the contract. Table 2.2 shows that this was the case on May 26, 2010. The settlement price on that day increased from \$1213.40 to \$1227.80 as the contract maturity month increased from June 2010 to June 2011. Markets where the futures price is an increasing function of the time to maturity are known as *normal markets*. Markets where the futures price decreases with the maturity of the futures contract are known as *inverted markets*.³

Table 2.2 shows that there was a normal market for crude oil on May 26, 2010. This is not always the case. For example, on October 15, 2007, oil futures prices were inverted. The November 2007, December 2007, January 2008, February 2008, March 2008, and April 2008 settlement prices were 86.13, 85.13, 84.25, 83.41, 82.69, and 82.05, respectively. Sometimes futures prices, perhaps because of seasonality, show a mixture of normal and inverted markets. For example, on May 26, 2010, the futures price of soybeans first decreased and then increased as the maturity of the contract increased.

³ The term *contango* is sometimes used to describe situations where the futures price is an increasing function of the maturity of the contract and the term *backwardation* is sometimes used to describe the situation where the futures price is a decreasing function of the maturity of the contract. Strictly speaking, as will be explained in Chapter 5, these terms refer to whether the price of the underlying asset is expected to increase or decrease over time.

2.7 DELIVERY

As mentioned earlier in this chapter, very few of the futures contracts that are entered into lead to delivery of the underlying asset. Most are closed out early. Nevertheless, it is the possibility of eventual delivery that determines the futures price. An understanding of delivery procedures is therefore important.

The period during which delivery can be made is defined by the exchange and varies from contract to contract. The decision on when to deliver is made by the party with the short position, whom we shall refer to as investor A. When investor A decides to deliver, investor A's broker issues a notice of intention to deliver to the exchange clearing house. This notice states how many contracts will be delivered and, in the case of commodities, also specifies where delivery will be made and what grade will be delivered. The exchange then chooses a party with a long position to accept delivery.

Suppose that the party on the other side of investor A's futures contract when it was entered into was investor B. It is important to realize that there is no reason to expect that it will be investor B who takes delivery. Investor B may well have closed out his or her position by trading with investor C, investor C may have closed out his or her position by trading with investor D, and so on. The usual rule chosen by the exchange is to pass the notice of intention to deliver on to the party with the oldest outstanding long position. Parties with long positions must accept delivery notices. However, if the notices are transferable, long investors have a short period of time, usually half an hour, to find another party with a long position that is prepared to accept the notice from them.

In the case of a commodity, taking delivery usually means accepting a warehouse receipt in return for immediate payment. The party taking delivery is then responsible for all warehousing costs. In the case of livestock futures, there may be costs associated with feeding and looking after the animals (see Business Snapshot 2.1). In the case of financial futures, delivery is usually made by wire transfer. For all contracts, the price paid is usually the most recent settlement price. If specified by the exchange, this price is adjusted for grade, location of delivery, and so on. The whole delivery procedure from the issuance of the notice of intention to deliver to the delivery itself generally takes about two to three days.

There are three critical days for a contract. These are the first notice day, the last notice day, and the last trading day. The *first notice day* is the first day on which a notice of intention to make delivery can be submitted to the exchange. The *last notice day* is the last such day. The *last trading day* is generally a few days before the last notice day. To avoid the risk of having to take delivery, an investor with a long position should close out his or her contracts prior to the first notice day.

Cash Settlement

Some financial futures, such as those on stock indices discussed in Chapter 3, are settled in cash because it is inconvenient or impossible to deliver the underlying asset. In the case of the futures contract on the S&P 500, for example, delivering the underlying asset would involve delivering a portfolio of 500 stocks. When a contract is settled in cash, all outstanding contracts are declared closed on a predetermined day. The final settlement price is set equal to the spot price of the underlying asset at either the opening or close of trading on that day. For example, in the S&P 500 futures contract traded by the

CME Group, the predetermined day is the third Friday of the delivery month and final settlement is at the opening price.

2.8 TYPES OF TRADERS AND TYPES OF ORDERS

There are two main types of traders executing trades: *futures commission merchants* (FCMs) and *locals*. FCMs are following the instructions of their clients and charge a commission for doing so; locals are trading on their own account.

Individuals taking positions, whether locals or the clients of FCMs, can be categorized as hedgers, speculators, or arbitrageurs, as discussed in Chapter 1. Speculators can be classified as scalpers, day traders, or position traders. *Scalpers* are watching for very short-term trends and attempt to profit from small changes in the contract price. They usually hold their positions for only a few minutes. *Day traders* hold their positions for less than one trading day. They are unwilling to take the risk that adverse news will occur overnight. *Position traders* hold their positions for much longer periods of time. They hope to make significant profits from major movements in the markets.

Orders

The simplest type of order placed with a broker is a *market order*. It is a request that a trade be carried out immediately at the best price available in the market. However, there are many other types of orders. We will consider those that are more commonly used.

A *limit order* specifies a particular price. The order can be executed only at this price or at one more favorable to the investor. Thus, if the limit price is \$30 for an investor wanting to buy, the order will be executed only at a price of \$30 or less. There is, of course, no guarantee that the order will be executed at all, because the limit price may never be reached.

A *stop order* or *stop-loss order* also specifies a particular price. The order is executed at the best available price once a bid or offer is made at that particular price or a less-favorable price. Suppose a stop order to sell at \$30 is issued when the market price is \$35. It becomes an order to sell when and if the price falls to \$30. In effect, a stop order becomes a market order as soon as the specified price has been hit. The purpose of a stop order is usually to close out a position if unfavorable price movements take place. It limits the loss that can be incurred.

A *stop-limit order* is a combination of a stop order and a limit order. The order becomes a limit order as soon as a bid or offer is made at a price equal to or less favorable than the stop price. Two prices must be specified in a stop-limit order: the stop price and the limit price. Suppose that at the time the market price is \$35, a stop-limit order to buy is issued with a stop price of \$40 and a limit price of \$41. As soon as there is a bid or offer at \$40, the stop-limit becomes a limit order at \$41. If the stop price and the limit price are the same, the order is sometimes called a *stop-and-limit order*.

A *market-if-touched* (MIT) order is executed at the best available price after a trade occurs at a specified price or at a price more favorable than the specified price. In effect, an MIT becomes a market order once the specified price has been hit. An MIT is also known as a *board order*. Consider an investor who has a long position in a futures contract and is issuing instructions that would lead to closing out the contract. A stop order is designed to place a limit on the loss that can occur in the event of unfavorable

price movements. By contrast, a market-if-touched order is designed to ensure that profits are taken if sufficiently favorable price movements occur.

A *discretionary order* or *market-not-held order* is traded as a market order except that execution may be delayed at the broker's discretion in an attempt to get a better price.

Some orders specify time conditions. Unless otherwise stated, an order is a day order and expires at the end of the trading day. A *time-of-day order* specifies a particular period of time during the day when the order can be executed. An *open order* or a *good-till-canceled order* is in effect until executed or until the end of trading in the particular contract. A *fill-or-kill order*, as its name implies, must be executed immediately on receipt or not at all.

2.9 REGULATION

Futures markets in the United States are currently regulated federally by the Commodity Futures Trading Commission (CFTC; www.cftc.gov), which was established in 1974. This body is responsible for licensing futures exchanges and approving contracts. All new contracts and changes to existing contracts must be approved by the CFTC. To be approved, the contract must have some useful economic purpose. Usually this means that it must serve the needs of hedgers as well as speculators.

The CFTC looks after the public interest. It is responsible for ensuring that prices are communicated to the public and that futures traders report their outstanding positions if they are above certain levels. The CFTC also licenses all individuals who offer their services to the public in futures trading. The backgrounds of these individuals are investigated, and there are minimum capital requirements. The CFTC deals with complaints brought by the public and ensures that disciplinary action is taken against individuals when appropriate. It has the authority to force exchanges to take disciplinary action against members who are in violation of exchange rules.

With the formation of the National Futures Association (NFA; www.nfa.futures.org) in 1982, some of responsibilities of the CFTC were shifted to the futures industry itself. The NFA is an organization of individuals who participate in the futures industry. Its objective is to prevent fraud and to ensure that the market operates in the best interests of the general public. It is authorized to monitor trading and take disciplinary action when appropriate. The agency has set up an efficient system for arbitrating disputes between individuals and its members.

From time to time, other bodies, such as the Securities and Exchange Commission (SEC; www.sec.gov), the Federal Reserve Board (www.federalreserve.gov), and the US Treasury Department (www.treas.gov), have claimed jurisdictional rights over some aspects of futures trading. These bodies are concerned with the effects of futures trading on the spot markets for securities such as stocks, Treasury bills, and Treasury bonds. The SEC currently has an effective veto over the approval of new stock or bond index futures contracts. However, the basic responsibility for all futures and options on futures rests with the CFTC.

Trading Irregularities

Most of the time futures markets operate efficiently and in the public interest. However, from time to time, trading irregularities do come to light. One type of trading

irregularity occurs when an investor group tries to “corner the market.”⁴ The investor group takes a huge long futures position and also tries to exercise some control over the supply of the underlying commodity. As the maturity of the futures contracts is approached, the investor group does not close out its position, so that the number of outstanding futures contracts may exceed the amount of the commodity available for delivery. The holders of short positions realize that they will find it difficult to deliver and become desperate to close out their positions. The result is a large rise in both futures and spot prices. Regulators usually deal with this type of abuse of the market by increasing margin requirements or imposing stricter position limits or prohibiting trades that increase a speculator’s open position or requiring market participants to close out their positions.

Other types of trading irregularity can involve the traders on the floor of the exchange. These received some publicity early in 1989, when it was announced that the FBI had carried out a two-year investigation, using undercover agents, of trading on the Chicago Board of Trade and the Chicago Mercantile Exchange. The investigation was initiated because of complaints filed by a large agricultural concern. The alleged offenses included overcharging customers, not paying customers the full proceeds of sales, and traders using their knowledge of customer orders to trade first for themselves (an offence known as *front running*).

2.10 ACCOUNTING AND TAX

The full details of the accounting and tax treatment of futures contracts are beyond the scope of this book. A trader who wants detailed information on this should consult experts. In this section we provide some general background information.

Accounting

Accounting standards require changes in the market value of a futures contract to be recognized when they occur unless the contract qualifies as a hedge. If the contract does qualify as a hedge, gains or losses are generally recognized for accounting purposes in the same period in which the gains or losses from the item being hedged are recognized. The latter treatment is referred to as *hedge accounting*.

Consider a company with a December year end. In September 2011 it buys a March 2012 corn futures contract and closes out the position at the end of February 2012. Suppose that the futures prices are 250 cents per bushel when the contract is entered into, 270 cents per bushel at the end of 2011, and 280 cents per bushel when the contract is closed out. The contract is for the delivery of 5,000 bushels. If the contract does not qualify as a hedge, the gains for accounting purposes are

$$5,000 \times (2.70 - 2.50) = \$1,000$$

in 2011 and

$$5,000 \times (2.80 - 2.70) = \$500$$

⁴ Possibly the best known example of this was the attempt by the Hunt brothers to corner the silver market in 1979–80. Between the middle of 1979 and the beginning of 1980, their activities led to a price rise from \$6 per ounce to \$50 per ounce.

in 2012. If the company is hedging the purchase of 5,000 bushels of corn in February 2012 so that the contract qualifies for hedge accounting, the entire gain of \$1,500 is realized in 2012 for accounting purposes.

The treatment of hedging gains and losses is sensible. If the company is hedging the purchase of 5,000 bushels of corn in February 2012, the effect of the futures contract is to ensure that the price paid is close to 250 cents per bushel. The accounting treatment reflects that this price is paid in 2012.

In June 1998, the Financial Accounting Standards Board issued Statement No. 133 (FAS 133), Accounting for Derivative Instruments and Hedging Activities. FAS 133 applies to all types of derivatives (including futures, forwards, swaps, and options). It requires all derivatives to be included on the balance sheet at fair market value.⁵ It increases disclosure requirements. It also gives companies far less latitude than previously in using hedge accounting. For hedge accounting to be used, the hedging instrument must be highly effective in offsetting exposures and an assessment of this effectiveness is required every three months. A similar standard IAS 39 has been issued by the International Accounting Standards Board.

Tax

Under the US tax rules, two key issues are the nature of a taxable gain or loss and the timing of the recognition of the gain or loss. Gains or losses are either classified as capital gains or losses or alternatively as part of ordinary income.

For a corporate taxpayer, capital gains are taxed at the same rate as ordinary income, and the ability to deduct losses is restricted. Capital losses are deductible only to the extent of capital gains. A corporation may carry back a capital loss for three years and carry it forward for up to five years. For a noncorporate taxpayer, short-term capital gains are taxed at the same rate as ordinary income, but long-term capital gains are subject to a maximum capital gains tax rate of 15%. (Long-term capital gains are gains from the sale of a capital asset held for longer than one year; short-term capital gains are the gains from the sale of a capital asset held one year or less.) For a noncorporate taxpayer, capital losses are deductible to the extent of capital gains plus ordinary income up to \$3,000 and can be carried forward indefinitely.

Generally, positions in futures contracts are treated as if they are closed out on the last day of the tax year. For the noncorporate taxpayer, this gives rise to capital gains and losses that are treated as if they were 60% long term and 40% short term without regard to the holding period. This is referred to as the “60/40” rule. A noncorporate taxpayer may elect to carry back for three years any net losses from the 60/40 rule to offset any gains recognized under the rule in the previous three years.

Hedging transactions are exempt from this rule. The definition of a hedge transaction for tax purposes is different from that for accounting purposes. The tax regulations define a hedging transaction as a transaction entered into in the normal course of business primarily for one of the following reasons:

1. To reduce the risk of price changes or currency fluctuations with respect to property that is held or to be held by the taxpayer for the purposes of producing ordinary income

⁵ Previously the attraction of derivatives in some situations was that they were “off-balance-sheet” items.

2. To reduce the risk of price or interest rate changes or currency fluctuations with respect to borrowings made by the taxpayer.

A hedging transaction must be clearly identified as such in the company's records. Gains or losses from hedging transactions are treated as ordinary income. The timing of the recognition of gains or losses from hedging transactions generally matches the timing of the recognition of income or expense associated with the transaction being hedged.

2.11 FORWARD vs. FUTURES CONTRACTS

The main differences between forward and futures contracts are summarized in Table 2.3. Both contracts are agreements to buy or sell an asset for a certain price at a certain future time. A forward contract is traded in the over-the-counter market and there is no standard contract size or standard delivery arrangements. A single delivery date is usually specified and the contract is usually held to the end of its life and then settled. A futures contract is a standardized contract traded on an exchange. A range of delivery dates is usually specified. It is settled daily and usually closed out prior to maturity.

Profits from Forward and Futures Contracts

Suppose that the sterling exchange rate for a 90-day forward contract is 1.5000 and that this rate is also the futures price for a contract that will be delivered in exactly 90 days. What is the difference between the gains and losses under the two contracts?

Under the forward contract, the whole gain or loss is realized at the end of the life of the contract. Under the futures contract, the gain or loss is realized day by day because of the daily settlement procedures. Suppose that investor A is long £1 million in a 90-day forward contract and investor B is long £1 million in 90-day futures contracts. (Because each futures contract is for the purchase or sale of £62,500, investor B must purchase a total of 16 contracts.) Assume that the spot exchange rate in 90 days proves to be 1.7000 dollars per pound. Investor A makes a gain of \$200,000 on the 90th day. Investor B makes the same gain—but spread out over the 90-day period. On some days investor B may realize a loss, whereas on other days he

Table 2.3 Comparison of forward and futures contracts.

| <i>Forward</i> | <i>Futures</i> |
|---|--|
| Private contract between two parties | Traded on an exchange |
| Not standardized | Standardized contract |
| Usually one specified delivery date | Range of delivery dates |
| Settled at end of contract | Settled daily |
| Delivery or final cash settlement usually takes place | Contract is usually closed out prior to maturity |
| Some credit risk | Virtually no credit risk |

or she makes a gain. However, in total, when losses are netted against gains, there is a gain of \$200,000 over the 90-day period.

Foreign Exchange Quotes

Both forward and futures contracts trade actively on foreign currencies. However, there is sometimes a difference in the way exchange rates are quoted in the two markets. For example, futures prices where one currency is the US dollar are always quoted as the number of US dollars per unit of the foreign currency or as the number of US cents per unit of the foreign currency. Forward prices are always quoted in the same way as spot prices. This means that, for the British pound, the euro, the Australian dollar, and the New Zealand dollar, the forward quotes show the number of US dollars per unit of the foreign currency and are directly comparable with futures quotes. For other major currencies, forward quotes show the number of units of the foreign currency per US dollar (USD). Consider the Canadian dollar (CAD). A futures price quote of 0.9500 USD per CAD corresponds to a forward price quote of 1.0526 CAD per USD ($1.0526 = 1/0.9500$).

SUMMARY

A very high proportion of the futures contracts that are traded do not lead to the delivery of the underlying asset. Traders usually enter into offsetting contracts to close out their positions before the delivery period is reached. However, it is the possibility of final delivery that drives the determination of the futures price. For each futures contract, there is a range of days during which delivery can be made and a well-defined delivery procedure. Some contracts, such as those on stock indices, are settled in cash rather than by delivery of the underlying asset.

The specification of contracts is an important activity for a futures exchange. The two sides to any contract must know what can be delivered, where delivery can take place, and when delivery can take place. They also need to know details on the trading hours, how prices will be quoted, maximum daily price movements, and so on. New contracts must be approved by the Commodity Futures Trading Commission before trading starts.

Margins are an important aspect of futures markets. An investor keeps a margin account with his or her broker. The account is adjusted daily to reflect gains or losses, and from time to time the broker may require the account to be topped up if adverse price movements have taken place. The broker either must be a clearing house member or must maintain a margin account with a clearing house member. Each clearing house member maintains a margin account with the exchange clearing house. The balance in the account is adjusted daily to reflect gains and losses on the business for which the clearing house member is responsible.

Information on futures prices is collected in a systematic way at exchanges and relayed within a matter of seconds to investors throughout the world. Many daily newspapers such as the *Wall Street Journal* carry a summary of the previous day's trading.

Forward contracts differ from futures contracts in a number of ways. Forward contracts are private arrangements between two parties, whereas futures contracts are

traded on exchanges. There is generally a single delivery date in a forward contract, whereas futures contracts frequently involve a range of such dates. Because they are not traded on exchanges, forward contracts do not need to be standardized. A forward contract is not usually settled until the end of its life, and most contracts do in fact lead to delivery of the underlying asset or a cash settlement at this time.

In the next few chapters we shall examine in more detail the ways in which forward and futures contracts can be used for hedging. We shall also look at how forward and futures prices are determined.

FURTHER READING

- Duffie, D., and H. Zhu. "Does a Central Clearing Counterparty Reduce Counterparty Risk?" Working Paper, Stanford University, 2010.
- Gastineau, G. L., D. J. Smith, and R. Todd. *Risk Management, Derivatives, and Financial Analysis under SFAS No. 133*. The Research Foundation of AIMR and Blackwell Series in Finance, 2001.
- Hull, J., "OTC Derivatives and Central Clearing: Can All Transactions Be Cleared," *Financial Stability Review*, 14 (July 2010): 71–80.
- Jones, F. J., and R. J. Teweles. In: *The Futures Game*, edited by B. Warwick, 3rd edn. New York: McGraw-Hill, 1998.
- Jorion, P. "Risk Management Lessons from Long-Term Capital Management," *European Financial Management*, 6, 3 (September 2000): 277–300.
- Kawaller, I. G., and P. D. Koch. "Meeting the Highly Effective Expectation Criterion for Hedge Accounting," *Journal of Derivatives*, 7, 4 (Summer 2000): 79–87.
- Lowenstein, R. *When Genius Failed: The Rise and Fall of Long-Term Capital Management*. New York: Random House, 2000.

Practice Questions (Answers in Solutions Manual)

- 2.1. Distinguish between the terms *open interest* and *trading volume*.
- 2.2. What is the difference between a *local* and a *futures commission merchant*?
- 2.3. Suppose that you enter into a short futures contract to sell July silver for \$17.20 per ounce. The size of the contract is 5,000 ounces. The initial margin is \$4,000, and the maintenance margin is \$3,000. What change in the futures price will lead to a margin call? What happens if you do not meet the margin call?
- 2.4. Suppose that in September 2012 a company takes a long position in a contract on May 2013 crude oil futures. It closes out its position in March 2013. The futures price (per barrel) is \$68.30 when it enters into the contract, \$70.50 when it closes out its position, and \$69.10 at the end of December 2012. One contract is for the delivery of 1,000 barrels. What is the company's total profit? When is it realized? How is it taxed if it is (a) a hedger and (b) a speculator? Assume that the company has a December 31 year-end.
- 2.5. What does a stop order to sell at \$2 mean? When might it be used? What does a limit order to sell at \$2 mean? When might it be used?
- 2.6. What is the difference between the operation of the margin accounts administered by a clearing house and those administered by a broker?

- 2.7. What differences exist in the way prices are quoted in the foreign exchange futures market, the foreign exchange spot market, and the foreign exchange forward market?
- 2.8. The party with a short position in a futures contract sometimes has options as to the precise asset that will be delivered, where delivery will take place, when delivery will take place, and so on. Do these options increase or decrease the futures price? Explain your reasoning.
- 2.9. What are the most important aspects of the design of a new futures contract?
- 2.10. Explain how margins protect investors against the possibility of default.
- 2.11. A trader buys two July futures contracts on orange juice. Each contract is for the delivery of 15,000 pounds. The current futures price is 160 cents per pound, the initial margin is \$6,000 per contract, and the maintenance margin is \$4,500 per contract. What price change would lead to a margin call? Under what circumstances could \$2,000 be withdrawn from the margin account?
- 2.12. Show that, if the futures price of a commodity is greater than the spot price during the delivery period, then there is an arbitrage opportunity. Does an arbitrage opportunity exist if the futures price is less than the spot price? Explain your answer.
- 2.13. Explain the difference between a market-if-touched order and a stop order.
- 2.14. Explain what a stop–limit order to sell at 20.30 with a limit of 20.10 means.
- 2.15. At the end of one day a clearing house member is long 100 contracts, and the settlement price is \$50,000 per contract. The original margin is \$2,000 per contract. On the following day the member becomes responsible for clearing an additional 20 long contracts, entered into at a price of \$51,000 per contract. The settlement price at the end of this day is \$50,200. How much does the member have to add to its margin account with the exchange clearing house?
- 2.16. On July 1, 2012, a Japanese company enters into a forward contract to buy \$1 million with yen on January 1, 2013. On September 1, 2012, it enters into a forward contract to sell \$1 million on January 1, 2013. Describe the profit or loss the company will make in yen as a function of the forward exchange rates on July 1, 2012, and September 1, 2012.
- 2.17. The forward price of the Swiss franc for delivery in 45 days is quoted as 1.1000. The futures price for a contract that will be delivered in 45 days is 0.9000. Explain these two quotes. Which is more favorable for an investor wanting to sell Swiss francs?
- 2.18. Suppose you call your broker and issue instructions to sell one July hogs contract. Describe what happens.
- 2.19. “Speculation in futures markets is pure gambling. It is not in the public interest to allow speculators to trade on a futures exchange.” Discuss this viewpoint.
- 2.20. Live cattle futures trade with June, August, October, December, February, and April maturities. Why do you think the open interest for the June contract is less than that for the August contract in Table 2.2?
- 2.21. What do you think would happen if an exchange started trading a contract in which the quality of the underlying asset was incompletely specified?
- 2.22. “When a futures contract is traded on the floor of the exchange, it may be the case that the open interest increases by one, stays the same, or decreases by one.” Explain this statement.

- 2.23. Suppose that, on October 24, 2012, a company sells one April 2013 live cattle futures contract. It closes out its position on January 21, 2013. The futures price (per pound) is 91.20 cents when it enters into the contract, 88.30 cents when it closes out its position, and 88.80 cents at the end of December 2012. One contract is for the delivery of 40,000 pounds of cattle. What is the total profit? How is it taxed if the company is (a) a hedger and (b) a speculator? Assume that the company has a December 31 year-end.
- 2.24. A cattle farmer expects to have 120,000 pounds of live cattle to sell in 3 months. The live cattle futures contract traded by the CME Group is for the delivery of 40,000 pounds of cattle. How can the farmer use the contract for hedging? From the farmer's viewpoint, what are the pros and cons of hedging?
- 2.25. It is July 2011. A mining company has just discovered a small deposit of gold. It will take 6 months to construct the mine. The gold will then be extracted on a more or less continuous basis for 1 year. Futures contracts on gold are available with delivery months every 2 months from August 2011 to December 2012. Each contract is for the delivery of 100 ounces. Discuss how the mining company might use futures markets for hedging.

Further Questions

- 2.26. Trader A enters into futures contracts to buy 1 million euros for 1.4 million dollars in three months. Trader B enters in a forward contract to do the same thing. The exchange rate (dollars per euro) declines sharply during the first two months and then increases for the third month to close at 1.4300. Ignoring daily settlement, what is the total profit of each trader? When the impact of daily settlement is taken into account, which trader has done better?
- 2.27. Explain what is meant by open interest. Why does the open interest usually decline during the month preceding the delivery month? On a particular day, there were 2,000 trades in a particular futures contract. This means that there were 2,000 buyers (going long) and 2,000 sellers (going short). Of the 2,000 buyers, 1,400 were closing out positions and 600 were entering into new positions. Of the 2,000 sellers, 1,200 were closing out positions and 800 were entering into new positions. What is the impact of the day's trading on open interest?
- 2.28. One orange juice futures contract is on 15,000 pounds of frozen concentrate. Suppose that in September 2011 a company sells a March 2013 orange juice futures contract for 120 cents per pound. In December 2011, the futures price is 140 cents; in December 2012, it is 110 cents; and in February 2013, it is closed out at 125 cents. The company has a December year end. What is the company's profit or loss on the contract? How is it realized? What is the accounting and tax treatment of the transaction if the company is classified as (a) a hedger and (b) a speculator?
- 2.29. A company enters into a short futures contract to sell 5,000 bushels of wheat for 450 cents per bushel. The initial margin is \$3,000 and the maintenance margin is \$2,000. What price change would lead to a margin call? Under what circumstances could \$1,500 be withdrawn from the margin account?
- 2.30. Suppose that there are no storage costs for crude oil and the interest rate for borrowing or lending is 5% per annum. How could you make money on May 26, 2010, by trading July 2010 and December 2010 contracts? Use Table 2.2.

- 2.31. What position is equivalent to a long forward contract to buy an asset at K on a certain date and a put option to sell it for K on that date.
- 2.32. The author's website (www.rotman.utoronto.ca/~hull/data) contains daily closing prices for crude oil and gold futures contracts. You are required to download the data and answer the following:
- How high do the maintenance margin levels for oil and gold have to be set so that there is a 1% chance that an investor with a balance slightly above the maintenance margin level on a particular day has a negative balance 2 days later? How high do they have to be for a 0.1% chance? Assume daily price changes are normally distributed with mean zero. Explain why the exchange might be interested in this calculation.
 - Imagine an investor who starts with a long position in the oil contract at the beginning of the period covered by the data and keeps the contract for the whole of the period of time covered by the data. Margin balances in excess of the initial margin are withdrawn. Use the maintenance margin you calculated in part (a) for a 1% risk level and assume that the maintenance margin is 75% of the initial margin. Calculate the number of margin calls and the number of times the investor has a negative margin balance. Assume that all margin calls are met in your calculations. Repeat the calculations for an investor who starts with a short position in the gold contract.



3 CHAPTER

Hedging Strategies Using Futures

Many of the participants in futures markets are hedgers. Their aim is to use futures markets to reduce a particular risk that they face. This risk might relate to fluctuations in the price of oil, a foreign exchange rate, the level of the stock market, or some other variable. A *perfect hedge* is one that completely eliminates the risk. Perfect hedges are rare. For the most part, therefore, a study of hedging using futures contracts is a study of the ways in which hedges can be constructed so that they perform as close to perfect as possible.

In this chapter we consider a number of general issues associated with the way hedges are set up. When is a short futures position appropriate? When is a long futures position appropriate? Which futures contract should be used? What is the optimal size of the futures position for reducing risk? At this stage, we restrict our attention to what might be termed *hedge-and-forget* strategies. We assume that no attempt is made to adjust the hedge once it has been put in place. The hedger simply takes a futures position at the beginning of the life of the hedge and closes out the position at the end of the life of the hedge. In Chapter 18 we will examine dynamic hedging strategies in which the hedge is monitored closely and frequent adjustments are made.

The chapter initially treats futures contracts as forward contracts (that is, it ignores daily settlement). Later it explains an adjustment known as “tailing” that takes account of the difference between futures and forwards.

3.1 BASIC PRINCIPLES

When an individual or company chooses to use futures markets to hedge a risk, the objective is usually to take a position that neutralizes the risk as far as possible. Consider a company that knows it will gain \$10,000 for each 1 cent increase in the price of a commodity over the next 3 months and lose \$10,000 for each 1 cent decrease in the price during the same period. To hedge, the company’s treasurer should take a short futures position that is designed to offset this risk. The futures position should lead to a loss of \$10,000 for each 1 cent increase in the price of the commodity over the 3 months and a gain of \$10,000 for each 1 cent decrease in the price during this period. If the price of the commodity goes down, the gain on the futures position offsets the loss on the rest of the company’s business. If the price of the commodity

goes up, the loss on the futures position is offset by the gain on the rest of the company's business.

Short Hedges

A *short hedge* is a hedge, such as the one just described, that involves a short position in futures contracts. A short hedge is appropriate when the hedger already owns an asset and expects to sell it at some time in the future. For example, a short hedge could be used by a farmer who owns some hogs and knows that they will be ready for sale at the local market in two months. A short hedge can also be used when an asset is not owned right now but will be owned at some time in the future. Consider, for example, a US exporter who knows that he or she will receive euros in 3 months. The exporter will realize a gain if the euro increases in value relative to the US dollar and will sustain a loss if the euro decreases in value relative to the US dollar. A short futures position leads to a loss if the euro increases in value and a gain if it decreases in value. It has the effect of offsetting the exporter's risk.

To provide a more detailed illustration of the operation of a short hedge in a specific situation, we assume that it is May 15 today and that an oil producer has just negotiated a contract to sell 1 million barrels of crude oil. It has been agreed that the price that will apply in the contract is the market price on August 15. The oil producer is therefore in the position where it will gain \$10,000 for each 1 cent increase in the price of oil over the next 3 months and lose \$10,000 for each 1 cent decrease in the price during this period. Suppose that on May 15 the spot price is \$80 per barrel and the crude oil futures price for August delivery is \$79 per barrel. Because each futures contract is for the delivery of 1,000 barrels, the company can hedge its exposure by shorting (i.e., selling) 1,000 futures contracts. If the oil producer closes out its position on August 15, the effect of the strategy should be to lock in a price close to \$79 per barrel.

To illustrate what might happen, suppose that the spot price on August 15 proves to be \$75 per barrel. The company realizes \$75 million for the oil under its sales contract. Because August is the delivery month for the futures contract, the futures price on August 15 should be very close to the spot price of \$75 on that date. The company therefore gains approximately

$$\$79 - \$75 = \$4$$

per barrel, or \$4 million in total from the short futures position. The total amount realized from both the futures position and the sales contract is therefore approximately \$79 per barrel, or \$79 million in total.

For an alternative outcome, suppose that the price of oil on August 15 proves to be \$85 per barrel. The company realizes \$85 per barrel for the oil and loses approximately

$$\$85 - \$79 = \$6$$

per barrel on the short futures position. Again, the total amount realized is approximately \$79 million. It is easy to see that in all cases the company ends up with approximately \$79 million.

Long Hedges

Hedges that involve taking a long position in a futures contract are known as *long hedges*. A long hedge is appropriate when a company knows it will have to purchase a certain asset in the future and wants to lock in a price now.

Suppose that it is now January 15. A copper fabricator knows it will require 100,000 pounds of copper on May 15 to meet a certain contract. The spot price of copper is 340 cents per pound, and the futures price for May delivery is 320 cents per pound. The fabricator can hedge its position by taking a long position in four futures contracts offered by the COMEX division of the CME Group and closing its position on May 15. Each contract is for the delivery of 25,000 pounds of copper. The strategy has the effect of locking in the price of the required copper at close to 320 cents per pound.

Suppose that the spot price of copper on May 15 proves to be 325 cents per pound. Because May is the delivery month for the futures contract, this should be very close to the futures price. The fabricator therefore gains approximately

$$100,000 \times (\$3.25 - \$3.20) = \$5,000$$

on the futures contracts. It pays $100,000 \times \$3.25 = \$325,000$ for the copper, making the net cost approximately $\$325,000 - \$5,000 = \$320,000$. For an alternative outcome, suppose that the spot price is 305 cents per pound on May 15. The fabricator then loses approximately

$$100,000 \times (\$3.20 - \$3.05) = \$15,000$$

on the futures contract and pays $100,000 \times \$3.05 = \$305,000$ for the copper. Again, the net cost is approximately $\$320,000$, or 320 cents per pound.

Note that, in this case, it is clearly better for the company to use futures contracts than to buy the copper on January 15 in the spot market. If it does the latter, it will pay 340 cents per pound instead of 320 cents per pound and will incur both interest costs and storage costs. For a company using copper on a regular basis, this disadvantage would be offset by the convenience of having the copper on hand.¹ However, for a company that knows it will not require the copper until May 15, the futures contract alternative is likely to be preferred.

The examples we have looked at assume that the futures position is closed out in the delivery month. The hedge has the same basic effect if delivery is allowed to happen. However, making or taking delivery can be costly and inconvenient. For this reason, delivery is not usually made even when the hedger keeps the futures contract until the delivery month. As will be discussed later, hedgers with long positions usually avoid any possibility of having to take delivery by closing out their positions before the delivery period.

We have also assumed in the two examples that there is no daily settlement. In practice, daily settlement does have a small effect on the performance of a hedge. As explained in Chapter 2, it means that the payoff from the futures contract is realized day by day throughout the life of the hedge rather than all at the end.

3.2 ARGUMENTS FOR AND AGAINST HEDGING

The arguments in favor of hedging are so obvious that they hardly need to be stated. Most companies are in the business of manufacturing, or retailing or wholesaling, or providing a service. They have no particular skills or expertise in predicting variables such as interest rates, exchange rates, and commodity prices. It makes sense for them to

¹ See Section 5.11 for a discussion of convenience yields.

hedge the risks associated with these variables as they become aware of them. The companies can then focus on their main activities—for which presumably they do have particular skills and expertise. By hedging, they avoid unpleasant surprises such as sharp rises in the price of a commodity that is being purchased.

In practice, many risks are left unhedged. In the rest of this section we will explore some of the reasons for this.

Hedging and Shareholders

One argument sometimes put forward is that the shareholders can, if they wish, do the hedging themselves. They do not need the company to do it for them. This argument is, however, open to question. It assumes that shareholders have as much information as the company's management about the risks faced by a company. In most instances, this is not the case. The argument also ignores commissions and other transactions costs. These are less expensive per dollar of hedging for large transactions than for small transactions. Hedging is therefore likely to be less expensive when carried out by the company than when it is carried out by individual shareholders. Indeed, the size of futures contracts makes hedging by individual shareholders impossible in many situations.

One thing that shareholders can do far more easily than a corporation is diversify risk. A shareholder with a well-diversified portfolio may be immune to many of the risks faced by a corporation. For example, in addition to holding shares in a company that uses copper, a well-diversified shareholder may hold shares in a copper producer, so that there is very little overall exposure to the price of copper. If companies are acting in the best interests of well-diversified shareholders, it can be argued that hedging is unnecessary in many situations. However, the extent to which managers are in practice influenced by this type of argument is open to question.

Hedging and Competitors

If hedging is not the norm in a certain industry, it may not make sense for one particular company to choose to be different from all others. Competitive pressures within the industry may be such that the prices of the goods and services produced by the industry fluctuate to reflect raw material costs, interest rates, exchange rates, and so on. A company that does not hedge can expect its profit margins to be roughly constant. However, a company that does hedge can expect its profit margins to fluctuate!

To illustrate this point, consider two manufacturers of gold jewelry, SafeandSure Company and TakeaChance Company. We assume that most companies in the industry do not hedge against movements in the price of gold and that TakeaChance Company is no exception. However, SafeandSure Company has decided to be different from its competitors and to use futures contracts to hedge its purchase of gold over the next 18 months. If the price of gold goes up, economic pressures will tend to lead to a corresponding increase in the wholesale price of jewelry, so that TakeaChance Company's gross profit margin is unaffected. By contrast, SafeandSure Company's profit margin will increase after the effects of the hedge have been taken into account. If the price of gold goes down, economic pressures will tend to lead to a corresponding decrease in the wholesale price of jewelry. Again, TakeaChance Company's profit margin is unaffected. However, SafeandSure Company's profit margin goes down. In extreme conditions,

Table 3.1 Danger in hedging when competitors do not hedge.

| <i>Change in gold price</i> | <i>Effect on price of gold jewelry</i> | <i>Effect on profits of TakeaChance Co.</i> | <i>Effect on profits of SafeandSure Co.</i> |
|-----------------------------|--|---|---|
| Increase | Increase | None | Increase |
| Decrease | Decrease | None | Decrease |

SafeandSure Company's profit margin could become negative as a result of the "hedging" carried out! The situation is summarized in Table 3.1.

This example emphasizes the importance of looking at the big picture when hedging. All the implications of price changes on a company's profitability should be taken into account in the design of a hedging strategy to protect against the price changes.

Hedging Can Lead to a Worse Outcome

It is important to realize that a hedge using futures contracts can result in a decrease or an increase in a company's profits relative to the position it would be in with no hedging. In the example involving the oil producer considered earlier, if the price of oil goes down, the company loses money on its sale of 1 million barrels of oil, and the futures position leads to an offsetting gain. The treasurer can be congratulated for having had the foresight to put the hedge in place. Clearly, the company is better off than it would be with no hedging. Other executives in the organization, it is hoped, will appreciate the contribution made by the treasurer. If the price of oil goes up, the company gains from its sale of the oil, and the futures position leads to an offsetting loss. The company is in a worse position than it would be with no hedging. Although the hedging decision was perfectly logical, the treasurer may in practice have a difficult time justifying it. Suppose that the price of oil at the end of the hedge is \$89, so that the company loses \$10 per barrel on the futures contract. We can imagine a conversation such as the following between the treasurer and the president:

President: This is terrible. We've lost \$10 million in the futures market in the space of three months. How could it happen? I want a full explanation.

Treasurer: The purpose of the futures contracts was to hedge our exposure to the price of oil, not to make a profit. Don't forget we made \$10 million from the favorable effect of the oil price increases on our business.

President: What's that got to do with it? That's like saying that we do not need to worry when our sales are down in California because they are up in New York.

Treasurer: If the price of oil had gone down . . .

President: I don't care what would have happened if the price of oil had gone down. The fact is that it went up. I really do not know what you were doing playing the futures markets like this. Our shareholders will expect us to have done particularly well this quarter. I'm going to have to explain to them that your actions reduced profits by \$10 million. I'm afraid this is going to mean no bonus for you this year.

Business Snapshot 3.1 Hedging by Gold Mining Companies

It is natural for a gold mining company to consider hedging against changes in the price of gold. Typically it takes several years to extract all the gold from a mine. Once a gold mining company decides to go ahead with production at a particular mine, it has a big exposure to the price of gold. Indeed a mine that looks profitable at the outset could become unprofitable if the price of gold plunges.

Gold mining companies are careful to explain their hedging strategies to potential shareholders. Some gold mining companies do not hedge. They tend to attract shareholders who buy gold stocks because they want to benefit when the price of gold increases and are prepared to accept the risk of a loss from a decrease in the price of gold. Other companies choose to hedge. They estimate the number of ounces of gold they will produce each month for the next few years and enter into short futures or forward contracts to lock in the price for all or part of this.

Suppose you are Goldman Sachs and are approached by a gold mining company that wants to sell you a large amount of gold in 1 year at a fixed price. How do you set the price and then hedge your risk? The answer is that you can hedge by borrowing the gold from a central bank, selling it immediately in the spot market, and investing the proceeds at the risk-free rate. At the end of the year, you buy the gold from the gold mining company and use it to repay the central bank. The fixed forward price you set for the gold reflects the risk-free rate you can earn and the lease rate you pay the central bank for borrowing the gold.

Treasurer: That's unfair. I was only . . .

President: Unfair! You are lucky not to be fired. You lost \$10 million.

Treasurer: It all depends on how you look at it . . .

It is easy to see why many treasurers are reluctant to hedge! Hedging reduces risk for the company. However, it may increase risk for the treasurer if others do not fully understand what is being done. The only real solution to this problem involves ensuring that all senior executives within the organization fully understand the nature of hedging before a hedging program is put in place. Ideally, hedging strategies are set by a company's board of directors and are clearly communicated to both the company's management and the shareholders. (See Business Snapshot 3.1 for a discussion of hedging by gold mining companies.)

3.3 BASIS RISK

The hedges in the examples considered so far have been almost too good to be true. The hedger was able to identify the precise date in the future when an asset would be bought or sold. The hedger was then able to use futures contracts to remove almost all the risk arising from the price of the asset on that date. In practice, hedging is often not quite as straightforward as this. Some of the reasons are as follows:

1. The asset whose price is to be hedged may not be exactly the same as the asset underlying the futures contract.

2. The hedger may be uncertain as to the exact date when the asset will be bought or sold.
3. The hedge may require the futures contract to be closed out before its delivery month.

These problems give rise to what is termed *basis risk*. This concept will now be explained.

The Basis

The *basis* in a hedging situation is as follows:²

$$\text{Basis} = \text{Spot price of asset to be hedged} - \text{Futures price of contract used}$$

If the asset to be hedged and the asset underlying the futures contract are the same, the basis should be zero at the expiration of the futures contract. Prior to expiration, the basis may be positive or negative. From Table 2.2, we see that, on May 26, 2010, the basis was negative for gold and positive for short maturity contracts on soybeans.

As time passes, the spot price and the futures price for a particular month do not necessarily change by the same amount. As a result, the basis changes. An increase in the basis is referred to as a *strengthening of the basis*; a decrease in the basis is referred to as a *weakening of the basis*. Figure 3.1 illustrates how a basis might change over time in a situation where the basis is positive prior to expiration of the futures contract.

To examine the nature of basis risk, we will use the following notation:

S_1 : Spot price at time t_1

S_2 : Spot price at time t_2


F_1 : Futures price at time t_1

F_2 : Futures price at time t_2

b_1 : Basis at time t_1

b_2 : Basis at time t_2 .

Figure 3.1 Variation of basis over time.



² This is the usual definition. However, the alternative definition $\text{Basis} = \text{Futures price} - \text{Spot price}$ is sometimes used, particularly when the futures contract is on a financial asset.

We will assume that a hedge is put in place at time t_1 and closed out at time t_2 . As an example, we will consider the case where the spot and futures prices at the time the hedge is initiated are \$2.50 and \$2.20, respectively, and that at the time the hedge is closed out they are \$2.00 and \$1.90, respectively. This means that $S_1 = 2.50$, $F_1 = 2.20$, $S_2 = 2.00$, and $F_2 = 1.90$.

From the definition of the basis, we have

$$b_1 = S_1 - F_1 \quad \text{and} \quad b_2 = S_2 - F_2$$

so that, in our example, $b_1 = 0.30$ and $b_2 = 0.10$.

Consider first the situation of a hedger who knows that the asset will be sold at time t_2 and takes a short futures position at time t_1 . The price realized for the asset is S_2 and the profit on the futures position is $F_1 - F_2$. The effective price that is obtained for the asset with hedging is therefore

$$S_2 + F_1 - F_2 = F_1 + b_2$$

In our example, this is \$2.30. The value of F_1 is known at time t_1 . If b_2 were also known at this time, a perfect hedge would result. The hedging risk is the uncertainty associated with b_2 and is known as *basis risk*. Consider next a situation where a company knows it will buy the asset at time t_2 and initiates a long hedge at time t_1 . The price paid for the asset is S_2 and the loss on the hedge is $F_1 - F_2$. The effective price that is paid with hedging is therefore

$$S_2 + F_1 - F_2 = F_1 + b_2$$

This is the same expression as before and is \$2.30 in the example. The value of F_1 is known at time t_1 , and the term b_2 represents basis risk.

Note that basis risk can lead to an improvement or a worsening of a hedger's position. Consider a short hedge. If the basis strengthens (i.e., increases) unexpectedly, the hedger's position improves; if the basis weakens (i.e., decreases) unexpectedly, the hedger's position worsens. For a long hedge, the reverse holds. If the basis strengthens unexpectedly, the hedger's position worsens; if the basis weakens unexpectedly, the hedger's position improves.

The asset that gives rise to the hedger's exposure is sometimes different from the asset underlying the futures contract that is used for hedging. This is known as cross hedging and is discussed in the next section. It leads to an increase in basis risk. Define S_2^* as the price of the asset underlying the futures contract at time t_2 . As before, S_2 is the price of the asset being hedged at time t_2 . By hedging, a company ensures that the price that will be paid (or received) for the asset is

$$S_2 + F_1 - F_2$$

This can be written as

$$F_1 + (S_2^* - F_2) + (S_2 - S_2^*)$$

The terms $S_2^* - F_2$ and $S_2 - S_2^*$ represent the two components of the basis. The $S_2^* - F_2$ term is the basis that would exist if the asset being hedged were the same as the asset underlying the futures contract. The $S_2 - S_2^*$ term is the basis arising from the difference between the two assets.

Choice of Contract

One key factor affecting basis risk is the choice of the futures contract to be used for hedging. This choice has two components:

1. The choice of the asset underlying the futures contract
2. The choice of the delivery month.

If the asset being hedged exactly matches an asset underlying a futures contract, the first choice is generally fairly easy. In other circumstances, it is necessary to carry out a careful analysis to determine which of the available futures contracts has futures prices that are most closely correlated with the price of the asset being hedged.

The choice of the delivery month is likely to be influenced by several factors. In the examples given earlier in this chapter, we assumed that, when the expiration of the hedge corresponds to a delivery month, the contract with that delivery month is chosen. In fact, a contract with a later delivery month is usually chosen in these circumstances. The reason is that futures prices are in some instances quite erratic during the delivery month. Moreover, a long hedger runs the risk of having to take delivery of the physical asset if the contract is held during the delivery month. Taking delivery can be expensive and inconvenient. (Long hedgers normally prefer to close out the futures contract and buy the asset from their usual suppliers.)

In general, basis risk increases as the time difference between the hedge expiration and the delivery month increases. A good rule of thumb is therefore to choose a delivery month that is as close as possible to, but later than, the expiration of the hedge. Suppose delivery months are March, June, September, and December for a futures contract on a particular asset. For hedge expirations in December, January, and February, the March contract will be chosen; for hedge expirations in March, April, and May, the June contract will be chosen; and so on. This rule of thumb assumes that there is sufficient liquidity in all contracts to meet the hedger's requirements. In practice, liquidity tends to be greatest in short-maturity futures contracts. Therefore, in some situations, the hedger may be inclined to use short-maturity contracts and roll them forward. This strategy is discussed later in the chapter.

Example 3.1

It is March 1. A US company expects to receive 50 million Japanese yen at the end of July. Yen futures contracts on the CME Group have delivery months of March, June, September, and December. One contract is for the delivery of 12.5 million yen. The company therefore shorts four September yen futures contracts on March 1. When the yen are received at the end of July, the company closes out its position. We suppose that the futures price on March 1 in cents per yen is 0.7800 and that the spot and futures prices when the contract is closed out are 0.7200 and 0.7250, respectively.

The gain on the futures contract is $0.7800 - 0.7250 = 0.0550$ cents per yen. The basis is $0.7200 - 0.7250 = -0.0050$ cents per yen when the contract is closed out. The effective price obtained in cents per yen is the final spot price plus the gain on the futures:

$$0.7200 + 0.0550 = 0.7750$$

This can also be written as the initial futures price plus the final basis:

$$0.7800 + (-0.0050) = 0.7750$$

The total amount received by the company for the 50 million yen is 50×0.00775 million dollars, or \$387,500.

Example 3.2

It is June 8 and a company knows that it will need to purchase 20,000 barrels of crude oil at some time in October or November. Oil futures contracts are currently traded for delivery every month on the NYMEX division of the CME Group and the contract size is 1,000 barrels. The company therefore decides to use the December contract for hedging and takes a long position in 20 December contracts. The futures price on June 8 is \$68.00 per barrel. The company finds that it is ready to purchase the crude oil on November 10. It therefore closes out its futures contract on that date. The spot price and futures price on November 10 are \$70.00 per barrel and \$69.10 per barrel.

The gain on the futures contract is $69.10 - 68.00 = \$1.10$ per barrel. The basis when the contract is closed out is $70.00 - 69.10 = \$0.90$ per barrel. The effective price paid (in dollars per barrel) is the final spot price less the gain on the futures, or

$$70.00 - 1.10 = 68.90$$

This can also be calculated as the initial futures price plus the final basis,

$$68.00 + 0.90 = 68.90$$

The total price paid is $68.90 \times 20,000 = \$1,378,000$.

3.4 CROSS HEDGING

In Examples 3.1 and 3.2, the asset underlying the futures contract was the same as the asset whose price is being hedged. *Cross hedging* occurs when the two assets are different. Consider, for example, an airline that is concerned about the future price of jet fuel. Because jet fuel futures are not actively traded, it might choose to use heating oil futures contracts to hedge its exposure.

The *hedge ratio* is the ratio of the size of the position taken in futures contracts to the size of the exposure. When the asset underlying the futures contract is the same as the asset being hedged, it is natural to use a hedge ratio of 1.0. This is the hedge ratio we have used in the examples considered so far. For instance, in Example 3.2, the hedger's exposure was on 20,000 barrels of oil, and futures contracts were entered into for the delivery of exactly this amount of oil.

When cross hedging is used, setting the hedge ratio equal to 1.0 is not always optimal. The hedger should choose a value for the hedge ratio that minimizes the variance of the value of the hedged position. We now consider how the hedger can do this.

Calculating the Minimum Variance Hedge Ratio

The minimum variance hedge ratio depends on the relationship between changes in the spot price and changes in the futures price. Define:

ΔS : Change in spot price, S , during a period of time equal to the life of the hedge

ΔF : Change in futures price, F , during a period of time equal to the life of the hedge.

We will denote the minimum variance hedge ratio by h^* . It can be shown that h^* is the slope of the best-fit line from a linear regression of ΔS against ΔF (see Figure 3.2). This result is intuitively reasonable. We would expect h^* to be the ratio of the average change in S for a particular change in F .

The formula for h^* is:

$$h^* = \rho \frac{\sigma_S}{\sigma_F} \quad (3.1)$$

where σ_S is the standard deviation of ΔS , σ_F is the standard deviation of ΔF , and ρ is the coefficient of correlation between the two.

Equation (3.1) shows that the optimal hedge ratio is the product of the coefficient of correlation between ΔS and ΔF and the ratio of the standard deviation of ΔS to the standard deviation of ΔF . Figure 3.3 shows how the variance of the value of the hedger's position depends on the hedge ratio chosen.

If $\rho = 1$ and $\sigma_F = \sigma_S$, the hedge ratio, h^* , is 1.0. This result is to be expected, because in this case the futures price mirrors the spot price perfectly. If $\rho = 1$ and $\sigma_F = 2\sigma_S$, the

Figure 3.2 Regression of change in spot price against change in futures price.





Figure 3.3 Dependence of variance of hedger's position on hedge ratio.

hedge ratio h^* is 0.5. This result is also as expected, because in this case the futures price always changes by twice as much as the spot price. The *hedge effectiveness* can be defined as the proportion of the variance that is eliminated by hedging. This is the R^2 from the regression of ΔS against ΔF and equals ρ^2 .

The parameters ρ , σ_F , and σ_S in equation (3.1) are usually estimated from historical data on ΔS and ΔF . (The implicit assumption is that the future will in some sense be like the past.) A number of equal nonoverlapping time intervals are chosen, and the values of ΔS and ΔF for each of the intervals are observed. Ideally, the length of each time interval is the same as the length of the time interval for which the hedge is in effect. In practice, this sometimes severely limits the number of observations that are available, and a shorter time interval is used.

Optimal Number of Contracts

To calculate the number of contracts that should be used in hedging, define:

Q_A : Size of position being hedged (units)

Q_F : Size of one futures contract (units)

N^* : Optimal number of futures contracts for hedging.

The futures contracts should be on h^*Q_A units of the asset. The number of futures contracts required is therefore given by

$$N^* = \frac{h^* Q_A}{Q_F} \quad (3.2)$$

Example 3.3 will show how the results in this section can be used by an airline hedging the purchase of jet fuel.³

³ Heating oil futures are more actively traded than jet fuel futures contracts. For an account of how Delta Airlines used heating oil futures to hedge its future purchases of jet fuel, see A. Ness, "Delta Wins on Fuel," *Risk*, June 2001: 8.

Example 3.3

An airline expects to purchase 2 million gallons of jet fuel in 1 month and decides to use heating oil futures for hedging. We suppose that Table 3.2 gives, for 15 successive months, data on the change, ΔS , in the jet fuel price per gallon and the corresponding change, ΔF , in the futures price for the contract on heating oil that would be used for hedging price changes during the month. In this case, the usual formulas for calculating standard deviations and correlations give $\sigma_F = 0.0313$, $\sigma_S = 0.0263$, and $\rho = 0.928$.

From equation (3.1), the minimum variance hedge ratio, h^* , is therefore

$$0.928 \times \frac{0.0263}{0.0313} = 0.7777$$

Each heating oil contract traded on NYMEX is on 42,000 gallons of heating oil. From equation (3.2), the optimal number of contracts is

$$\frac{0.7777 \times 2,000,000}{42,000} = 37.03$$

or, rounding to the nearest whole number, 37.

Table 3.2 Data to calculate minimum variance hedge ratio when heating oil futures contract is used to hedge purchase of jet fuel.

| <i>Month</i> <i>i</i> | <i>Change in</i> <i>heating oil futures</i> <i>price per gallon</i> <i>(= ΔF)</i> | <i>Change in</i> <i>jet fuel</i> <i>price per gallon</i> <i>(= ΔS)</i> |
|--------------------------|--|---|
| 1 | 0.021 | 0.029 |
| 2 | 0.035 | 0.020 |
| 3 | -0.046 | -0.044 |
| 4 | 0.001 | 0.008 |
| 5 | 0.044 | 0.026 |
| 6 | -0.029 | -0.019 |
| 7 | -0.026 | -0.010 |
| 8 | -0.029 | -0.007 |
| 9 | 0.048 | 0.043 |
| 10 | -0.006 | 0.011 |
| 11 | -0.036 | -0.036 |
| 12 | -0.011 | -0.018 |
| 13 | 0.019 | 0.009 |
| 14 | -0.027 | -0.032 |
| 15 | 0.029 | 0.023 |

Tailing the Hedge

When futures are used for hedging, a small adjustment, known as *tailing the hedge*, can be made to allow for the impact of daily settlement. In practice this means that equation (3.2) becomes⁴

$$N^* = \frac{h^* V_A}{V_F} \quad (3.3)$$

where V_A is the dollar value of the position being hedged and V_F is the dollar value of one futures contract (the futures price times Q_F). Suppose that in Example 3.3 the spot price and the futures price are 1.94 and 1.99 dollars per gallon, respectively. Then $V_A = 2,000,000 \times 1.94 = 3,880,000$ while $V_F = 42,000 \times 1.99 = 83,580$, so that the optimal number of contracts is

$$\frac{0.7777 \times 3,880,000}{83,580} = 36.10$$

If we round this to the nearest whole number, the optimal number of contracts is now 36 rather than 37. The effect of tailing the hedge is to multiply the hedge ratio in equation (3.2) by the ratio of the spot price to the futures price. Theoretically, the futures position used for hedging should then be adjusted as the spot price and futures price change, but in practice this usually makes very little difference.

If forward contracts rather than futures contracts are used, there is no daily settlement and equation (3.2) should be used.

3.5 STOCK INDEX FUTURES

We now move on to consider stock index futures and how they are used to hedge or manage exposures to equity prices.

A *stock index* tracks changes in the value of a hypothetical portfolio of stocks. The weight of a stock in the portfolio at a particular time equals the proportion of the hypothetical portfolio invested in the stock at that time. The percentage increase in the stock index over a small interval of time is set equal to the percentage increase in the value of the hypothetical portfolio. Dividends are usually not included in the calculation so that the index tracks the capital gain/loss from investing in the portfolio.⁵

If the hypothetical portfolio of stocks remains fixed, the weights assigned to individual stocks in the portfolio do not remain fixed. When the price of one particular stock in the portfolio rises more sharply than others, more weight is automatically given to that stock. Sometimes indices are constructed from a hypothetical portfolio consisting of one of each of a number of stocks. The weights assigned to the stocks are then proportional to their market prices, with adjustments being made when there are stock splits. Other indices are constructed so that weights are proportional to market capitalization (stock price \times number of shares outstanding). The underlying portfolio is then automatically adjusted to reflect stock splits, stock dividends, and new equity issues.

⁴ See Problem 5.23 for an explanation of equation (3.3).

⁵ An exception to this is a *total return index*. This is calculated by assuming that dividends on the hypothetical portfolio are reinvested in the portfolio.

Stock Indices

Table 3.3 shows futures prices for contracts on a number of different stock indices on May 26, 2010.

The *Dow Jones Industrial Average* is based on a portfolio consisting of 30 blue-chip stocks in the United States. The weights given to the stocks are proportional to their prices. The CME Group trades two futures contracts on the index. The one shown is on \$10 times the index. The other (the Mini DJ Industrial Average) is on \$5 times the index.

The *Standard & Poor's 500 (S&P 500) Index* is based on a portfolio of 500 different stocks: 400 industrials, 40 utilities, 20 transportation companies, and 40 financial institutions. The weights of the stocks in the portfolio at any given time are proportional to their market capitalizations. The stocks are those of large publicly held companies that trade on NYSE Euronext or Nasdaq OMX. The CME Group trades two futures contracts on the S&P 500. The one shown is on \$250 times the index; the other (the Mini S&P 500 contract) is on \$50 times the index.

The *Nasdaq-100* is based on 100 stocks using the National Association of Securities Dealers Automatic Quotations Service. The CME Group trades two futures contracts. The one shown is on \$100 times the index; the other (the Mini Nasdaq-100 contract) is on \$20 times the index.

The *Russell 1000 Index* is an index of the prices of the 1,000 largest capitalization stocks in the United States. The *US Dollar Index* is a trade-weighted index of the values of six foreign currencies (the euro, yen, pound, Canadian dollar, Swedish krona, and Swiss franc).

As mentioned in Chapter 2, futures contracts on stock indices are settled in cash, not by delivery of the underlying asset. All contracts are marked to market to either the opening price or the closing price of the index on the last trading day, and the positions are then deemed to be closed. For example, contracts on the S&P 500 are closed out at the opening price of the S&P 500 index on the third Friday of the delivery month.

Table 3.3 Index futures quotes as reported by exchanges on May 26, 2010.

| | <i>Open</i> | <i>High</i> | <i>Low</i> | <i>Settlement</i> | <i>Change</i> | <i>Volume</i> | <i>Open interest</i> |
|---|-------------|-------------|------------|-------------------|---------------|---------------|----------------------|
| Dow Jones Industrial Average, \$10 times index (CME Group) | | | | | | | |
| June 2010 | 10080 | 10160 | 9921 | 9921 | −104 | 533 | 12,353 |
| Sept. 2010 | 10080 | 10085 | 9864 | 9864 | −104 | 2 | 225 |
| S&P 500, \$250 times index (CME Group) | | | | | | | |
| June 2010 | 1080.0 | 1089.5 | 1060.5 | 1061.2 | −11.8 | 43,076 | 296,397 |
| Sept. 2010 | 1084.0 | 1085.5 | 1057.5 | 1057.0 | −11.9 | 6,850 | 26,966 |
| Dec. 2010 | 1074.0 | 1081.9 | 1052.9 | 1053.1 | −11.8 | 7 | 4,326 |
| Nasdaq-100, \$100 times index (CME Group) | | | | | | | |
| June 2010 | 1827.0 | 1850.0 | 1788.0 | 1791.5 | −24.0 | 2,350 | 20,674 |
| Russell 1000, \$100 times index (ICE) | | | | | | | |
| June 2010 | 595.4 | 601.0 | 586.4 | 585.7 | −5.8 | 1,214 | 19,275 |
| US Dollar Index, \$1000 times index (ICE) | | | | | | | |
| June 2010 | 86.725 | 87.480 | 86.625 | 87.244 | 0.356 | 37,321 | 35,401 |
| Sept. 2010 | 87.270 | 87.800 | 87.050 | 87.584 | 0.356 | 74 | 2,533 |

Hedging an Equity Portfolio

Stock index futures can be used to hedge a well-diversified equity portfolio. Define:

V_A : Current value of the portfolio

V_F : Current value of one futures contract (the futures price times the contract size).

If the portfolio mirrors the index, the optimal hedge ratio, h^* , equals 1.0 and equation (3.3) shows that the number of futures contracts that should be shorted is

$$N^* = \frac{V_A}{V_F} \quad (3.4)$$

Suppose, for example, that a portfolio worth \$5,050,000 mirrors the S&P 500. The index futures price is 1,010 and each futures contract is on \$250 times the index. In this case $V_A = 5,050,000$ and $V_F = 1,010 \times 250 = 252,500$, so that 20 contracts should be shorted to hedge the portfolio.

When the portfolio does not exactly mirror the index, we can use the capital asset pricing model (see the appendix to this chapter). The parameter beta (β) from the capital asset pricing model is the slope of the best-fit line obtained when excess return on the portfolio over the risk-free rate is regressed against the excess return of the index over the risk-free rate. When $\beta = 1.0$, the return on the portfolio tends to mirror the return on the index; when $\beta = 2.0$, the excess return on the portfolio tends to be twice as great as the excess return on the index; when $\beta = 0.5$, it tends to be half as great; and so on.

A portfolio with a β of 2.0 is twice as sensitive to movements in the index as a portfolio with a beta 1.0. It is therefore necessary to use twice as many contracts to hedge the portfolio. Similarly, a portfolio with a beta of 0.5 is half as sensitive to market movements as a portfolio with a beta of 1.0 and we should use half as many contracts to hedge it. In general,

$$N^* = \beta \frac{V_A}{V_F} \quad (3.5)$$

This formula assumes that the maturity of the futures contract is close to the maturity of the hedge.

Comparing equation (3.5) with equation (3.3), we see that they imply $h^* = \beta$. This is not surprising. The hedge ratio h^* is the slope of the best-fit line when changes in the portfolio are regressed against changes in the futures price of the index. Beta (β) is the slope of the best-fit line when the return from the portfolio is regressed against the return for the index.

We illustrate that this formula gives good results by extending our earlier example. Suppose that a futures contract with 4 months to maturity is used to hedge the value of a portfolio over the next 3 months in the following situation:

Value of S&P 500 index = 1,000

S&P 500 futures price = 1,010

Value of portfolio = \$5,050,000

Risk-free interest rate = 4% per annum

Dividend yield on index = 1% per annum

Beta of portfolio = 1.5.

One futures contract is for delivery of \$250 times the index. As before, $V_F = 250 \times 1,010 = 252,500$. From equation (3.5), the number of futures contracts that should be shorted to hedge the portfolio is

$$1.5 \times \frac{5,050,000}{252,500} = 30$$

Suppose the index turns out to be 900 in 3 months and the futures price is 902. The gain from the short futures position is then

$$30 \times (1010 - 902) \times 250 = \$810,000$$

The loss on the index is 10%. The index pays a dividend of 1% per annum, or 0.25% per 3 months. When dividends are taken into account, an investor in the index would therefore earn -9.75% over the 3-month period. Because the portfolio has a β of 1.5, the capital asset pricing model gives

$$\begin{aligned} \text{Expected return on portfolio} - \text{Risk-free interest rate} \\ = 1.5 \times (\text{Return on index} - \text{Risk-free interest rate}) \end{aligned}$$

The risk-free interest rate is approximately 1% per 3 months. It follows that the expected return (%) on the portfolio during the 3 months when the 3-month return on the index is -9.75% is

$$1.0 + [1.5 \times (-9.75 - 1.0)] = -15.125$$

The expected value of the portfolio (inclusive of dividends) at the end of the 3 months is therefore

$$\$5,050,000 \times (1 - 0.15125) = \$4,286,187$$

It follows that the expected value of the hedger's position, including the gain on the hedge, is

$$\$4,286,187 + \$810,000 = \$5,096,187$$

Table 3.4 summarizes these calculations together with similar calculations for other values of the index at maturity. It can be seen that the total expected value of the hedger's position in 3 months is almost independent of the value of the index.

The only thing we have not covered in this example is the relationship between futures prices and spot prices. We will see in Chapter 5 that the 1,010 assumed for the futures price today is roughly what we would expect given the interest rate and dividend we are assuming. The same is true of the futures prices in 3 months shown in Table 3.4.⁶

Reasons for Hedging an Equity Portfolio

Table 3.4 shows that the hedging procedure results in a value for the hedger's position at the end of the 3-month period being about 1% higher than at the beginning of the 3-month period. There is no surprise here. The risk-free rate is 4% per annum, or 1% per 3 months. The hedge results in the investor's position growing at the risk-free rate.

⁶ The calculations in Table 3.4 assume that the dividend yield on the index is predictable, the risk-free interest rate remains constant, and the return on the index over the 3-month period is perfectly correlated with the return on the portfolio. In practice, these assumptions do not hold perfectly, and the hedge works rather less well than is indicated by Table 3.4.

Table 3.4 Performance of stock index hedge.

| | | | | | |
|---|-----------|-----------|-----------|-----------|-----------|
| Value of index in three months: | 900 | 950 | 1,000 | 1,050 | 1,100 |
| Futures price of index today: | 1,010 | 1,010 | 1,010 | 1,010 | 1,010 |
| Futures price of index in three months: | 902 | 952 | 1,003 | 1,053 | 1,103 |
| Gain on futures position (\$): | 810,000 | 435,000 | 52,500 | -322,500 | -697,500 |
| Return on market: | -9.750% | -4.750% | 0.250% | 5.250% | 10.250% |
| Expected return on portfolio: | -15.125% | -7.625% | -0.125% | 7.375% | 14.875% |
| Expected portfolio value in three months including dividends (\$): | 4,286,187 | 4,664,937 | 5,043,687 | 5,422,437 | 5,801,187 |
| Total value of position in three months (\$): | 5,096,187 | 5,099,937 | 5,096,187 | 5,099,937 | 5,103,687 |

It is natural to ask why the hedger should go to the trouble of using futures contracts. To earn the risk-free interest rate, the hedger can simply sell the portfolio and invest the proceeds in risk-free instruments such as Treasury bills.

One answer to this question is that hedging can be justified if the hedger feels that the stocks in the portfolio have been chosen well. In these circumstances, the hedger might be very uncertain about the performance of the market as a whole, but confident that the stocks in the portfolio will outperform the market (after appropriate adjustments have been made for the beta of the portfolio). A hedge using index futures removes the risk arising from market moves and leaves the hedger exposed only to the performance of the portfolio relative to the market. This will be discussed further shortly. Another reason for hedging may be that the hedger is planning to hold a portfolio for a long period of time and requires short-term protection in an uncertain market situation. The alternative strategy of selling the portfolio and buying it back later might involve unacceptably high transaction costs.

Changing the Beta of a Portfolio

In the example in Table 3.4, the beta of the hedger's portfolio is reduced to zero so that the hedger's expected return is almost independent of the performance of the index. Sometimes futures contracts are used to change the beta of a portfolio to some value other than zero. Continuing with our earlier example:

$$\begin{aligned} \text{S\&P 500 index} &= 1,000 \\ \text{S\&P 500 futures price} &= 1,010 \\ \text{Value of portfolio} &= \$5,050,000 \\ \text{Beta of portfolio} &= 1.5 \end{aligned}$$

As before, $V_F = 250 \times 1,010 = 252,500$ and a complete hedge requires

$$1.5 \times \frac{5,050,000}{252,500} = 30$$

contracts to be shorted. To reduce the beta of the portfolio from 1.5 to 0.75, the

number of contracts shorted should be 15 rather than 30; to increase the beta of the portfolio to 2.0, a long position in 10 contracts should be taken; and so on. In general, to change the beta of the portfolio from β to β^* , where $\beta > \beta^*$, a short position in

$$(\beta - \beta^*) \frac{V_A}{V_F}$$

contracts is required. When $\beta < \beta^*$, a long position in

$$(\beta^* - \beta) \frac{V_A}{V_F}$$

contracts is required.

Locking in the Benefits of Stock Picking

Suppose you consider yourself to be good at picking stocks that will outperform the market. You own a single stock or a small portfolio of stocks. You do not know how well the market will perform over the next few months, but you are confident that your portfolio will do better than the market. What should you do?

You should short $\beta V_A / V_F$ index futures contracts, where β is the beta of your portfolio, V_A is the total value of the portfolio, and V_F is the current value of one index futures contract. If your portfolio performs better than a well-diversified portfolio with the same beta, you will then make money.

Consider an investor who in April holds 20,000 IBM shares, each worth \$100. The investor feels that the market will be very volatile over the next three months but that IBM has a good chance of outperforming the market. The investor decides to use the August futures contract on the S&P 500 to hedge the market's return during the three-month period. The β of IBM is estimated at 1.1. Suppose that the current futures price for the August contract on the S&P 500 is 900. Each contract is for delivery of \$250 times the index. In this case, $V_A = 20,000 \times 100 = 2,000,000$ and $V_F = 900 \times 250 = 225,000$. The number of contracts that should be shorted is therefore

$$1.1 \times \frac{2,000,000}{225,000} = 9.78$$

Rounding to the nearest integer, the investor shorts 10 contracts, closing out the position in July. Suppose IBM falls to \$90 and the futures price of the S&P 500 falls to 750. The investor loses $20,000 \times (\$100 - \$90) = \$200,000$ on IBM, while gaining $10 \times 250 \times (900 - 750) = \$375,000$ on the futures contracts.

The overall gain to the investor in this case is \$175,000 because IBM did not go down by as much as a well-diversified portfolio with a β of 1.1. If the market had gone up and IBM went up by more than a portfolio with a β of 1.1 (as expected by the investor), then a profit would be made in this case as well.

3.6 STACK AND ROLL

Sometimes the expiration date of the hedge is later than the delivery dates of all the futures contracts that can be used. The hedger must then roll the hedge forward by closing out one futures contract and taking the same position in a futures contract with

Table 3.5 Data for the example on rolling oil hedge forward.

| <i>Date</i> | <i>Apr. 2011</i> | <i>Sept. 2011</i> | <i>Feb. 2012</i> | <i>June 2012</i> |
|-------------------------|------------------|-------------------|------------------|------------------|
| Oct. 2011 futures price | 68.20 | 67.40 | | |
| Mar. 2012 futures price | | 67.00 | 66.50 | |
| July 2012 futures price | | | 66.30 | 65.90 |
| Spot price | 69.00 | | | 66.00 |

a later delivery date. Hedges can be rolled forward many times. The procedure is known as *stack and roll*. Consider a company that wishes to use a short hedge to reduce the risk associated with the price to be received for an asset at time T . If there are futures contracts 1, 2, 3, ..., n (not all necessarily in existence at the present time) with progressively later delivery dates, the company can use the following strategy:

- Time t_1 : Short futures contract 1
- Time t_2 : Close out futures contract 1
Short futures contract 2
- Time t_3 : Close out futures contract 2
Short futures contract 3
- ⋮
- Time t_n : Close out futures contract $n - 1$
Short futures contract n
- Time T : Close out futures contract n .

Suppose that in April 2011 a company realizes that it will have 100,000 barrels of oil to sell in June 2012 and decides to hedge its risk with a hedge ratio of 1.0. (In this example, we do not make the “tailing” adjustment described in Section 3.4.) The current spot price is \$69. Although futures contracts are traded with maturities stretching several years into the future, we suppose that only the first six delivery months have sufficient liquidity to meet the company’s needs. The company therefore shorts 100 October 2011 contracts. In September 2011, it rolls the hedge forward into the March 2012 contract. In February 2012, it rolls the hedge forward again into the July 2012 contract.

One possible outcome is shown in Table 3.5. The October 2011 contract is shorted at \$68.20 per barrel and closed out at \$67.40 per barrel for a profit of \$0.80 per barrel; the March 2012 contract is shorted at \$67.00 per barrel and closed out at \$66.50 per barrel for a profit of \$0.50 per barrel. The July 2012 contract is shorted at \$66.30 per barrel and closed out at \$65.90 per barrel for a profit of \$0.40 per barrel. The final spot price is \$66.

The dollar gain per barrel of oil from the short futures contracts is

$$(68.20 - 67.40) + (67.00 - 66.50) + (66.30 - 65.90) = 1.70$$

The oil price declined from \$69 to \$66. Receiving only \$1.70 per barrel compensation for a price decline of \$3.00 may appear unsatisfactory. However, we cannot expect total compensation for a price decline when futures prices are below spot prices. The best we

Business Snapshot 3.2 Metallgesellschaft: Hedging Gone Awry

Sometimes rolling hedges forward can lead to cash flow pressures. The problem was illustrated dramatically by the activities of a German company, Metallgesellschaft (MG), in the early 1990s.

MG sold a huge volume of 5- to 10-year heating oil and gasoline fixed-price supply contracts to its customers at 6 to 8 cents above market prices. It hedged its exposure with long positions in short-dated futures contracts that were rolled forward. As it turned out, the price of oil fell and there were margin calls on the futures positions. Considerable short-term cash flow pressures were placed on MG. The members of MG who devised the hedging strategy argued that these short-term cash outflows were offset by positive cash flows that would ultimately be realized on the long-term fixed-price contracts. However, the company's senior management and its bankers became concerned about the huge cash drain. As a result, the company closed out all the hedge positions and agreed with its customers that the fixed-price contracts would be abandoned. The outcome was a loss to MG of \$1.33 billion.

can hope for is to lock in the futures price that would apply to a June 2012 contract if it were actively traded.

In practice, a company usually has an exposure every month to the underlying asset and uses a 1-month futures contract for hedging because it is the most liquid. Initially it enters into ("stacks") sufficient contracts to cover its exposure to the end of its hedging horizon. One month later, it closes out all the contracts and "rolls" them into new 1-month contracts to cover its new exposure, and so on.

As described in Business Snapshot 3.2, a German company, Metallgesellschaft, followed this strategy in the early 1990s to hedge contracts it had entered into to supply commodities at a fixed price. It ran into difficulties because the prices of the commodities declined so that there were immediate cash outflows on the futures and the expectation of eventual gains on the contracts. This mismatch between the timing of the cash flows on hedge and the timing of the cash flows from the position being hedged led to liquidity problems that could not be handled. The moral of the story is that potential liquidity problems should always be considered when a hedging strategy is being planned.

SUMMARY

This chapter has discussed various ways in which a company can take a position in futures contracts to offset an exposure to the price of an asset. If the exposure is such that the company gains when the price of the asset increases and loses when the price of the asset decreases, a short hedge is appropriate. If the exposure is the other way round (i.e., the company gains when the price of the asset decreases and loses when the price of the asset increases), a long hedge is appropriate.

Hedging is a way of reducing risk. As such, it should be welcomed by most executives. In reality, there are a number of theoretical and practical reasons why companies do not hedge. On a theoretical level, we can argue that shareholders, by holding well-diversified portfolios, can eliminate many of the risks faced by a company.

They do not require the company to hedge these risks. On a practical level, a company may find that it is increasing rather than decreasing risk by hedging if none of its competitors does so. Also, a treasurer may fear criticism from other executives if the company makes a gain from movements in the price of the underlying asset and a loss on the hedge.

An important concept in hedging is basis risk. The basis is the difference between the spot price of an asset and its futures price. Basis risk arises from uncertainty as to what the basis will be at maturity of the hedge.

The hedge ratio is the ratio of the size of the position taken in futures contracts to the size of the exposure. It is not always optimal to use a hedge ratio of 1.0. If the hedger wishes to minimize the variance of a position, a hedge ratio different from 1.0 may be appropriate. The optimal hedge ratio is the slope of the best-fit line obtained when changes in the spot price are regressed against changes in the futures price.

Stock index futures can be used to hedge the systematic risk in an equity portfolio. The number of futures contracts required is the beta of the portfolio multiplied by the ratio of the value of the portfolio to the value of one futures contract. Stock index futures can also be used to change the beta of a portfolio without changing the stocks that make up the portfolio.

When there is no liquid futures contract that matures later than the expiration of the hedge, a strategy known as stack and roll may be appropriate. This involves entering into a sequence of futures contracts. When the first futures contract is near expiration, it is closed out and the hedger enters into a second contract with a later delivery month. When the second contract is close to expiration, it is closed out and the hedger enters into a third contract with a later delivery month; and so on. The result of all this is the creation of a long-dated futures contract by trading a series of short-dated contracts.

FURTHER READING

- Allayannis, G., and J. Weston. "The Use of Foreign Currency Derivatives and Firm Market Value," *Review of Financial Studies*, 14, 1 (Spring 2001): 243–76.
- Bodnar, G. M., G. S. Hayt, and R. C. Marston. "1998 Wharton Survey of Financial Risk Management by U.S. Non-Financial Firms," *Financial Management*, 2, 4 (1998): 70–91.
- Brown, G. W. "Managing Foreign Exchange Risk with Derivatives," *Journal of Financial Economics*, 60 (2001): 401–48.
- Culp, C. and M. H. Miller. "Metallgesellschaft and the Economics of Synthetic Storage," *Journal of Applied Corporate Finance*, 7, 4 (Winter 1995): 62–76.
- Campbell, J. Y., K. Serfaty-de Medeiros, and L. M. Viceira. "Global Currency Hedging," *Journal of Finance*, 65, 1 (February 2010): 87–121.
- Cotter, J., and J. Hanly. "Hedging: Scaling and the Investor Horizon," *Journal of Risk*, 12, 2 (Winter 2009/2010): 49–77.
- Ederington, L. H. "The Hedging Performance of the New Futures Market," *Journal of Finance*, 34 (March 1979): 157–70.
- Edwards, F. R. and M. S. Canter. "The Collapse of Metallgesellschaft: Unhedgeable Risks, Poor Hedging Strategy, or Just Bad Luck?" *Journal of Applied Corporate Finance*, 8, 1 (Spring 1995): 86–105.
- Geczy, C., B. A. Minton, and C. Schrand. "Why Firms Use Currency Derivatives," *Journal of Finance*, 52, 4 (1997): 1323–54.

- Graham, J. R. and C. W. Smith, Jr. "Tax Incentives to Hedge," *Journal of Finance*, 54, 6 (1999): 2241–62.
- Haushalter, G. D. "Financing Policy, Basis Risk, and Corporate Hedging: Evidence from Oil and Gas Producers," *Journal of Finance*, 55, 1 (2000): 107–52.
- Mello, A. S. and J. E. Parsons. "Hedging and Liquidity," *Review of Financial Studies*, 13 (Spring 2000): 127–53.
- Neuberger, A. J. "Hedging Long-Term Exposures with Multiple Short-Term Futures Contracts," *Review of Financial Studies*, 12 (1999): 429–59.
- Petersen, M. A. and S. R. Thiagarajan, "Risk Management and Hedging: With and Without Derivatives," *Financial Management*, 29, 4 (Winter 2000): 5–30.
- Rendleman, R. "A Reconciliation of Potentially Conflicting Approaches to Hedging with Futures," *Advances in Futures and Options*, 6 (1993): 81–92.
- Stulz, R. M. "Optimal Hedging Policies," *Journal of Financial and Quantitative Analysis*, 19 (June 1984): 127–40.
- Tufano, P. "Who Manages Risk? An Empirical Examination of Risk Management Practices in the Gold Mining Industry," *Journal of Finance*, 51, 4 (1996): 1097–1138.
- Tufano, P. "The Determinants of Stock Price Exposure: Financial Engineering and the Gold Mining Industry," *Journal of Finance*, 53, 3 (1998): 1015–52.

Practice Questions (Answers in Solutions Manual)

- 3.1. Under what circumstances are (a) a short hedge and (b) a long hedge appropriate?
- 3.2. Explain what is meant by *basis risk* when futures contracts are used for hedging.
- 3.3. Explain what is meant by a *perfect hedge*. Does a perfect hedge always lead to a better outcome than an imperfect hedge? Explain your answer.
- 3.4. Under what circumstances does a minimum variance hedge portfolio lead to no hedging at all?
- 3.5. Give three reasons why the treasurer of a company might not hedge the company's exposure to a particular risk.
- 3.6. Suppose that the standard deviation of quarterly changes in the prices of a commodity is \$0.65, the standard deviation of quarterly changes in a futures price on the commodity is \$0.81, and the coefficient of correlation between the two changes is 0.8. What is the optimal hedge ratio for a 3-month contract? What does it mean?
- 3.7. A company has a \$20 million portfolio with a beta of 1.2. It would like to use futures contracts on the S&P 500 to hedge its risk. The index futures price is currently standing at 1080, and each contract is for delivery of \$250 times the index. What is the hedge that minimizes risk? What should the company do if it wants to reduce the beta of the portfolio to 0.6?
- 3.8. In the corn futures contract traded on an exchange, the following delivery months are available: March, May, July, September, and December. Which of the available contracts should be used for hedging when the expiration of the hedge is in (a) June, (b) July, and (c) January.
- 3.9. Does a perfect hedge always succeed in locking in the current spot price of an asset for a future transaction? Explain your answer.

- 3.10. Explain why a short hedger's position improves when the basis strengthens unexpectedly and worsens when the basis weakens unexpectedly.
- 3.11. Imagine you are the treasurer of a Japanese company exporting electronic equipment to the United States. Discuss how you would design a foreign exchange hedging strategy and the arguments you would use to sell the strategy to your fellow executives.
- 3.12. Suppose that in Example 3.2 of Section 3.3 the company decides to use a hedge ratio of 0.8. How does the decision affect the way in which the hedge is implemented and the result?
- 3.13. "If the minimum variance hedge ratio is calculated as 1.0, the hedge must be perfect." Is this statement true? Explain your answer.
- 3.14. "If there is no basis risk, the minimum variance hedge ratio is always 1.0." Is this statement true? Explain your answer.
- 3.15. "For an asset where futures prices are usually less than spot prices, long hedges are likely to be particularly attractive." Explain this statement.
- 3.16. The standard deviation of monthly changes in the spot price of live cattle is (in cents per pound) 1.2. The standard deviation of monthly changes in the futures price of live cattle for the closest contract is 1.4. The correlation between the futures price changes and the spot price changes is 0.7. It is now October 15. A beef producer is committed to purchasing 200,000 pounds of live cattle on November 15. The producer wants to use the December live cattle futures contracts to hedge its risk. Each contract is for the delivery of 40,000 pounds of cattle. What strategy should the beef producer follow?
- 3.17. A corn farmer argues "I do not use futures contracts for hedging. My real risk is not the price of corn. It is that my whole crop gets wiped out by the weather." Discuss this viewpoint. Should the farmer estimate his or her expected production of corn and hedge to try to lock in a price for expected production?
- 3.18. On July 1, an investor holds 50,000 shares of a certain stock. The market price is \$30 per share. The investor is interested in hedging against movements in the market over the next month and decides to use the September Mini S&P 500 futures contract. The index futures price is currently 1,500 and one contract is for delivery of \$50 times the index. The beta of the stock is 1.3. What strategy should the investor follow? Under what circumstances will it be profitable?
- 3.19. Suppose that in Table 3.5 the company decides to use a hedge ratio of 1.5. How does the decision affect the way the hedge is implemented and the result?
- 3.20. A futures contract is used for hedging. Explain why the daily settlement of the contract can give rise to cash-flow problems.
- 3.21. An airline executive has argued: "There is no point in our using oil futures. There is just as much chance that the price of oil in the future will be less than the futures price as there is that it will be greater than this price." Discuss the executive's viewpoint.
- 3.22. Suppose that the 1-year gold lease rate is 1.5% and the 1-year risk-free rate is 5.0%. Both rates are compounded annually. Use the discussion in Business Snapshot 3.1 to calculate the maximum 1-year gold forward price Goldman Sachs should quote to the gold-mining company when the spot price is \$1,200.
- 3.23. The expected return on the S&P 500 is 12% and the risk-free rate is 5%. What is the expected return on an investment with a beta of (a) 0.2, (b) 0.5, and (c) 1.4?

Further Questions

- 3.24. A company wishes to hedge its exposure to a new fuel whose price changes have a 0.6 correlation with gasoline futures price changes. The company will lose \$1 million for each 1 cent increase in the price per gallon of the new fuel over the next three months. The new fuel's price changes have a standard deviation that is 50% greater than price changes in gasoline futures prices. If gasoline futures are used to hedge the exposure, what should the hedge ratio be? What is the company's exposure measured in gallons of the new fuel? What position, measured in gallons, should the company take in gasoline futures? How many gasoline futures contracts should be traded? Each contract is on 42,000 gallons.
- 3.25. A portfolio manager has maintained an actively managed portfolio with a beta of 0.2. During the last year, the risk-free rate was 5% and equities performed very badly providing a return of -30% . The portfolio manager produced a return of -10% and claims that in the circumstances it was a good performance. Discuss this claim.
- 3.26. The following table gives data on monthly changes in the spot price and the futures price for a certain commodity. Use the data to calculate a minimum variance hedge ratio.

| | | | | | |
|----------------------|-------|-------|-------|-------|-------|
| Spot price change | +0.50 | +0.61 | -0.22 | -0.35 | +0.79 |
| Futures price change | +0.56 | +0.63 | -0.12 | -0.44 | +0.60 |
| Spot price change | +0.04 | +0.15 | +0.70 | -0.51 | -0.41 |
| Futures price change | -0.06 | +0.01 | +0.80 | -0.56 | -0.46 |

- 3.27. It is July 16. A company has a portfolio of stocks worth \$100 million. The beta of the portfolio is 1.2. The company would like to use the CME December futures contract on the S&P 500 to change the beta of the portfolio to 0.5 during the period July 16 to November 16. The index futures price is currently 1,000 and each contract is on \$250 times the index.
- What position should the company take?
 - Suppose that the company changes its mind and decides to increase the beta of the portfolio from 1.2 to 1.5. What position in futures contracts should it take?
- 3.28. A fund manager has a portfolio worth \$50 million with a beta of 0.87. The manager is concerned about the performance of the market over the next 2 months and plans to use 3-month futures contracts on the S&P 500 to hedge the risk. The current level of the index is 1,250, one contract is on 250 times the index, the risk-free rate is 6% per annum, and the dividend yield on the index is 3% per annum. The current 3-month futures price is 1,259.
- What position should the fund manager take to hedge all exposure to the market over the next 2 months?
 - Calculate the effect of your strategy on the fund manager's returns if the index in 2 months is 1,000, 1,100, 1,200, 1,300, and 1,400. Assume that the 1-month futures price is 0.25% higher than the index level at this time.
- 3.29. It is now October 2010. A company anticipates that it will purchase 1 million pounds of copper in each of February 2011, August 2011, February 2012, and August 2012. The company has decided to use the futures contracts traded in the COMEX division of the CME Group to hedge its risk. One contract is for the delivery of 25,000 pounds of copper. The initial margin is \$2,000 per contract and the maintenance margin is \$1,500 per contract. The company's policy is to hedge 80% of its exposure. Contracts with maturities up to 13 months into the future are considered to have sufficient liquidity to meet the

company's needs. Devise a hedging strategy for the company. (Do not make the "tailing" adjustment described in Section 3.4.)

Assume the market prices (in cents per pound) today and at future dates are as in the following table. What is the impact of the strategy you propose on the price the company pays for copper? What is the initial margin requirement in October 2010? Is the company subject to any margin calls?

| <i>Date</i> | <i>Oct. 2010</i> | <i>Feb. 2011</i> | <i>Aug. 2011</i> | <i>Feb. 2012</i> | <i>Aug. 2012</i> |
|--------------------------|------------------|------------------|------------------|------------------|------------------|
| Spot price | 372.00 | 369.00 | 365.00 | 377.00 | 388.00 |
| Mar. 2011 futures price | 372.30 | 369.10 | | | |
| Sept. 2011 futures price | 372.80 | 370.20 | 364.80 | | |
| Mar. 2012 futures price | | 370.70 | 364.30 | 376.70 | |
| Sept. 2012 futures price | | | 364.20 | 376.50 | 388.20 |

APPENDIX

CAPITAL ASSET PRICING MODEL

The capital asset pricing model (CAPM) is a model that can be used to calculate the expected return from an asset during a period in terms of the risk of the return. The risk in the return from an asset is divided into two parts. *Systematic risk* is risk related to the return from the market as a whole and cannot be diversified away. *Nonsystematic risk* is risk that is unique to the asset and can be diversified away by choosing a large portfolio of different assets. CAPM argues that the return should depend only on systematic risk. The CAPM formula is⁷

$$\text{Expected return on asset} = R_F + \beta(R_M - R_F) \quad (3A.1)$$

where R_M is the return on the portfolio of all available investments, R_F is the return on a risk-free investment, and β (the Greek letter beta) is a parameter measuring systematic risk.

The return from the portfolio of all available investments, R_M , is referred to as the *return on the market* and is usually approximated as the return on a well-diversified stock index such as the S&P 500. The beta (β) of an asset is a measure of the sensitivity of its returns to returns from the market. It can be estimated from historical data as the slope obtained when the excess return on the asset over the risk-free rate is regressed against the excess return on the market over the risk-free rate. When $\beta = 0$, an asset's returns are not sensitive to returns from the market. In this case, it has no systematic risk and equation (3A.1) shows that its expected return is the risk-free rate; when $\beta = 0.5$, the excess return on the asset over the risk-free rate is on average half of the excess return of the market over the risk-free rate; when $\beta = 1$, the expected return on the asset equals to the return on the market; and so on.

Suppose that the risk-free rate R_F is 5% and the return on the market is 13%. Equation (3A.1) shows that, when the beta of an asset is zero, its expected return is 5%. When $\beta = 0.75$, its expected return is $0.05 + 0.75 \times (0.13 - 0.05) = 0.11$, or 11%.

The derivation of CAPM requires a number of assumptions.⁸ In particular:

1. Investors care only about the expected return and standard deviation of the return from an asset.
2. The returns from two assets are correlated with each other only because of their correlation with the return from the market. This is equivalent to assuming that there is only one factor driving returns.
3. Investors focus on returns over a single period and that period is the the same for all investors.
4. Investors can borrow and lend at the same risk-free rate.
5. Tax does not influence investment decisions.
6. All investors make the same estimates of expected returns, standard deviations of returns, and correlations between returns.

⁷ If the return on the market is not known, R_M is replaced by the expected value of R_M in this formula.

⁸ For details on the derivation, see, for example, J. Hull, *Risk Management and Financial Institutions*, 2nd edn. Upper Saddle River, NJ: Pearson, 2010, Chap. 1.

These assumptions are at best only approximately true. Nevertheless CAPM has proved to be a useful tool for portfolio managers and is often used as a benchmark for assessing their performance.

When the asset is an individual stock, the expected return given by equation (3A.1) is not a particularly good predictor of the actual return. But, when the asset is a well-diversified portfolio of stocks, it is a much better predictor. As a result, the equation

$$\text{Return on diversified portfolio} = R_F + \beta(R_M - R_F)$$

can be used as a basis for hedging a diversified portfolio, as described in Section 3.5. The β in the equation is the beta of the portfolio. It can be calculated as the weighted average of the betas of the stocks in the portfolio.



4 CHAPTER

Interest Rates

Interest rates are a factor in the valuation of virtually all derivatives and will feature prominently in much of the material that will be presented in the rest of this book. This chapter deals with some fundamental issues concerned with the way interest rates are measured and analyzed. It explains the compounding frequency used to define an interest rate and the meaning of continuously compounded interest rates, which are used extensively in the analysis of derivatives. It covers zero rates, par yields, and yield curves, discusses bond pricing, and outlines a “bootstrap” procedure commonly used by a derivatives trading desk to calculate zero-coupon Treasury interest rates. It also covers forward rates and forward rate agreements and reviews different theories of the term structure of interest rates. Finally, it explains the use of duration and convexity measures to determine the sensitivity of bond prices to interest rate changes.

Chapter 6 will cover interest rate futures and show how the duration measure can be used when interest rate exposures are hedged. For ease of exposition, day count conventions will be ignored throughout this chapter. The nature of these conventions and their impact on calculations will be discussed in Chapters 6 and 7.

4.1 TYPES OF RATES

An interest rate in a particular situation defines the amount of money a borrower promises to pay the lender. For any given currency, many different types of interest rates are regularly quoted. These include mortgage rates, deposit rates, prime borrowing rates, and so on. The interest rate applicable in a situation depends on the credit risk. This is the risk that there will be a default by the borrower of funds, so that the interest and principal are not paid to the lender as promised. The higher the credit risk, the higher the interest rate that is promised by the borrower.

Treasury Rates

Treasury rates are the rates an investor earns on Treasury bills and Treasury bonds. These are the instruments used by a government to borrow in its own currency. Japanese Treasury rates are the rates at which the Japanese government borrows in yen; US Treasury rates are the rates at which the US government borrows in US dollars; and so on. It is usually assumed that there is no chance that a government will

default on an obligation denominated in its own currency. Treasury rates are therefore totally risk-free rates in the sense that an investor who buys a Treasury bill or Treasury bond is certain that interest and principal payments will be made as promised.

LIBOR

LIBOR is short for *London Interbank Offered Rate*. It is a reference interest rate, produced once a day by the British Bankers' Association, and is designed to reflect the rate of interest at which banks are prepared to make large wholesale deposits with other banks. LIBOR is quoted in all major currencies for maturities up to 12 months: 1-month LIBOR is the rate at which 1-month deposits are offered, 3-month LIBOR is the rate at which 3-month deposits are offered, and so on.

A deposit with a bank can be regarded as a loan to that bank. A bank must therefore satisfy certain creditworthiness criteria in order to be able to receive deposits from another bank at LIBOR. Typically it must have a AA credit rating.¹

A rate closely related to LIBOR is LIBID. This is the *London Interbank Bid Rate* and is the rate at which banks will accept deposits from other banks. At any specified time, there is a small spread between LIBID and LIBOR rates (with LIBOR higher than LIBID). The rates themselves are determined by active trading between banks and adjust so that the supply of funds in the interbank market equals the demand for funds in that market. For example, if more banks want to borrow US dollars for 3 months than lend US dollars for 3 months, the 3-month US LIBID and LIBOR rates will increase. Similarly, if more banks want to lend 3-month funds than borrow these funds, the 3-month LIBID and LIBOR rates will decrease. This interbank market is known as the *Eurocurrency market*. It is outside the control of any one government.

Repo Rates

Sometimes trading activities are funded with a *repurchase agreement*, *repo*. This is a contract where an investment dealer who owns securities agrees to sell them to another company now and buy them back later at a slightly higher price. The other company is providing a loan to the investment dealer. The difference between the price at which the securities are sold and the price at which they are repurchased is the interest it earns. The interest rate is referred to as the *repo rate*. If structured carefully, the loan involves very little credit risk. If the borrower does not honor the agreement, the lending company simply keeps the securities. If the lending company does not keep to its side of the agreement, the original owner of the securities keeps the cash. The most common type of repo is an *overnight repo*, in which the agreement is renegotiated each day. However, longer-term arrangements, known as *term repos*, are sometimes used.

The Risk-Free Rate

The “risk-free rate” is used extensively in the evaluation of derivatives. It might be thought that derivatives traders would use the interest rates implied by Treasury bills

¹ As explained in Chapter 23, the best credit rating given to a company by rating agencies S&P and Fitch is AAA. The second best is AA. The corresponding ratings from Moody's are Aaa and Aa, respectively.

Business Snapshot 4.1 What Is the Risk-Free Rate?

Derivatives dealers argue that the interest rates implied by Treasury bills and Treasury bonds are artificially low because:

1. Treasury bills and Treasury bonds must be purchased by financial institutions to fulfill a variety of regulatory requirements. This increases demand for these Treasury instruments driving the price up and the yield down.
2. The amount of capital a bank is required to hold to support an investment in Treasury bills and bonds is substantially smaller than the capital required to support a similar investment in other instruments with very low risk.
3. In the United States, Treasury instruments are given a favorable tax treatment compared with most other fixed-income investments because they are not taxed at the state level.

Traditionally derivatives dealers have assumed that LIBOR rates are risk-free and, for ease of exposition, this what we will do in this book. But LIBOR rates are not totally risk-free. Following the credit crisis that started in 2007, many dealers switched to using overnight indexed swap rates as risk-free rates. These rates are discussed in Section 7.8.

and bonds as risk-free rates. In fact, they do not do this. As indicated in Business Snapshot 4.1, there are a number of tax and regulatory issues that cause Treasury rates to be artificially low.

Financial institutions have traditionally used LIBOR rates as risk-free rates. For a AA-rated financial institution LIBOR is the short-term opportunity cost of capital. The financial institution can borrow short-term funds at the LIBOR quotes of other financial institutions and can lend funds to other financial institutions at its own LIBOR quotes. LIBOR rates are not totally free of credit risk. For example, when funds are lent at 3-month LIBOR to a AA-rated financial institution, there is a small chance that it will default during the 3 months. However, they are close to risk-free in normal market conditions. LIBOR rates have maturities up to 1 year. As we will explain later, traders have traditionally used Eurodollar futures and interest rate swaps to extend the risk-free LIBOR yield curve beyond 1 year.

The credit crisis that started in 2007 caused many derivatives dealers to critically review their practices. This is because banks became very reluctant to lend to each other during the crisis and LIBOR rates soared. Many dealers have now switched to using the overnight indexed swap (OIS) rate as a proxy for the risk-free rate. This rate will be explained in Section 7.8. It is closer to risk-free than LIBOR.

4.2 MEASURING INTEREST RATES

A statement by a bank that the interest rate on one-year deposits is 10% per annum sounds straightforward and unambiguous. In fact, its precise meaning depends on the way the interest rate is measured.

Table 4.1 Effect of the compounding frequency on the value of \$100 at the end of 1 year when the interest rate is 10% per annum.

| <i>Compounding frequency</i> | <i>Value of \$100 at end of year (\$)</i> |
|------------------------------|---|
| Annually ($m = 1$) | 110.00 |
| Semiannually ($m = 2$) | 110.25 |
| Quarterly ($m = 4$) | 110.38 |
| Monthly ($m = 12$) | 110.47 |
| Weekly ($m = 52$) | 110.51 |
| Daily ($m = 365$) | 110.52 |

If the interest rate is measured with annual compounding, the bank's statement that the interest rate is 10% means that \$100 grows to

$$\$100 \times 1.1 = \$110$$

at the end of 1 year. When the interest rate is measured with semiannual compounding, it means that 5% is earned every 6 months, with the interest being reinvested. In this case \$100 grows to

$$\$100 \times 1.05 \times 1.05 = \$110.25$$

at the end of 1 year. When the interest rate is measured with quarterly compounding, the bank's statement means that 2.5% is earned every 3 months, with the interest being reinvested. The \$100 then grows to

$$\$100 \times 1.025^4 = \$110.38$$

at the end of 1 year. Table 4.1 shows the effect of increasing the compounding frequency further.

The compounding frequency defines the units in which an interest rate is measured. A rate expressed with one compounding frequency can be converted into an equivalent rate with a different compounding frequency. For example, from Table 4.1 we see that 10.25% with annual compounding is equivalent to 10% with semiannual compounding. We can think of the difference between one compounding frequency and another to be analogous to the difference between kilometers and miles. They are two different units of measurement.

To generalize our results, suppose that an amount A is invested for n years at an interest rate of R per annum. If the rate is compounded once per annum, the terminal value of the investment is

$$A(1 + R)^n$$

If the rate is compounded m times per annum, the terminal value of the investment is

$$A \left(1 + \frac{R}{m} \right)^{mn} \quad (4.1)$$

When $m = 1$, the rate is sometimes referred to as the *equivalent annual interest rate*.

Continuous Compounding

The limit as the compounding frequency, m , tends to infinity is known as *continuous compounding*.² With continuous compounding, it can be shown that an amount A invested for n years at rate R grows to

$$Ae^{Rn} \quad (4.2)$$

where $e = 2.71828$. The exponential function, e^x , is built into most calculators, so the computation of the expression in equation (4.2) presents no problems. In the example in Table 4.1, $A = 100$, $n = 1$, and $R = 0.1$, so that the value to which A grows with continuous compounding is

$$100e^{0.1} = \$110.52$$

This is (to two decimal places) the same as the value with daily compounding. For most practical purposes, continuous compounding can be thought of as being equivalent to daily compounding. Compounding a sum of money at a continuously compounded rate R for n years involves multiplying it by e^{Rn} . Discounting it at a continuously compounded rate R for n years involves multiplying by e^{-Rn} .

In this book, interest rates will be measured with continuous compounding except where stated otherwise. Readers used to working with interest rates that are measured with annual, semiannual, or some other compounding frequency may find this a little strange at first. However, continuously compounded interest rates are used to such a great extent in pricing derivatives that it makes sense to get used to working with them now.

Suppose that R_c is a rate of interest with continuous compounding and R_m is the equivalent rate with compounding m times per annum. From the results in equations (4.1) and (4.2), we have

$$Ae^{R_c n} = A \left(1 + \frac{R_m}{m} \right)^{mn}$$

or

$$e^{R_c} = \left(1 + \frac{R_m}{m} \right)^m$$

This means that

$$R_c = m \ln \left(1 + \frac{R_m}{m} \right) \quad (4.3)$$

and

$$R_m = m(e^{R_c/m} - 1) \quad (4.4)$$

These equations can be used to convert a rate with a compounding frequency of m times per annum to a continuously compounded rate and vice versa. The natural logarithm function $\ln x$, which is built into most calculators, is the *inverse* of the exponential function, so that, if $y = \ln x$, then $x = e^y$.

Example 4.1

Consider an interest rate that is quoted as 10% per annum with semiannual compounding. From equation (4.3) with $m = 2$ and $R_m = 0.1$, the equivalent rate

² Actuaries sometimes refer to a continuously compounded rate as the *force of interest*.

with continuous compounding is

$$2 \ln \left(1 + \frac{0.1}{2} \right) = 0.09758$$

or 9.758% per annum.

Example 4.2

Suppose that a lender quotes the interest rate on loans as 8% per annum with continuous compounding, and that interest is actually paid quarterly. From equation (4.4) with $m = 4$ and $R_c = 0.08$, the equivalent rate with quarterly compounding is

$$4 \times (e^{0.08/4} - 1) = 0.0808$$

or 8.08% per annum. This means that on a \$1,000 loan, interest payments of \$20.20 would be required each quarter.

4.3 ZERO RATES

The n -year zero-coupon interest rate is the rate of interest earned on an investment that starts today and lasts for n years. All the interest and principal is realized at the end of n years. There are no intermediate payments. The n -year zero-coupon interest rate is sometimes also referred to as the n -year *spot rate*, the n -year *zero rate*, or just the n -year zero. Suppose a 5-year zero rate with continuous compounding is quoted as 5% per annum. This means that \$100, if invested for 5 years, grows to

$$100 \times e^{0.05 \times 5} = 128.40$$

Most of the interest rates we observe directly in the market are not pure zero rates. Consider a 5-year government bond that provides a 6% coupon. The price of this bond does not by itself determine the 5-year Treasury zero rate because some of the return on the bond is realized in the form of coupons prior to the end of year 5. Later in this chapter we will discuss how we can determine Treasury zero rates from the market prices of coupon-bearing bonds.

4.4 BOND PRICING

Most bonds pay coupons to the holder periodically. The bond's principal (which is also known as its par value or face value) is paid at the end of its life. The theoretical price of a bond can be calculated as the present value of all the cash flows that will be received by the owner of the bond. Sometimes bond traders use the same discount rate for all the cash flows underlying a bond, but a more accurate approach is to use a different zero rate for each cash flow.

To illustrate this, consider the situation where Treasury zero rates, measured with continuous compounding, are as in Table 4.2. (We explain later how these can be calculated.) Suppose that a 2-year Treasury bond with a principal of \$100 provides coupons at the rate of 6% per annum semiannually. To calculate the present value of the first coupon of \$3, we discount it at 5.0% for 6 months; to calculate the present

Table 4.2 Treasury zero rates.

| <i>Maturity (years)</i> | <i>Zero rate (%) (continuously compounded)</i> |
|-------------------------|--|
| 0.5 | 5.0 |
| 1.0 | 5.8 |
| 1.5 | 6.4 |
| 2.0 | 6.8 |

value of the second coupon of \$3, we discount it at 5.8% for 1 year; and so on. Therefore the theoretical price of the bond is

$$3e^{-0.05 \times 0.5} + 3e^{-0.058 \times 1.0} + 3e^{-0.064 \times 1.5} + 103e^{-0.068 \times 2.0} = 98.39$$

or \$98.39.

Bond Yield

A bond's yield is the single discount rate that, when applied to all cash flows, gives a bond price equal to its market price. Suppose that the theoretical price of the bond we have been considering, \$98.39, is also its market value (i.e., the market's price of the bond is in exact agreement with the data in Table 4.2). If y is the yield on the bond, expressed with continuous compounding, it must be true that

$$3e^{-y \times 0.5} + 3e^{-y \times 1.0} + 3e^{-y \times 1.5} + 103e^{-y \times 2.0} = 98.39$$

This equation can be solved using an iterative ("trial and error") procedure to give $y = 6.76\%$.³

Par Yield

The *par yield* for a certain bond maturity is the coupon rate that causes the bond price to equal its par value. (The par value is the same as the principal value.) Usually the bond is assumed to provide semiannual coupons. Suppose that the coupon on a 2-year bond in our example is c per annum (or $\frac{1}{2}c$ per 6 months). Using the zero rates in Table 4.2, the value of the bond is equal to its par value of 100 when

$$\frac{c}{2}e^{-0.05 \times 0.5} + \frac{c}{2}e^{-0.058 \times 1.0} + \frac{c}{2}e^{-0.064 \times 1.5} + \left(100 + \frac{c}{2}\right)e^{-0.068 \times 2.0} = 100$$

This equation can be solved in a straightforward way to give $c = 6.87$. The 2-year par yield is therefore 6.87% per annum. This has semiannual compounding because payments are assumed to be made every 6 months. With continuous compounding, the rate is 6.75% per annum.

More generally, if d is the present value of \$1 received at the maturity of the bond, A is the value of an annuity that pays one dollar on each coupon payment date, and m

³ One way of solving nonlinear equations of the form $f(y) = 0$, such as this one, is to use the Newton–Raphson method. We start with an estimate y_0 of the solution and produce successively better estimates y_1, y_2, y_3, \dots using the formula $y_{i+1} = y_i - f(y_i)/f'(y_i)$, where $f'(y)$ denotes the derivative of f with respect to y .

is the number of coupon payments per year, then the par yield c must satisfy

$$100 = A \frac{c}{m} + 100d$$

so that

$$c = \frac{(100 - 100d)m}{A}$$

In our example, $m = 2$, $d = e^{-0.068 \times 2} = 0.87284$, and

$$A = e^{-0.05 \times 0.5} + e^{-0.058 \times 1.0} + e^{-0.064 \times 1.5} + e^{-0.068 \times 2.0} = 3.70027$$

The formula confirms that the par yield is 6.87% per annum.

4.5 DETERMINING TREASURY ZERO RATES

One way of determining Treasury zero rates such as those in Table 4.2 is to observe the yields on “strips.” These are zero-coupon bonds that are synthetically created by traders when they sell coupons on a Treasury bond separately from the principal.

Another way to determine Treasury zero rates is from Treasury bills and coupon-bearing bonds. The most popular approach is known as the *bootstrap method*. To illustrate the nature of the method, consider the data in Table 4.3 on the prices of five bonds. Because the first three bonds pay no coupons, the zero rates corresponding to the maturities of these bonds can easily be calculated. The 3-month bond has the effect of turning an investment of 97.5 into 100 in 3 months. The continuously compounded 3-month rate R is therefore given by solving

$$100 = 97.5e^{R \times 0.25}$$

It is 10.127% per annum. The 6-month continuously compounded rate is similarly given by solving

$$100 = 94.9e^{R \times 0.5}$$

It is 10.469% per annum. Similarly, the 1-year rate with continuous compounding is given by solving

$$100 = 90e^{R \times 1.0}$$

It is 10.536% per annum.

Table 4.3 Data for bootstrap method.

| Bond principal (\$) | Time to maturity (years) | Annual coupon* (\$) | Bond price (\$) |
|------------------------|-----------------------------|------------------------|--------------------|
| 100 | 0.25 | 0 | 97.5 |
| 100 | 0.50 | 0 | 94.9 |
| 100 | 1.00 | 0 | 90.0 |
| 100 | 1.50 | 8 | 96.0 |
| 100 | 2.00 | 12 | 101.6 |

* Half the stated coupon is assumed to be paid every 6 months.

The fourth bond lasts 1.5 years. The payments are as follows:

6 months: \$4

1 year: \$4

1.5 years: \$104.

From our earlier calculations, we know that the discount rate for the payment at the end of 6 months is 10.469% and that the discount rate for the payment at the end of 1 year is 10.536%. We also know that the bond's price, \$96, must equal the present value of all the payments received by the bondholder. Suppose the 1.5-year zero rate is denoted by R . It follows that

$$4e^{-0.10469 \times 0.5} + 4e^{-0.10536 \times 1.0} + 104e^{-R \times 1.5} = 96$$

This reduces to

$$e^{-1.5R} = 0.85196$$

or

$$R = -\frac{\ln(0.85196)}{1.5} = 0.10681$$

The 1.5-year zero rate is therefore 10.681%. This is the only zero rate that is consistent with the 6-month rate, 1-year rate, and the data in Table 4.3.

The 2-year zero rate can be calculated similarly from the 6-month, 1-year, and 1.5-year zero rates, and the information on the last bond in Table 4.3. If R is the 2-year zero rate, then

$$6e^{-0.10469 \times 0.5} + 6e^{-0.10536 \times 1.0} + 6e^{-0.10681 \times 1.5} + 106e^{-R \times 2.0} = 101.6$$


This gives $R = 0.10808$, or 10.808%.

The rates we have calculated are summarized in Table 4.4. A chart showing the zero rate as a function of maturity is known as the *zero curve*. A common assumption is that the zero curve is linear between the points determined using the bootstrap method. (This means that the 1.25-year zero rate is $0.5 \times 10.536 + 0.5 \times 10.681 = 10.6085\%$ in our example.) It is also usually assumed that the zero curve is horizontal prior to the first point and horizontal beyond the last point. Figure 4.1 shows the zero curve for our data using these assumptions. By using longer maturity bonds, the zero curve would be more accurately determined beyond 2 years.

In practice, we do not usually have bonds with maturities equal to exactly 1.5 years, 2 years, 2.5 years, and so on. The approach often used by analysts is to interpolate

Table 4.4 Continuously compounded zero rates determined from data in Table 4.3.

| <i>Maturity</i> (years) | <i>Zero rate (%)</i> (continuously compounded) |
|----------------------------|---|
| 0.25 | 10.127 |
| 0.50 | 10.469 |
| 1.00 | 10.536 |
| 1.50 | 10.681 |
| 2.00 | 10.808 |

Figure 4.1 Zero rates given by the bootstrap method.

between the bond price data before it is used to calculate the zero curve. For example, if they know that a 2.3-year bond with a coupon of 6% sells for 98 and a 2.7-year bond with a coupon of 6.5% sells for 99, it might be assumed that a 2.5-year bond with a coupon of 6.25% would sell for 98.5.

4.6 FORWARD RATES

Forward interest rates are the rates of interest implied by current zero rates for periods of time in the future. To illustrate how they are calculated, we suppose that LIBOR zero rates are as shown in the second column of Table 4.5. (As we shall see in Chapter 7, LIBOR zero rates are calculated in a similar way to the Treasury zero rates calculated in the previous section.) The rates are assumed to be continuously compounded. Thus, the 3% per annum rate for 1 year means that, in return for an investment of \$100 today, an amount $100e^{0.03 \times 1} = \$103.05$ is received in 1 year; the 4% per annum rate for 2 years means that, in return for an investment of \$100 today, an amount $100e^{0.04 \times 2} = \$108.33$ is received in 2 years; and so on.

The forward interest rate in Table 4.5 for year 2 is 5% per annum. This is the rate of interest that is implied by the zero rates for the period of time between the end of the first year and the end of the second year. It can be calculated from the 1-year zero interest rate of 3% per annum and the 2-year zero interest rate of 4% per annum. It is the rate of interest for year 2 that, when combined with 3% per annum for year 1, gives 4% overall for the 2 years. To show that the correct answer is 5% per annum, suppose that \$100 is invested. A rate of 3% for the first year and 5% for the second year gives

$$100e^{0.03 \times 1} e^{0.05 \times 1} = \$108.33$$

at the end of the second year. A rate of 4% per annum for 2 years gives

$$100e^{0.04 \times 2}$$

which is also \$108.33. This example illustrates the general result that when interest rates are continuously compounded and rates in successive time periods are combined, the

overall equivalent rate is simply the average rate during the whole period. In our example, 3% for the first year and 5% for the second year average to 4% over the 2 years. The result is only approximately true when the rates are not continuously compounded.

The forward rate for year 3 is the rate of interest that is implied by a 4% per annum 2-year zero rate and a 4.6% per annum 3-year zero rate. It is 5.8% per annum. The reason is that an investment for 2 years at 4% per annum combined with an investment for one year at 5.8% per annum gives an overall average return for the three years of 4.6% per annum. The other forward rates can be calculated similarly and are shown in the third column of the table. In general, if R_1 and R_2 are the zero rates for maturities T_1 and T_2 , respectively, and R_F is the forward interest rate for the period of time between T_1 and T_2 , then

$$R_F = \frac{R_2 T_2 - R_1 T_1}{T_2 - T_1} \quad (4.5)$$

To illustrate this formula, consider the calculation of the year-4 forward rate from the data in Table 4.5: $T_1 = 3$, $T_2 = 4$, $R_1 = 0.046$, and $R_2 = 0.05$, and the formula gives $R_F = 0.062$.

Equation (4.5) can be written as

$$R_F = R_2 + (R_2 - R_1) \frac{T_1}{T_2 - T_1} \quad (4.6)$$

This shows that if the zero curve is upward sloping between T_1 and T_2 , so that $R_2 > R_1$, then $R_F > R_2$ (i.e., the forward rate for a period of time ending at T_2 is greater than the T_2 zero rate). Similarly, if the zero curve is downward sloping with $R_2 < R_1$, then $R_F < R_2$ (i.e., the forward rate is less than the T_2 zero rate). Taking limits as T_2 approaches T_1 in equation (4.6) and letting the common value of the two be T , we obtain

$$R_F = R + T \frac{\partial R}{\partial T}$$

where R is the zero rate for a maturity of T . The value of R_F obtained in this way is known as the *instantaneous forward rate* for a maturity of T . This is the forward rate that is applicable to a very short future time period that begins at time T . Define $P(0, T)$

Table 4.5 Calculation of forward LIBOR rates.

| Year (n) | Zero rate for an n -year investment (% per annum) | Forward rate for n th year (% per annum) |
|--------------|---|--|
| 1 | 3.0 | |
| 2 | 4.0 | 5.0 |
| 3 | 4.6 | 5.8 |
| 4 | 5.0 | 6.2 |
| 5 | 5.3 | 6.5 |

as the price of a zero-coupon bond maturing at time T . Because $P(0, T) = e^{-RT}$, the equation for the instantaneous forward rate can also be written as

$$R_F = -\frac{\partial}{\partial T} \ln P(0, T)$$

By borrowing and lending at LIBOR, a large financial institution can lock in LIBOR forward rates. Suppose LIBOR zero rates are as in Table 4.5. It can borrow \$100 at 3% for 1 year and invest the money at 4% for 2 years, the result is a cash outflow of $100e^{0.03 \times 1} = \$103.05$ at the end of year 1 and an inflow of $100e^{0.04 \times 2} = \$108.33$ at the end of year 2. Since $108.33 = 103.05e^{0.05}$, a return equal to the forward rate (5%) is earned on \$103.05 during the second year. Alternatively, it can borrow \$100 for four years at 5% and invest it for three years at 4.6%. The result is a cash inflow of $100e^{0.046 \times 3} = \114.80 at the end of the third year and a cash outflow of $100e^{0.05 \times 4} = \$122.14$ at the end of the fourth year. Since $122.14 = 114.80e^{0.062}$, money is being borrowed for the fourth year at the forward rate of 6.2%.

If a large investor thinks that rates in the future will be different from today's forward rates there are many trading strategies that the investor will find attractive (see Business Snapshot 4.2). One of these involves entering into a contract known as a *forward rate agreement*. We will now discuss how this contract works and how it is valued.

4.7 FORWARD RATE AGREEMENTS

A forward rate agreement (FRA) is an over-the-counter agreement designed to ensure that a certain interest rate will apply to either borrowing or lending a certain principal during a specified future period of time. The assumption underlying the contract is that the borrowing or lending would normally be done at LIBOR.

Consider a forward rate agreement where company X is agreeing to lend money to company Y for the period of time between T_1 and T_2 . Define:

R_K : The rate of interest agreed to in the FRA

R_F : The forward LIBOR interest rate for the period between times T_1 and T_2 , calculated today⁴

R_M : The actual LIBOR interest rate observed in the market at time T_1 for the period between times T_1 and T_2

L : The principal underlying the contract.

We will depart from our usual assumption of continuous compounding and assume that the rates R_K , R_F , and R_M are all measured with a compounding frequency reflecting the length of the period to which they apply. This means that if $T_2 - T_1 = 0.5$, they are expressed with semiannual compounding; if $T_2 - T_1 = 0.25$, they are expressed with quarterly compounding; and so on. (This assumption corresponds to the usual market practice for FRAs.)

Normally company X would earn R_M from the LIBOR loan. The FRA means that it will earn R_K . The extra interest rate (which may be negative) that it earns as a result of

⁴ LIBOR forward rates are calculated as described in Section 4.6 from the LIBOR/swap zero curve. The latter is determined in the way described in Section 7.6.

Business Snapshot 4.2 Orange County's Yield Curve Plays

Suppose a large investor can borrow or lend at the rates given in Table 4.5 and thinks that 1-year interest rates will not change much over the next 5 years. The investor can borrow 1-year funds and invest for 5-years. The 1-year borrowings can be rolled over for further 1-year periods at the end of the first, second, third, and fourth years. If interest rates do stay about the same, this strategy will yield a profit of about 2.3% per year, because interest will be received at 5.3% and paid at 3%. This type of trading strategy is known as a *yield curve play*. The investor is speculating that rates in the future will be quite different from the forward rates observed in the market today. (In our example, forward rates observed in the market today for future 1-year periods are 5%, 5.8%, 6.2%, and 6.5%.)

Robert Citron, the Treasurer at Orange County, used yield curve plays similar to the one we have just described very successfully in 1992 and 1993. The profit from Mr. Citron's trades became an important contributor to Orange County's budget and he was re-elected. (No one listened to his opponent in the election who said his trading strategy was too risky.)

In 1994 Mr. Citron expanded his yield curve plays. He invested heavily in *inverse floaters*. These pay a rate of interest equal to a fixed rate of interest minus a floating rate. He also leveraged his position by borrowing in the repo market. If short-term interest rates had remained the same or declined he would have continued to do well. As it happened, interest rates rose sharply during 1994. On December 1, 1994, Orange County announced that its investment portfolio had lost \$1.5 billion and several days later it filed for bankruptcy protection.

entering into the FRA is $R_K - R_M$. The interest rate is set at time T_1 and paid at time T_2 . The extra interest rate therefore leads to a cash flow to company X at time T_2 of

$$L(R_K - R_M)(T_2 - T_1) \quad (4.7)$$

Similarly there is a cash flow to company Y at time T_2 of

$$L(R_M - R_K)(T_2 - T_1) \quad (4.8)$$

From equations (4.7) and (4.8), we see that there is another interpretation of the FRA. It is an agreement where company X will receive interest on the principal between T_1 and T_2 at the fixed rate of R_K and pay interest at the realized LIBOR rate of R_M . Company Y will pay interest on the principal between T_1 and T_2 at the fixed rate of R_K and receive interest at R_M .

Usually FRAs are settled at time T_1 rather than T_2 . The payoff must then be discounted from time T_2 to T_1 . For company X, the payoff at time T_1 is

$$\frac{L(R_K - R_M)(T_2 - T_1)}{1 + R_M(T_2 - T_1)}$$

and, for company Y, the payoff at time T_1 is

$$\frac{L(R_M - R_K)(T_2 - T_1)}{1 + R_M(T_2 - T_1)}$$

Example 4.3

Suppose that a company enters into an FRA that is designed to ensure it will receive a fixed rate of 4% on a principal of \$100 million for a 3-month period starting in 3 years. The FRA is an exchange where LIBOR is paid and 4% is received for the 3-month period. If 3-month LIBOR proves to be 4.5% for the 3-month period the cash flow to the lender will be

$$100,000,000 \times (0.04 - 0.045) \times 0.25 = -\$125,000$$

at the 3.25-year point. This is equivalent to a cash flow of

$$-\frac{125,000}{1 + 0.045 \times 0.25} = -\$123,609$$

at the 3-year point. The cash flow to the party on the opposite side of the transaction will be +\$125,000 at the 3.25-year point or +\$123,609 at the 3-year point. (All interest rates quoted in this example are expressed with quarterly compounding.)

Valuation

To value an FRA we first note that it is always worth zero when $R_K = R_F$.⁵ This is because, as noted in Section 4.6, a large financial institution can at no cost lock in the forward rate for a future time period. For example, it can ensure that it earns the forward rate for the time period between years 2 and 3 by borrowing a certain amount of money for 2 years and investing it for 3 years. Similarly, it can ensure that it pays the forward rate for the time period between years 2 and 3 by borrowing for a certain amount of money for 3 years and investing it for 2 years.

Compare two FRAs. The first promises that the LIBOR forward rate R_F will be received on a principal of L between times T_1 and T_2 ; the second promises that R_K will be received on the same principal between the same two dates. The two contracts are the same except for the interest payments received at time T_2 . The excess of the value of the second contract over the first is, therefore, the present value of the difference between these interest payments, or

$$L(R_K - R_F)(T_2 - T_1)e^{-R_2T_2}$$

where R_2 is the continuously compounded riskless zero rate for a maturity T_2 .⁶ Because the value of the first FRA, where R_F is received, is zero, the value of the second FRA, where R_K is received, is

$$V_{\text{FRA}} = L(R_K - R_F)(T_2 - T_1)e^{-R_2T_2} \quad (4.9)$$

Similarly, the value of an FRA where R_K is paid is

$$V_{\text{FRA}} = L(R_F - R_K)(T_2 - T_1)e^{-R_2T_2} \quad (4.10)$$

By comparing equations (4.7) and (4.9), or equations (4.8) and (4.10), we see that an

⁵ It is usually the case that R_K is set equal to R_F when the FRA is first initiated.

⁶ Note that R_K , R_M , and R_F are expressed with a compounding frequency corresponding to $T_2 - T_1$, whereas R_2 is expressed with continuous compounding.

FRA can be valued if we:

1. Calculate the payoff on the assumption that forward rates are realized (that is, on the assumption that $R_M = R_F$).
2. Discount this payoff at the risk-free rate.

We will use this result when we value swaps (which are portfolios of FRAs) in Chapter 7.

Example 4.4

Suppose that LIBOR zero and forward rates are as in Table 4.5. Consider an FRA where we will receive a rate of 6%, measured with annual compounding, and pay LIBOR on a principal of \$100 million between the end of year 1 and the end of year 2. In this case, the forward rate is 5% with continuous compounding or 5.127% with annual compounding. From equation (4.9), it follows that the value of the FRA is

$$100,000,000 \times (0.06 - 0.05127)e^{-0.04 \times 2} = \$805,800$$

4.8 DURATION

The *duration* of a bond, as its name implies, is a measure of how long on average the holder of the bond has to wait before receiving cash payments. A zero-coupon bond that lasts n years has a duration of n years. However, a coupon-bearing bond lasting n years has a duration of less than n years, because the holder receives some of the cash payments prior to year n .

Suppose that a bond provides the holder with cash flows c_i at time t_i ($1 \leq i \leq n$). The bond price B and bond yield y (continuously compounded) are related by

$$B = \sum_{i=1}^n c_i e^{-yt_i} \quad (4.11)$$

The duration of the bond, D , is defined as

$$D = \frac{\sum_{i=1}^n t_i c_i e^{-yt_i}}{B} \quad (4.12)$$

This can be written

$$D = \sum_{i=1}^n t_i \left[\frac{c_i e^{-yt_i}}{B} \right]$$

The term in square brackets is the ratio of the present value of the cash flow at time t_i to the bond price. The bond price is the present value of all payments. The duration is therefore a weighted average of the times when payments are made, with the weight applied to time t_i being equal to the proportion of the bond's total present value provided by the cash flow at time t_i . The sum of the weights is 1.0. Note that for the purposes of the definition of duration all discounting is done at the bond yield rate of interest, y . (We do not use a different zero rate for each cash flow in the way described in Section 4.4.)

When a small change Δy in the yield is considered, it is approximately true that

$$\Delta B = \frac{dB}{dy} \Delta y \quad (4.13)$$

From equation (4.11), this becomes

$$\Delta B = -\Delta y \sum_{i=1}^n c_i t_i e^{-y t_i} \quad (4.14)$$

(Note that there is a negative relationship between B and y . When bond yields increase, bond prices decrease. When bond yields decrease, bond prices increase.) From equations (4.12) and (4.14), the key duration relationship is obtained:

$$\Delta B = -BD \Delta y \quad (4.15)$$

This can be written

$$\frac{\Delta B}{B} = -D \Delta y \quad (4.16)$$

Equation (4.16) is an approximate relationship between percentage changes in a bond price and changes in its yield. It is easy to use and is the reason why duration, which was first suggested by Macaulay in 1938, has become such a popular measure.

Consider a 3-year 10% coupon bond with a face value of \$100. Suppose that the yield on the bond is 12% per annum with continuous compounding. This means that $y = 0.12$. Coupon payments of \$5 are made every 6 months. Table 4.6 shows the calculations necessary to determine the bond's duration. The present values of the bond's cash flows, using the yield as the discount rate, are shown in column 3 (e.g., the present value of the first cash flow is $5e^{-0.12 \times 0.5} = 4.709$). The sum of the numbers in column 3 gives the bond's price as 94.213. The weights are calculated by dividing the numbers in column 3 by 94.213. The sum of the numbers in column 5 gives the duration as 2.653 years.

Small changes in interest rates are often measured in *basis points*. As mentioned earlier, a basis point is 0.01% per annum. The following example investigates the accuracy of the duration relationship in equation (4.15).

Table 4.6 Calculation of duration.

| <i>Time</i> (years) | <i>Cash flow</i> (\$) | <i>Present value</i> | <i>Weight</i> | <i>Time × weight</i> |
|------------------------|--------------------------|----------------------|---------------|----------------------|
| 0.5 | 5 | 4.709 | 0.050 | 0.025 |
| 1.0 | 5 | 4.435 | 0.047 | 0.047 |
| 1.5 | 5 | 4.176 | 0.044 | 0.066 |
| 2.0 | 5 | 3.933 | 0.042 | 0.083 |
| 2.5 | 5 | 3.704 | 0.039 | 0.098 |
| 3.0 | 105 | 73.256 | 0.778 | 2.333 |
| <i>Total:</i> | 130 | 94.213 | 1.000 | 2.653 |

Example 4.5

For the bond in Table 4.6, the bond price, B , is 94.213 and the duration, D , is 2.653, so that equation (4.15) gives

$$\Delta B = -94.213 \times 2.653 \times \Delta y$$

or

$$\Delta B = -249.95 \times \Delta y$$

When the yield on the bond increases by 10 basis points ($= 0.1\%$), $\Delta y = +0.001$. The duration relationship predicts that $\Delta B = -249.95 \times 0.001 = -0.250$, so that the bond price goes down to $94.213 - 0.250 = 93.963$. How accurate is this? Valuing the bond in terms of its yield in the usual way, we find that, when the bond yield increases by 10 basis points to 12.1%, the bond price is

$$5e^{-0.121 \times 0.5} + 5e^{-0.121 \times 1.0} + 5e^{-0.121 \times 1.5} + 5e^{-0.121 \times 2.0} + 5e^{-0.121 \times 2.5} + 105e^{-0.121 \times 3.0} = 93.963$$

which is (to three decimal places) the same as that predicted by the duration relationship.

Modified Duration

The preceding analysis is based on the assumption that y is expressed with continuous compounding. If y is expressed with annual compounding, it can be shown that the approximate relationship in equation (4.15) becomes

$$\Delta B = -\frac{BD \Delta y}{1 + y}$$

More generally, if y is expressed with a compounding frequency of m times per year, then

$$\Delta B = -\frac{BD \Delta y}{1 + y/m}$$

A variable D^* , defined by

$$D^* = \frac{D}{1 + y/m}$$

is sometimes referred to as the bond's *modified duration*. It allows the duration relationship to be simplified to

$$\Delta B = -BD^* \Delta y \quad (4.17)$$

when y is expressed with a compounding frequency of m times per year. The following example investigates the accuracy of the modified duration relationship.

Example 4.6

The bond in Table 4.6 has a price of 94.213 and a duration of 2.653. The yield, expressed with semiannual compounding is 12.3673%. The modified duration, D^* , is given by

$$D^* = \frac{2.653}{1 + 0.123673/2} = 2.499$$

From equation (4.17),

$$\Delta B = -94.213 \times 2.4985 \times \Delta y$$

or

$$\Delta B = -235.39 \times \Delta y$$

When the yield (semiannually compounded) increases by 10 basis points ($= 0.1\%$), we have $\Delta y = +0.001$. The duration relationship predicts that we expect ΔB to be $-235.39 \times 0.001 = -0.235$, so that the bond price goes down to $94.213 - 0.235 = 93.978$. How accurate is this? An exact calculation similar to that in the previous example shows that, when the bond yield (semiannually compounded) increases by 10 basis points to 12.4673%, the bond price becomes 93.978. This shows that the modified duration calculation gives good accuracy for small yield changes.

Another term that is sometimes used is *dollar duration*. This is the product of modified duration and bond price, so that $\Delta B = -D^{**} \Delta y$, where D^{**} is dollar duration.

Bond Portfolios

The duration, D , of a bond portfolio can be defined as a weighted average of the durations of the individual bonds in the portfolio, with the weights being proportional to the bond prices. Equations (4.15) to (4.17) then apply, with B being defined as the value of the bond portfolio. They estimate the change in the value of the bond portfolio for a small change Δy in the yields of all the bonds.

It is important to realize that, when duration is used for bond portfolios, there is an implicit assumption that the yields of all bonds will change by approximately the same amount. When the bonds have widely differing maturities, this happens only when there is a parallel shift in the zero-coupon yield curve. We should therefore interpret equations (4.15) to (4.17) as providing estimates of the impact on the price of a bond portfolio of a small parallel shift, Δy , in the zero curve.

By choosing a portfolio so that the duration of assets equals the duration of liabilities (i.e., the net duration is zero), a financial institution eliminates its exposure to small parallel shifts in the yield curve. But it is still exposed to shifts that are either large or nonparallel.


4.9 CONVEXITY

The duration relationship applies only to small changes in yields. This is illustrated in Figure 4.2, which shows the relationship between the percentage change in value and change in yield for two bond portfolios having the same duration. The gradients of the two curves are the same at the origin. This means that both bond portfolios change in value by the same percentage for small yield changes and is consistent with equation (4.16). For large yield changes, the portfolios behave differently. Portfolio X has more curvature in its relationship with yields than portfolio Y . A factor known as *convexity* measures this curvature and can be used to improve the relationship in equation (4.16).

A measure of convexity is

$$C = \frac{1}{B} \frac{d^2 B}{dy^2} = \frac{\sum_{i=1}^n c_i t_i^2 e^{-y t_i}}{B}$$

Figure 4.2 Two bond portfolios with the same duration.



From Taylor series expansions, we obtain a more accurate expression than equation (4.13), given by

$$\Delta B = \frac{dB}{dy} \Delta y + \frac{1}{2} \frac{d^2 B}{dy^2} \Delta y^2$$

This leads to

$$\frac{\Delta B}{B} = -D \Delta y + \frac{1}{2} C (\Delta y)^2$$

For a portfolio with a particular duration, the convexity of a bond portfolio tends to be greatest when the portfolio provides payments evenly over a long period of time. It is least when the payments are concentrated around one particular point in time. By choosing a portfolio of assets and liabilities with a net duration of zero and a net convexity of zero, a financial institution can make itself immune to relatively large parallel shifts in the zero curve. However, it is still exposed to nonparallel shifts.

4.10 THEORIES OF THE TERM STRUCTURE OF INTEREST RATES

It is natural to ask what determines the shape of the zero curve. Why is it sometimes downward sloping, sometimes upward sloping, and sometimes partly upward sloping and partly downward sloping? A number of different theories have been proposed. The simplest is *expectations theory*, which conjectures that long-term interest rates should reflect expected future short-term interest rates. More precisely, it argues that a forward interest rate corresponding to a certain future period is equal to the expected future zero interest rate for that period. Another idea, *market segmentation theory*, conjectures that

there need be no relationship between short-, medium-, and long-term interest rates. Under the theory, a major investor such as a large pension fund invests in bonds of a certain maturity and does not readily switch from one maturity to another. The short-term interest rate is determined by supply and demand in the short-term bond market; the medium-term interest rate is determined by supply and demand in the medium-term bond market; and so on.

The theory that is most appealing is *liquidity preference theory*. The basic assumption underlying the theory is that investors prefer to preserve their liquidity and invest funds for short periods of time. Borrowers, on the other hand, usually prefer to borrow at fixed rates for long periods of time. This leads to a situation in which forward rates are greater than expected future zero rates. The theory is also consistent with the empirical result that yield curves tend to be upward sloping more often than they are downward sloping.

The Management of Net Interest Income

To understand liquidity preference theory, it is useful to consider the interest rate risk faced by banks when they take deposits and make loans. The *net interest income* of the bank is the excess of the interest received over the interest paid and needs to be carefully managed.

Consider a simple situation where a bank offers consumers a one-year and a five-year deposit rate as well as a one-year and five-year mortgage rate. The rates are shown in Table 4.7. We make the simplifying assumption that the expected one-year interest rate for future time periods to equal the one-year rates prevailing in the market today. Loosely speaking this means that the market considers interest rate increases to be just as likely as interest rate decreases. As a result, the rates in Table 4.7 are “fair” in that they reflect the market’s expectations (i.e., they correspond to expectations theory). Investing money for one year and reinvesting for four further one-year periods give the same expected overall return as a single five-year investment. Similarly, borrowing money for one year and refinancing each year for the next four years leads to the same expected financing costs as a single five-year loan.

Suppose you have money to deposit and agree with the prevailing view that interest rate increases are just as likely as interest rate decreases. Would you choose to deposit your money for one year at 3% per annum or for five years at 3% per annum? The chances are that you would choose one year because this gives you more financial flexibility. It ties up your funds for a shorter period of time.

Now suppose that you want a mortgage. Again you agree with the prevailing view that interest rate increases are just as likely as interest rate decreases. Would you choose a one-year mortgage at 6% or a five-year mortgage at 6%? The chances are that you would choose a five-year mortgage because it fixes your borrowing rate for the next five years and subjects you to less refinancing risk.

When the bank posts the rates shown in Table 4.7, it is likely to find that the majority

Table 4.7 Example of rates offered by a bank to its customers.

| <i>Maturity (years)</i> | <i>Deposit rate</i> | <i>Mortgage rate</i> |
|-------------------------|---------------------|----------------------|
| 1 | 3% | 6% |
| 5 | 3% | 6% |

Table 4.8 Five-year rates are increased in an attempt to match maturities of assets and liabilities.

| <i>Maturity (years)</i> | <i>Deposit rate</i> | <i>Mortgage rate</i> |
|-------------------------|---------------------|----------------------|
| 1 | 3% | 6% |
| 5 | 4% | 7% |

of its customers opt for one-year deposits and five-year mortgages. This creates an asset/liability mismatch for the bank and subjects it to risks. There is no problem if interest rates fall. The bank will find itself financing the five-year 6% loans with deposits that cost less than 3% in the future and net interest income will increase. However, if rates rise, the deposits that are financing these 6% loans will cost more than 3% in the future and net interest income will decline. A 3% rise in interest rates would reduce the net interest income to zero.

It is the job of the asset/liability management group to ensure that the maturities of the assets on which interest is earned and the maturities of the liabilities on which interest is paid are matched. One way it can do this is by increasing the five-year rate on both deposits and mortgages. For example, it could move to the situation in Table 4.8 where the five-year deposit rate is 4% and the five-year mortgage rate 7%. This would make five-year deposits relatively more attractive and one-year mortgages relatively more attractive. Some customers who chose one-year deposits when the rates were as in Table 4.7 will switch to five-year deposits in the Table 4.8 situation. Some customers who chose five-year mortgages when the rates were as in Table 4.7 will choose one-year mortgages. This may lead to the maturities of assets and liabilities being matched. If there is still an imbalance with depositors tending to choose a one-year maturity and borrowers a five-year maturity, five-year deposit and mortgage rates could be increased even further. Eventually the imbalance will disappear.

The net result of all banks behaving in the way we have just described is liquidity preference theory. Long-term rates tend to be higher than those that would be predicted by expected future short-term rates. The yield curve is upward sloping most of the time. It is downward sloping only when the market expects a steep decline in short-term rates.

Many banks now have sophisticated systems for monitoring the decisions being made by customers so that, when they detect small differences between the maturities of the assets and liabilities being chosen by customers they can fine tune the rates they offer. Sometimes derivatives such as interest rate swaps (which will be discussed in Chapter 7) are also used to manage their exposure. The result of all this is that net interest income is usually very stable. This has not always been the case. In the United States, the failure of Savings and Loan companies in the 1980s and the failure of Continental Illinois in 1984 were to a large extent a result of the fact that they did not match the maturities of assets and liabilities. Both failures proved to be very expensive for US taxpayers.

Liquidity

In addition to creating problems in the way that has been described, a portfolio where maturities are mismatched can lead to liquidity problems. Consider a financial institution that funds 5-year fixed rate loans with wholesale deposits that last only 3 months. It might recognize its exposure to rising interest rates and hedge its interest rate risk.

Business Snapshot 4.3 Liquidity and the 2007–2009 Financial Crisis

During the credit crisis that started in July 2007 there was a “flight to quality,” where financial institutions and investors looked for safe investments and were less inclined than before to take credit risks. Financial institutions that relied on short-term funding experienced liquidity problems. One example is Northern Rock in the United Kingdom, which chose to finance much of its mortgage portfolio with wholesale deposits, some lasting only 3 months. Starting in September 2007, the depositors became nervous and refused to roll over the funding they were providing to Northern Rock, i.e., at the end of a 3-month period they would refuse to deposit their funds for a further 3-month period. As a result, Northern Rock was unable to finance its assets. It was taken over by the UK government in early 2008. In the US, financial institutions such as Bear Stearns and Lehman Brothers experienced similar liquidity problems because they had chosen to fund part of their operations with short-term funds.

(One way of doing this is by using interest rate swaps, as mentioned earlier.) However, it still has a liquidity risk. Wholesale depositors may, for some reason, lose confidence in the financial institution and refuse to continue to provide the financial institution with short-term funding. The financial institution, even if it has adequate equity capital, will then experience a severe liquidity problem that could lead to its downfall. As described in Business Snapshot 4.3, these types of liquidity problems were the root cause of some of the failures of financial institutions during the crisis that started in 2007.

SUMMARY

Two important interest rates for derivative traders are Treasury rates and LIBOR rates. Treasury rates are the rates paid by a government on borrowings in its own currency. LIBOR rates are short-term lending rates offered by banks in the interbank market. Derivatives traders have traditionally assumed that the LIBOR rate is the short-term risk-free rate at which funds can be borrowed or lent.

The compounding frequency used for an interest rate defines the units in which it is measured. The difference between an annually compounded rate and a quarterly compounded rate is analogous to the difference between a distance measured in miles and a distance measured in kilometers. Traders frequently use continuous compounding when analyzing the value of options and more complex derivatives.

Many different types of interest rates are quoted in financial markets and calculated by analysts. The n -year zero or spot rate is the rate applicable to an investment lasting for n years when all of the return is realized at the end. The par yield on a bond of a certain maturity is the coupon rate that causes the bond to sell for its par value. Forward rates are the rates applicable to future periods of time implied by today's zero rates.

The method most commonly used to calculate zero rates is known as the bootstrap method. It involves starting with short-term instruments and moving progressively to longer-term instruments, making sure that the zero rates calculated at each stage are

consistent with the prices of the instruments. It is used daily by trading desks to calculate a Treasury zero-rate curve.

A forward rate agreement (FRA) is an over-the-counter agreement that the LIBOR rate will be exchanged for a specified interest rate during a specified future period of time. An FRA can be valued by assuming that forward LIBOR rates are realized and discounting the resulting payoff.

An important concept in interest rate markets is *duration*. Duration measures the sensitivity of the value of a bond portfolio to a small parallel shift in the zero-coupon yield curve. Specifically,

$$\Delta B = -BD \Delta y$$

where B is the value of the bond portfolio, D is the duration of the portfolio, Δy is the size of a small parallel shift in the zero curve, and ΔB is the resultant effect on the value of the bond portfolio.

Liquidity preference theory can be used to explain the interest rate term structures that are observed in practice. The theory argues that most entities like to borrow long and lend short. To match the maturities of borrowers and lenders, it is necessary for financial institutions to raise long-term rates so that forward interest rates are higher than expected future spot interest rates.

FURTHER READING

Allen, S. L., and A. D. Kleinstein. *Valuing Fixed-Income Investments and Derivative Securities: Cash Flow Analysis and Calculations*. New York: New York Institute of Finance, 1991.

Fabozzi, F. J. *Fixed-Income Mathematics: Analytical and Statistical Techniques*, 4th edn. New York: McGraw-Hill, 2006.

Fabozzi, F. J. *Duration, Convexity, and Other Bond Risk Measures*. Frank J. Fabozzi Assoc., 1999.

Grinblatt, M., and F. A. Longstaff. "Financial Innovation and the Role of Derivatives Securities: An Empirical Analysis of the Treasury Strips Program," *Journal of Finance*, 55, 3 (2000): 1415–36.

Jorion, P. *Big Bets Gone Bad: Derivatives and Bankruptcy in Orange County*. New York: Academic Press, 1995.

Stigum, M., and F. L. Robinson. *Money Markets and Bond Calculations*. Chicago: Irwin, 1996.

Practice Questions (Answers in Solutions Manual)

- 4.1. A bank quotes an interest rate of 14% per annum with quarterly compounding. What is the equivalent rate with (a) continuous compounding and (b) annual compounding?
- 4.2. What is meant by LIBOR and LIBID. Which is higher?
- 4.3. The 6-month and 1-year zero rates are both 10% per annum. For a bond that has a life of 18 months and pays a coupon of 8% per annum (with semiannual payments and one having just been made), the yield is 10.4% per annum. What is the bond's price? What is the 18-month zero rate? All rates are quoted with semiannual compounding.
- 4.4. An investor receives \$1,100 in one year in return for an investment of \$1,000 now. Calculate the percentage return per annum with:
 - (a) Annual compounding
 - (b) Semiannual compounding

- (c) Monthly compounding
 (d) Continuous compounding.

4.5. Suppose that zero interest rates with continuous compounding are as follows:

| <i>Maturity (months)</i> | <i>Rate (% per annum)</i> |
|--------------------------|---------------------------|
| 3 | 8.0 |
| 6 | 8.2 |
| 9 | 8.4 |
| 12 | 8.5 |
| 15 | 8.6 |
| 18 | 8.7 |

Calculate forward interest rates for the second, third, fourth, fifth, and sixth quarters.

- 4.6. Assuming that zero rates are as in Problem 4.5, what is the value of an FRA that enables the holder to earn 9.5% for a 3-month period starting in 1 year on a principal of \$1,000,000? The interest rate is expressed with quarterly compounding.
- 4.7. The term structure of interest rates is upward-sloping. Put the following in order of magnitude:
 (a) The 5-year zero rate
 (b) The yield on a 5-year coupon-bearing bond
 (c) The forward rate corresponding to the period between 4.75 and 5 years in the future.
 What is the answer when the term structure of interest rates is downward-sloping?
- 4.8. What does duration tell you about the sensitivity of a bond portfolio to interest rates. What are the limitations of the duration measure?
- 4.9. What rate of interest with continuous compounding is equivalent to 15% per annum with monthly compounding?
- 4.10. A deposit account pays 12% per annum with continuous compounding, but interest is actually paid quarterly. How much interest will be paid each quarter on a \$10,000 deposit?
- 4.11. Suppose that 6-month, 12-month, 18-month, 24-month, and 30-month zero rates are, respectively, 4%, 4.2%, 4.4%, 4.6%, and 4.8% per annum, with continuous compounding. Estimate the cash price of a bond with a face value of 100 that will mature in 30 months and pays a coupon of 4% per annum semiannually.
- 4.12. A 3-year bond provides a coupon of 8% semiannually and has a cash price of 104. What is the bond's yield?
- 4.13. Suppose that the 6-month, 12-month, 18-month, and 24-month zero rates are 5%, 6%, 6.5%, and 7%, respectively. What is the 2-year par yield?
- 4.14. Suppose that zero interest rates with continuous compounding are as follows:

| <i>Maturity (years)</i> | <i>Rate (% per annum)</i> |
|-------------------------|---------------------------|
| 1 | 2.0 |
| 2 | 3.0 |
| 3 | 3.7 |
| 4 | 4.2 |
| 5 | 4.5 |

Calculate forward interest rates for the second, third, fourth, and fifth years.

- 4.15. Use the rates in Problem 4.14 to value an FRA where you will pay 5% (compounded annually) for the third year on \$1 million.
- 4.16. A 10-year 8% coupon bond currently sells for \$90. A 10-year 4% coupon bond currently sells for \$80. What is the 10-year zero rate? (*Hint*: Consider taking a long position in two of the 4% coupon bonds and a short position in one of the 8% coupon bonds.)
- 4.17. Explain carefully why liquidity preference theory is consistent with the observation that the term structure of interest rates tends to be upward-sloping more often than it is downward-sloping.
- 4.18. “When the zero curve is upward-sloping, the zero rate for a particular maturity is greater than the par yield for that maturity. When the zero curve is downward-sloping the reverse is true.” Explain why this is so.
- 4.19. Why are US Treasury rates significantly lower than other rates that are close to risk-free?
- 4.20. Why does a loan in the repo market involve very little credit risk?
- 4.21. Explain why an FRA is equivalent to the exchange of a floating rate of interest for a fixed rate of interest.
- 4.22. A 5-year bond with a yield of 11% (continuously compounded) pays an 8% coupon at the end of each year.
- (a) What is the bond’s price?
 - (b) What is the bond’s duration?
 - (c) Use the duration to calculate the effect on the bond’s price of a 0.2% decrease in its yield.
 - (d) Recalculate the bond’s price on the basis of a 10.8% per annum yield and verify that the result is in agreement with your answer to (c).
- 4.23. The cash prices of 6-month and 1-year Treasury bills are 94.0 and 89.0. A 1.5-year bond that will pay coupons of \$4 every 6 months currently sells for \$94.84. A 2-year bond that will pay coupons of \$5 every 6 months currently sells for \$97.12. Calculate the 6-month, 1-year, 1.5-year, and 2-year zero rates.
- 4.24. “An interest rate swap where 6-month LIBOR is exchanged for a fixed rate of 5% on a principal of \$100 million for 5 years involves a known cash flow and a portfolio of nine FRAs.” Explain this statement.

Further Questions

- 4.25. A five-year bond provides a coupon of 5% per annum payable semiannually. Its price is 104. What is the bond’s yield? You may find Excel’s Solver useful.
- 4.26. Suppose that LIBOR rates for maturities of one, two, three, four, five, and six months are 2.6%, 2.9%, 3.1%, 3.2%, 3.25%, and 3.3% with continuous compounding. What are the forward rates for future one-month periods?
- 4.27. A bank can borrow or lend at LIBOR. The two-month LIBOR rate is 0.28% per annum with continuous compounding. Assuming that interest rates cannot be negative, what is the arbitrage opportunity if the three-month LIBOR rate is 0.1% per year with continuous compounding. How low can the three-month LIBOR rate become without an arbitrage opportunity being created?

- 4.28. A bank can borrow or lend at LIBOR. Suppose that the six-month rate is 5% and the nine-month rate is 6%. The rate that can be locked in for the period between six months and nine months using an FRA is 7%. What arbitrage opportunities are open to the bank? All rates are continuously compounded.
- 4.29. An interest rate is quoted as 5% per annum with semiannual compounding. What is the equivalent rate with (a) annual compounding, (b) monthly compounding, and (c) continuous compounding.
- 4.30. The 6-month, 12-month, 18-month, and 24-month zero rates are 4%, 4.5%, 4.75%, and 5%, with semiannual compounding.
- (a) What are the rates with continuous compounding?
- (b) What is the forward rate for the 6-month period beginning in 18 months?
- (c) What is the value of an FRA that promises to pay you 6% (compounded semi-annually) on a principal of \$1 million for the 6-month period starting in 18 months?
- 4.31. What is the 2-year par yield when the zero rates are as in Problem 4.30? What is the yield on a 2-year bond that pays a coupon equal to the par yield?
- 4.32. The following table gives the prices of bonds:

| <i>Bond principal</i> (\$) | <i>Time to maturity</i> (years) | <i>Annual coupon*</i> (\$) | <i>Bond price</i> (\$) |
|-------------------------------|------------------------------------|-------------------------------|---------------------------|
| 100 | 0.50 | 0.0 | 98 |
| 100 | 1.00 | 0.0 | 95 |
| 100 | 1.50 | 6.2 | 101 |
| 100 | 2.00 | 8.0 | 104 |

* Half the stated coupon is assumed to be paid every six months.

- (a) Calculate zero rates for maturities of 6 months, 12 months, 18 months, and 24 months.
- (b) What are the forward rates for the following periods: 6 months to 12 months, 12 months to 18 months, and 18 months to 24 months?
- (c) What are the 6-month, 12-month, 18-month, and 24-month par yields for bonds that provide semiannual coupon payments?
- (d) Estimate the price and yield of a 2-year bond providing a semiannual coupon of 7% per annum.
- 4.33. Portfolio A consists of a 1-year zero-coupon bond with a face value of \$2,000 and a 10-year zero-coupon bond with a face value of \$6,000. Portfolio B consists of a 5.95-year zero-coupon bond with a face value of \$5,000. The current yield on all bonds is 10% per annum.
- (a) Show that both portfolios have the same duration.
- (b) Show that the percentage changes in the values of the two portfolios for a 0.1% per annum increase in yields are the same.
- (c) What are the percentage changes in the values of the two portfolios for a 5% per annum increase in yields?



5 CHAPTER

Determination of Forward and Futures Prices

In this chapter we examine how forward prices and futures prices are related to the spot price of the underlying asset. Forward contracts are easier to analyze than futures contracts because there is no daily settlement—only a single payment at maturity. We therefore start this chapter by considering the relationship between the forward price and the spot price. Luckily it can be shown that the forward price and futures price of an asset are usually very close when the maturities of the two contracts are the same. This is convenient because it means that results obtained for forwards are usually also true for futures.

In the first part of the chapter we derive some important general results on the relationship between forward (or futures) prices and spot prices. We then use the results to examine the relationship between futures prices and spot prices for contracts on stock indices, foreign exchange, and commodities. We will consider interest rate futures contracts in the next chapter.

5.1 INVESTMENT ASSETS vs. CONSUMPTION ASSETS

When considering forward and futures contracts, it is important to distinguish between investment assets and consumption assets. An *investment asset* is an asset that is held for investment purposes by significant numbers of investors. Stocks and bonds are clearly investment assets. Gold and silver are also examples of investment assets. Note that investment assets do not have to be held exclusively for investment. (Silver, for example, has a number of industrial uses.) However, they do have to satisfy the requirement that they are held by significant numbers of investors solely for investment. A *consumption asset* is an asset that is held primarily for consumption. It is not usually held for investment. Examples of consumption assets are commodities such as copper, oil, and pork bellies.

As we shall see later in this chapter, we can use arbitrage arguments to determine the forward and futures prices of an investment asset from its spot price and other observable market variables. We cannot do this for consumption assets.

5.2 SHORT SELLING

Some of the arbitrage strategies presented in this chapter involve *short selling*. This trade, usually simply referred to as “shorting”, involves selling an asset that is not owned. It is something that is possible for some—but not all—investment assets. We will illustrate how it works by considering a short sale of shares of a stock.

Suppose an investor instructs a broker to short 500 IBM shares. The broker will carry out the instructions by borrowing the shares from another client and selling them in the market in the usual way. The investor can maintain the short position for as long as desired, provided there are always shares for the broker to borrow. At some stage, however, the investor will close out the position by purchasing 500 IBM shares. These are then replaced in the account of the client from which the shares were borrowed. The investor takes a profit if the stock price has declined and a loss if it has risen. If at any time while the contract is open the broker is not able to borrow shares, the investor is forced to close out the position, even if not ready to do so. Sometimes a fee is charged for lending shares or other securities to the party doing the shorting.

An investor with a short position must pay to the broker any income, such as dividends or interest, that would normally be received on the securities that have been shorted. The broker will transfer this income to the account of the client from whom the securities have been borrowed. Consider the position of an investor who shorts 500 shares in April when the price per share is \$120 and closes out the position by buying them back in July when the price per share is \$100. Suppose that a dividend of \$1 per share is paid in May. The investor receives $500 \times \$120 = \$60,000$ in April when the short position is initiated. The dividend leads to a payment by the investor of $500 \times \$1 = \500 in May. The investor also pays $500 \times \$100 = \$50,000$ for shares when the position is closed out in July. The net gain, therefore, is

$$\$60,000 - \$500 - \$50,000 = \$9,500$$

assuming there is no fee for borrowing the shares. Table 5.1 illustrates this example and shows that the cash flows from the short sale are the mirror image of the cash flows from purchasing the shares in April and selling them in July. (Again, this assumes no borrowing fee.)

Table 5.1 Cash flows from short sale and purchase of shares.

| <i>Purchase of shares</i> | |
|--|-----------------------|
| April: Purchase 500 shares for \$120 | −\$60,000 |
| May: Receive dividend | +\$500 |
| July: Sell 500 shares for \$100 per share | +\$50,000 |
| | Net profit = −\$9,500 |
| <i>Short sale of shares</i> | |
| April: Borrow 500 shares and sell them for \$120 | +\$60,000 |
| May: Pay dividend | −\$500 |
| July: Buy 500 shares for \$100 per share | −\$50,000 |
| Replace borrowed shares to close short position | |
| | Net profit = +\$9,500 |

The investor is required to maintain a *margin account* with the broker. The margin account consists of cash or marketable securities deposited by the investor with the broker to guarantee that the investor will not walk away from the short position if the share price increases. It is similar to the margin account discussed in Chapter 2 for futures contracts. An initial margin is required and if there are adverse movements (i.e., increases) in the price of the asset that is being shorted, additional margin may be required. If the additional margin is not provided, the short position is closed out. The margin account does not represent a cost to the investor. This is because interest is usually paid on the balance in margin accounts and, if the interest rate offered is unacceptable, marketable securities such as Treasury bills can be used to meet margin requirements. The proceeds of the sale of the asset belong to the investor and normally form part of the initial margin.

From time to time regulations are changed on short selling. In 1938, the uptick rule was introduced. This allowed shares to be shorted only on an “uptick”—that is, when the most recent movement in the share price was an increase. The SEC abolished the uptick rule in July 2007, but introduced an “alternative uptick” rule in February 2010. Under this rule, when the price of a stock has decreased by more than 10% in one day, there are restrictions on short selling for that day and the next. These restrictions are that the stock can be shorted only at a price that is higher than the best current bid price. Occasionally there are temporary bans on short selling. This happened in a number of countries in 2008 because it was considered that short selling contributed to the high market volatility that was being experienced.

5.3 ASSUMPTIONS AND NOTATION

In this chapter we will assume that the following are all true for some market participants:

1. The market participants are subject to no transaction costs when they trade.
2. The market participants are subject to the same tax rate on all net trading profits.
3. The market participants can borrow money at the same risk-free rate of interest as they can lend money.
4. The market participants take advantage of arbitrage opportunities as they occur.

Note that we do not require these assumptions to be true for all market participants. All that we require is that they be true—or at least approximately true—for a few key market participants such as large derivatives dealers. It is the trading activities of these key market participants and their eagerness to take advantage of arbitrage opportunities as they occur that determine the relationship between forward and spot prices.

The following notation will be used throughout this chapter:

- T : Time until delivery date in a forward or futures contract (in years)
- S_0 : Price of the asset underlying the forward or futures contract today
- F_0 : Forward or futures price today
- r : Zero-coupon risk-free rate of interest per annum, expressed with continuous compounding, for an investment maturing at the delivery date (i.e., in T years).

The risk-free rate r is the rate at which money is borrowed or lent when there is no credit risk, so that the money is certain to be repaid. As discussed in Chapter 4, participants in derivatives markets have traditionally assumed that LIBOR rates rather than Treasury rates are the relevant risk-free rates.

5.4 FORWARD PRICE FOR AN INVESTMENT ASSET

The easiest forward contract to value is one written on an investment asset that provides the holder with no income. Non-dividend-paying stocks and zero-coupon bonds are examples of such investment assets.

Consider a long forward contract to purchase a non-dividend-paying stock in 3 months.¹ Assume the current stock price is \$40 and the 3-month risk-free interest rate is 5% per annum.

Suppose first that the forward price is relatively high at \$43. An arbitrageur can borrow \$40 at the risk-free interest rate of 5% per annum, buy one share, and short a forward contract to sell one share in 3 months. At the end of the 3 months, the arbitrageur delivers the share and receives \$43. The sum of money required to pay off the loan is

$$40e^{0.05 \times 3/12} = \$40.50$$

By following this strategy, the arbitrageur locks in a profit of $\$43.00 - \$40.50 = \$2.50$ at the end of the 3-month period.

Suppose next that the forward price is relatively low at \$39. An arbitrageur can short one share, invest the proceeds of the short sale at 5% per annum for 3 months, and take a long position in a 3-month forward contract. The proceeds of the short sale grow to $40e^{0.05 \times 3/12}$, or \$40.50 in 3 months. At the end of the 3 months, the arbitrageur pays \$39, takes delivery of the share under the terms of the forward contract, and uses it to close out the short position. A net gain of

$$\$40.50 - \$39.00 = \$1.50$$

is therefore made at the end of the 3 months. The two trading strategies we have considered are summarized in Table 5.2.

Under what circumstances do arbitrage opportunities such as those in Table 5.2 not exist? The first arbitrage works when the forward price is greater than \$40.50. The second arbitrage works when the forward price is less than \$40.50. We deduce that for there to be no arbitrage the forward price must be exactly \$40.50.

A Generalization

To generalize this example, we consider a forward contract on an investment asset with price S_0 that provides no income. Using our notation, T is the time to maturity, r is the risk-free rate, and F_0 is the forward price. The relationship between F_0 and S_0 is

$$F_0 = S_0 e^{rT} \tag{5.1}$$

¹ Forward contracts on individual stocks do not often arise in practice. However, they form useful examples for developing our ideas. Futures on individual stocks started trading in the United States in November 2002.

Table 5.2 Arbitrage opportunities when forward price is out of line with spot price for asset providing no income. (Asset price = \$40; interest rate = 5%; maturity of forward contract = 3 months.)

| <i>Forward Price = \$43</i> | <i>Forward Price = \$39</i> |
|--|---|
| <i>Action now:</i> | <i>Action now:</i> |
| Borrow \$40 at 5% for 3 months | Short 1 unit of asset to realize \$40 |
| Buy one unit of asset | Invest \$40 at 5% for 3 months |
| Enter into forward contract to sell asset in 3 months for \$43 | Enter into a forward contract to buy asset in 3 months for \$39 |
| <i>Action in 3 months:</i> | <i>Action in 3 months:</i> |
| Sell asset for \$43 | Buy asset for \$39 |
| Use \$40.50 to repay loan with interest | Close short position |
| | Receive \$40.50 from investment |
| Profit realized = \$2.50 | Profit realized = \$1.50 |

If $F_0 > S_0e^{rT}$, arbitrageurs can buy the asset and short forward contracts on the asset. If $F_0 < S_0e^{rT}$, they can short the asset and enter into long forward contracts on it.² In our example, $S_0 = 40$, $r = 0.05$, and $T = 0.25$, so that equation (5.1) gives

$$F_0 = 40e^{0.05 \times 0.25} = \$40.50$$

which is in agreement with our earlier calculations.

A long forward contract and a spot purchase both lead to the asset being owned at time T . The forward price is higher than the spot price because of the cost of financing the spot purchase of the asset during the life of the forward contract. This point was overlooked by Kidder Peabody in 1994, much to its cost (see Business Snapshot 5.1).

Example 5.1

Consider a 4-month forward contract to buy a zero-coupon bond that will mature 1 year from today. (This means that the bond will have 8 months to go when the forward contract matures.) The current price of the bond is \$930. We assume that the 4-month risk-free rate of interest (continuously compounded) is 6% per annum. Because zero-coupon bonds provide no income, we can use equation (5.1) with $T = 4/12$, $r = 0.06$, and $S_0 = 930$. The forward price, F_0 , is given by

$$F_0 = 930e^{0.06 \times 4/12} = \$948.79$$

This would be the delivery price in a contract negotiated today.

² For another way of seeing that equation (5.1) is correct, consider the following strategy: buy one unit of the asset and enter into a short forward contract to sell it for F_0 at time T . This costs S_0 and is certain to lead to a cash inflow of F_0 at time T . Therefore S_0 must equal the present value of F_0 ; that is, $S_0 = F_0e^{-rT}$, or equivalently $F_0 = S_0e^{rT}$.

Business Snapshot 5.1 Kidder Peabody's Embarrassing Mistake

Investment banks have developed a way of creating a zero-coupon bond, called a *strip*, from a coupon-bearing Treasury bond by selling each of the cash flows underlying the coupon-bearing bond as a separate security. Joseph Jett, a trader working for Kidder Peabody, had a relatively simple trading strategy. He would buy strips and sell them in the forward market. As equation (5.1) shows, the forward price of a security providing no income is always higher than the spot price. Suppose, for example, that the 3-month interest rate is 4% per annum and the spot price of a strip is \$70. The 3-month forward price of the strip is $70e^{0.04 \times 3/12} = \70.70 .

Kidder Peabody's computer system reported a profit on each of Jett's trades equal to the excess of the forward price over the spot price (\$0.70 in our example). In fact this profit was nothing more than the cost of financing the purchase of the strip. But, by rolling his contracts forward, Jett was able to prevent this cost from accruing to him.

The result was that the system reported a profit of \$100 million on Jett's trading (and Jett received a big bonus) when in fact there was a loss in the region of \$350 million. This shows that even large financial institutions can get relatively simple things wrong!

What If Short Sales Are Not Possible?

Short sales are not possible for all investment assets and sometimes a fee is charged for borrowing assets. As it happens, this does not matter. To derive equation (5.1), we do not need to be able to short the asset. All that we require is that there be a significant number of people who hold the asset purely for investment (and by definition this is always true of an investment asset). If the forward price is too low, they will find it attractive to sell the asset and take a long position in a forward contract.

Suppose that the underlying asset has no storage costs or income. If $F_0 > S_0e^{rT}$, an investor can adopt the following strategy:

1. Borrow S_0 dollars at an interest rate r for T years.
2. Buy 1 unit of the asset.
3. Short a forward contract on 1 unit of the asset.

At time T , the asset is sold for F_0 . An amount S_0e^{rT} is required to repay the loan at this time and the investor makes a profit of $F_0 - S_0e^{rT}$.

Suppose next that $F_0 < S_0e^{rT}$. In this case, an investor who owns the asset can:

1. Sell the asset for S_0 .
2. Invest the proceeds at interest rate r for time T .
3. Take a long position in a forward contract on 1 unit of the asset.

At time T , the cash invested has grown to S_0e^{rT} . The asset is repurchased for F_0 and the investor makes a profit of $S_0e^{rT} - F_0$ relative to the position the investor would have been in if the asset had been kept.

As in the non-dividend-paying stock example considered earlier, we can expect the forward price to adjust so that neither of the two arbitrage opportunities we have considered exists. This means that the relationship in equation (5.1) must hold.

5.5 KNOWN INCOME

In this section we consider a forward contract on an investment asset that will provide a perfectly predictable cash income to the holder. Examples are stocks paying known dividends and coupon-bearing bonds. We adopt the same approach as in the previous section. We first look at a numerical example and then review the formal arguments.

Consider a long forward contract to purchase a coupon-bearing bond whose current price is \$900. We will suppose that the forward contract matures in 9 months. We will also suppose that a coupon payment of \$40 is expected after 4 months. We assume that the 4-month and 9-month risk-free interest rates (continuously compounded) are, respectively, 3% and 4% per annum.

Suppose first that the forward price is relatively high at \$910. An arbitrageur can borrow \$900 to buy the bond and short a forward contract. The coupon payment has a present value of $40e^{-0.03 \times 4/12} = \39.60 . Of the \$900, \$39.60 is therefore borrowed at 3% per annum for 4 months so that it can be repaid with the coupon payment. The remaining \$860.40 is borrowed at 4% per annum for 9 months. The amount owing at the end of the 9-month period is $860.40e^{0.04 \times 0.75} = \886.60 . A sum of \$910 is received for the bond under the terms of the forward contract. The arbitrageur therefore makes a net profit of

$$910.00 - 886.60 = \$23.40$$

Suppose next that the forward price is relatively low at \$870. An investor can short the bond and enter into a long forward contract. Of the \$900 realized from shorting the bond, \$39.60 is invested for 4 months at 3% per annum so that it grows into an amount sufficient to pay the coupon on the bond. The remaining \$860.40 is invested for 9 months at 4% per annum and grows to \$886.60. Under the terms of the forward contract, \$870 is paid to buy the bond and the short position is closed out. The investor therefore gains

$$886.60 - 870 = \$16.60$$

The two strategies we have considered are summarized in Table 5.3.³ The first strategy in Table 5.3 produces a profit when the forward price is greater than \$886.60, whereas the second strategy produces a profit when the forward price is less than \$886.60. It follows that if there are no arbitrage opportunities then the forward price must be \$886.60.

A Generalization

We can generalize from this example to argue that, when an investment asset will provide income with a present value of I during the life of a forward contract, we have

$$F_0 = (S_0 - I)e^{rT} \quad (5.2)$$

In our example, $S_0 = 900.00$, $I = 40e^{-0.03 \times 4/12} = 39.60$, $r = 0.04$, and $T = 0.75$, so that

$$F_0 = (900.00 - 39.60)e^{0.04 \times 0.75} = \$886.60$$

³ If shorting the bond is not possible, investors who already own the bond will sell it and buy a forward contract on the bond increasing the value of their position by \$16.60. This is similar to the strategy we described for gold in Section 5.4.

Table 5.3 Arbitrage opportunities when 9-month forward price is out of line with spot price for asset providing known cash income. (Asset price = \$900; income of \$40 occurs at 4 months; 4-month and 9-month rates are, respectively, 3% and 4% per annum.)

| <i>Forward price = \$910</i> | <i>Forward price = \$870</i> |
|--|--|
| <i>Action now:</i> | <i>Action now:</i> |
| Borrow \$900: \$39.60 for 4 months and \$860.40 for 9 months | Short 1 unit of asset to realize \$900 Invest \$39.60 for 4 months and \$860.40 for 9 months |
| Buy 1 unit of asset | Enter into a forward contract to buy asset in 9 months for \$870 |
| Enter into forward contract to sell asset in 9 months for \$910 | |
| <i>Action in 4 months:</i> | <i>Action in 4 months:</i> |
| Receive \$40 of income on asset | Receive \$40 from 4-month investment |
| Use \$40 to repay first loan with interest | Pay income of \$40 on asset |
| <i>Action in 9 months:</i> | <i>Action in 9 months:</i> |
| Sell asset for \$910 | Receive \$886.60 from 9-month investment |
| Use \$886.60 to repay second loan with interest | Buy asset for \$870 Close out short position |
| Profit realized = \$23.40 | Profit realized = \$16.60 |

This is in agreement with our earlier calculation. Equation (5.2) applies to any investment asset that provides a known cash income.

If $F_0 > (S_0 - I)e^{rT}$, an arbitrageur can lock in a profit by buying the asset and shorting a forward contract on the asset; if $F_0 < (S_0 - I)e^{rT}$, an arbitrageur can lock in a profit by shorting the asset and taking a long position in a forward contract. If short sales are not possible, investors who own the asset will find it profitable to sell the asset and enter into long forward contracts.⁴

Example 5.2

Consider a 10-month forward contract on a stock when the stock price is \$50. We assume that the risk-free rate of interest (continuously compounded) is 8% per annum for all maturities. We also assume that dividends of \$0.75 per share are expected after 3 months, 6 months, and 9 months. The present value of the dividends, I , is

$$I = 0.75e^{-0.08 \times 3/12} + 0.75e^{-0.08 \times 6/12} + 0.75e^{-0.08 \times 9/12} = 2.162$$

The variable T is 10 months, so that the forward price, F_0 , from equation (5.2), is given by

$$F_0 = (50 - 2.162)e^{0.08 \times 10/12} = \$51.14$$

⁴ For another way of seeing that equation (5.2) is correct, consider the following strategy: buy one unit of the asset and enter into a short forward contract to sell it for F_0 at time T . This costs S_0 and is certain to lead to a cash inflow of F_0 at time T and an income with a present value of I . The initial outflow is S_0 . The present value of the inflows is $I + F_0e^{-rT}$. Hence, $S_0 = I + F_0e^{-rT}$, or equivalently $F_0 = (S_0 - I)e^{rT}$.

If the forward price were less than this, an arbitrageur would short the stock and buy forward contracts. If the forward price were greater than this, an arbitrageur would short forward contracts and buy the stock in the spot market.

5.6 KNOWN YIELD

We now consider the situation where the asset underlying a forward contract provides a known yield rather than a known cash income. This means that the income is known when expressed as a percentage of the asset's price at the time the income is paid. Suppose that an asset is expected to provide a yield of 5% per annum. This could mean that income is paid once a year and is equal to 5% of the asset price at the time it is paid, in which case the yield would be 5% with annual compounding. Alternatively, it could mean that income is paid twice a year and is equal to 2.5% of the asset price at the time it is paid, in which case the yield would be 5% per annum with semiannual compounding. In Section 4.2 we explained that we will normally measure interest rates with continuous compounding. Similarly, we will normally measure yields with continuous compounding. Formulas for translating a yield measured with one compounding frequency to a yield measured with another compounding frequency are the same as those given for interest rates in Section 4.2.

Define q as the average yield per annum on an asset during the life of a forward contract with continuous compounding. It can be shown (see Problem 5.20) that

$$F_0 = S_0 e^{(r-q)T} \quad (5.3)$$

Example 5.3

Consider a 6-month forward contract on an asset that is expected to provide income equal to 2% of the asset price once during a 6-month period. The risk-free rate of interest (with continuous compounding) is 10% per annum. The asset price is \$25. In this case, $S_0 = 25$, $r = 0.10$, and $T = 0.5$. The yield is 4% per annum with semiannual compounding. From equation (4.3), this is 3.96% per annum with continuous compounding. It follows that $q = 0.0396$, so that from equation (5.3) the forward price, F_0 , is given by

$$F_0 = 25e^{(0.10-0.0396) \times 0.5} = \$25.77$$

5.7 VALUING FORWARD CONTRACTS

The value of a forward contract at the time it is first entered into is zero. At a later stage, it may prove to have a positive or negative value. It is important for banks and other financial institutions to value the contract each day. (This is referred to as marking to market the contract.) Using the notation introduced earlier, we suppose K is the delivery price for a contract that was negotiated some time ago, the delivery date is T years from today, and r is the T -year risk-free interest rate. The variable F_0 is the forward price that would be applicable if we negotiated the contract today. In addition, we define f to be the value of forward contract today.

It is important to be clear about the meaning of the variables F_0 , K , and f . At the beginning of the life of the forward contract, the delivery price, K , is set equal to the

forward price and the value of the contract, f , is 0. As time passes, K stays the same (because it is part of the definition of the contract), but the forward price changes and the value of the contract becomes either positive or negative.

A general result, applicable to all long forward contracts (both those on investment assets and those on consumption assets), is

$$f = (F_0 - K)e^{-rT} \quad (5.4)$$

To see why equation (5.4) is correct, we use an argument analogous to the one we used for forward rate agreements in Section 4.7. We compare a long forward contract that has a delivery price of F_0 with an otherwise identical long forward contract that has a delivery price of K . The difference between the two is only in the amount that will be paid for the underlying asset at time T . Under the first contract, this amount is F_0 ; under the second contract, it is K . A cash outflow difference of $F_0 - K$ at time T translates to a difference of $(F_0 - K)e^{-rT}$ today. The contract with a delivery price F_0 is therefore less valuable than the contract with delivery price K by an amount $(F_0 - K)e^{-rT}$. The value of the contract that has a delivery price of F_0 is by definition zero. It follows that the value of the contract with a delivery price of K is $(F_0 - K)e^{-rT}$. This proves equation (5.4). Similarly, the value of a short forward contract with delivery price K is

$$(K - F_0)e^{-rT}$$

Example 5.4

A long forward contract on a non-dividend-paying stock was entered into some time ago. It currently has 6 months to maturity. The risk-free rate of interest (with continuous compounding) is 10% per annum, the stock price is \$25, and the delivery price is \$24. In this case, $S_0 = 25$, $r = 0.10$, $T = 0.5$, and $K = 24$. From equation (5.1), the 6-month forward price, F_0 , is given by

$$F_0 = 25e^{0.1 \times 0.5} = \$26.28$$

From equation (5.4), the value of the forward contract is

$$f = (26.28 - 24)e^{-0.1 \times 0.5} = \$2.17$$

Equation (5.4) shows that we can value a long forward contract on an asset by making the assumption that the price of the asset at the maturity of the forward contract equals the forward price F_0 . To see this, note that when we make that assumption, a long forward contract provides a payoff at time T of $F_0 - K$. This has a present value of $(F_0 - K)e^{-rT}$, which is the value of f in equation (5.4). Similarly, we can value a short forward contract on the asset by assuming that the current forward price of the asset is realized. These results are analogous to the result in Section 4.7 that we can value a forward rate agreement on the assumption that forward rates are realized.

Using equation (5.4) in conjunction with equation (5.1) gives the following expression for the value of a forward contract on an investment asset that provides no income

$$f = S_0 - Ke^{-rT} \quad (5.5)$$

Similarly, using equation (5.4) in conjunction with equation (5.2) gives the following

Business Snapshot 5.2 A Systems Error?

A foreign exchange trader working for a bank enters into a long forward contract to buy 1 million pounds sterling at an exchange rate of 1.5000 in 3 months. At the same time, another trader on the next desk takes a long position in 16 contracts for 3-month futures on sterling. The futures price is 1.5000 and each contract is on 62,500 pounds. The positions taken by the forward and futures traders are therefore the same. Within minutes of the positions being taken the forward and the futures prices both increase to 1.5040. The bank's systems show that the futures trader has made a profit of \$4,000, while the forward trader has made a profit of only \$3,900. The forward trader immediately calls the bank's systems department to complain. Does the forward trader have a valid complaint?

The answer is no! The daily settlement of futures contracts ensures that the futures trader realizes an almost immediate profit corresponding to the increase in the futures price. If the forward trader closed out the position by entering into a short contract at 1.5040, the forward trader would have contracted to buy 1 million pounds at 1.5000 in 3 months and sell 1 million pounds at 1.5040 in 3 months. This would lead to a \$4,000 profit—but in 3 months, not today. The forward trader's profit is the present value of \$4,000. This is consistent with equation (5.4).

The forward trader can gain some consolation from the fact that gains and losses are treated symmetrically. If the forward/futures prices dropped to 1.4960 instead of rising to 1.5040, then the futures trader would take a loss of \$4,000 while the forward trader would take a loss of only \$3,900.

expression for the value of a long forward contract on an investment asset that provides a known income with present value I :

$$f = S_0 - I - Ke^{-rT} \quad (5.6)$$

Finally, using equation (5.4) in conjunction with equation (5.3) gives the following expression for the value of a long forward contract on an investment asset that provides a known yield at rate q :

$$f = S_0e^{-qT} - Ke^{-rT} \quad (5.7)$$

When a futures price changes, the gain or loss on a futures contract is calculated as the change in the futures price multiplied by the size of the position. This gain is realized almost immediately because of the way futures contracts are settled daily. Equation (5.4) shows that, when a forward price changes, the gain or loss is the present value of the change in the forward price multiplied by the size of the position. The difference between the gain/loss on forward and futures contracts can cause confusion on a foreign exchange trading desk (see Business Snapshot 5.2).

5.8 ARE FORWARD PRICES AND FUTURES PRICES EQUAL?

Technical Note 24 at www.rotman.utoronto.ca/~hull/TechnicalNotes provides an arbitrage argument to show that, when the short-term risk-free interest rate is constant,

the forward price for a contract with a certain delivery date is in theory the same as the futures price for a contract with that delivery date. The argument in the appendix can be extended to cover situations where the interest rate is a known function of time.

When interest rates vary unpredictably (as they do in the real world), forward and futures prices are in theory no longer the same. We can get a sense of the nature of the relationship by considering the situation where the price of the underlying asset, S , is strongly positively correlated with interest rates. When S increases, an investor who holds a long futures position makes an immediate gain because of the daily settlement procedure. The positive correlation indicates that it is likely that interest rates have also increased. The gain will therefore tend to be invested at a higher than average rate of interest. Similarly, when S decreases, the investor will incur an immediate loss. This loss will tend to be financed at a lower than average rate of interest. An investor holding a forward contract rather than a futures contract is not affected in this way by interest rate movements. It follows that a long futures contract will be slightly more attractive than a similar long forward contract. Hence, when S is strongly positively correlated with interest rates, futures prices will tend to be slightly higher than forward prices. When S is strongly negatively correlated with interest rates, a similar argument shows that forward prices will tend to be slightly higher than futures prices.

The theoretical differences between forward and futures prices for contracts that last only a few months are in most circumstances sufficiently small to be ignored. In practice, there are a number of factors not reflected in theoretical models that may cause forward and futures prices to be different. These include taxes, transactions costs, and the treatment of margins. The risk that the counterparty will default may be less in the case of a futures contract because of the role of the exchange clearinghouse. Also, in some instances, futures contracts are more liquid and easier to trade than forward contracts. Despite all these points, for most purposes it is reasonable to assume that forward and futures prices are the same. This is the assumption we will usually make in this book. We will use the symbol F_0 to represent both the futures price and the forward price of an asset today.

One exception to the rule that futures and forward contracts can be assumed to be the same concerns Eurodollar futures. This will be discussed in Section 6.3.

5.9 FUTURES PRICES OF STOCK INDICES

We introduced futures on stock indices in Section 3.5 and showed how a stock index futures contract is a useful tool in managing equity portfolios. Table 3.3 shows futures prices for a number of different indices. We are now in a position to consider how index futures prices are determined.

A stock index can usually be regarded as the price of an investment asset that pays dividends.⁵ The investment asset is the portfolio of stocks underlying the index, and the dividends paid by the investment asset are the dividends that would be received by the holder of this portfolio. It is usually assumed that the dividends provide a known yield rather than a known cash income. If q is the dividend yield rate, equation (5.3) gives the futures price, F_0 , as

$$F_0 = S_0 e^{(r-q)T} \quad (5.8)$$

⁵ Occasionally this is not the case: see Business Snapshot 5.3.

Business Snapshot 5.3 The CME Nikkei 225 Futures Contract

The arguments in this chapter on how index futures prices are determined require that the index be the value of an investment asset. This means that it must be the value of a portfolio of assets that can be traded. The asset underlying the Chicago Mercantile Exchange's futures contract on the Nikkei 225 Index does not qualify, and the reason why is quite subtle. Suppose S is the value of the Nikkei 225 Index. This is the value of a portfolio of 225 Japanese stocks measured in yen. The variable underlying the CME futures contract on the Nikkei 225 has a *dollar value* of $5S$. In other words, the futures contract takes a variable that is measured in yen and treats it as though it is dollars.

We cannot invest in a portfolio whose value will always be $5S$ dollars. The best we can do is to invest in one that is always worth $5S$ yen or in one that is always worth $5QS$ dollars, where Q is the dollar value of 1 yen. The variable $5S$ dollars is not, therefore, the price of an investment asset and equation (5.8) does not apply.

CME's Nikkei 225 futures contract is an example of a *quanto*. A quanto is a derivative where the underlying asset is measured in one currency and the payoff is in another currency. Quantos will be discussed further in Chapter 29.

This shows that the futures price increases at rate $r - q$ with the maturity of the futures contract. In Table 3.3, the December futures settlement price of the S&P 500 is about 0.76% less than the June settlement price. This indicates that, on May 26, 2010, the short-term risk-free rate r was less than the dividend yield q by about 1.52% per year.

Example 5.5

Consider a 3-month futures contract on an index. Suppose that the stocks underlying the index provide a dividend yield of 1% per annum, that the current value of the index is 1,300, and that the continuously compounded risk-free interest rate is 5% per annum. In this case, $r = 0.05$, $S_0 = 1,300$, $T = 0.25$, and $q = 0.01$. Hence, the futures price, F_0 , is given by

$$F_0 = 1,300e^{(0.05-0.01) \times 0.25} = \$1,313.07$$

In practice, the dividend yield on the portfolio underlying an index varies week by week throughout the year. For example, a large proportion of the dividends on the NYSE stocks are paid in the first week of February, May, August, and November each year. The chosen value of q should represent the average annualized dividend yield during the life of the contract. The dividends used for estimating q should be those for which the ex-dividend date is during the life of the futures contract.

Index Arbitrage

If $F_0 > S_0e^{(r-q)T}$, profits can be made by buying the stocks underlying the index at the spot price (i.e., for immediate delivery) and shorting futures contracts. If $F_0 < S_0e^{(r-q)T}$, profits can be made by doing the reverse—that is, shorting or selling the stocks underlying the index and taking a long position in futures contracts. These strategies are known as *index arbitrage*. When $F_0 < S_0e^{(r-q)T}$, index arbitrage is often done by a pension fund that owns an indexed portfolio of stocks. When $F_0 > S_0e^{(r-q)T}$, it might be

Business Snapshot 5.4 Index Arbitrage in October 1987

To do index arbitrage, a trader must be able to trade both the index futures contract and the portfolio of stocks underlying the index very quickly at the prices quoted in the market. In normal market conditions this is possible using program trading, and the relationship in equation (5.8) holds well. Examples of days when the market was anything but normal are October 19 and 20 of 1987. On what is termed “Black Monday,” October 19, 1987, the market fell by more than 20%, and the 604 million shares traded on the New York Stock Exchange easily exceeded all previous records. The exchange’s systems were overloaded, and orders placed to buy or sell shares on that day could be delayed by up to two hours before being executed.

For most of October 19, 1987, futures prices were at a significant discount to the underlying index. For example, at the close of trading the S&P 500 Index was at 225.06 (down 57.88 on the day), whereas the futures price for December delivery on the S&P 500 was 201.50 (down 80.75 on the day). This was largely because the delays in processing orders made index arbitrage impossible. On the next day, Tuesday, October 20, 1987, the New York Stock Exchange placed temporary restrictions on the way in which program trading could be done. This also made index arbitrage very difficult and the breakdown of the traditional linkage between stock indices and stock index futures continued. At one point the futures price for the December contract was 18% less than the S&P 500 Index. However, after a few days the market returned to normal, and the activities of arbitrageurs ensured that equation (5.8) governed the relationship between futures and spot prices of indices.


done by a corporation holding short-term money market investments. For indices involving many stocks, index arbitrage is sometimes accomplished by trading a relatively small representative sample of stocks whose movements closely mirror those of the index. Usually index arbitrage is implemented through *program trading*. This involves using a computer system to generate the trades.

Most of the time the activities of arbitrageurs ensure that equation (5.8) holds, but occasionally arbitrage is impossible and the futures price does get out of line with the spot price (see Business Snapshot 5.4).

5.10 FORWARD AND FUTURES CONTRACTS ON CURRENCIES

We now move on to consider forward and futures foreign currency contracts from the perspective of a US investor. The underlying asset is one unit of the foreign currency. We will therefore define the variable S_0 as the current spot price in US dollars of one unit of the foreign currency and F_0 as the forward or futures price in US dollars of one unit of the foreign currency. This is consistent with the way we have defined S_0 and F_0 for other assets underlying forward and futures contracts. However, as mentioned in Section 2.11, it does not necessarily correspond to the way spot and forward exchange rates are quoted. For major exchange rates other than the British pound, euro, Australian dollar, and New Zealand dollar, a spot or forward exchange rate is normally quoted as the number of units of the currency that are equivalent to one US dollar.

Figure 5.1 Two ways of converting 1,000 units of a foreign currency to dollars at time T . Here, S_0 is spot exchange rate, F_0 is forward exchange rate, and r and r_f are the dollar and foreign risk-free rates.



A foreign currency has the property that the holder of the currency can earn interest at the risk-free interest rate prevailing in the foreign country. For example, the holder can invest the currency in a foreign-denominated bond. We define r_f as the value of the foreign risk-free interest rate when money is invested for time T . The variable r is the US dollar risk-free rate when money is invested for this period of time.

The relationship between F_0 and S_0 is

$$F_0 = S_0 e^{(r-r_f)T} \tag{5.9}$$

This is the well-known interest rate parity relationship from international finance. The reason it is true is illustrated in Figure 5.1. Suppose that an individual starts with 1,000 units of the foreign currency. There are two ways it can be converted to dollars at time T . One is by investing it for T years at r_f and entering into a forward contract to sell the proceeds for dollars at time T . This generates $1,000e^{r_f T} F_0$ dollars. The other is by exchanging the foreign currency for dollars in the spot market and investing the proceeds for T years at rate r . This generates $1,000S_0e^{rT}$ dollars. In the absence of arbitrage opportunities, the two strategies must give the same result. Hence,

$$1,000e^{r_f T} F_0 = 1,000S_0e^{rT}$$

so that

$$F_0 = S_0 e^{(r-r_f)T}$$

Example 5.6

Suppose that the 2-year interest rates in Australia and the United States are 5% and 7%, respectively, and the spot exchange rate between the Australian dollar (AUD) and the US dollar (USD) is 0.6200 USD per AUD. From equation (5.9), the 2-year forward exchange rate should be

$$0.62e^{(0.07-0.05)\times 2} = 0.6453$$

Suppose first that the 2-year forward exchange rate is less than this, say 0.6300. An arbitrageur can:

1. Borrow 1,000 AUD at 5% per annum for 2 years, convert to 620 USD and invest the USD at 7% (both rates are continuously compounded).
2. Enter into a forward contract to buy 1,105.17 AUD for $1,105.17 \times 0.63 = 696.26$ USD.

The 620 USD that are invested at 7% grow to $620e^{0.07 \times 2} = 713.17$ USD in 2 years. Of this, 696.26 USD are used to purchase 1,105.17 AUD under the terms of the forward contract. This is exactly enough to repay principal and interest on the 1,000 AUD that are borrowed ($1,000e^{0.05 \times 2} = 1,105.17$). The strategy therefore gives rise to a riskless profit of $713.17 - 696.26 = 16.91$ USD. (If this does not sound very exciting, consider following a similar strategy where you borrow 100 million AUD!)

Suppose next that the 2-year forward rate is 0.6600 (greater than the 0.6453 value given by equation (5.9)). An arbitrageur can:

1. Borrow 1,000 USD at 7% per annum for 2 years, convert to $1,000/0.6200 = 1,612.90$ AUD, and invest the AUD at 5%.
2. Enter into a forward contract to sell 1,782.53 AUD for $1,782.53 \times 0.66 = 1,176.47$ USD.

The 1,612.90 AUD that are invested at 5% grow to $1,612.90e^{0.05 \times 2} = 1,782.53$ AUD in 2 years. The forward contract has the effect of converting this to 1,176.47 USD. The amount needed to payoff the USD borrowings is $1,000e^{0.07 \times 2} = 1,150.27$ USD. The strategy therefore gives rise to a riskless profit of $1,176.47 - 1,150.27 = 26.20$ USD.

Table 5.4 shows currency futures quotes on May 26, 2010. The quotes are US dollars per unit of the foreign currency. (In the case of the Japanese yen, the quote is US dollars per 100 yen; for the Mexican peso, it is US dollars per 10 pesos.) This is the usual quotation convention for futures contracts. Equation (5.9) applies with r equal to the US risk-free rate and r_f equal to the foreign risk-free rate.

On May 26, 2010, short-term interest rates on the Japanese yen, British pound, Swiss franc, and euro were lower than the short-term interest rate on the US dollar. This corresponds to the $r > r_f$ situation and explains why futures prices for these currencies increase with maturity in Table 5.4. For the Australian dollar, Canadian dollar, and Mexican peso, short-term interest rates were higher than in the United States. This corresponds to the $r_f > r$ situation and explains why the futures prices of these currencies decrease with maturity.

Example 5.7

In Table 5.4, the September settlement price for the Australian dollar is about 1% lower than the June settlement price. This indicates that the futures prices are increasing at about 4% per year with maturity. From equation (5.9) this is an estimate of the amount by which short-term Australian LIBOR interest rates exceeded short-term US LIBOR interest rates on May 26, 2010.

Table 5.4 Currency futures quotes as reported by exchanges on May 26, 2010.

| | <i>Open</i> | <i>High</i> | <i>Low</i> | <i>Settlement</i> | <i>Change</i> | <i>Volume</i> | <i>Open interest</i> |
|--|-------------|-------------|------------|-------------------|---------------|---------------|----------------------|
| Australian dollar, \$100,000, USD per AUD (CME Group) | | | | | | | |
| June 2010 | 0.8266 | 0.8373 | 0.8171 | 0.8236 | 0.0062 | 146,968 | 101,448 |
| Sept. 2010 | 0.8165 | 0.8285 | 0.8090 | 0.8152 | 0.0059 | 860 | 4,650 |
| British pound, \$62,500, USD per GBP (CME Group) | | | | | | | |
| June 2010 | 1.4429 | 1.4446 | 1.4330 | 1.4411 | 0.0046 | 105,256 | 140,369 |
| Sep 2010 | 1.4432 | 1.4450 | 1.4339 | 1.4416 | 0.0044 | 1,448 | 10,811 |
| Canadian Dollar, \$100,000, USD per CAD (CME Group) | | | | | | | |
| June 2010 | 0.9384 | 0.9452 | 0.9305 | 0.9393 | 0.0097 | 126,564 | 111,697 |
| Sept. 2010 | 0.9378 | 0.9449 | 0.9309 | 0.9392 | 0.0094 | 2,264 | 8,647 |
| Euro, 125,000 euros, USD per EUR (CME Group) | | | | | | | |
| June 2010 | 1.2371 | 1.2380 | 1.2170 | 1.2201 | -0.0117 | 400,948 | 267,552 |
| Sept. 2010 | 1.2388 | 1.2388 | 1.2186 | 1.2216 | -0.0118 | 4,702 | 13,939 |
| Japanese Yen, 12,500,000 yen, USD per 100 yen (CME Group) | | | | | | | |
| June 2010 | 1.1073 | 1.1136 | 1.1031 | 1.1108 | 0.0009 | 172,240 | 135,113 |
| Sept. 2010 | 1.1100 | 1.1156 | 1.1053 | 1.1129 | 0.0005 | 2,098 | 5,506 |
| Mexican Peso, MXN500,000, USD per 10MXN (CME Group) | | | | | | | |
| June 2010 | 0.76800 | 0.77175 | 0.76000 | 0.76375 | 0.00225 | 37,776 | 84,207 |
| Sept. 2010 | 0.76375 | 0.76375 | 0.75275 | 0.75625 | 0.00225 | 107 | 727 |
| Swiss Franc, CHF125,000, USD per CHF (CME Group) | | | | | | | |
| June 2010 | 0.8661 | 0.8688 | 0.8613 | 0.8629 | -0.0012 | 68,960 | 46,212 |
| Sept. 2010 | 0.8693 | 0.8713 | 0.8644 | 0.8657 | -0.0017 | 1,817 | 1,938 |

A Foreign Currency as an Asset Providing a Known Yield

Equation (5.9) is identical to equation (5.3) with q replaced by r_f . This is not a coincidence. A foreign currency can be regarded as an investment asset paying a known yield. The yield is the risk-free rate of interest in the foreign currency.

To understand this, we note that the value of interest paid in a foreign currency depends on the value of the foreign currency. Suppose that the interest rate on British pounds is 5% per annum. To a US investor the British pound provides an income equal to 5% of the value of the British pound per annum. In other words it is an asset that provides a yield of 5% per annum.

5.11 FUTURES ON COMMODITIES

We now move on to consider futures contracts on commodities. First we look at the futures prices of commodities that are investment assets such as gold and silver.⁶ We then go on to examine the futures prices of consumption assets.

⁶ Recall that, for an asset to be an investment asset, it need not be held solely for investment purposes. What is required is that some individuals hold it for investment purposes and that these individuals be prepared to sell their holdings and go long forward contracts, if the latter look more attractive. This explains why silver, although it has significant industrial uses, is an investment asset.

Income and Storage Costs

As explained in Business Snapshot 3.1, the hedging strategies of gold producers leads to a requirement on the part of investment banks to borrow gold. Gold owners such as central banks charge interest in the form of what is known as the *gold lease rate* when they lend gold. The same is true of silver. Gold and silver can therefore provide income to the holder. Like other commodities they also have storage costs.

Equation (5.1) shows that, in the absence of storage costs and income, the forward price of a commodity that is an investment asset is given by

$$F_0 = S_0 e^{rT} \quad (5.10)$$

Storage costs can be treated as negative income. If U is the present value of all the storage costs, net of income, during the life of a forward contract, it follows from equation (5.2) that

$$F_0 = (S_0 + U)e^{rT} \quad (5.11)$$

Example 5.8

Consider a 1-year futures contract on an investment asset that provides no income. It costs \$2 per unit to store the asset, with the payment being made at the end of the year. Assume that the spot price is \$450 per unit and the risk-free rate is 7% per annum for all maturities. This corresponds to $r = 0.07$, $S_0 = 450$, $T = 1$, and

$$U = 2e^{-0.07 \times 1} = 1.865$$

From equation (5.11), the theoretical futures price, F_0 , is given by

$$F_0 = (450 + 1.865)e^{0.07 \times 1} = \$484.63$$

If the actual futures price is greater than 484.63, an arbitrageur can buy the asset and short 1-year futures contracts to lock in a profit. If the actual futures price is less than 484.63, an investor who already owns the asset can improve the return by selling the asset and buying futures contracts.

If the storage costs (net of income) incurred at any time are proportional to the price of the commodity, they can be treated as negative yield. In this case, from equation (5.3),

$$F_0 = S_0 e^{(r+u)T} \quad (5.12)$$

where u denotes the storage costs per annum as a proportion of the spot price net of any yield earned on the asset.

Consumption Commodities

Commodities that are consumption assets rather than investment assets usually provide no income, but can be subject to significant storage costs. We now review the arbitrage strategies used to determine futures prices from spot prices carefully.⁷

⁷ For some commodities the spot price depends on the delivery location. We assume that the delivery location for spot and futures are the same.

Suppose that, instead of equation (5.11), we have

$$F_0 > (S_0 + U)e^{rT} \quad (5.13)$$

To take advantage of this opportunity, an arbitrageur can implement the following strategy:

1. Borrow an amount $S_0 + U$ at the risk-free rate and use it to purchase one unit of the commodity and to pay storage costs.
2. Short a futures contract on one unit of the commodity.

If we regard the futures contract as a forward contract, so that there is no daily settlement, this strategy leads to a profit of $F_0 - (S_0 + U)e^{rT}$ at time T . There is no problem in implementing the strategy for any commodity. However, as arbitrageurs do so, there will be a tendency for S_0 to increase and F_0 to decrease until equation (5.13) is no longer true. We conclude that equation (5.13) cannot hold for any significant length of time.

Suppose next that

$$F_0 < (S_0 + U)e^{rT} \quad (5.14)$$

When the commodity is an investment asset, we can argue that many investors hold the commodity solely for investment. When they observe the inequality in equation (5.14), they will find it profitable to do the following:

1. Sell the commodity, save the storage costs, and invest the proceeds at the risk-free interest rate.
2. Take a long position in a futures contract.

The result is a riskless profit at maturity of $(S_0 + U)e^{rT} - F_0$ relative to the position the investors would have been in if they had held the commodity. It follows that equation (5.14) cannot hold for long. Because neither equation (5.13) nor (5.14) can hold for long, we must have $F_0 = (S_0 + U)e^{rT}$.

This argument cannot be used for a commodity that is a consumption asset rather than an investment asset. Individuals and companies who own a consumption commodity usually plan to use it in some way. They are reluctant to sell the commodity in the spot market and buy forward or futures contracts, because forward and futures contracts cannot be used in a manufacturing process or consumed in some other way. There is therefore nothing to stop equation (5.14) from holding, and all we can assert for a consumption commodity is

$$F_0 \leq (S_0 + U)e^{rT} \quad (5.15)$$

If storage costs are expressed as a proportion u of the spot price, the equivalent result is

$$F_0 \leq S_0 e^{(r+u)T} \quad (5.16)$$

Convenience Yields

We do not necessarily have equality in equations (5.15) and (5.16) because users of a consumption commodity may feel that ownership of the physical commodity provides benefits that are not obtained by holders of futures contracts. For example, an oil

refiner is unlikely to regard a futures contract on crude oil in the same way as crude oil held in inventory. The crude oil in inventory can be an input to the refining process, whereas a futures contract cannot be used for this purpose. In general, ownership of the physical asset enables a manufacturer to keep a production process running and perhaps profit from temporary local shortages. A futures contract does not do the same. The benefits from holding the physical asset are sometimes referred to as the *convenience yield* provided by the commodity. If the dollar amount of storage costs is known and has a present value U , then the convenience yield y is defined such that

$$F_0 e^{yT} = (S_0 + U) e^{rT}$$

If the storage costs per unit are a constant proportion, u , of the spot price, then y is defined so that

$$F_0 e^{yT} = S_0 e^{(r+u)T}$$

or

$$F_0 = S_0 e^{(r+u-y)T} \quad (5.17)$$

The convenience yield simply measures the extent to which the left-hand side is less than the right-hand side in equation (5.15) or (5.16). For investment assets the convenience yield must be zero; otherwise, there are arbitrage opportunities. Table 2.2 in Chapter 2 shows that, on May 26, 2010, the futures price of soybeans decreased as the maturity of the contract increased from July 2010 to November 2010. This pattern suggests that the convenience yield, y , is greater than $r + u$.

The convenience yield reflects the market's expectations concerning the future availability of the commodity. The greater the possibility that shortages will occur, the higher the convenience yield. If users of the commodity have high inventories, there is very little chance of shortages in the near future and the convenience yield tends to be low. If inventories are low, shortages are more likely and the convenience yield is usually higher.

5.12 THE COST OF CARRY

The relationship between futures prices and spot prices can be summarized in terms of the *cost of carry*. This measures the storage cost plus the interest that is paid to finance the asset less the income earned on the asset. For a non-dividend-paying stock, the cost of carry is r , because there are no storage costs and no income is earned; for a stock index, it is $r - q$, because income is earned at rate q on the asset. For a currency, it is $r - r_f$; for a commodity that provides income at rate q and requires storage costs at rate u , it is $r - q + u$; and so on.

Define the cost of carry as c . For an investment asset, the futures price is

$$F_0 = S_0 e^{cT} \quad (5.18)$$

For a consumption asset, it is

$$F_0 = S_0 e^{(c-y)T} \quad (5.19)$$

where y is the convenience yield.

5.13 DELIVERY OPTIONS

Whereas a forward contract normally specifies that delivery is to take place on a particular day, a futures contract often allows the party with the short position to choose to deliver at any time during a certain period. (Typically the party has to give a few days' notice of its intention to deliver.) The choice introduces a complication into the determination of futures prices. Should the maturity of the futures contract be assumed to be the beginning, middle, or end of the delivery period? Even though most futures contracts are closed out prior to maturity, it is important to know when delivery would have taken place in order to calculate the theoretical futures price.

If the futures price is an increasing function of the time to maturity, it can be seen from equation (5.19) that $c > y$, so that the benefits from holding the asset (including convenience yield and net of storage costs) are less than the risk-free rate. It is usually optimal in such a case for the party with the short position to deliver as early as possible, because the interest earned on the cash received outweighs the benefits of holding the asset. As a rule, futures prices in these circumstances should be calculated on the basis that delivery will take place at the beginning of the delivery period. If futures prices are decreasing as time to maturity increases ($c < y$), the reverse is true. It is then usually optimal for the party with the short position to deliver as late as possible, and futures prices should, as a rule, be calculated on this assumption.

5.14 FUTURES PRICES AND EXPECTED FUTURE SPOT PRICES

We refer to the market's average opinion about what the spot price of an asset will be at a certain future time as the *expected spot price* of the asset at that time. Suppose that it is now June and the September futures price of corn is 350 cents. It is interesting to ask what the expected spot price of corn in September is. Is it less than 350 cents, greater than 350 cents, or exactly equal to 350 cents? As illustrated in Figure 2.1, the futures price converges to the spot price at maturity. If the expected spot price is less than 350 cents, the market must be expecting the September futures price to decline, so that traders with short positions gain and traders with long positions lose. If the expected spot price is greater than 350 cents, the reverse must be true. The market must be expecting the September futures price to increase, so that traders with long positions gain while those with short positions lose.

Keynes and Hicks

Economists John Maynard Keynes and John Hicks argued that, if hedgers tend to hold short positions and speculators tend to hold long positions, the futures price of an asset will be below the expected spot price.⁸ This is because speculators require compensation for the risks they are bearing. They will trade only if they can expect to make money on average. Hedgers will lose money on average, but they are likely to be prepared to accept this because the futures contract reduces their risks. If hedgers tend to hold long positions while speculators hold short positions, Keynes and Hicks argued that the futures price will be above the expected spot price for a similar reason.

⁸ See: J. M. Keynes, *A Treatise on Money*. London: Macmillan, 1930; and J. R. Hicks, *Value and Capital*. Oxford: Clarendon Press, 1939.

Risk and Return

The modern approach to explaining the relationship between futures prices and expected spot prices is based on the relationship between risk and expected return in the economy. In general, the higher the risk of an investment, the higher the expected return demanded by an investor. The capital asset pricing model, which is explained in the appendix to Chapter 3, shows that there are two types of risk in the economy: systematic and nonsystematic. Nonsystematic risk should not be important to an investor. It can be almost completely eliminated by holding a well-diversified portfolio. An investor should not therefore require a higher expected return for bearing nonsystematic risk. Systematic risk, in contrast, cannot be diversified away. It arises from a correlation between returns from the investment and returns from the whole stock market. An investor generally requires a higher expected return than the risk-free interest rate for bearing positive amounts of systematic risk. Also, an investor is prepared to accept a lower expected return than the risk-free interest rate when the systematic risk in an investment is negative.

The Risk in a Futures Position

Let us consider a speculator who takes a long position in a futures contract that lasts for T years in the hope that the spot price of the asset will be above the futures price at the end of the life of the futures contract. We ignore daily settlement and assume that the futures contract can be treated as a forward contract. We suppose that the speculator puts the present value of the futures price into a risk-free investment while simultaneously taking a long futures position. The proceeds of the risk-free investment are used to buy the asset on the delivery date. The asset is then immediately sold for its market price. The cash flows to the speculator are as follows:

Today: $-F_0e^{-rT}$

End of futures contract: $+S_T$

where F_0 is the futures price today, S_T is the price of the asset at time T at the end of the futures contract, and r is the risk-free return on funds invested for time T .

How do we value this investment? The discount rate we should use for the expected cash flow at time T equals an investor's required return on the investment. Suppose that k is an investor's required return for this investment. The present value of this investment is

$$-F_0e^{-rT} + E(S_T)e^{-kT}$$

where E denotes expected value. We can assume that all investments in securities markets are priced so that they have zero net present value. This means that

$$-F_0e^{-rT} + E(S_T)e^{-kT} = 0$$

or

$$F_0 = E(S_T)e^{(r-k)T} \quad (5.20)$$

As we have just discussed, the returns investors require on an investment depend on its systematic risk. The investment we have been considering is in essence an investment in the asset underlying the futures contract. If the returns from this asset are uncorrelated

Table 5.5 Relationship between futures price and expected future spot price.

| <i>Underlying asset</i> | <i>Relationship of expected return k from asset to risk-free rate r</i> | <i>Relationship of futures price F to expected future spot price $E(S_T)$</i> |
|--------------------------|---|---|
| No systematic risk | $k = r$ | $F_0 = E(S_T)$ |
| Positive systematic risk | $k > r$ | $F_0 < E(S_T)$ |
| Negative systematic risk | $k < r$ | $F_0 > E(S_T)$ |

with the stock market, the correct discount rate to use is the risk-free rate r , so we should set $k = r$. Equation (5.20) then gives

$$F_0 = E(S_T)$$

This shows that the futures price is an unbiased estimate of the expected future spot price when the return from the underlying asset is uncorrelated with the stock market.

If the return from the asset is positively correlated with the stock market, $k > r$ and equation (5.20) leads to $F_0 < E(S_T)$. This shows that, when the asset underlying the futures contract has positive systematic risk, we should expect the futures price to understate the expected future spot price. An example of an asset that has positive systematic risk is a stock index. The expected return of investors on the stocks underlying an index is generally more than the risk-free rate, r . The dividends provide a return of q . The expected increase in the index must therefore be more than $r - q$. Equation (5.8) is therefore consistent with the prediction that the futures price understates the expected future stock price for a stock index.

If the return from the asset is negatively correlated with the stock market, $k < r$ and equation (5.20) gives $F_0 > E(S_T)$. This shows that, when the asset underlying the futures contract has negative systematic risk, we should expect the futures price to overstate the expected future spot price.

These results are summarized in Table 5.5.

Normal Backwardation and Contango

When the futures price is below the expected future spot price, the situation is known as *normal backwardation*; and when the futures price is above the expected future spot price, the situation is known as *contango*. However, it should be noted that sometimes these terms are used to refer to whether the futures price is below or above the current spot price, rather than the expected future spot price.

SUMMARY

For most purposes, the futures price of a contract with a certain delivery date can be considered to be the same as the forward price for a contract with the same delivery date. It can be shown that in theory the two should be exactly the same when interest rates are perfectly predictable.

Table 5.6 Summary of results for a contract with time to maturity T on an investment asset with price S_0 when the risk-free interest rate for a T -year period is r .

| <i>Asset</i> | <i>Forward/futures price</i> | <i>Value of long forward contract with delivery price K</i> |
|--|------------------------------|--|
| Provides no income: | $S_0 e^{rT}$ | $S_0 - K e^{-rT}$ |
| Provides known income with present value I : | $(S_0 - I) e^{rT}$ | $S_0 - I - K e^{-rT}$ |
| Provides known yield q : | $S_0 e^{(r-q)T}$ | $S_0 e^{-qT} - K e^{-rT}$ |

For the purposes of understanding futures (or forward) prices, it is convenient to divide futures contracts into two categories: those in which the underlying asset is held for investment by a significant number of investors and those in which the underlying asset is held primarily for consumption purposes.

In the case of investment assets, we have considered three different situations:

1. The asset provides no income.
2. The asset provides a known dollar income.
3. The asset provides a known yield.

The results are summarized in Table 5.6. They enable futures prices to be obtained for contracts on stock indices, currencies, gold, and silver. Storage costs can be treated as negative income.

In the case of consumption assets, it is not possible to obtain the futures price as a function of the spot price and other observable variables. Here the parameter known as the asset's convenience yield becomes important. It measures the extent to which users of the commodity feel that ownership of the physical asset provides benefits that are not obtained by the holders of the futures contract. These benefits may include the ability to profit from temporary local shortages or the ability to keep a production process running. We can obtain an upper bound for the futures price of consumption assets using arbitrage arguments, but we cannot nail down an equality relationship between futures and spot prices.

The concept of cost of carry is sometimes useful. The cost of carry is the storage cost of the underlying asset plus the cost of financing it minus the income received from it. In the case of investment assets, the futures price is greater than the spot price by an amount reflecting the cost of carry. In the case of consumption assets, the futures price is greater than the spot price by an amount reflecting the cost of carry net of the convenience yield.

If we assume the capital asset pricing model is true, the relationship between the futures price and the expected future spot price depends on whether the return on the asset is positively or negatively correlated with the return on the stock market. Positive correlation will tend to lead to a futures price lower than the expected future spot price, whereas negative correlation will tend to lead to a futures price higher than the expected future spot price. Only when the correlation is zero will the theoretical futures price be equal to the expected future spot price.

FURTHER READING

- Cox, J. C., J. E. Ingersoll, and S. A. Ross. "The Relation between Forward Prices and Futures Prices," *Journal of Financial Economics*, 9 (December 1981): 321–46.
- Ghon, R. S. and R. P. Chang. "Intra-day Arbitrage in Foreign Exchange and Eurocurrency Markets," *Journal of Finance*, 47, 1 (1992): 363–380.
- Jarrow, R. A., and G. S. Oldfield. "Forward Contracts and Futures Contracts," *Journal of Financial Economics*, 9 (December 1981): 373–82.
- Kane, E. J. "Market Incompleteness and Divergences between Forward and Futures Interest Rates," *Journal of Finance*, 35 (May 1980): 221–34.
- Pindyck, R. S. "Inventories and the Short-Run Dynamics of Commodity Prices," *Rand Journal of Economics*, 25, 1 (1994): 141–159.
- Richard, S., and S. Sundaresan. "A Continuous-Time Model of Forward and Futures Prices in a Multigood Economy," *Journal of Financial Economics*, 9 (December 1981): 347–72.
- Routledge, B. R., D. J. Seppi, and C. S. Spatt. "Equilibrium Forward Curves for Commodities," *Journal of Finance*, 55, 3 (2000) 1297–1338.

Practice Questions (Answers in Solutions Manual)

- 5.1. Explain what happens when an investor shorts a certain share.
- 5.2. What is the difference between the forward price and the value of a forward contract?
- 5.3. Suppose that you enter into a 6-month forward contract on a non-dividend-paying stock when the stock price is \$30 and the risk-free interest rate (with continuous compounding) is 12% per annum. What is the forward price?
- 5.4. A stock index currently stands at 350. The risk-free interest rate is 8% per annum (with continuous compounding) and the dividend yield on the index is 4% per annum. What should the futures price for a 4-month contract be?
- 5.5. Explain carefully why the futures price of gold can be calculated from its spot price and other observable variables whereas the futures price of copper cannot.
- 5.6. Explain carefully the meaning of the terms *convenience yield* and *cost of carry*. What is the relationship between futures price, spot price, convenience yield, and cost of carry?
- 5.7. Explain why a foreign currency can be treated as an asset providing a known yield.
- 5.8. Is the futures price of a stock index greater than or less than the expected future value of the index? Explain your answer.
- 5.9. A 1-year long forward contract on a non-dividend-paying stock is entered into when the stock price is \$40 and the risk-free rate of interest is 10% per annum with continuous compounding.
 - (a) What are the forward price and the initial value of the forward contract?
 - (b) Six months later, the price of the stock is \$45 and the risk-free interest rate is still 10%. What are the forward price and the value of the forward contract?
- 5.10. The risk-free rate of interest is 7% per annum with continuous compounding, and the dividend yield on a stock index is 3.2% per annum. The current value of the index is 150. What is the 6-month futures price?

- 5.11. Assume that the risk-free interest rate is 9% per annum with continuous compounding and that the dividend yield on a stock index varies throughout the year. In February, May, August, and November, dividends are paid at a rate of 5% per annum. In other months, dividends are paid at a rate of 2% per annum. Suppose that the value of the index on July 31 is 1,300. What is the futures price for a contract deliverable in December 31 of the same year?
- 5.12. Suppose that the risk-free interest rate is 10% per annum with continuous compounding and that the dividend yield on a stock index is 4% per annum. The index is standing at 400, and the futures price for a contract deliverable in four months is 405. What arbitrage opportunities does this create?
- 5.13. Estimate the difference between short-term interest rates in Mexico and the United States on May 26, 2010, from the information in Table 5.4.
- 5.14. The 2-month interest rates in Switzerland and the United States are, respectively, 2% and 5% per annum with continuous compounding. The spot price of the Swiss franc is \$0.8000. The futures price for a contract deliverable in 2 months is \$0.8100. What arbitrage opportunities does this create?
- 5.15. The spot price of silver is \$15 per ounce. The storage costs are \$0.24 per ounce per year payable quarterly in advance. Assuming that interest rates are 10% per annum for all maturities, calculate the futures price of silver for delivery in 9 months.
- 5.16. Suppose that F_1 and F_2 are two futures contracts on the same commodity with times to maturity, t_1 and t_2 , where $t_2 > t_1$. Prove that

$$F_2 \leq F_1 e^{r(t_2-t_1)}$$

where r is the interest rate (assumed constant) and there are no storage costs. For the purposes of this problem, assume that a futures contract is the same as a forward contract.

- 5.17. When a known future cash outflow in a foreign currency is hedged by a company using a forward contract, there is no foreign exchange risk. When it is hedged using futures contracts, the daily settlement process does leave the company exposed to some risk. Explain the nature of this risk. In particular, consider whether the company is better off using a futures contract or a forward contract when:
- The value of the foreign currency falls rapidly during the life of the contract.
 - The value of the foreign currency rises rapidly during the life of the contract.
 - The value of the foreign currency first rises and then falls back to its initial value.
 - The value of the foreign currency first falls and then rises back to its initial value.
- Assume that the forward price equals the futures price.
- 5.18. It is sometimes argued that a forward exchange rate is an unbiased predictor of future exchange rates. Under what circumstances is this so?
- 5.19. Show that the growth rate in an index futures price equals the excess return on the portfolio underlying the index over the risk-free rate. Assume that the risk-free interest rate and the dividend yield are constant.
- 5.20. Show that equation (5.3) is true by considering an investment in the asset combined with a short position in a futures contract. Assume that all income from the asset is reinvested in the asset. Use an argument similar to that in footnotes 2 and 4 and explain in detail what an arbitrageur would do if equation (5.3) did not hold.

- 5.21. Explain carefully what is meant by the expected price of a commodity on a particular future date. Suppose that the futures price for crude oil declines with the maturity of the contract at the rate of 2% per year. Assume that speculators tend to be short crude oil futures and hedgers tend to be long. What does the Keynes and Hicks argument imply about the expected future price of oil?
- 5.22. The Value Line Index is designed to reflect changes in the value of a portfolio of over 1,600 equally weighted stocks. Prior to March 9, 1988, the change in the index from one day to the next was calculated as the *geometric* average of the changes in the prices of the stocks underlying the index. In these circumstances, does equation (5.8) correctly relate the futures price of the index to its cash price? If not, does the equation overstate or understate the futures price?
- 5.23. A US company is interested in using the futures contracts traded by the CME Group to hedge its Australian dollar exposure. Define r as the interest rate (all maturities) on the US dollar and r_f as the interest rate (all maturities) on the Australian dollar. Assume that r and r_f are constant and that the company uses a contract expiring at time T to hedge an exposure at time t ($T > t$).
- Show that the optimal hedge ratio is $e^{(r_f - r)(T - t)}$.
 - Show that, when t is 1 day, the optimal hedge ratio is almost exactly S_0/F_0 , where S_0 is the current spot price of the currency and F_0 is the current futures price of the currency for the contract maturing at time T .
 - Show that the company can take account of the daily settlement of futures contracts for a hedge that lasts longer than 1 day by adjusting the hedge ratio so that it always equals the spot price of the currency divided by the futures price of the currency.

Further Questions

- 5.24. An index is 1,200. The three-month risk-free rate is 3% per annum and the dividend yield over the next three months is 1.2% per annum. The six-month risk-free rate is 3.5% per annum and the dividend yield over the next six months is 1% per annum. Estimate the futures price of the index for three-month and six-month contracts. All interest rates and dividend yields are continuously compounded.
- 5.25. The current USD/euro exchange rate is 1.4000 dollar per euro. The six-month forward exchange rate is 1.3950. The six-month USD interest rate is 1% per annum continuously compounded. Estimate the six-month euro interest rate.
- 5.26. The spot price of oil is \$80 per barrel and the cost of storing a barrel of oil for one year is \$3, payable at the end of the year. The risk-free interest rate is 5% per annum continuously compounded. What is an upper bound for the one-year futures price of oil?
- 5.27. A stock is expected to pay a dividend of \$1 per share in 2 months and in 5 months. The stock price is \$50, and the risk-free rate of interest is 8% per annum with continuous compounding for all maturities. An investor has just taken a short position in a 6-month forward contract on the stock.
- What are the forward price and the initial value of the forward contract?
 - Three months later, the price of the stock is \$48 and the risk-free rate of interest is still 8% per annum. What are the forward price and the value of the short position in the forward contract?

- 5.28. A bank offers a corporate client a choice between borrowing cash at 11% per annum and borrowing gold at 2% per annum. (If gold is borrowed, interest must be repaid in gold. Thus, 100 ounces borrowed today would require 102 ounces to be repaid in 1 year.) The risk-free interest rate is 9.25% per annum, and storage costs are 0.5% per annum. Discuss whether the rate of interest on the gold loan is too high or too low in relation to the rate of interest on the cash loan. The interest rates on the two loans are expressed with annual compounding. The risk-free interest rate and storage costs are expressed with continuous compounding.
- 5.29. A company that is uncertain about the exact date when it will pay or receive a foreign currency may try to negotiate with its bank a forward contract that specifies a period during which delivery can be made. The company wants to reserve the right to choose the exact delivery date to fit in with its own cash flows. Put yourself in the position of the bank. How would you price the product that the company wants?
- 5.30. A trader owns gold as part of a long-term investment portfolio. The trader can buy gold for \$1,250 per ounce and sell it for \$1,249 per ounce. The trader can borrow funds at 6% per year and invest funds at 5.5% per year (both interest rates are expressed with annual compounding). For what range of 1-year forward prices of gold does the trader have no arbitrage opportunities? Assume there is no bid–offer spread for forward prices.
- 5.31. A company enters into a forward contract with a bank to sell a foreign currency for K_1 at time T_1 . The exchange rate at time T_1 proves to be $S_1 (> K_1)$. The company asks the bank if it can roll the contract forward until time $T_2 (> T_1)$ rather than settle at time T_1 . The bank agrees to a new delivery price, K_2 . Explain how K_2 should be calculated.



6 CHAPTER

Interest Rate Futures

So far we have covered futures contracts on commodities, stock indices, and foreign currencies. We have seen how they work, how they are used for hedging, and how futures prices are determined. We now move on to consider interest rate futures.

This chapter explains the popular Treasury bond and Eurodollar futures contracts that trade in the United States. Many of the other interest rate futures contracts throughout the world have been modeled on these contracts. The chapter also shows how interest rate futures contracts, when used in conjunction with the duration measure introduced in Chapter 4, can be used to hedge a company's exposure to interest rate movements.

6.1 DAY COUNT AND QUOTATION CONVENTIONS

As a preliminary to the material in this chapter, we consider the day count and quotation conventions that apply to bonds and other instruments dependent on the interest rate.

Day Counts

The day count defines the way in which interest accrues over time. Generally, we know the interest earned over some reference period (e.g., the time between coupon payments on a bond), and we are interested in calculating the interest earned over some other period.

The day count convention is usually expressed as X/Y . When we are calculating the interest earned between two dates, X defines the way in which the number of days between the two dates is calculated, and Y defines the way in which the total number of days in the reference period is measured. The interest earned between the two dates is

$$\frac{\text{Number of days between dates}}{\text{Number of days in reference period}} \times \text{Interest earned in reference period}$$

Three day count conventions that are commonly used in the United States are:

1. Actual/actual (in period)
2. 30/360
3. Actual/360

Business Snapshot 6.1 Day Counts Can Be Deceptive

Between February 28 and March 1, 2013, you have a choice between owning a US government bond and a US corporate bond. They pay the same coupon and have the same quoted price. Assuming no risk of default, which would you prefer?

It sounds as though you should be indifferent, but in fact you should have a marked preference for the corporate bond. Under the 30/360 day count convention used for corporate bonds, there are 3 days between February 28, 2013, and March 1, 2013. Under the actual/actual (in period) day count convention used for government bonds, there is only 1 day. You would earn approximately three times as much interest by holding the corporate bond!

The actual/actual (in period) day count is used for Treasury bonds in the United States. This means that the interest earned between two dates is based on the ratio of the actual days elapsed to the actual number of days in the period between coupon payments. Assume that the bond principal is \$100, coupon payment dates are March 1 and September 1, and the coupon rate is 8% per annum. (This means that \$4 of interest is paid on each of March 1 and September 1.) Suppose that we wish to calculate the interest earned between March 1 and July 3. The reference period is from March 1 to September 1. There are 184 (actual) days in the reference period, and interest of \$4 is earned during the period. There are 124 (actual) days between March 1 and July 3. The interest earned between March 1 and July 3 is therefore

$$\frac{124}{184} \times 4 = 2.6957$$

The 30/360 day count is used for corporate and municipal bonds in the United States. This means that we assume 30 days per month and 360 days per year when carrying out calculations. With the 30/360 day count, the total number of days between March 1 and September 1 is 180. The total number of days between March 1 and July 3 is $(4 \times 30) + 2 = 122$. In a corporate bond with the same terms as the Treasury bond just considered, the interest earned between March 1 and July 3 would therefore be

$$\frac{122}{180} \times 4 = 2.7111$$

As shown in Business Snapshot 6.1, sometimes the 30/360 day count convention has surprising consequences.

The actual/360 day count is used for money market instruments in the United States. This indicates that the reference period is 360 days. The interest earned during part of a year is calculated by dividing the actual number of elapsed days by 360 and multiplying by the rate. The interest earned in 90 days is therefore exactly one-fourth of the quoted rate, and the interest earned in a whole year of 365 days is 365/360 times the quoted rate.

Conventions vary from country to country and from instrument to instrument. For example, money market instruments are quoted on an actual/365 basis in Australia, Canada, and New Zealand. LIBOR is quoted on an actual/360 for all currencies except sterling, for which it is quoted on an actual/365 basis. Euro-denominated and sterling bonds are usually quoted on an actual/actual basis.

Price Quotations of US Treasury Bills

The prices of money market instruments are sometimes quoted using a *discount rate*. This is the interest earned as a percentage of the final face value rather than as a percentage of the initial price paid for the instrument. An example is Treasury bills in the United States. If the price of a 91-day Treasury bill is quoted as 8, this means that the rate of interest earned is 8% of the face value per 360 days. Suppose that the face value is \$100. Interest of \$2.0222 ($= \$100 \times 0.08 \times 91/360$) is earned over the 91-day life. This corresponds to a true rate of interest of $2.0222/(100 - 2.0222) = 2.064\%$ for the 91-day period. In general, the relationship between the cash price and quoted price of a Treasury bill in the United States is

$$P = \frac{360}{n}(100 - Y)$$

where P is the quoted price, Y is the cash price, and n is the remaining life of the Treasury bill measured in calendar days.

Price Quotations of US Treasury Bonds

Treasury bond prices in the United States are quoted in dollars and thirty-seconds of a dollar. The quoted price is for a bond with a face value of \$100. Thus, a quote of 90-05 indicates that the quoted price for a bond with a face value of \$100,000 is \$90,156.25.

The quoted price, which traders refer to as the *clean price*, is not the same as the cash price paid by the purchaser of the bond, which is referred to by traders as the *dirty price*. In general,

$$\text{Cash price} = \text{Quoted price} + \text{Accrued interest since last coupon date}$$

To illustrate this formula, suppose that it is March 5, 2010, and the bond under consideration is an 11% coupon bond maturing on July 10, 2018, with a quoted price of 95-16 or \$95.50. Because coupons are paid semiannually on government bonds (and the final coupon is at maturity), the most recent coupon date is January 10, 2010, and the next coupon date is July 10, 2010. The number of days between January 10, 2010, and March 5, 2010, is 54, whereas the number of days between January 10, 2010, and July 10, 2010, is 181. On a bond with \$100 face value, the coupon payment is \$5.50 on January 10 and July 10. The accrued interest on March 5, 2010, is the share of the July 10 coupon accruing to the bondholder on March 5, 2010. Because actual/actual in period is used for Treasury bonds in the United States, this is

$$\frac{54}{181} \times \$5.50 = \$1.64$$

The cash price per \$100 face value for the bond is therefore

$$\$95.50 + \$1.64 = \$97.14$$

Thus, the cash price of a \$100,000 bond is \$97,140.

6.2 TREASURY BOND FUTURES

Table 6.1 shows interest rate futures quotes on May 26, 2010. One of the most popular long-term interest rate futures contracts is the Treasury bond futures contract traded by the CME Group. In this contract, any government bond that has more than 15 years to maturity on the first day of the delivery month and is not callable within 15 years from that day can be delivered. As will be explained later in this section, the exchange has developed a procedure for adjusting the price received by the party with the short position according to the particular bond it chooses to deliver.

The 10-year, 5-year, and 2-year Treasury note futures contract in the United States are also very popular. In the 10-year Treasury note futures contract, any government bond (or note) with a maturity between $6\frac{1}{2}$ and 10 years can be delivered. In the 5-year Treasury note futures contract, the bond delivered must have a remaining life between 4.167 and 5.25 years; in the 2-year contract, the remaining life must be between 1.75 and 5.25 years.

The remaining discussion in this section focuses on the Treasury bond futures. The Treasury note futures traded in the United States and many other futures contracts in the rest of the world are designed in a similar way to the Treasury bond futures, so that many of the points we will make are applicable to these contracts as well.

Table 6.1 Interest rate futures quotes as reported by exchanges on May 26, 2010.

| | <i>Open</i> | <i>High</i> | <i>Low</i> | <i>Settlement</i> | <i>Change</i> | <i>Volume</i> | <i>Open interest</i> |
|--|-------------|-------------|------------|-------------------|---------------|---------------|----------------------|
| Treasury Bonds \$100,000 (CME Group) | | | | | | | |
| June 2010 | 125-000 | 125-090 | 123-280 | 124-150 | -25.0 | 691,927 | 462,946 |
| Sept. 2010 | 124-170 | 124-290 | 123-130 | 124-010 | -28.0 | 240,475 | 294,151 |
| Treasury Notes 10 Year \$100,000 (CME Group) | | | | | | | |
| June 2010 | 121-180 | 121-230 | 120-245 | 121-050 | -16.5 | 2,139,365 | 1,085,236 |
| Sept. 2010 | 120-230 | 120-295 | 119-300 | 120-105 | -18.0 | 657,677 | 754,551 |
| Treasury Notes 5 Year \$100,000 (CME Group) | | | | | | | |
| June 2010 | 117-260 | 117-287 | 117-082 | 117-157 | -12.0 | 1,008,580 | 710,630 |
| Sept. 2010 | 116-312 | 117-035 | 116-140 | 116-217 | -12.5 | 360,038 | 353,682 |
| Treasury Notes 2 Year \$100,000 (CME Group) | | | | | | | |
| June 2010 | 109-102 | 109-110 | 109-050 | 109-080 | -2.5 | 595,833 | 642,470 |
| Sept. 2010 | 109-002 | 109-010 | 108-272 | 108-302 | -2.7 | 282,058 | 343,447 |
| 30-day Fed Funds Rate \$5,000,000 (CME Group) | | | | | | | |
| May 2010 | 99.7950 | 99.7975 | 99.7925 | 99.7925 | 0.0000 | 2,440 | 63,402 |
| Nov. 2010 | 99.7150 | 99.7250 | 99.7000 | 99.7150 | 0.0000 | 8,440 | 83,036 |
| Eurodollar \$1,000,000 (CME Group) | | | | | | | |
| June 2010 | 99.3400 | 99.3400 | 99.3050 | 99.3100 | 0.0425 | 370,183 | 1,110,424 |
| Sept. 2010 | 99.1150 | 99.1150 | 99.0500 | 99.0500 | 0.0250 | 693,097 | 1,107,562 |
| Dec. 2010 | 98.9700 | 98.9700 | 98.8950 | 98.8950 | -0.0050 | 500,388 | 1,065,630 |
| Dec. 2011 | 98.2200 | 98.2250 | 98.1750 | 98.2050 | -0.0450 | 222,979 | 530,952 |
| Dec. 2012 | 97.2950 | 97.3350 | 97.2750 | 97.3350 | -0.0650 | 61,428 | 141,428 |
| Dec. 2015 | 95.5100 | 95.5800 | 95.5100 | 95.5800 | -0.0400 | 1,403 | 10,380 |

Quotes

Treasury bond and Treasury note futures contracts are quoted in dollars and thirty-seconds of a dollar per \$100 face value. This is similar to the way bonds and notes are quoted in the spot market. In Table 6.1, the settlement price of the June 2010 Treasury bond futures contract is specified as 124-150. This means $124\frac{15.0}{32}$, or 124.46875. The settlement price of the September 2010 10-year Treasury note futures contract is quoted as 120-105. This means $120\frac{10.5}{32}$, or 120.328125. The settlement price of the June 2010 5-year Treasury bond price is quoted as 117-157. This means $117\frac{15.75}{32}$, or 117.492188. Finally, the settlement price of the September 2010 2-year Treasury note futures contract is quoted as 108-302. This means $108\frac{30.25}{32}$, or 108.945313.

Conversion Factors

As mentioned, the Treasury bond futures contract allows the party with the short position to choose to deliver any bond that has a maturity of more than 15 years and is not callable within 15 years. When a particular bond is delivered, a parameter known as its *conversion factor* defines the price received for the bond by the party with the short position. The applicable quoted price is the product of the conversion factor and the most recent settlement price for the futures contract. Taking accrued interest into account (see Section 6.1), the cash received for each \$100 face value of the bond delivered is

$$(\text{Most recent settlement price} \times \text{Conversion factor}) + \text{Accrued interest}$$

Each contract is for the delivery of \$100,000 face value of bonds. Suppose that the most recent settlement price is 90-00, the conversion factor for the bond delivered is 1.3800, and the accrued interest on this bond at the time of delivery is \$3 per \$100 face value. The cash received by the party with the short position (and paid by the party with the long position) is then

$$(1.3800 \times 90.00) + 3.00 = \$127.20$$

per \$100 face value. A party with the short position in one contract would deliver bonds with a face value of \$100,000 and receive \$127,200.

The conversion factor for a bond is set equal to the quoted price the bond would have per dollar of principal on the first day of the delivery month on the assumption that the interest rate for all maturities equals 6% per annum (with semiannual compounding). The bond maturity and the times to the coupon payment dates are rounded down to the nearest 3 months for the purposes of the calculation. The practice enables the exchange to produce comprehensive tables. If, after rounding, the bond lasts for an exact number of 6-month periods, the first coupon is assumed to be paid in 6 months. If, after rounding, the bond does not last for an exact number of 6-month periods (i.e., there are an extra 3 months), the first coupon is assumed to be paid after 3 months and accrued interest is subtracted.

As a first example of these rules, consider a 10% coupon bond with 20 years and 2 months to maturity. For the purposes of calculating the conversion factor, the bond is assumed to have exactly 20 years to maturity. The first coupon payment is assumed to be made after 6 months. Coupon payments are then assumed to be made at 6-month intervals until the end of the 20 years when the principal payment is made. Assume that the face value is \$100. When the discount rate is 6% per annum with semiannual

compounding (or 3% per 6 months), the value of the bond is

$$\sum_{i=1}^{40} \frac{5}{1.03^i} + \frac{100}{1.03^{40}} = \$146.23$$

Dividing by the face value gives a conversion factor of 1.4623.

As a second example of the rules, consider an 8% coupon bond with 18 years and 4 months to maturity. For the purposes of calculating the conversion factor, the bond is assumed to have exactly 18 years and 3 months to maturity. Discounting all the payments back to a point in time 3 months from today at 6% per annum (compounded semi-annually) gives a value of

$$4 + \sum_{i=1}^{36} \frac{4}{1.03^i} + \frac{100}{1.03^{36}} = \$125.83$$

The interest rate for a 3-month period is $\sqrt{1.03} - 1$, or 1.4889%. Hence, discounting back to the present gives the bond's value as $125.83/1.014889 = \$123.99$. Subtracting the accrued interest of 2.0, this becomes \$121.99. The conversion factor is therefore 1.2199.

Cheapest-to-Deliver Bond

At any given time during the delivery month, there are many bonds that can be delivered in the Treasury bond futures contract. These vary widely as far as coupon and maturity are concerned. The party with the short position can choose which of the available bonds is "cheapest" to deliver. Because the party with the short position receives

$$(\text{Most recent settlement price} \times \text{Conversion factor}) + \text{Accrued interest}$$

and the cost of purchasing a bond is

$$\text{Quoted bond price} + \text{Accrued interest}$$

the cheapest-to-deliver bond is the one for which

$$\text{Quoted bond price} - (\text{Most recent settlement price} \times \text{Conversion factor})$$

is least. Once the party with the short position has decided to deliver, it can determine the cheapest-to-deliver bond by examining each of the deliverable bonds in turn.

Example 6.1

The party with the short position has decided to deliver and is trying to choose between the three bonds in the table below. Assume the most recent settlement price is 93-08, or 93.25.

| <i>Bond</i> | <i>Quoted bond price (\$)</i> | <i>Conversion factor</i> |
|-------------|-------------------------------|--------------------------|
| 1 | 99.50 | 1.0382 |
| 2 | 143.50 | 1.5188 |
| 3 | 119.75 | 1.2615 |

Business Snapshot 6.2 The Wild Card Play

Trading in the CME Group's Treasury bond futures contract ceases at 2:00 p.m. Chicago time. However, Treasury bonds themselves continue trading in the spot market until 4:00 p.m. Furthermore, a trader with a short futures position has until 8:00 p.m. to issue to the clearinghouse a notice of intention to deliver. If the notice is issued, the invoice price is calculated on the basis of the settlement price that day. This is the price at which trading was conducted just before the closing bell at 2:00 p.m.

This practice gives rise to an option known as the *wild card play*. If bond prices decline after 2:00 p.m. on the first day of the delivery month, the party with the short position can issue a notice of intention to deliver at, say, 3:45 p.m. and proceed to buy bonds in the spot market for delivery at the 2:00 p.m. futures price. If the bond price does not decline, the party with the short position keeps the position open and waits until the next day when the same strategy can be used.

As with the other options open to the party with the short position, the wild card play is not free. Its value is reflected in the futures price, which is lower than it would be without the option.

The cost of delivering each of the bonds is as follows:

$$\text{Bond 1: } 99.50 - (93.25 \times 1.0382) = \$2.69$$

$$\text{Bond 2: } 143.50 - (93.25 \times 1.5188) = \$1.87$$

$$\text{Bond 3: } 119.75 - (93.25 \times 1.2615) = \$2.12$$

The cheapest-to-deliver bond is Bond 2.

A number of factors determine the cheapest-to-deliver bond. When bond yields are in excess of 6%, the conversion factor system tends to favor the delivery of low-coupon long-maturity bonds. When yields are less than 6%, the system tends to favor the delivery of high-coupon short-maturity bonds. Also, when the yield curve is upward-sloping, there is a tendency for bonds with a long time to maturity to be favored, whereas when it is downward-sloping, there is a tendency for bonds with a short time to maturity to be delivered.

In addition to the cheapest-to-deliver bond option, the party with a short position has an option known as the wild card play. This is described in Business Snapshot 6.2.

Determining the Futures Price

An exact theoretical futures price for the Treasury bond contract is difficult to determine because the short party's options concerned with the timing of delivery and choice of the bond that is delivered cannot easily be valued. However, if we assume that both the cheapest-to-deliver bond and the delivery date are known, the Treasury bond futures contract is a futures contract on a traded security (the bond) that provides the holder with known income.¹ Equation (5.2) then shows that the futures price, F_0 , is related to the spot price, S_0 , by

$$F_0 = (S_0 - I)e^{rT} \quad (6.1)$$

¹ In practice, for the purposes of estimating the cheapest-to-deliver bond, analysts usually assume that zero rates at the maturity of the futures contract will equal today's forward rates.

where I is the present value of the coupons during the life of the futures contract, T is the time until the futures contract matures, and r is the risk-free interest rate applicable to a time period of length T .

Example 6.2

Suppose that, in a Treasury bond futures contract, it is known that the cheapest-to-deliver bond will be a 12% coupon bond with a conversion factor of 1.6000. Suppose also that it is known that delivery will take place in 270 days. Coupons are payable semiannually on the bond. As illustrated in Figure 6.1, the last coupon date was 60 days ago, the next coupon date is in 122 days, and the coupon date thereafter is in 305 days. The term structure is flat, and the rate of interest (with continuous compounding) is 10% per annum. Assume that the current quoted bond price is \$115. The cash price of the bond is obtained by adding to this quoted price the proportion of the next coupon payment that accrues to the holder. The cash price is therefore

$$115 + \frac{60}{60 + 122} \times 6 = 116.978$$

A coupon of \$6 will be received after 122 days (= 0.3342 years). The present value of this is

$$6e^{-0.1 \times 0.3342} = 5.803$$

The futures contract lasts for 270 days (= 0.7397 years). The cash futures price, if the contract were written on the 12% bond, would therefore be

$$(116.978 - 5.803)e^{0.1 \times 0.7397} = 119.711$$


At delivery, there are 148 days of accrued interest. The quoted futures price, if the contract were written on the 12% bond, is calculated by subtracting the accrued interest

$$119.711 - 6 \times \frac{148}{148 + 35} = 114.859$$

From the definition of the conversion factor, 1.6000 standard bonds are considered equivalent to each 12% bond. The quoted futures price should therefore be

$$\frac{114.859}{1.6000} = 71.79$$

Figure 6.1 Time chart for Example 6.2.



6.3 EURODOLLAR FUTURES

The most popular interest rate futures contract in the United States is the three-month Eurodollar futures contract traded by the CME Group. A Eurodollar is a dollar deposited in a U.S. or foreign bank outside the United States. The Eurodollar interest rate is the rate of interest earned on Eurodollars deposited by one bank with another bank. It is essentially the same as the London Interbank Offered Rate (LIBOR) introduced in Chapter 4.

A three-month Eurodollar futures contract is a futures contract on the interest that will be paid (by someone who borrows at the Eurodollar interest rate) on \$1 million for a future three-month period. It allows a trader to speculate on a future three-month interest rate or to hedge an exposure to a future three-month interest rate. Eurodollar futures contracts have maturities in March, June, September, and December for up to 10 years into the future. This means that in 2010 a trader can use Eurodollar futures to take a position on what interest rates will be as far into the future as 2020. Short-maturity contracts trade for months other than March, June, September, and December.

To understand how Eurodollar futures contracts work, consider the June 2010 contract in Table 6.1. The quoted settlement price on May 26, 2010, is 99.3100. The contract ends on the third Wednesday of the delivery month. In the case of this contract, the third Wednesday of the delivery month is June 16, 2010. The contract is settled daily in the usual way until that date. On June 16, 2010, the settlement price is set equal to $100 - R$, where R is the actual three-month Eurodollar interest rate on that day, expressed with quarterly compounding and an actual/360 day count convention. Thus, if the three-month Eurodollar interest rate on June 16, 2010, turned out to be 0.75% (actual/360 with quarterly compounding), the final settlement price would be 99.2500. Once a final settlement has taken place, all contracts are declared closed.

The contract is designed so that a one-basis-point ($= 0.01$) move in the futures quote corresponds to a gain or loss of \$25 per contract. When a Eurodollar futures quote increases by one basis point, a trader who is long one contract gains \$25 and a trader who is short one contract loses \$25. Similarly, when the quote decreases by one basis point a trader who is long one contract loses \$25 and a trader who is short one contract gains \$25. Suppose, for example, a settlement price changes from 99.3100 to 99.2700. Traders with long positions lose $4 \times 25 = \$100$ per contract; traders with short positions gain \$100 per contract. A one-basis-point change in the futures quote corresponds to a 0.01% change in the futures interest rate. This in turn leads to a

$$1,000,000 \times 0.0001 \times 0.25 = 25$$

or \$25 change in the interest that will be earned on \$1 million in three months. The \$25 per basis point rule is therefore consistent with the point made earlier that the contract locks in an interest rate on \$1 million for three months.

The futures quote is 100 minus the futures interest rate, an investor who is long gains when interest rates fall and one who is short gains when interest rates rise. Table 6.2 shows a possible set of outcomes for the June 2010 contract in Table 6.1.

The contract price is defined as

$$10,000 \times [100 - 0.25 \times (100 - Q)] \quad (6.2)$$

where Q is the quote. Thus, the settlement price of 99.3100 for the June 2010 contract

Table 6.2 Possible sequence of prices for June 2010 Eurodollar futures contract.

| <i>Date</i> | <i>Futures price</i> | <i>Change</i> | <i>Gain per contract (\$)</i> |
|---------------|----------------------|---------------|-------------------------------|
| May 26, 2010 | 99.3100 | | |
| May 27, 2010 | 99.2700 | -0.0400 | -100 |
| May 28, 2010 | 99.3200 | +0.0500 | +125 |
| ⋮ | ⋮ | ⋮ | ⋮ |
| June 16, 2010 | 99.5300 | +0.0600 | +150 |
| <i>Total</i> | | +0.2200 | +550 |

in Table 6.1 corresponds to a contract price of

$$10,000 \times [100 - 0.25 \times (100 - 99.3100)] = \$998,275$$

In Table 6.2, the final contract price is

$$10,000 \times [100 - 0.25 \times (100 - 99.5300)] = \$998,825$$

and the difference between the initial and final contract price is \$550. This is consistent with the gain calculated in Table 6.2 using the “\$25 per one-basis-point move” rule.

Example 6.3

An investor wants to lock in the interest rate for a three-month period beginning September 19, 2012, on a principal of \$100 million. The September 2012 Eurodollar futures quote is 96.50, indicating that the investor can lock in an interest rate of $100 - 96.5$ or 3.5% per annum. The investor hedges by buying 100 contracts. Suppose that on September 19, 2012, the three-month Eurodollar rate turns out to be 2.6%. The final settlement in the contract is then at a price of 97.40. The investor gains

$$100 \times 25 \times (9,740 - 9,650) = 225,000$$

or \$225,000 on the Eurodollar futures contracts. The interest earned on the three-month investment is

$$100,000,000 \times 0.25 \times 0.026 = 650,000$$

or \$650,000. The gain on the Eurodollar futures brings this up to \$875,000, which is what the interest would be at 3.5% ($100,000,000 \times 0.25 \times 0.035 = 875,000$).

It appears that the futures trade has the effect of exactly locking an interest rate of 3.5% in all circumstances. In fact, the hedge is less than perfect because (a) futures contracts are settled daily (not all at the end) and (b) the final settlement in the futures contract happens on September 19, 2012, whereas the interest payment on the investment is three months later. One way of adjusting for the second point is to reduce the size of the hedge to reflect the difference between funds received on September 19, 2012, and funds received three months later. In this case, we would assume an interest rate of 3.5% for the three-month period and multiply the number of contracts by $1/(1 + 0.035 \times 0.25) = 0.9913$. This would lead to 99 rather than 100 contracts being purchased.

Example 6.3 shows how Eurodollar futures contracts can be used by an investor who wants to hedge the interest that will be earned during a future three-month period starting on September 19, 2012. Note that the timing of the cash flows from the hedge does not line up exactly with the timing of the interest cash flows. This is because the futures contract is settled daily. Also, the final settlement is on September 19, 2012, whereas interest payments on the investment are received three months after September 19, 2012. As indicated in the example, a small adjustment can be made to the hedge position in an attempt to allow for this second point.

Table 6.1 shows that the first year of the interest rate term structure in the U.S. was upward sloping on August 4, 2009. The futures rate for a three-month period beginning June 16, 2010, was 0.69%; for a three-month period beginning September 15, 2010, it was 0.95%; for a three-month period beginning December 15, 2010, it was 1.105%; and for a three-month period beginning December 16, 2015, it was 4.42%.

Other contracts similar to the CME Group's Eurodollar futures contracts trade on interest rates in other countries. The CME Group trades Euroyen contracts. The London International Financial Futures and Options Exchange (part of Euronext) trades three-month Euribor contracts (i.e., contracts on the three-month LIBOR rate for the euro) and three-month Euroswiss futures.

Forward vs. Futures Interest Rates

The Eurodollar futures contract is similar to a forward rate agreement (FRA: see Section 4.7) in that it locks in an interest rate for a future period. For short maturities (up to a year or so), the Eurodollar futures interest rate can be assumed to be the same as the corresponding forward interest rate. For longer-dated contracts, differences between the contracts become important. Compare a Eurodollar futures contract on an interest rate for the period between times T_1 and T_2 with an FRA for the same period. The Eurodollar futures contract is settled daily. The final settlement is at time T_1 and reflects the realized interest rate for the period between times T_1 and T_2 . By contrast the FRA is not settled daily and the final settlement reflecting the realized interest rate between times T_1 and T_2 is made at time T_2 .²

There are therefore two differences between a Eurodollar futures contract and an FRA. These are:

1. The difference between a Eurodollar futures contract and a similar contract where there is no daily settlement. The latter is a forward contract where a payoff equal to the difference between the forward interest rate and the realized interest rate is paid at time T_1 .
2. The difference between a forward contract where there is settlement at time T_1 and a forward contract where there is settlement at time T_2 .

These two components to the difference between the contracts cause some confusion in practice. Both decrease the forward rate relative to the futures rate, but for long-dated contracts the reduction caused by the second difference is much smaller than that caused by the first. The reason why the first difference (daily settlement) decreases the forward rate follows from the arguments in Section 5.8. Suppose you have a

² As mentioned in Section 4.7, settlement may occur at time T_1 , but it is then equal to the present value of what the forward contract payoff would be at time T_2 .

contract where the payoff is $R_M - R_F$ at time T_1 , where R_F is a predetermined rate for the period between T_1 and T_2 , and R_M is the realized rate for this period, and you have the option to switch to daily settlement. In this case daily settlement tends to lead to cash inflows when rates are high and cash outflows when rates are low. You would therefore find switching to daily settlement to be attractive because you tend to have more money in your margin account when rates are high. As a result the market would therefore set R_F higher for the daily settlement alternative (reducing your cumulative expected payoff). To put this the other way round, switching from daily settlement to settlement at time T_1 reduces R_F .

To understand the reason why the second difference reduces the forward rate, suppose that the payoff of $R_M - R_F$ is at time T_2 instead of T_1 (as it is for a regular FRA). If R_M is high, the payoff is positive. Because rates are high, the cost to you of having the payoff that you receive at time T_2 rather than time T_1 is relatively high. If R_M is low, the payoff is negative. Because rates are low, the benefit to you of having the payoff you make at time T_2 rather than time T_1 is relatively low. Overall you would rather have the payoff at time T_1 . If it is at time T_2 rather than T_1 , you must be compensated by a reduction in R_F .³

Convexity Adjustment

Analysts make what is known as a *convexity adjustment* to account for the total difference between the two rates. One popular adjustment is⁴

$$\text{Forward rate} = \text{Futures rate} - \frac{1}{2}\sigma^2 T_1 T_2 \quad (6.3)$$

where, as above, T_1 is the time to maturity of the futures contract and T_2 is the time to the maturity of the rate underlying the futures contract. The variable σ is the standard deviation of the change in the short-term interest rate in 1 year. Both rates are expressed with continuous compounding.⁵

Example 6.4

Consider the situation where $\sigma = 0.012$ and we wish to calculate the forward rate when the 8-year Eurodollar futures price quote is 94. In this case $T_1 = 8$, $T_2 = 8.25$, and the convexity adjustment is

$$\frac{1}{2} \times 0.012^2 \times 8 \times 8.25 = 0.00475$$

or 0.475% (47.5 basis points). The futures rate is 6% per annum on an actual/360 basis with quarterly compounding. This corresponds to 1.5% per 90 days or an annual rate of $(365/90) \ln 1.015 = 6.038\%$ with continuous compounding and an actual/365 day count. The estimate of the forward rate given by equation (6.3), therefore, is $6.038 - 0.475 = 5.563\%$ per annum with continuous compounding. The table below shows how the size of the adjustment increases with the time to maturity.

³ Quantifying the effect of this type of timing difference on the value of a derivative is discussed further in Chapter 29.

⁴ See Technical Note 1 at www.rotman.utoronto.ca/~hull/TechnicalNotes for a proof of this.

⁵ This formula is based on the Ho–Lee interest rate model, which will be discussed in Chapter 30. See T. S. Y. Ho and S.-B. Lee, “Term structure movements and pricing interest rate contingent claims,” *Journal of Finance*, 41 (December 1986), 1011–29.

| <i>Maturity of futures (years)</i> | <i>Convexity adjustments (basis points)</i> |
|--|---|
| 2 | 3.2 |
| 4 | 12.2 |
| 6 | 27.0 |
| 8 | 47.5 |
| 10 | 73.8 |

We can see from this table that the size of the adjustment is roughly proportional to the square of the time to maturity of the futures contract. For example, when the maturity doubles from 2 to 4 years, the size of the convexity quadruples.

Using Eurodollar Futures to Extend the LIBOR Zero Curve

The LIBOR zero curve out to 1 year is determined by the 1-month, 3-month, 6-month, and 12-month LIBOR rates. Once the convexity adjustment just described has been made, Eurodollar futures are often used to extend the zero curve. Suppose that the i th Eurodollar futures contract matures at time T_i ($i = 1, 2, \dots$). It is usually assumed that the forward interest rate calculated from the i th futures contract applies to the period T_i to T_{i+1} . (In practice this is close to true.) This enables a bootstrap procedure to be used to determine zero rates. Suppose that F_i is the forward rate calculated from the i th Eurodollar futures contract and R_i is the zero rate for a maturity T_i . From equation (4.5),

$$F_i = \frac{R_{i+1}T_{i+1} - R_iT_i}{T_{i+1} - T_i}$$

so that

$$R_{i+1} = \frac{F_i(T_{i+1} - T_i) + R_iT_i}{T_{i+1}} \quad (6.4)$$

Other Euro rates such as Euroswiss, Euroyen, and Euribor are used in a similar way.

Example 6.5

The 400-day LIBOR zero rate has been calculated as 4.80% with continuous compounding and, from Eurodollar futures quotes, it has been calculated that (a) the forward rate for a 90-day period beginning in 400 days is 5.30% with continuous compounding, (b) the forward rate for a 90-day period beginning in 491 days is 5.50% with continuous compounding, and (c) the forward rate for a 90-day period beginning in 589 days is 5.60% with continuous compounding. We can use equation (6.4) to obtain the 491-day rate as

$$\frac{0.053 \times 91 + 0.048 \times 400}{491} = 0.04893$$

or 4.893%. Similarly we can use the second forward rate to obtain the 589-day rate as

$$\frac{0.055 \times 98 + 0.04893 \times 491}{589} = 0.04994$$

or 4.994%. The next forward rate of 5.60% would be used to determine the zero curve out to the maturity of the next Eurodollar futures contract. (Note that, even though the rate underlying the Eurodollar futures contract is a 90-day rate, it is assumed to apply to the 91 or 98 days elapsing between Eurodollar contract maturities.)

6.4 DURATION-BASED HEDGING STRATEGIES USING FUTURES

We discussed duration in Section 4.8. Consider the situation where a position in an asset that is interest rate dependent, such as a bond portfolio or a money market security, is being hedged using an interest rate futures contract. Define:

V_F : Contract price for one interest rate futures contract

D_F : Duration of the asset underlying the futures contract at the maturity of the futures contract

P : Forward value of the portfolio being hedged at the maturity of the hedge (in practice, this is usually assumed to be the same as the value of the portfolio today)

D_P : Duration of the portfolio at the maturity of the hedge

If we assume that the change in the yield, Δy , is the same for all maturities, which means that only parallel shifts in the yield curve can occur, it is approximately true that

$$\Delta P = -PD_P \Delta y$$

It is also approximately true that

$$\Delta V_F = -V_F D_F \Delta y$$

The number of contracts required to hedge against an uncertain Δy , therefore, is

$$N^* = \frac{PD_P}{V_F D_F} \quad (6.5)$$

This is the *duration-based hedge ratio*. It is sometimes also called the *price sensitivity hedge ratio*.⁶ Using it has the effect of making the duration of the entire position zero.

When the hedging instrument is a Treasury bond futures contract, the hedger must base D_F on an assumption that one particular bond will be delivered. This means that the hedger must estimate which of the available bonds is likely to be cheapest to deliver at the time the hedge is put in place. If, subsequently, the interest rate environment changes so that it looks as though a different bond will be cheapest to deliver, then the hedge has to be adjusted and as a result its performance may be worse than anticipated.

When hedges are constructed using interest rate futures, it is important to bear in mind that interest rates and futures prices move in opposite directions. When interest rates go up, an interest rate futures price goes down. When interest rates go down, the reverse happens, and the interest rate futures price goes up. Thus, a company in a position to lose money if interest rates drop should hedge by taking a long futures

⁶ For a more detailed discussion of equation (6.5), see R.J. Rendleman, "Duration-Based Hedging with Treasury Bond Futures," *Journal of Fixed Income* 9, 1 (June 1999): 84–91.

position. Similarly, a company in a position to lose money if interest rates rise should hedge by taking a short futures position.

The hedger tries to choose the futures contract so that the duration of the underlying asset is as close as possible to the duration of the asset being hedged. Eurodollar futures tend to be used for exposures to short-term interest rates, whereas Treasury bond and Treasury note futures contracts are used for exposures to longer-term rates.

Example 6.6

It is August 2 and a fund manager with \$10 million invested in government bonds is concerned that interest rates are expected to be highly volatile over the next 3 months. The fund manager decides to use the December T-bond futures contract to hedge the value of the portfolio. The current futures price is 93-02, or 93.0625. Because each contract is for the delivery of \$100,000 face value of bonds, the futures contract price is \$93,062.50.

Suppose that the duration of the bond portfolio in 3 months will be 6.80 years. The cheapest-to-deliver bond in the T-bond contract is expected to be a 20-year 12% per annum coupon bond. The yield on this bond is currently 8.80% per annum, and the duration will be 9.20 years at maturity of the futures contract.

The fund manager requires a short position in T-bond futures to hedge the bond portfolio. If interest rates go up, a gain will be made on the short futures position, but a loss will be made on the bond portfolio. If interest rates decrease, a loss will be made on the short position, but there will be a gain on the bond portfolio. The number of bond futures contracts that should be shorted can be calculated from equation (6.5) as

$$\frac{10,000,000}{93,062.50} \times \frac{6.80}{9.20} = 79.42$$

To the nearest whole number, the portfolio manager should short 79 contracts.

6.5 HEDGING PORTFOLIOS OF ASSETS AND LIABILITIES

Financial institutions sometimes attempt to hedge themselves against interest rate risk by ensuring that the average duration of their assets equals the average duration of their liabilities. (The liabilities can be regarded as short positions in bonds.) This strategy is known as *duration matching* or *portfolio immunization*. When implemented, it ensures that a small parallel shift in interest rates will have little effect on the value of the portfolio of assets and liabilities. The gain (loss) on the assets should offset the loss (gain) on the liabilities.

Duration matching does not immunize a portfolio against nonparallel shifts in the zero curve. This is a weakness of the approach. In practice, short-term rates are usually more volatile than, and are not perfectly correlated with, long-term rates. Sometimes it even happens that short- and long-term rates move in opposite directions to each other. Duration matching is therefore only a first step and financial institutions have developed other tools to help them manage their interest rate exposure. See Business Snapshot 6.3.

Business Snapshot 6.3 Asset–Liability Management by Banks

The asset–liability management (ALM) committees of banks now monitor their exposure to interest rates very carefully. Matching the durations of assets and liabilities is sometimes a first step, but this does not protect a bank against non-parallel shifts in the yield curve. A popular approach is known as *GAP management*. This involves dividing the zero-coupon yield curve into segments, known as *buckets*. The first bucket might be 0 to 1 month, the second 1 to 3 months, and so on. The ALM committee then investigates the effect on the value of the bank's portfolio of the zero rates corresponding to one bucket changing while those corresponding to all other buckets stay the same.

If there is a mismatch, corrective action is usually taken. This can involve changing deposit and lending rates in the way described in Section 4.10. Alternatively, tools such as swaps, FRAs, bond futures, Eurodollar futures, and other interest rate derivatives can be used.

SUMMARY

Two very popular interest rate contracts are the Treasury bond and Eurodollar futures contracts that trade in the United States. In the Treasury bond futures contracts, the party with the short position has a number of interesting delivery options:

1. Delivery can be made on any day during the delivery month.
2. There are a number of alternative bonds that can be delivered.
3. On any day during the delivery month, the notice of intention to deliver at the 2:00 p.m. settlement price can be made any time up to 8:00 p.m.

These options all tend to reduce the futures price.

The Eurodollar futures contract is a contract on the 3-month rate on the third Wednesday of the delivery month. Eurodollar futures are frequently used to estimate LIBOR forward rates for the purpose of constructing a LIBOR zero curve. When long-dated contracts are used in this way, it is important to make what is termed a convexity adjustment to allow for the marking to market in the futures contract.

The concept of duration is important in hedging interest rate risk. It enables a hedger to assess the sensitivity of a bond portfolio to small parallel shifts in the yield curve. It also enables the hedger to assess the sensitivity of an interest rate futures price to small changes in the yield curve. The number of futures contracts necessary to protect the bond portfolio against small parallel shifts in the yield curve can therefore be calculated.

The key assumption underlying duration-based hedging is that all interest rates change by the same amount. This means that only parallel shifts in the term structure are allowed for. In practice, short-term interest rates are generally more volatile than are long-term interest rates, and hedge performance is liable to be poor if the duration of the bond underlying the futures contract differs markedly from the duration of the asset being hedged.

FURTHER READING

- Burghardt, G., and W. Hoskins. "The Convexity Bias in Eurodollar Futures," *Risk*, 8, 3 (1995): 63–70.
- Duffie, D. "Debt Management and Interest Rate Risk," in W. Beaver and G. Parker (eds.), *Risk Management: Challenges and Solutions*. New York: McGraw-Hill, 1994.
- Fabozzi, F. J. *Duration, Convexity, and Other Bond Risk Measures*. Frank Fabozzi Assoc., 1999.
- Grinblatt, M., and N. Jegadeesh. "The Relative Price of Eurodollar Futures and Forward Contracts," *Journal of Finance*, 51, 4 (September 1996): 1499–1522.

Practice Questions (Answers in Solutions Manual)

- 6.1. A US Treasury bond pays a 7% coupon on January 7 and July 7. How much interest accrues per \$100 of principal to the bondholder between July 7, 2011, and August 9, 2011? How would your answer be different if it were a corporate bond?
- 6.2. It is January 9, 2013. The price of a Treasury bond with a 12% coupon that matures on October 12, 2020, is quoted as 102-07. What is the cash price?
- 6.3. How is the conversion factor of a bond calculated by the CME Group? How is it used?
- 6.4. A Eurodollar futures price changes from 96.76 to 96.82. What is the gain or loss to an investor who is long two contracts?
- 6.5. What is the purpose of the convexity adjustment made to Eurodollar futures rates? Why is the convexity adjustment necessary?
- 6.6. The 350-day LIBOR rate is 3% with continuous compounding and the forward rate calculated from a Eurodollar futures contract that matures in 350 days is 3.2% with continuous compounding. Estimate the 440-day zero rate.
- 6.7. It is January 30. You are managing a bond portfolio worth \$6 million. The duration of the portfolio in 6 months will be 8.2 years. The September Treasury bond futures price is currently 108-15, and the cheapest-to-deliver bond will have a duration of 7.6 years in September. How should you hedge against changes in interest rates over the next 6 months?
- 6.8. The price of a 90-day Treasury bill is quoted as 10.00. What continuously compounded return (on an actual/365 basis) does an investor earn on the Treasury bill for the 90-day period?
- 6.9. It is May 5, 2011. The quoted price of a government bond with a 12% coupon that matures on July 27, 2014, is 110-17. What is the cash price?
- 6.10. Suppose that the Treasury bond futures price is 101-12. Which of the following four bonds is cheapest to deliver?

| <i>Bond</i> | <i>Price</i> | <i>Conversion factor</i> |
|-------------|--------------|--------------------------|
| 1 | 125-05 | 1.2131 |
| 2 | 142-15 | 1.3792 |
| 3 | 115-31 | 1.1149 |
| 4 | 144-02 | 1.4026 |

- 6.11. It is July 30, 2013. The cheapest-to-deliver bond in a September 2013 Treasury bond futures contract is a 13% coupon bond, and delivery is expected to be made on September 30, 2013. Coupon payments on the bond are made on February 4 and August 4 each year. The term structure is flat, and the rate of interest with semiannual compounding is 12% per annum. The conversion factor for the bond is 1.5. The current quoted bond price is \$110. Calculate the quoted futures price for the contract.
- 6.12. An investor is looking for arbitrage opportunities in the Treasury bond futures market. What complications are created by the fact that the party with a short position can choose to deliver any bond with a maturity of over 15 years?
- 6.13. Suppose that the 9-month LIBOR interest rate is 8% per annum and the 6-month LIBOR interest rate is 7.5% per annum (both with actual/365 and continuous compounding). Estimate the 3-month Eurodollar futures price quote for a contract maturing in 6 months.
- 6.14. Suppose that the 300-day LIBOR zero rate is 4% and Eurodollar quotes for contracts maturing in 300, 398, and 489 days are 95.83, 95.62, and 95.48. Calculate 398-day and 489-day LIBOR zero rates. Assume no difference between forward and futures rates for the purposes of your calculations.
- 6.15. Suppose that a bond portfolio with a duration of 12 years is hedged using a futures contract in which the underlying asset has a duration of 4 years. What is likely to be the impact on the hedge of the fact that the 12-year rate is less volatile than the 4-year rate?
- 6.16. Suppose that it is February 20 and a treasurer realizes that on July 17 the company will have to issue \$5 million of commercial paper with a maturity of 180 days. If the paper were issued today, the company would realize \$4,820,000. (In other words, the company would receive \$4,820,000 for its paper and have to redeem it at \$5,000,000 in 180 days' time.) The September Eurodollar futures price is quoted as 92.00. How should the treasurer hedge the company's exposure?
- 6.17. On August 1, a portfolio manager has a bond portfolio worth \$10 million. The duration of the portfolio in October will be 7.1 years. The December Treasury bond futures price is currently 91-12 and the cheapest-to-deliver bond will have a duration of 8.8 years at maturity. How should the portfolio manager immunize the portfolio against changes in interest rates over the next 2 months?
- 6.18. How can the portfolio manager change the duration of the portfolio to 3.0 years in Problem 6.17?
- 6.19. Between October 30, 2012, and November 1, 2012, you have a choice between owning a US government bond paying a 12% coupon and a US corporate bond paying a 12% coupon. Consider carefully the day count conventions discussed in this chapter and decide which of the two bonds you would prefer to own. Ignore the risk of default.
- 6.20. Suppose that a Eurodollar futures quote is 88 for a contract maturing in 60 days. What is the LIBOR forward rate for the 60- to 150-day period? Ignore the difference between futures and forwards for the purposes of this question.
- 6.21. The 3-month Eurodollar futures price for a contract maturing in 6 years is quoted as 95.20. The standard deviation of the change in the short-term interest rate in 1 year is 1.1%. Estimate the forward LIBOR interest rate for the period between 6.00 and 6.25 years in the future.
- 6.22. Explain why the forward interest rate is less than the corresponding futures interest rate calculated from a Eurodollar futures contract.

Further Questions

- 6.23. The December Eurodollar futures contract is quoted as 98.40 and a company plans to borrow \$8 million for three months starting in December at LIBOR plus 0.5%.
- (a) What rate can the company lock in by using the Eurodollar futures contract?
 - (b) What position should the company take in the contracts?
 - (c) If the actual three-month rate turns out to be 1.3%, what is the final settlement price on the futures contracts.
- 6.24. A Eurodollar futures quote for the period between 5.1 and 5.35 years in the future is 97.1. The standard deviation of the change in the short-term interest rate in one year is 1.4%. Estimate the forward interest rate in an FRA.
- 6.25. It is March 10, 2011. The cheapest-to-deliver bond in a December 2011 Treasury bond futures contract is an 8% coupon bond, and delivery is expected to be made on December 31, 2011. Coupon payments on the bond are made on March 1 and September 1 each year. The rate of interest with continuous compounding is 5% per annum for all maturities. The conversion factor for the bond is 1.2191. The current quoted bond price is \$137. Calculate the quoted futures price for the contract.
- 6.26. Assume that a bank can borrow or lend money at the same interest rate in the LIBOR market. The 90-day rate is 10% per annum, and the 180-day rate is 10.2% per annum, both expressed with continuous compounding and actual/actual day count. The Eurodollar futures price for a contract maturing in 91 days is quoted as 89.5. What arbitrage opportunities are open to the bank?
- 6.27. A Canadian company wishes to create a Canadian LIBOR futures contract from a US Eurodollar futures contract and forward contracts on foreign exchange. Using an example, explain how the company should proceed. For the purposes of this problem, assume that a futures contract is the same as a forward contract.
- 6.28. The futures price for the June 2011 CBOT bond futures contract is 118-23.
- (a) Calculate the conversion factor for a bond maturing on January 1, 2027, paying a coupon of 10%.
 - (b) Calculate the conversion factor for a bond maturing on October 1, 2032, paying a coupon of 7%.
 - (c) Suppose that the quoted prices of the bonds in (a) and (b) are 169.00 and 136.00, respectively. Which bond is cheaper to deliver?
 - (d) Assuming that the cheapest-to-deliver bond is actually delivered on June 25, 2011, what is the cash price received for the bond?
- 6.29. A portfolio manager plans to use a Treasury bond futures contract to hedge a bond portfolio over the next 3 months. The portfolio is worth \$100 million and will have a duration of 4.0 years in 3 months. The futures price is 122, and each futures contract is on \$100,000 of bonds. The bond that is expected to be cheapest to deliver will have a duration of 9.0 years at the maturity of the futures contract. What position in futures contracts is required?
- (a) What adjustments to the hedge are necessary if after 1 month the bond that is expected to be cheapest to deliver changes to one with a duration of 7 years?
 - (b) Suppose that all rates increase over the next 3 months, but long-term rates increase less than short-term and medium-term rates. What is the effect of this on the performance of the hedge?

7

CHAPTER



Swaps

The first swap contracts were negotiated in the early 1980s. Since then the market has seen phenomenal growth. Swaps now occupy a position of central importance in derivatives markets.

A swap is an over-the-counter agreement between two companies to exchange cash flows in the future. The agreement defines the dates when the cash flows are to be paid and the way in which they are to be calculated. Usually the calculation of the cash flows involves the future value of an interest rate, an exchange rate, or other market variable.

A forward contract can be viewed as a simple example of a swap. Suppose it is March 1, 2012, and a company enters into a forward contract to buy 100 ounces of gold for \$1,200 per ounce in 1 year. The company can sell the gold in 1 year as soon as it is received. The forward contract is therefore equivalent to a swap where the company agrees that on March 1, 2012, it will pay \$120,000 and receive 100 S , where S is the market price of 1 ounce of gold on that date.

Whereas a forward contract is equivalent to the exchange of cash flows on just one future date, swaps typically lead to cash flow exchanges on several future dates. In this chapter we examine how swaps are designed, how they are used, and how they are valued. Most of this chapter focuses on two popular swaps: plain vanilla interest rate swaps and fixed-for-fixed currency swaps. Other types of swaps are briefly reviewed at the end of the chapter and discussed in more detail in Chapter 32.

7.1 MECHANICS OF INTEREST RATE SWAPS

The most common type of swap is a “plain vanilla” interest rate swap. In this swap a company agrees to pay cash flows equal to interest at a predetermined fixed rate on a notional principal for a predetermined number of years. In return, it receives interest at a floating rate on the same notional principal for the same period of time.

LIBOR

The floating rate in most interest rate swap agreements is the London Interbank Offered Rate (LIBOR). We introduced this in Chapter 4. It is the rate of interest at which a bank is prepared to deposit money with other banks that have a AA credit rating. One-month, three-month, six-month, and 12-month LIBOR are quoted in all major currencies.

Just as prime is often the reference rate of interest for floating-rate loans in the domestic financial market, LIBOR is a reference rate of interest for loans in international financial markets. To understand how it is used, consider a 5-year bond with a rate of interest specified as 6-month LIBOR plus 0.5% per annum. The life of the bond is divided into 10 periods, each 6 months in length. For each period, the rate of interest is set at 0.5% per annum above the 6-month LIBOR rate at the beginning of the period. Interest is paid at the end of the period.

Illustration


Consider a hypothetical 3-year swap initiated on March 5, 2012, between Microsoft and Intel. We suppose Microsoft agrees to pay Intel an interest rate of 5% per annum on a principal of \$100 million, and in return Intel agrees to pay Microsoft the 6-month LIBOR rate on the same principal. Microsoft is the *fixed-rate payer*; Intel is the *floating-rate payer*. We assume the agreement specifies that payments are to be exchanged every 6 months and that the 5% interest rate is quoted with semiannual compounding. This swap is represented diagrammatically in Figure 7.1.

The first exchange of payments would take place on September 5, 2012, 6 months after the initiation of the agreement. Microsoft would pay Intel \$2.5 million. This is the interest on the \$100 million principal for 6 months at 5%. Intel would pay Microsoft interest on the \$100 million principal at the 6-month LIBOR rate prevailing 6 months prior to September 5, 2012—that is, on March 5, 2012. Suppose that the 6-month LIBOR rate on March 5, 2012, is 4.2%. Intel pays Microsoft $0.5 \times 0.042 \times \$100 = \$2.1$ million.¹ Note that there is no uncertainty about this first exchange of payments because it is determined by the LIBOR rate at the time the contract is entered into.

The second exchange of payments would take place on March 5, 2013, a year after the initiation of the agreement. Microsoft would pay \$2.5 million to Intel. Intel would pay interest on the \$100 million principal to Microsoft at the 6-month LIBOR rate prevailing 6 months prior to March 5, 2013—that is, on September 5, 2012. Suppose that the 6-month LIBOR rate on September 5, 2012, is 4.8%. Intel pays $0.5 \times 0.048 \times \$100 = \$2.4$ million to Microsoft.

In total, there are six exchanges of payment on the swap. The fixed payments are always \$2.5 million. The floating-rate payments on a payment date are calculated using the 6-month LIBOR rate prevailing 6 months before the payment date. An interest rate swap is generally structured so that one side remits the difference between the two payments to the other side. In our example, Microsoft would pay Intel \$0.4 million (= \$2.5 million – \$2.1 million) on September 5, 2012, and \$0.1 million (= \$2.5 million – \$2.4 million) on March 5, 2013.

Figure 7.1 Interest rate swap between Microsoft and Intel.



¹ The calculations here are simplified in that they ignore day count conventions. This point is discussed in more detail later in the chapter.

Table 7.1 Cash flows (millions of dollars) to Microsoft in a \$100 million 3-year interest rate swap when a fixed rate of 5% is paid and LIBOR is received.

| <i>Date</i> | <i>Six-month LIBOR rate (%)</i> | <i>Floating cash flow received</i> | <i>Fixed cash flow paid</i> | <i>Net cash flow</i> |
|---------------|---------------------------------|------------------------------------|-----------------------------|----------------------|
| Mar. 5, 2012 | 4.20 | | | |
| Sept. 5, 2012 | 4.80 | +2.10 | -2.50 | -0.40 |
| Mar. 5, 2013 | 5.30 | +2.40 | -2.50 | -0.10 |
| Sept. 5, 2013 | 5.50 | +2.65 | -2.50 | +0.15 |
| Mar. 5, 2014 | 5.60 | +2.75 | -2.50 | +0.25 |
| Sept. 5, 2014 | 5.90 | +2.80 | -2.50 | +0.30 |
| Mar. 5, 2015 | | +2.95 | -2.50 | +0.45 |

Table 7.1 provides a complete example of the payments made under the swap for one particular set of 6-month LIBOR rates. The table shows the swap cash flows from the perspective of Microsoft. Note that the \$100 million principal is used only for the calculation of interest payments. The principal itself is not exchanged. For this reason it is termed the *notional principal*, or just the *notional*.

If the principal were exchanged at the end of the life of the swap, the nature of the deal would not be changed in any way. The principal is the same for both the fixed and floating payments. Exchanging \$100 million for \$100 million at the end of the life of the swap is a transaction that would have no financial value to either Microsoft or Intel. Table 7.2 shows the cash flows in Table 7.1 with a final exchange of principal added in. This provides an interesting way of viewing the swap. The cash flows in the third column of this table are the cash flows from a long position in a floating-rate bond. The cash flows in the fourth column of the table are the cash flows from a short position in a fixed-rate bond. The table shows that the swap can be regarded as the exchange of a fixed-rate bond for a floating-rate bond. Microsoft, whose position is described by Table 7.2, is long a floating-rate bond and short a fixed-rate bond. Intel is long a fixed-rate bond and short a floating-rate bond.

Table 7.2 Cash flows (millions of dollars) from Table 7.1 when there is a final exchange of principal.

| <i>Date</i> | <i>Six-month LIBOR rate (%)</i> | <i>Floating cash flow received</i> | <i>Fixed cash flow paid</i> | <i>Net cash flow</i> |
|---------------|---------------------------------|------------------------------------|-----------------------------|----------------------|
| Mar. 5, 2012 | 4.20 | | | |
| Sept. 5, 2012 | 4.80 | +2.10 | -2.50 | -0.40 |
| Mar. 5, 2013 | 5.30 | +2.40 | -2.50 | -0.10 |
| Sept. 5, 2013 | 5.50 | +2.65 | -2.50 | +0.15 |
| Mar. 5, 2014 | 5.60 | +2.75 | -2.50 | +0.25 |
| Sept. 5, 2014 | 5.90 | +2.80 | -2.50 | +0.30 |
| Mar. 5, 2015 | | +102.95 | -102.50 | +0.45 |

This characterization of the cash flows in the swap helps to explain why the floating rate in the swap is set 6 months before it is paid. On a floating-rate bond, interest is generally set at the beginning of the period to which it will apply and is paid at the end of the period. The calculation of the floating-rate payments in a “plain vanilla” interest rate swap such as the one in Table 7.2 reflects this.

Using the Swap to Transform a Liability

For Microsoft, the swap could be used to transform a floating-rate loan into a fixed-rate loan. Suppose that Microsoft has arranged to borrow \$100 million at LIBOR plus 10 basis points. (One basis point is one-hundredth of 1%, so the rate is LIBOR plus 0.1%.) After Microsoft has entered into the swap, it has the following three sets of cash flows:

1. It pays LIBOR plus 0.1% to its outside lenders.
2. It receives LIBOR under the terms of the swap.
3. It pays 5% under the terms of the swap.

These three sets of cash flows net out to an interest rate payment of 5.1%. Thus, for Microsoft, the swap could have the effect of transforming borrowings at a floating rate of LIBOR plus 10 basis points into borrowings at a fixed rate of 5.1%.

For Intel, the swap could have the effect of transforming a fixed-rate loan into a floating-rate loan. Suppose that Intel has a 3-year \$100 million loan outstanding on which it pays 5.2%. After it has entered into the swap, it has the following three sets of cash flows:

1. It pays 5.2% to its outside lenders.
2. It pays LIBOR under the terms of the swap.
3. It receives 5% under the terms of the swap.

These three sets of cash flows net out to an interest rate payment of LIBOR plus 0.2% (or LIBOR plus 20 basis points). Thus, for Intel, the swap could have the effect of transforming borrowings at a fixed rate of 5.2% into borrowings at a floating rate of LIBOR plus 20 basis points. These potential uses of the swap by Intel and Microsoft are illustrated in Figure 7.2.

Using the Swap to Transform an Asset

Swaps can also be used to transform the nature of an asset. Consider Microsoft in our example. The swap could have the effect of transforming an asset earning a fixed rate of interest into an asset earning a floating rate of interest. Suppose that Microsoft owns \$100 million in bonds that will provide interest at 4.7% per annum over the next 3 years.

Figure 7.2 Microsoft and Intel use the swap to transform a liability.





Figure 7.3 Microsoft and Intel use the swap to transform an asset.

After Microsoft has entered into the swap, it has the following three sets of cash flows:

1. It receives 4.7% on the bonds.
2. It receives LIBOR under the terms of the swap.
3. It pays 5% under the terms of the swap.

These three sets of cash flows net out to an interest rate inflow of LIBOR minus 30 basis points. Thus, one possible use of the swap for Microsoft is to transform an asset earning 4.7% into an asset earning LIBOR minus 30 basis points.

Next, consider Intel. The swap could have the effect of transforming an asset earning a floating rate of interest into an asset earning a fixed rate of interest. Suppose that Intel has an investment of \$100 million that yields LIBOR minus 20 basis points. After it has entered into the swap, it has the following three sets of cash flows:

1. It receives LIBOR minus 20 basis points on its investment.
2. It pays LIBOR under the terms of the swap.
3. It receives 5% under the terms of the swap.

These three sets of cash flows net out to an interest rate inflow of 4.8%. Thus, one possible use of the swap for Intel is to transform an asset earning LIBOR minus 20 basis points into an asset earning 4.8%. These potential uses of the swap by Intel and Microsoft are illustrated in Figure 7.3.

Role of Financial Intermediary

Usually two nonfinancial companies such as Intel and Microsoft do not get in touch directly to arrange a swap in the way indicated in Figures 7.2 and 7.3. They each deal with a financial intermediary such as a bank or other financial institution. “Plain vanilla” fixed-for-floating swaps on US interest rates are usually structured so that the financial institution earns about 3 or 4 basis points (0.03% or 0.04%) on a pair of offsetting transactions.

Figure 7.4 shows what the role of the financial institution might be in the situation in Figure 7.2. The financial institution enters into two offsetting swap transactions with


Figure 7.4 Interest rate swap from Figure 7.2 when financial institution is involved.

Figure 7.5 Interest rate swap from Figure 7.3 when financial institution is involved.



Intel and Microsoft. Assuming that both companies honor their obligations, the financial institution is certain to make a profit of 0.03% (3 basis points) per year multiplied by the notional principal of \$100 million. This amounts to \$30,000 per year for the 3-year period. Microsoft ends up borrowing at 5.115% (instead of 5.1%, as in Figure 7.2), and Intel ends up borrowing at LIBOR plus 21.5 basis points (instead of at LIBOR plus 20 basis points, as in Figure 7.2).

Figure 7.5 illustrates the role of the financial institution in the situation in Figure 7.3. The swap is the same as before and the financial institution is certain to make a profit of 3 basis points if neither company defaults. Microsoft ends up earning LIBOR minus 31.5 basis points (instead of LIBOR minus 30 basis points, as in Figure 7.3), and Intel ends up earning 4.785% (instead of 4.8%, as in Figure 7.3).

Note that in each case the financial institution has two separate contracts: one with Intel and the other with Microsoft. In most instances, Intel will not even know that the financial institution has entered into an offsetting swap with Microsoft, and vice versa. If one of the companies defaults, the financial institution still has to honor its agreement with the other company. The 3-basis-point spread earned by the financial institution is partly to compensate it for the risk that one of the two companies will default on the swap payments.

Market Makers

In practice, it is unlikely that two companies will contact a financial institution at the same time and want to take opposite positions in exactly the same swap. For this reason, many large financial institutions act as market makers for swaps. This means that they are prepared to enter into a swap without having an offsetting swap with another counterparty.² Market makers must carefully quantify and hedge the risks they are taking. Bonds, forward rate agreements, and interest rate futures are examples of the instruments that can be used for hedging by swap market makers. Table 7.3 shows quotes for plain vanilla US dollar swaps that might be posted by a market maker.³ As mentioned earlier, the bid–offer spread is 3 to 4 basis points. The average of the bid and offer fixed rates is known as the *swap rate*. This is shown in the final column of Table 7.3.

Consider a new swap where the fixed rate equals the current swap rate. We can reasonably assume that the value of this swap is zero. (Why else would a market maker choose bid–offer quotes centered on the swap rate?) In Table 7.2 we saw that a swap can

² This is sometimes referred to as *warehousing* swaps.

³ The standard swap in the United States is one where fixed payments made every 6 months are exchanged for floating LIBOR payments made every 3 months. In Table 7.1 we assumed that fixed and floating payments are exchanged every 6 months. The fixed rate should be almost exactly the same in both cases.

Table 7.3 Bid and offer fixed rates in the swap market and swap rates (percent per annum).

| <i>Maturity (years)</i> | <i>Bid</i> | <i>Offer</i> | <i>Swap rate</i> |
|-------------------------|------------|--------------|------------------|
| 2 | 6.03 | 6.06 | 6.045 |
| 3 | 6.21 | 6.24 | 6.225 |
| 4 | 6.35 | 6.39 | 6.370 |
| 5 | 6.47 | 6.51 | 6.490 |
| 7 | 6.65 | 6.68 | 6.665 |
| 10 | 6.83 | 6.87 | 6.850 |

be characterized as the difference between a fixed-rate bond and a floating-rate bond. Define:

B_{fix} : Value of fixed-rate bond underlying the swap we are considering

B_{fl} : Value of floating-rate bond underlying the swap we are considering

Since the swap is worth zero, it follows that

$$B_{\text{fix}} = B_{\text{fl}} \quad (7.1)$$

We will use this result later in the chapter when discussing how the LIBOR/swap zero curve is determined.

7.2 DAY COUNT ISSUES

We discussed day count conventions in Section 6.1. The day count conventions affect payments on a swap, and some of the numbers calculated in the examples we have given do not exactly reflect these day count conventions. Consider, for example, the 6-month LIBOR payments in Table 7.1. Because it is a US money market rate, 6-month LIBOR is quoted on an actual/360 basis. The first floating payment in Table 7.1, based on the LIBOR rate of 4.2%, is shown as \$2.10 million. Because there are 184 days between March 5, 2012, and September 5, 2012, it should be

$$100 \times 0.042 \times \frac{184}{360} = \$2.1467 \text{ million}$$

In general, a LIBOR-based floating-rate cash flow on a swap payment date is calculated as $LRn/360$, where L is the principal, R is the relevant LIBOR rate, and n is the number of days since the last payment date.

The fixed rate that is paid in a swap transaction is similarly quoted with a particular day count basis being specified. As a result, the fixed payments may not be exactly equal on each payment date. The fixed rate is usually quoted as actual/365 or 30/360. It is not therefore directly comparable with LIBOR because it applies to a full year. To make the rates approximately comparable, either the 6-month LIBOR rate must be multiplied by 365/360 or the fixed rate must be multiplied by 360/365.

For clarity of exposition, we will ignore day count issues in the calculations in the rest of this chapter.

| Business Snapshot 7.1 Extract from Hypothetical Swap Confirmation | |
|--|---|
| Trade date: | 27-February-2012 |
| Effective date: | 5-March-2012 |
| Business day convention (all dates): | Following business day |
| Holiday calendar: | US |
| Termination date: | 5-March-2015 |
| <i>Fixed amounts</i> | |
| Fixed-rate payer: | Microsoft |
| Fixed-rate notional principal: | USD 100 million |
| Fixed rate: | 5.015% per annum |
| Fixed-rate day count convention: | Actual/365 |
| Fixed-rate payment dates: | Each 5-March and 5-September, commencing 5-September-2012, up to and including 5-March-2015 |
| <i>Floating amounts</i> | |
| Floating-rate payer: | Goldman Sachs |
| Floating-rate notional principal: | USD 100 million |
| Floating rate: | USD 6-month LIBOR |
| Floating-rate day count convention: | Actual/360 |
| Floating-rate payment dates: | Each 5-March and 5-September, commencing 5-September-2012, up to and including 5-March-2015 |

7.3 CONFIRMATIONS

A *confirmation* is the legal agreement underlying a swap and is signed by representatives of the two parties. The drafting of confirmations has been facilitated by the work of the International Swaps and Derivatives Association (ISDA; www.isda.org) in New York. This organization has produced a number of Master Agreements that consist of clauses defining in some detail the terminology used in swap agreements, what happens in the event of default by either side, and so on. Master Agreements cover all outstanding transactions between two parties. In Business Snapshot 7.1, we show a possible extract from the confirmation for the swap shown in Figure 7.4 between Microsoft and a financial institution (assumed here to be Goldman Sachs). The full confirmation might state that the provisions of an ISDA Master Agreement apply.

The confirmation specifies that the following business day convention is to be used and that the US calendar determines which days are business days and which days are holidays. This means that, if a payment date falls on a weekend or a US holiday, the payment is made on the next business day.⁴

⁴ Another business day convention that is sometimes specified is the *modified following* business day convention, which is the same as the following business day convention except that, when the next business day falls in a different month from the specified day, the payment is made on the immediately preceding business day. *Preceding* and *modified preceding* business day conventions are defined analogously.

7.4 THE COMPARATIVE-ADVANTAGE ARGUMENT

An explanation commonly put forward to explain the popularity of swaps concerns comparative advantages. Consider the use of an interest rate swap to transform a liability. Some companies, it is argued, have a comparative advantage when borrowing in fixed-rate markets, whereas other companies have a comparative advantage in floating-rate markets. To obtain a new loan, it makes sense for a company to go to the market where it has a comparative advantage. As a result, the company may borrow fixed when it wants floating, or borrow floating when it wants fixed. The swap is used to transform a fixed-rate loan into a floating-rate loan, and vice versa.

Suppose that two companies, AAACorp and BBBCorp, both wish to borrow \$10 million for 5 years and have been offered the rates shown in Table 7.4. AAACorp has a AAA credit rating; BBBCorp has a BBB credit rating.⁵ We assume that BBBCorp wants to borrow at a fixed rate of interest, whereas AAACorp wants to borrow at a floating rate of interest linked to 6-month LIBOR. Because it has a worse credit rating than AAACorp, BBBCorp pays a higher rate of interest than AAACorp in both fixed and floating markets.

A key feature of the rates offered to AAACorp and BBBCorp is that the difference between the two fixed rates is greater than the difference between the two floating rates. BBBCorp pays 1.2% more than AAACorp in fixed-rate markets and only 0.7% more than AAACorp in floating-rate markets. BBBCorp appears to have a comparative advantage in floating-rate markets, whereas AAACorp appears to have a comparative advantage in fixed-rate markets.⁶ It is this apparent anomaly that can lead to a swap being negotiated. AAACorp borrows fixed-rate funds at 4% per annum. BBBCorp borrows floating-rate funds at LIBOR plus 0.6% per annum. They then enter into a swap agreement to ensure that AAACorp ends up with floating-rate funds and BBBCorp ends up with fixed-rate funds.

To understand how this swap might work, we first assume that AAACorp and BBBCorp get in touch with each other directly. The sort of swap they might negotiate is shown in Figure 7.6. This is similar to our example in Figure 7.2. AAACorp agrees to pay BBBCorp interest at 6-month LIBOR on \$10 million. In return, BBBCorp agrees to pay AAACorp interest at a fixed rate of 4.35% per annum on \$10 million.


Table 7.4 Borrowing rates that provide a basis for the comparative-advantage argument.

| | <i>Fixed</i> | <i>Floating</i> |
|---------|--------------|----------------------|
| AAACorp | 4.0% | 6-month LIBOR – 0.1% |
| BBBCorp | 5.2% | 6-month LIBOR + 0.6% |

⁵ The credit ratings assigned to companies by S&P and Fitch (in order of decreasing creditworthiness) are AAA, AA, A, BBB, BB, B, CCC, CC, and C. The corresponding ratings assigned by Moody's are Aaa, Aa, A, Baa, Ba, B, Caa, Ca, and C, respectively.

⁶ Note that BBBCorp's comparative advantage in floating-rate markets does not imply that BBBCorp pays less than AAACorp in this market. It means that the extra amount that BBBCorp pays over the amount paid by AAACorp is less in this market. One of my students summarized the situation as follows: "AAACorp pays more less in fixed-rate markets; BBBCorp pays less more in floating-rate markets."

Figure 7.6 Swap agreement between AAACorp and BBBCorp when rates in Table 7.4 apply.



AAACorp has three sets of interest rate cash flows:

1. It pays 4% per annum to outside lenders.
2. It receives 4.35% per annum from BBBCorp.
3. It pays LIBOR to BBBCorp.

The net effect of the three cash flows is that AAACorp pays LIBOR minus 0.35% per annum. This is 0.25% per annum less than it would pay if it went directly to floating-rate markets. BBBCorp also has three sets of interest rate cash flows:


1. It pays LIBOR + 0.6% per annum to outside lenders.
2. It receives LIBOR from AAACorp.
3. It pays 4.35% per annum to AAACorp.

The net effect of the three cash flows is that BBBCorp pays 4.95% per annum. This is 0.25% per annum less than it would pay if it went directly to fixed-rate markets.

In this example, the swap has been structured so that the net gain to both sides is the same, 0.25%. This need not be the case. However, the total apparent gain from this type of interest rate swap arrangement is always $a - b$, where a is the difference between the interest rates facing the two companies in fixed-rate markets, and b is the difference between the interest rates facing the two companies in floating-rate markets. In this case, $a = 1.2\%$ and $b = 0.7\%$, so that the total gain is 0.5%.

If AAACorp and BBBCorp did not deal directly with each other and used a financial institution, an arrangement such as that shown in Figure 7.7 might result. (This is similar to the example in Figure 7.4.) In this case, AAACorp ends up borrowing at LIBOR minus 0.33%, BBBCorp ends up borrowing at 4.97%, and the financial institution earns a spread of 4 basis points per year. The gain to AAACorp is 0.23%; the gain to BBBCorp is 0.23%; and the gain to the financial institution is 0.04%. The total gain to all three parties is 0.50% as before.

Figure 7.7 Swap agreement between AAACorp and BBBCorp when rates in Table 7.4 apply and a financial intermediary is involved.



Criticism of the Argument

The comparative-advantage argument we have just outlined for explaining the attractiveness of interest rate swaps is open to question. Why in Table 7.4 should the spreads between the rates offered to AAACorp and BBBCorp be different in fixed and floating markets? Now that the swap market has been in existence for some time, we might reasonably expect these types of differences to have been arbitrated away.

The reason that spread differentials appear to exist is due to the nature of the contracts available to companies in fixed and floating markets. The 4.0% and 5.2% rates available to AAACorp and BBBCorp in fixed-rate markets are 5-year rates (e.g., the rates at which the companies can issue 5-year fixed-rate bonds). The $\text{LIBOR} - 0.1\%$ and $\text{LIBOR} + 0.6\%$ rates available to AAACorp and BBBCorp in floating-rate markets are 6-month rates. In the floating-rate market, the lender usually has the opportunity to review the floating rates every 6 months. If the creditworthiness of AAACorp or BBBCorp has declined, the lender has the option of increasing the spread over LIBOR that is charged. In extreme circumstances, the lender can refuse to roll over the loan at all. The providers of fixed-rate financing do not have the option to change the terms of the loan in this way.⁷

The spreads between the rates offered to AAACorp and BBBCorp are a reflection of the extent to which BBBCorp is more likely than AAACorp to default. During the next 6 months, there is very little chance that either AAACorp or BBBCorp will default. As we look further ahead, the probability of a default by a company with a relatively low credit rating (such as BBBCorp) is liable to increase faster than the probability of a default by a company with a relatively high credit rating (such as AAACorp). This is why the spread between the 5-year rates is greater than the spread between the 6-month rates.

After negotiating a floating-rate loan at $\text{LIBOR} + 0.6\%$ and entering into the swap shown in Figure 7.7, BBBCorp appears to obtain a fixed-rate loan at 4.97%. The arguments just presented show that this is not really the case. In practice, the rate paid is 4.97% only if BBBCorp can continue to borrow floating-rate funds at a spread of 0.6% over LIBOR. If, for example, the credit rating of BBBCorp declines so that the floating-rate loan is rolled over at $\text{LIBOR} + 1.6\%$, the rate paid by BBBCorp increases to 5.97%. The market expects that BBBCorp's spread over 6-month LIBOR will on average rise during the swap's life. BBBCorp's expected average borrowing rate when it enters into the swap is therefore greater than 4.97%.

The swap in Figure 7.7 locks in $\text{LIBOR} - 0.33\%$ for AAACorp for the whole of the next 5 years, not just for the next 6 months. This appears to be a good deal for AAACorp. The downside is that it is bearing the risk of a default by the financial institution. If it borrowed floating-rate funds in the usual way, it would not be bearing this risk.

7.5 THE NATURE OF SWAP RATES

At this stage it is appropriate to examine the nature of swap rates and the relationship between swap and LIBOR markets. We explained in Section 4.1 that LIBOR is the rate of interest at which AA-rated banks borrow for periods between 1 and 12 months from other banks. Also, as indicated in Table 7.3, a swap rate is the average of (a) the fixed rate that a

⁷ If the floating-rate loans are structured so that the spread over LIBOR is guaranteed in advance regardless of changes in credit rating, the spread differentials disappear.

swap market maker is prepared to pay in exchange for receiving LIBOR (its bid rate) and (b) the fixed rate that it is prepared to receive in return for paying LIBOR (its offer rate).

Like LIBOR rates, swap rates are not risk-free lending rates. However, they are close to risk-free. A financial institution can earn the 5-year swap rate on a certain principal by doing the following:

1. Lend the principal for the first 6 months to a AA borrower and then relend it for successive 6-month periods to other AA borrowers; and
2. Enter into a swap to exchange the LIBOR income for the 5-year swap rate.

This shows that the 5-year swap rate is an interest rate with a credit risk corresponding to the situation where 10 consecutive 6-month LIBOR loans to AA companies are made. Similarly the 7-year swap rate is an interest rate with a credit risk corresponding to the situation where 14 consecutive 6-month LIBOR loans to AA companies are made. Swap rates of other maturities can be interpreted analogously.

Note that 5-year swap rates are less than 5-year AA borrowing rates. It is much more attractive to lend money for successive 6-month periods to borrowers who are always AA at the beginning of the periods than to lend it to one borrower for the whole 5 years when all we can be sure of is that the borrower is AA at the beginning of the 5 years.

7.6 DETERMINING LIBOR/SWAP ZERO RATES

We explained in Section 4.1 that derivatives traders have traditionally used LIBOR rates as proxies for risk-free rates when valuing derivatives. One problem with LIBOR rates is that direct observations are possible only for maturities out to 12 months. As described in Section 6.3, one way of extending the LIBOR zero curve beyond 12 months is to use Eurodollar futures. Typically Eurodollar futures are used to produce a LIBOR zero curve out to 2 years—and sometimes out to as far as 5 years. Traders then use swap rates to extend the LIBOR zero curve further. The resulting zero curve is sometimes referred to as the LIBOR zero curve and sometimes as the swap zero curve. To avoid any confusion, we will refer to it as the LIBOR/swap zero curve. We will now describe how swap rates are used in the determination of the LIBOR/swap zero curve.

The first point to note is that the value of a newly issued floating-rate bond that pays 6-month LIBOR is always equal to its principal value (or par value) when the LIBOR/swap zero curve is used for discounting.⁸ The reason is that the bond provides a rate of interest of LIBOR, and LIBOR is the discount rate. The interest on the bond exactly matches the discount rate, and as a result the bond is fairly priced at par.

In equation (7.1), we showed that for a newly issued swap where the fixed rate equals the swap rate, $B_{\text{fix}} = B_{\text{fl}}$. We have just argued that B_{fl} equals the notional principal. It follows that B_{fix} also equals the swap's notional principal. Swap rates therefore define a set of par yield bonds. For example, from the swap rates in Table 7.3, we can deduce that the 2-year LIBOR/swap par yield is 6.045%, the 3-year LIBOR/swap par yield is 6.225%, and so on.⁹

⁸ The same is of course true of a newly issued bond that pays 1-month, 3-month, or 12-month LIBOR.

⁹ Analysts frequently interpolate between swap rates before calculating the zero curve, so that they have swap rates for maturities at 6-month intervals. For example, for the data in Table 7.3 the 2.5-year swap rate would be assumed to be 6.135%; the 7.5-year swap rate would be assumed to be 6.696%; and so on.

Section 4.5 showed how the bootstrap method can be used to determine the Treasury zero curve from Treasury bond prices. It can be used with swap rates in a similar way to extend the LIBOR/swap zero curve.

Example 7.1

Suppose that the 6-month, 12-month, and 18-month LIBOR/swap zero rates have been determined as 4%, 4.5%, and 4.8% with continuous compounding and that the 2-year swap rate (for a swap where payments are made semiannually) is 5%. This 5% swap rate means that a bond with a principal of \$100 and a semiannual coupon of 5% per annum sells for par. It follows that, if R is the 2-year zero rate, then

$$2.5e^{-0.04 \times 0.5} + 2.5e^{-0.045 \times 1.0} + 2.5e^{-0.048 \times 1.5} + 102.5e^{-2R} = 100$$

Solving this, we obtain $R = 4.953\%$. (Note that this calculation is simplified in that it does not take the swap's day count conventions and holiday calendars into account. See Section 7.2.)

7.7 VALUATION OF INTEREST RATE SWAPS

We now move on to discuss the valuation of interest rate swaps. An interest rate swap is worth close to zero when it is first initiated. After it has been in existence for some time, its value may be positive or negative. There are two valuation approaches. The first regards the swap as the difference between two bonds; the second regards it as a portfolio of FRAs.

Valuation in Terms of Bond Prices

Principal payments are not exchanged in an interest rate swap. However, as illustrated in Table 7.2, we can assume that principal payments are both received and paid at the end of the swap without changing its value. By doing this, we find that, from the point of view of the floating-rate payer, a swap can be regarded as a long position in a fixed-rate bond and a short position in a floating-rate bond, so that

$$V_{\text{swap}} = B_{\text{fix}} - B_{\text{fl}}$$

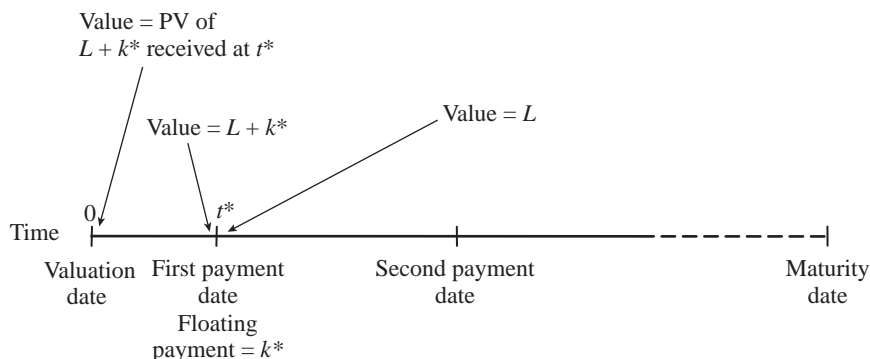
where V_{swap} is the value of the swap, B_{fl} is the value of the floating-rate bond (corresponding to payments that are made), and B_{fix} is the value of the fixed-rate bond (corresponding to payments that are received). Similarly, from the point of view of the fixed-rate payer, a swap is a long position in a floating-rate bond and a short position in a fixed-rate bond, so that the value of the swap is

$$V_{\text{swap}} = B_{\text{fl}} - B_{\text{fix}}$$

The value of the fixed rate bond, B_{fix} , can be determined as described in Section 4.4. To value the floating-rate bond, we note that the bond is worth the notional principal immediately after an interest payment. This is because at this time the bond is a "fair deal" where the borrower pays LIBOR for each subsequent accrual period.

Suppose that the notional principal is L , the next exchange of payments is at time t^* , and the floating payment that will be made at time t^* (which was determined at the last

Figure 7.8 Valuation of floating-rate bond when bond principal is L and next payment is k^* at t^*



payment date) is k^* . Immediately after the payment $B_{\text{fl}} = L$ as just explained. It follows that immediately before the payment $B_{\text{fl}} = L + k^*$. The floating-rate bond can therefore be regarded as an instrument providing a single cash flow of $L + k^*$ at time t^* . Discounting this, the value of the floating-rate bond today is $(L + k^*)e^{-r^*t^*}$, where r^* is the LIBOR/swap zero rate for a maturity of t^* . This argument is illustrated in Figure 7.8.

Example 7.2

Suppose that a financial institution has agreed to pay 6-month LIBOR and receive 8% per annum (with semiannual compounding) on a notional principal of \$100 million. The swap has a remaining life of 1.25 years. The LIBOR rates with continuous compounding for 3-month, 9-month, and 15-month maturities are 10%, 10.5%, and 11%, respectively. The 6-month LIBOR rate at the last payment date was 10.2% (with semiannual compounding).

The calculations for valuing the swap in terms of bonds are summarized in Table 7.5. The fixed-rate bond has cash flows of 4, 4, and 104 on the three payment dates. The discount factors for these cash flows are, respectively, $e^{-0.1 \times 0.25}$, $e^{-0.105 \times 0.75}$, and $e^{-0.11 \times 1.25}$ and are shown in the fourth column of Table 7.5. The table shows that the value of the fixed-rate bond (in millions of dollars) is 98.238.

Table 7.5 Valuing a swap in terms of bonds (\$ millions). Here, B_{fix} is fixed-rate bond underlying the swap, and B_{fl} is floating-rate bond underlying the swap.

| Time | B_{fix} cash flow | B_{fl} cash flow | Discount factor | Present value B_{fix} cash flow | Present value B_{fl} cash flow |
|--------|-------------------------------|------------------------------|--------------------|---|--|
| 0.25 | 4.0 | 105.100 | 0.9753 | 3.901 | 102.505 |
| 0.75 | 4.0 | | 0.9243 | 3.697 | |
| 1.25 | 104.0 | | 0.8715 | 90.640 | |
| Total: | | | | 98.238 | 102.505 |

In this example, $L = \$100$ million, $k^* = 0.5 \times 0.102 \times 100 = \5.1 million, and $t^* = 0.25$, so that the floating-rate bond can be valued as though it produces a cash flow of \$105.1 million in 3 months. The table shows that the value of the floating bond (in millions of dollars) is $105.100 \times 0.9753 = 102.505$.

The value of the swap is the difference between the two bond prices:

$$V_{\text{swap}} = 98.238 - 102.505 = -4.267$$

or -4.267 million dollars.

If the financial institution had been in the opposite position of paying fixed and receiving floating, the value of the swap would be $+\$4.267$ million. Note that these calculations do not take account of day count conventions and holiday calendars.

Valuation in Terms of FRAs

A swap can be characterized as a portfolio of forward rate agreements. Consider the swap between Microsoft and Intel in Figure 7.1. The swap is a 3-year deal entered into on March 5, 2012, with semiannual payments. The first exchange of payments is known at the time the swap is negotiated. The other five exchanges can be regarded as FRAs. The exchange on March 5, 2013, is an FRA where interest at 5% is exchanged for interest at the 6-month rate observed in the market on September 5, 2012; the exchange on September 5, 2013, is an FRA where interest at 5% is exchanged for interest at the 6-month rate observed in the market on March 5, 2013; and so on.

As shown at the end of Section 4.7, an FRA can be valued by assuming that forward interest rates are realized. Because it is nothing more than a portfolio of forward rate agreements, a plain vanilla interest rate swap can also be valued by making the assumption that forward interest rates are realized. The procedure is as follows:

1. Use the LIBOR/swap zero curve to calculate forward rates for each of the LIBOR rates that will determine swap cash flows.
2. Calculate swap cash flows on the assumption that the LIBOR rates will equal the forward rates.
3. Discount these swap cash flows (using the LIBOR/swap zero curve) to obtain the swap value.

Example 7.3

Consider again the situation in Example 7.2. Under the terms of the swap, a financial institution has agreed to pay 6-month LIBOR and receive 8% per annum (with semiannual compounding) on a notional principal of \$100 million. The swap has a remaining life of 1.25 years. The LIBOR rates with continuous compounding for 3-month, 9-month, and 15-month maturities are 10%, 10.5%, and 11%, respectively. The 6-month LIBOR rate at the last payment date was 10.2% (with semiannual compounding).

The calculations are summarized in Table 7.6. The first row of the table shows the cash flows that will be exchanged in 3 months. These have already been determined. The fixed rate of 8% will lead to a cash inflow of $100 \times 0.08 \times 0.5 = \4 million. The floating rate of 10.2% (which was set 3 months ago) will lead to a cash outflow of $100 \times 0.102 \times 0.5 = \5.1 million. The second row of the table shows the cash flows

Table 7.6 Valuing swap in terms of FRAs (\$ millions). Floating cash flows are calculated by assuming that forward rates will be realized.

| <i>Time</i> | <i>Fixed cash flow</i> | <i>Floating cash flow</i> | <i>Net cash flow</i> | <i>Discount factor</i> | <i>Present value of net cash flow</i> |
|---------------|------------------------|---------------------------|----------------------|------------------------|---------------------------------------|
| 0.25 | 4.0 | −5.100 | −1.100 | 0.9753 | −1.073 |
| 0.75 | 4.0 | −5.522 | −1.522 | 0.9243 | −1.407 |
| 1.25 | 4.0 | −6.051 | −2.051 | 0.8715 | −1.787 |
| <i>Total:</i> | | | | | −4.267 |

that will be exchanged in 9 months assuming that forward rates are realized. The cash inflow is \$4.0 million as before. To calculate the cash outflow, we must first calculate the forward rate corresponding to the period between 3 and 9 months. From equation (4.5), this is

$$\frac{0.105 \times 0.75 - 0.10 \times 0.25}{0.5} = 0.1075$$

or 10.75% with continuous compounding. From equation (4.4), the forward rate becomes 11.044% with semiannual compounding. The cash outflow is therefore $100 \times 0.11044 \times 0.5 = \5.522 million. The third row similarly shows the cash flows that will be exchanged in 15 months assuming that forward rates are realized. The discount factors for the three payment dates are, respectively,

$$e^{-0.1 \times 0.25}, \quad e^{-0.105 \times 0.75}, \quad e^{-0.11 \times 1.25}$$

The present value of the exchange in three months is $-\$1.073$ million. The values of the FRAs corresponding to the exchanges in 9 months and 15 months are $-\$1.407$ and $-\$1.787$ million, respectively. The total value of the swap is $-\$4.267$ million. This is in agreement with the value we calculated in Example 7.2 by decomposing the swap into bonds.

A swap is worth close to zero initially. This means that at the outset of a swap the sum of the values of the FRAs underlying the swap is close to zero. It does not mean that the value of each individual FRA is close to zero. In general, some FRAs will have positive values whereas others have negative values.

Consider the FRAs underlying the swap between Microsoft and Intel in Figure 7.1:


Value of FRA to Microsoft > 0 when forward interest rate $> 5.0\%$

Value of FRA to Microsoft $= 0$ when forward interest rate $= 5.0\%$

Value of FRA to Microsoft < 0 when forward interest rate $< 5.0\%$.

Suppose that the term structure of interest rates is upward-sloping at the time the swap is negotiated. This means that the forward interest rates increase as the maturity of the FRA increases. Since the sum of the values of the FRAs is close to zero, the forward interest rate must be less than 5.0% for the early payment dates and greater than 5.0% for the later payment dates. The value to Microsoft of the FRAs corresponding to early payment dates is therefore negative, whereas the value of the FRAs corresponding to later payment dates is positive. If the term structure of interest rates is downward-

Figure 7.9 Valuing of forward rate agreements underlying a swap as a function of maturity. In (a) the term structure of interest rates is upward-sloping and we receive fixed, or it is downward-sloping and we receive floating; in (b) the term structure of interest rates is upward-sloping and we receive floating, or it is downward-sloping and we receive fixed.



sloping at the time the swap is negotiated, the reverse is true. The impact of the shape of the term structure of interest rates on the values of the forward contracts underlying a swap is illustrated in Figure 7.9.

7.8 OVERNIGHT INDEXED SWAPS

Before leaving interest rate swaps, we discuss overnight indexed swaps. Since their introduction in the 1990s, they have become popular in all the major currencies. Their use arises from the fact that banks satisfy their liquidity needs at the end of each day by borrowing from and lending at an overnight rate. This rate is often a rate targeted by the central bank to influence monetary policy. In the United States, the rate is called the Fed Funds rate.

An overnight indexed swap (OIS) is a swap where a fixed rate for a period (e.g., 1 month, 3 months, 1 year, or 2 years) is exchanged for the geometric average of the overnight rates during the period. If during a certain period a bank borrows

funds at the overnight rate (rolling the loan forward each day), then its effective interest rate is the geometric average of the overnight interest rates. Similarly, if it lends money at the overnight interest rate every day, the effective rate of interest that it earns is the geometric average of the overnight interest rates. An OIS therefore allows overnight borrowing or lending to be swapped for borrowing or lending at a fixed rate. The fixed rate in an OIS is referred to as the overnight indexed swap rate.

A bank (Bank A) can engage in the following transactions:

1. Borrow \$100 million in the overnight market for 3 months, rolling the loan forward each night
2. Lend the \$100 million for 3 months at LIBOR to another bank (Bank B)
3. Use an OIS to exchange the overnight borrowings for fixed-rate borrowings.

This will lead to Bank A receiving the 3-month LIBOR rate and paying the 3-month overnight indexed swap rate. We might therefore expect the 3-month overnight indexed swap rate to equal the 3-month LIBOR rate. However, it is generally lower. This is because Bank A requires some compensation for the risk it is taking that Bank B will default on the LIBOR loan.

The excess of the 3-month LIBOR rate over the 3-month overnight indexed swap rate is known as the LIBOR–OIS spread. It is used as a measure of stress in financial markets. In normal market conditions, it is about 10 basis points. However, it rose sharply during the 2007–2009 credit crisis because banks became less willing to lend to each other. In October 2008, the spread spiked to an all time high of 364 basis points. By a year later, it had returned to more normal levels. It rose to over 30 basis points in June 2010 as a result of concerns about the financial health of Greece and a few other European countries.

The OIS rate is increasingly being regarded as a better proxy for the risk-free rate than LIBOR.


7.9 CURRENCY SWAPS

Another popular type of swap is known as a *currency swap*. In its simplest form, this involves exchanging principal and interest payments in one currency for principal and interest payments in another.

A currency swap agreement requires the principal to be specified in each of the two currencies. The principal amounts are usually exchanged at the beginning and at the end of the life of the swap. Usually the principal amounts are chosen to be approximately equivalent using the exchange rate at the swap's initiation. When they are exchanged at the end of the life of the swap, their values may be quite different.

Illustration

Consider a hypothetical 5-year currency swap agreement between IBM and British Petroleum entered into on February 1, 2011. We suppose that IBM pays a fixed rate of interest of 5% in sterling and receives a fixed rate of interest of 6% in dollars from British Petroleum. Interest rate payments are made once a year and the principal amounts are \$18 million and £10 million. This is termed a *fixed-for-fixed* currency swap because the interest rate in each currency is at a fixed rate. The swap is shown in Figure 7.10. Initially, the principal amounts flow in the opposite direction to the arrows in Figure 7.10. The

Figure 7.10 A currency swap.

interest payments during the life of the swap and the final principal payment flow in the same direction as the arrows. Thus, at the outset of the swap, IBM pays \$18 million and receives £10 million. Each year during the life of the swap contract, IBM receives \$1.08 million (= 6% of \$18 million) and pays £0.50 million (= 5% of £10 million). At the end of the life of the swap, it pays a principal of £10 million and receives a principal of \$18 million. These cash flows are shown in Table 7.7.

Use of a Currency Swap to Transform Liabilities and Assets

A swap such as the one just considered can be used to transform borrowings in one currency to borrowings in another. Suppose that IBM can issue \$18 million of US-dollar-denominated bonds at 6% interest. The swap has the effect of transforming this transaction into one where IBM has borrowed £10 million at 5% interest. The initial exchange of principal converts the proceeds of the bond issue from US dollars to sterling. The subsequent exchanges in the swap have the effect of swapping the interest and principal payments from dollars to sterling.

The swap can also be used to transform the nature of assets. Suppose that IBM can invest £10 million in the UK to yield 5% per annum for the next 5 years, but feels that the US dollar will strengthen against sterling and prefers a US-dollar-denominated investment. The swap has the effect of transforming the UK investment into a \$18 million investment in the US yielding 6%.

Comparative Advantage

Currency swaps can be motivated by comparative advantage. To illustrate this, we consider another hypothetical example. Suppose the 5-year fixed-rate borrowing costs to General Electric and Qantas Airways in US dollars (USD) and Australian dollars

Table 7.7 Cash flows to IBM in currency swap.

| <i>Date</i> | <i>Dollar cash flow (millions)</i> | <i>Sterling cash flow (millions)</i> |
|------------------|------------------------------------|--------------------------------------|
| February 1, 2011 | -18.00 | +10.00 |
| February 1, 2012 | +1.08 | -0.50 |
| February 1, 2013 | +1.08 | -0.50 |
| February 1, 2014 | +1.08 | -0.50 |
| February 1, 2015 | +1.08 | -0.50 |
| February 1, 2016 | +19.08 | -10.50 |

Table 7.8 Borrowing rates providing basis for currency swap.

| | USD* | AUD* |
|------------------|------|------|
| General Electric | 5.0% | 7.6% |
| Qantas Airways | 7.0% | 8.0% |

* Quoted rates have been adjusted to reflect the differential impact of taxes.

(AUD) are as shown in Table 7.8. The data in the table suggest that Australian rates are higher than USD interest rates, and also that General Electric is more creditworthy than Qantas Airways, because it is offered a more favorable rate of interest in both currencies. From the viewpoint of a swap trader, the interesting aspect of Table 7.8 is that the spreads between the rates paid by General Electric and Qantas Airways in the two markets are not the same. Qantas Airways pays 2% more than General Electric in the US dollar market and only 0.4% more than General Electric in the AUD market.

This situation is analogous to that in Table 7.4. General Electric has a comparative advantage in the USD market, whereas Qantas Airways has a comparative advantage in the AUD market. In Table 7.4, where a plain vanilla interest rate swap was considered, we argued that comparative advantages are largely illusory. Here we are comparing the rates offered in two different currencies, and it is more likely that the comparative advantages are genuine. One possible source of comparative advantage is tax. General Electric's position might be such that USD borrowings lead to lower taxes on its worldwide income than AUD borrowings. Qantas Airways' position might be the reverse. (Note that we assume that the interest rates shown in Table 7.8 have been adjusted to reflect these types of tax advantages.)

We suppose that General Electric wants to borrow 20 million AUD and Qantas Airways wants to borrow 15 million USD and that the current exchange rate (USD per AUD) is 0.7500. This creates a perfect situation for a currency swap. General Electric and Qantas Airways each borrow in the market where they have a comparative advantage; that is, General Electric borrows USD whereas Qantas Airways borrows AUD. They then use a currency swap to transform General Electric's loan into an AUD loan and Qantas Airways' loan into a USD loan.

As already mentioned, the difference between the USD interest rates is 2%, whereas the difference between the AUD interest rates is 0.4%. By analogy with the interest rate swap case, we expect the total gain to all parties to be $2.0 - 0.4 = 1.6\%$ per annum.

There are several ways in which the swap can be arranged. Figure 7.11 shows one way swaps might be entered into with a financial institution. General Electric borrows USD and Qantas Airways borrows AUD. The effect of the swap is to transform the USD

Figure 7.11 A currency swap motivated by comparative advantage.

Figure 7.12 Alternative arrangement for currency swap: Qantas Airways bears some foreign exchange risk.



interest rate of 5% per annum to an AUD interest rate of 6.9% per annum for General Electric. As a result, General Electric is 0.7% per annum better off than it would be if it went directly to AUD markets. Similarly, Qantas exchanges an AUD loan at 8% per annum for a USD loan at 6.3% per annum and ends up 0.7% per annum better off than it would be if it went directly to USD markets. The financial institution gains 1.3% per annum on its USD cash flows and loses 1.1% per annum on its AUD flows. If we ignore the difference between the two currencies, the financial institution makes a net gain of 0.2% per annum. As predicted, the total gain to all parties is 1.6% per annum.


Each year the financial institution makes a gain of USD 195,000 (= 1.3% of 15 million) and incurs a loss of AUD 220,000 (= 1.1% of 20 million). The financial institution can avoid any foreign exchange risk by buying AUD 220,000 per annum in the forward market for each year of the life of the swap, thus locking in a net gain in USD.

It is possible to redesign the swap so that the financial institution makes a 0.2% spread in USD. Figures 7.12 and 7.13 present two alternatives. These alternatives are unlikely to be used in practice because they do not lead to General Electric and Qantas being free of foreign exchange risk.¹⁰ In Figure 7.12, Qantas bears some foreign exchange risk because it pays 1.1% per annum in AUD and pays 5.2% per annum in USD. In Figure 7.13, General Electric bears some foreign exchange risk because it receives 1.1% per annum in USD and pays 8% per annum in AUD.

7.10 VALUATION OF CURRENCY SWAPS

Like interest rate swaps, fixed-for-fixed currency swaps can be decomposed into either the difference between two bonds or a portfolio of forward contracts.

Figure 7.13 Alternative arrangement for currency swap: General Electric bears some foreign exchange risk.



¹⁰ Usually it makes sense for the financial institution to bear the foreign exchange risk, because it is in the best position to hedge the risk.

Valuation in Terms of Bond Prices

If we define V_{swap} as the value in US dollars of an outstanding swap where dollars are received and a foreign currency is paid, then

$$V_{\text{swap}} = B_D - S_0 B_F$$

where B_F is the value, measured in the foreign currency, of the bond defined by the foreign cash flows on the swap and B_D is the value of the bond defined by the domestic cash flows on the swap, and S_0 is the spot exchange rate (expressed as number of dollars per unit of foreign currency). The value of a swap can therefore be determined from LIBOR rates in the two currencies, the term structure of interest rates in the domestic currency, and the spot exchange rate.

Similarly, the value of a swap where the foreign currency is received and dollars are paid is

$$V_{\text{swap}} = S_0 B_F - B_D$$

Example 7.4

Suppose that the term structure of LIBOR/swap interest rates is flat in both Japan and the United States. The Japanese rate is 4% per annum and the US rate is 9% per annum (both with continuous compounding). Some time ago a financial institution has entered into a currency swap in which it receives 5% per annum in yen and pays 8% per annum in dollars once a year. The principals in the two currencies are \$10 million and 1,200 million yen. The swap will last for another 3 years, and the current exchange rate is 110 yen = \$1.

The calculations are summarized in Table 7.9. In this case the cash flows from the dollar bond underlying the swap are as shown in the second column. The present value of the cash flows using the dollar discount rate of 9% are shown in the third column. The cash flows from the yen bond underlying the swap are shown in the fourth column of the table. The present value of the cash flows using the yen discount rate of 4% are shown in the final column of the table.

The value of the dollar bond, B_D , is 9.6439 million dollars. The value of the yen bond is 1230.55 million yen. The value of the swap in dollars is therefore

$$\frac{1,230.55}{110} - 9.6439 = 1.5430 \text{ million}$$

Table 7.9 Valuation of currency swap in terms of bonds. (All amounts in millions.)

| <i>Time</i> | <i>Cash flows on dollar bond (\$)</i> | <i>Present value (\$)</i> | <i>Cash flows on yen bond (yen)</i> | <i>Present value (yen)</i> |
|---------------|---|-------------------------------|---|--------------------------------|
| 1 | 0.8 | 0.7311 | 60 | 57.65 |
| 2 | 0.8 | 0.6682 | 60 | 55.39 |
| 3 | 0.8 | 0.6107 | 60 | 53.22 |
| 3 | 10.0 | 7.6338 | 1,200 | 1,064.30 |
| <i>Total:</i> | | 9.6439 | | 1,230.55 |

Valuation as Portfolio of Forward Contracts

Each exchange of payments in a fixed-for-fixed currency swap is a forward foreign exchange contract. In Section 5.7, forward foreign exchange contracts were valued by assuming that forward exchange rates are realized. The same assumption can therefore be made for a currency swap.

Example 7.5

Consider again the situation in Example 7.4. The LIBOR/swap term structure of interest rates is flat in both Japan and the United States. The Japanese rate is 4% per annum and the US rate is 9% per annum (both with continuous compounding). Some time ago a financial institution has entered into a currency swap in which it receives 5% per annum in yen and pays 8% per annum in dollars once a year. The principals in the two currencies are \$10 million and 1,200 million yen. The swap will last for another 3 years, and the current exchange rate is 110 yen = \$1.

The calculations are summarized in Table 7.10. The financial institution pays $0.08 \times 10 = \$0.8$ million dollars and receives $1,200 \times 0.05 = 60$ million yen each year. In addition, the dollar principal of \$10 million is paid and the yen principal of 1,200 is received at the end of year 3. The current spot rate is 0.009091 dollar per yen. In this case $r = 9\%$ and $r_f = 4\%$, so that, from equation (5.9), the 1-year forward rate is

$$0.009091 e^{(0.09-0.04) \times 1} = 0.009557$$

The 2- and 3-year forward rates in Table 7.10 are calculated similarly. The forward contracts underlying the swap can be valued by assuming that the forward rates are realized. If the 1-year forward rate is realized, the yen cash flow in year 1 is worth $60 \times 0.009557 = 0.5734$ million dollars and the net cash flow at the end of year 1 is $0.8 - 0.5734 = -0.2266$ million dollars. This has a present value of

$$-0.2266 e^{-0.09 \times 1} = -0.2071$$

million dollars. This is the value of forward contract corresponding to the exchange of cash flows at the end of year 1. The value of the other forward contracts are calculated similarly. As shown in Table 7.10, the total value of the forward contracts is \$1.5430 million. This agrees with the value calculated for the swap in Example 7.4 by decomposing it into bonds.

Table 7.10 Valuation of currency swap as a portfolio of forward contracts.
(All amounts in millions.)

| <i>Time</i> | <i>Dollar cash flow</i> | <i>Yen cash flow</i> | <i>Forward exchange rate</i> | <i>Dollar value of yen cash flow</i> | <i>Net cash flow (\$)</i> | <i>Present value</i> |
|---------------|-------------------------|----------------------|------------------------------|--------------------------------------|---------------------------|----------------------|
| 1 | -0.8 | 60 | 0.009557 | 0.5734 | -0.2266 | -0.2071 |
| 2 | -0.8 | 60 | 0.010047 | 0.6028 | -0.1972 | -0.1647 |
| 3 | -0.8 | 60 | 0.010562 | 0.6337 | -0.1663 | -0.1269 |
| 3 | -10.0 | 1200 | 0.010562 | 12.6746 | +2.6746 | 2.0417 |
| <i>Total:</i> | | | | | | 1.5430 |

The value of a currency swap is normally close to zero initially. If the two principals are worth the same at the start of the swap, the value of the swap is also close to zero immediately after the initial exchange of principal. However, as in the case of interest rate swaps, this does not mean that each of the individual forward contracts underlying the swap has a value close to zero. It can be shown that, when interest rates in two currencies are significantly different, the payer of the currency with the high interest rate is in the position where the forward contracts corresponding to the early exchanges of cash flows have negative values, and the forward contract corresponding to final exchange of principals has a positive value. The payer of the currency with the low interest rate is in the opposite position; that is, the forward contracts corresponding to the early exchanges of cash flows have positive values, while that corresponding to the final exchange has a negative value. These results are important when the credit risk in the swap is being evaluated.

7.11 CREDIT RISK


Contracts such as swaps that are private arrangements between two companies entail credit risks. Consider a financial institution that has entered into offsetting contracts with two companies (see Figure 7.4, 7.5, or 7.7). If neither party defaults, the financial institution remains fully hedged. A decline in the value of one contract will always be offset by an increase in the value of the other contract. However, there is a chance that one party will get into financial difficulties and default. The financial institution then still has to honor the contract it has with the other party.

Suppose that, some time after the initiation of the contracts in Figure 7.4, the contract with Microsoft has a positive value to the financial institution, whereas the contract with Intel has a negative value. If Microsoft defaults, the financial institution is liable to lose the whole of the positive value it has in this contract. To maintain a hedged position, it would have to find a third party willing to take Microsoft's position. To induce the third party to take the position, the financial institution would have to pay the third party an amount roughly equal to the value of its contract with Microsoft prior to the default.

A financial institution clearly has credit-risk exposure from a swap when the value of the swap to the financial institution is positive. What happens when this value is negative and the counterparty gets into financial difficulties? In theory, the financial institution could realize a windfall gain, because a default would lead to it getting rid of a liability. In practice, it is likely that the counterparty would choose to sell the contract to a third party or rearrange its affairs in some way so that its positive value in the contract is not lost. The most realistic assumption for the financial institution is therefore as follows. If the counterparty goes bankrupt, there will be a loss if the value of the swap to the financial institution is positive, and there will be no effect on the financial institution's position if the value of the swap to the financial institution is negative. This situation is summarized in Figure 7.14.

In swaps, it is sometimes the case that the early exchanges of cash flows have positive values and the later exchanges have negative values. (This would be true in Figure 7.9a and in a currency swap where the interest paid is lower than the interest received.) These swaps are likely to have negative values for most of their lives and therefore entail less credit risk than swaps where the reverse is true.

Figure 7.14 The credit exposure in a swap.



Potential losses from defaults on a swap are much less than the potential losses from defaults on a loan with the same principal. This is because the value of the swap is usually only a small fraction of the value of the loan. Potential losses from defaults on a currency swap are greater than on an interest rate swap. The reason is that, because principal amounts in two different currencies are exchanged at the end of the life of a currency swap, a currency swap is liable to have a greater value at the time of a default than an interest rate swap.

It is important to distinguish between the credit risk and market risk to a financial institution in any contract. As discussed earlier, the credit risk arises from the possibility of a default by the counterparty when the value of the contract to the financial institution is positive. The market risk arises from the possibility that market variables such as interest rates and exchange rates will move in such a way that the value of a contract to the financial institution becomes negative. Market risks can be hedged relatively easily by entering into offsetting contracts; credit risks are less easy to hedge.

One of the more bizarre stories in swap markets is outlined in Business Snapshot 7.2. It concerns the British Local Authority Hammersmith and Fulham and shows that, in addition to bearing market risk and credit risk, banks trading swaps also sometimes bear legal risk.

Clearing Houses

As explained in Business Snapshot 2.3, regulators are concerned about the potential for credit risk in the over-the-counter market to cause systemic risk. The volume of trading between financial institutions is huge. A default by one financial institution can lead to losses by other financial institutions. As a result, some of these financial institutions may also default, creating yet more losses for other financial institutions, more defaults, and so on. It was concerns about systemic risk that led governments to bail out financial institutions during the crisis that started in July 2007. To reduce systemic risk, governments have, since the crisis, introduced legislation requiring that clearing houses be used for many swaps and other derivatives. The way this works was discussed in Section 2.5. The clearing house acts as an intermediary between the two sides in a transaction. It requires an initial margin and variation margins in the same way that these are required for futures contracts.

Business Snapshot 7.2 The Hammersmith and Fulham Story

Between 1987 to 1989 the London Borough of Hammersmith and Fulham in the UK entered into about 600 interest rate swaps and related instruments with a total notional principal of about 6 billion pounds. The transactions appear to have been entered into for speculative rather than hedging purposes. The two employees of Hammersmith and Fulham responsible for the trades had only a sketchy understanding of the risks they were taking and how the products they were trading worked.

By 1989, because of movements in sterling interest rates, Hammersmith and Fulham had lost several hundred million pounds on the swaps. To the banks on the other side of the transactions, the swaps were worth several hundred million pounds. The banks were concerned about credit risk. They had entered into offsetting swaps to hedge their interest rate risks. If Hammersmith and Fulham defaulted, the banks would still have to honor their obligations on the offsetting swaps and would take a huge loss.

What happened was something a little different from a default. Hammersmith and Fulham's auditor asked to have the transactions declared void because Hammersmith and Fulham did not have the authority to enter into the transactions. The British courts agreed. The case was appealed and went all the way to the House of Lords, Britain's highest court. The final decision was that Hammersmith and Fulham did not have the authority to enter into the swaps, but that they ought to have the authority to do so in the future for risk-management purposes. Needless to say, banks were furious that their contracts were overturned in this way by the courts.

7.12 OTHER TYPES OF SWAPS

In this chapter, we have covered interest rate swaps where LIBOR is exchanged for a fixed rate of interest and currency swaps where a fixed rate of interest in one currency is exchanged for a fixed rate of interest in another currency. Many other types of swaps are traded. We will discuss some of them in detail in Chapters 24, 29, and 32. At this stage, we will provide an overview.

Variations on the Standard Interest Rate Swap

In fixed-for-floating interest rate swaps, LIBOR is the most common reference floating interest rate. In the examples in this chapter, the tenor (i.e., payment frequency) of LIBOR has been 6 months, but swaps where the tenor of LIBOR is 1 month, 3 months, and 12 months trade regularly. The tenor on the floating side does not have to match the tenor on the fixed side. (Indeed, as pointed out in footnote 3, the standard interest rate swap in the United States is one where there are quarterly LIBOR payments and semiannual fixed payments.) LIBOR is the most common floating rate, but others such as the commercial paper (CP) rate are occasionally used. Sometimes floating-for-floating interest rates swaps are negotiated. For example, the 3-month CP rate plus 10 basis points might be exchanged for 3-month LIBOR with both being applied to the same principal. (This deal would allow a company to hedge its exposure when assets and liabilities are subject to different floating rates.)

The principal in a swap agreement can be varied throughout the term of the swap to meet the needs of a counterparty. In an *amortizing swap*, the principal reduces in a predetermined way. (This might be designed to correspond to the amortization schedule on a loan.) In a *step-up swap*, the principal increases in a predetermined way. (This might be designed to correspond to drawdowns on a loan agreement.) Deferred swaps or *forward swaps*, where the parties do not begin to exchange interest payments until some future date, can also be arranged. Sometimes swaps are negotiated where the principal to which the fixed payments are applied is different from the principal to which the floating payments are applied.

A *constant maturity swap* (CMS swap) is an agreement to exchange a LIBOR rate for a swap rate. An example would be an agreement to exchange 6-month LIBOR applied to a certain principal for the 10-year swap rate applied to the same principal every 6 months for the next 5 years. A *constant maturity Treasury swap* (CMT swap) is a similar agreement to exchange a LIBOR rate for a particular Treasury rate (e.g., the 10-year Treasury rate).

In a *compounding swap*, interest on one or both sides is compounded forward to the end of the life of the swap according to preagreed rules and there is only one payment date at the end of the life of the swap. In a *LIBOR-in arrears* swap, the LIBOR rate observed on a payment date is used to calculate the payment on that date. (As explained in Section 7.1, in a standard deal the LIBOR rate observed on one payment date is used to determine the payment on the next payment date.) In an *accrual swap*, the interest on one side of the swap accrues only when the floating reference rate is in a certain range.

Other Currency Swaps

In this chapter we have considered fixed-for-fixed currency swaps. Another type of swap is a fixed-for-floating currency swap, whereby a floating rate (usually LIBOR) in one currency is exchanged for a fixed rate in another currency. This is a combination of a fixed-for-floating interest rate swap and a fixed-for-fixed currency swap and is known as a *cross-currency interest rate swap*. A further type of currency swap is a *floating-for-floating currency swap*, where a floating rate in one currency is exchanged for a floating rate in another currency.

Sometimes a rate observed in one currency is applied to a principal amount in another currency. One such deal might be where 3-month LIBOR observed in the United States is exchanged for 3-month LIBOR in Britain, with both rates being applied to a principal of 10 million British pounds. This type of swap is referred to as a *diff swap* or a *quanto*.

Equity Swaps

An *equity swap* is an agreement to exchange the total return (dividends and capital gains) realized on an equity index for either a fixed or a floating rate of interest. For example, the total return on the S&P 500 in successive 6-month periods might be exchanged for LIBOR, with both being applied to the same principal. Equity swaps can be used by portfolio managers to convert returns from a fixed or floating investment to the returns from investing in an equity index, and vice versa.

Options

Sometimes there are options embedded in a swap agreement. For example, in an *extendable swap*, one party has the option to extend the life of the swap beyond the specified period. In a *puttable swap*, one party has the option to terminate the swap early. Options on swaps, or *swaptions*, are also available. These provide one party with the right at a future time to enter into a swap where a predetermined fixed rate is exchanged for floating.

Commodity Swaps, Volatility Swaps, and Other Exotic Instruments

Commodity swaps are in essence a series of forward contracts on a commodity with different maturity dates and the same delivery prices. In a *volatility swap* there are a series of time periods. At the end of each period, one side pays a preagreed volatility, while the other side pays the historical volatility realized during the period. Both volatilities are multiplied by the same notional principal in calculating payments.

Swaps are limited only by the imagination of financial engineers and the desire of corporate treasurers and fund managers for exotic structures. In Chapter 32, we will describe the famous 5/30 swap entered into between Procter and Gamble and Bankers Trust, where payments depended in a complex way on the 30-day commercial paper rate, a 30-year Treasury bond price, and the yield on a 5-year Treasury bond.

SUMMARY

The two most common types of swaps are interest rate swaps and currency swaps. In an interest rate swap, one party agrees to pay the other party interest at a fixed rate on a notional principal for a number of years. In return, it receives interest at a floating rate on the same notional principal for the same period of time. In a currency swap, one party agrees to pay interest on a principal amount in one currency. In return, it receives interest on a principal amount in another currency.

Principal amounts are not usually exchanged in an interest rate swap. In a currency swap, principal amounts are usually exchanged at both the beginning and the end of the life of the swap. For a party paying interest in the foreign currency, the foreign principal is received, and the domestic principal is paid at the beginning of the swap's life. At the end of the swap's life, the foreign principal is paid and the domestic principal is received.

An interest rate swap can be used to transform a floating-rate loan into a fixed-rate loan, or vice versa. It can also be used to transform a floating-rate investment to a fixed-rate investment, or vice versa. A currency swap can be used to transform a loan in one currency into a loan in another currency. It can also be used to transform an investment denominated in one currency into an investment denominated in another currency.

There are two ways of valuing interest rate and currency swaps. In the first, the swap is decomposed into a long position in one bond and a short position in another bond. In the second it is regarded as a portfolio of forward contracts.

When a financial institution enters into a pair of offsetting swaps with different counterparties, it is exposed to credit risk. If one of the counterparties defaults when the financial institution has positive value in its swap with that counterparty, the financial institution loses money because it still has to honor its swap agreement with the other counterparty.

FURTHER READING

- Baz, J., and M. Pascutti. "Alternative Swap Contracts Analysis and Pricing," *Journal of Derivatives*, (Winter 1996): 7–21.
- Brown, K. C., and D.J. Smith. *Interest Rate and Currency Swaps: A Tutorial*. Association for Investment Management and Research, 1996.
- Cooper, I., and A. Mello. "The Default Risk in Interest Rate Swaps," *Journal of Finance*, 46, 2 (1991): 597–620.
- Dattatreya, R. E., and K. Hotta. *Advanced Interest Rate and Currency Swaps: State-of-the-Art Products, Strategies, and Risk Management Applications*. Irwin, 1993.
- Flavell, R. *Swaps and Other Instruments*. Chichester: Wiley, 2002.
- Gupta, A., and M. G. Subrahmanyam. "An Empirical Examination of the Convexity Bias in the Pricing of Interest Rate Swaps," *Journal of Financial Economics*, 55, 2 (2000): 239–79.
- Litzenberger, R. H. "Swaps: Plain and Fanciful," *Journal of Finance*, 47, 3 (1992): 831–50.
- Minton, B. A. "An Empirical Examination of the Basic Valuation Models for Interest Rate Swaps," *Journal of Financial Economics*, 44, 2 (1997): 251–77.
- Sun, T., S. Sundaresan, and C. Wang. "Interest Rate Swaps: An Empirical Investigation," *Journal of Financial Economics*, 34, 1 (1993): 77–99.
- Titman, S. "Interest Rate Swaps and Corporate Financing Choices," *Journal of Finance*, 47, 4 (1992): 1503–16.

Practice Questions (Answers in Solutions Manual)

- 7.1. Companies A and B have been offered the following rates per annum on a \$20 million 5-year loan:

| | <i>Fixed rate</i> | <i>Floating rate</i> |
|------------|-------------------|----------------------|
| Company A: | 5.0% | LIBOR + 0.1% |
| Company B: | 6.4% | LIBOR + 0.6% |

Company A requires a floating-rate loan; company B requires a fixed-rate loan. Design a swap that will net a bank, acting as intermediary, 0.1% per annum and that will appear equally attractive to both companies.

- 7.2. Company X wishes to borrow US dollars at a fixed rate of interest. Company Y wishes to borrow Japanese yen at a fixed rate of interest. The amounts required by the two companies are roughly the same at the current exchange rate. The companies are subject to the following interest rates, which have been adjusted to reflect the impact of taxes:

| | <i>Yen</i> | <i>Dollars</i> |
|------------|------------|----------------|
| Company X: | 5.0% | 9.6% |
| Company Y: | 6.5% | 10.0% |

Design a swap that will net a bank, acting as intermediary, 50 basis points per annum. Make the swap equally attractive to the two companies and ensure that all foreign exchange risk is assumed by the bank.

- 7.3. A \$100 million interest rate swap has a remaining life of 10 months. Under the terms of the swap, 6-month LIBOR is exchanged for 7% per annum (compounded semiannually). The average of the bid–offer rate being exchanged for 6-month LIBOR in swaps of all

maturities is currently 5% per annum with continuous compounding. The 6-month LIBOR rate was 4.6% per annum 2 months ago. What is the current value of the swap to the party paying floating? What is its value to the party paying fixed?

- 7.4. Explain what a swap rate is. What is the relationship between swap rates and par yields?
- 7.5. A currency swap has a remaining life of 15 months. It involves exchanging interest at 10% on £20 million for interest at 6% on \$30 million once a year. The term structure of interest rates in both the United Kingdom and the United States is currently flat, and if the swap were negotiated today the interest rates exchanged would be 4% in dollars and 7% in sterling. All interest rates are quoted with annual compounding. The current exchange rate (dollars per pound sterling) is 1.8500. What is the value of the swap to the party paying sterling? What is the value of the swap to the party paying dollars?
- 7.6. Explain the difference between the credit risk and the market risk in a financial contract.
- 7.7. A corporate treasurer tells you that he has just negotiated a 5-year loan at a competitive fixed rate of interest of 5.2%. The treasurer explains that he achieved the 5.2% rate by borrowing at 6-month LIBOR plus 150 basis points and swapping LIBOR for 3.7%. He goes on to say that this was possible because his company has a comparative advantage in the floating-rate market. What has the treasurer overlooked?
- 7.8. Explain why a bank is subject to credit risk when it enters into two offsetting swap contracts.
- 7.9. Companies X and Y have been offered the following rates per annum on a \$5 million 10-year investment:

| | <i>Fixed rate</i> | <i>Floating rate</i> |
|------------|-------------------|----------------------|
| Company X: | 8.0% | LIBOR |
| Company Y: | 8.8% | LIBOR |

Company X requires a fixed-rate investment; company Y requires a floating-rate investment. Design a swap that will net a bank, acting as intermediary, 0.2% per annum and will appear equally attractive to X and Y.

- 7.10. A financial institution has entered into an interest rate swap with company X. Under the terms of the swap, it receives 10% per annum and pays 6-month LIBOR on a principal of \$10 million for 5 years. Payments are made every 6 months. Suppose that company X defaults on the sixth payment date (at the end of year 3) when the interest rate (with semiannual compounding) is 8% per annum for all maturities. What is the loss to the financial institution? Assume that 6-month LIBOR was 9% per annum halfway through year 3.
- 7.11. Companies A and B face the following interest rates (adjusted for the differential impact of taxes):

| | <i>Company A</i> | <i>Company B</i> |
|--------------------------------|------------------|------------------|
| US dollars (floating rate): | LIBOR + 0.5% | LIBOR + 1.0% |
| Canadian dollars (fixed rate): | 5.0% | 6.5% |

Assume that A wants to borrow US dollars at a floating rate of interest and B wants to borrow Canadian dollars at a fixed rate of interest. A financial institution is planning to

- arrange a swap and requires a 50-basis-point spread. If the swap is to appear equally attractive to A and B, what rates of interest will A and B end up paying?
- 7.12. A financial institution has entered into a 10-year currency swap with company Y. Under the terms of the swap, the financial institution receives interest at 3% per annum in Swiss francs and pays interest at 8% per annum in US dollars. Interest payments are exchanged once a year. The principal amounts are 7 million dollars and 10 million francs. Suppose that company Y declares bankruptcy at the end of year 6, when the exchange rate is \$0.80 per franc. What is the cost to the financial institution? Assume that, at the end of year 6, the interest rate is 3% per annum in Swiss francs and 8% per annum in US dollars for all maturities. All interest rates are quoted with annual compounding.
 - 7.13. After it hedges its foreign exchange risk using forward contracts, is the financial institution's average spread in Figure 7.11 likely to be greater than or less than 20 basis points? Explain your answer.
 - 7.14. "Companies with high credit risks are the ones that cannot access fixed-rate markets directly. They are the companies that are most likely to be paying fixed and receiving floating in an interest rate swap." Assume that this statement is true. Do you think it increases or decreases the risk of a financial institution's swap portfolio? Assume that companies are most likely to default when interest rates are high.
 - 7.15. Why is the expected loss from a default on a swap less than the expected loss from the default on a loan with the same principal?
 - 7.16. A bank finds that its assets are not matched with its liabilities. It is taking floating-rate deposits and making fixed-rate loans. How can swaps be used to offset the risk?
 - 7.17. Explain how you would value a swap that is the exchange of a floating rate in one currency for a fixed rate in another currency.
 - 7.18. The LIBOR zero curve is flat at 5% (continuously compounded) out to 1.5 years. Swap rates for 2- and 3-year semiannual pay swaps are 5.4% and 5.6%, respectively. Estimate the LIBOR zero rates for maturities of 2.0, 2.5, and 3.0 years. (Assume that the 2.5-year swap rate is the average of the 2- and 3-year swap rates.)

Further Questions

- 7.19. (a) Company A has been offered the rates shown in Table 7.3. It can borrow for 3 years at 6.45%. What floating rate can it swap this fixed rate into?
(b) Company B has been offered the rates shown in Table 7.3. It can borrow for 5 years at LIBOR plus 75 basis points. What fixed rate can it swap this floating rate into?
- 7.20. (a) Company X has been offered the rates shown in Table 7.3. It can invest for 4 years at 5.5%. What floating rate can it swap this fixed rate into?
(b) Company Y has been offered the rates shown in Table 7.3. It can invest for 10 years at LIBOR minus 50 basis points. What fixed rate can it swap this floating rate into?
- 7.21. The 1-year LIBOR rate is 10% with annual compounding. A bank trades swaps where a fixed rate of interest is exchanged for 12-month LIBOR with payments being exchanged annually. The 2- and 3-year swap rates (expressed with annual compounding) are 11% and 12% per annum. Estimate the 2- and 3-year LIBOR zero rates.

- 7.22. Company A, a British manufacturer, wishes to borrow US dollars at a fixed rate of interest. Company B, a US multinational, wishes to borrow sterling at a fixed rate of interest. They have been quoted the following rates per annum (adjusted for differential tax effects):

| | <i>Sterling</i> | <i>US dollars</i> |
|-----------|-----------------|-------------------|
| Company A | 11.0% | 7.0% |
| Company B | 10.6% | 6.2% |

- Design a swap that will net a bank, acting as intermediary, 10 basis points per annum and that will produce a gain of 15 basis points per annum for each of the two companies.
- 7.23. Under the terms of an interest rate swap, a financial institution has agreed to pay 10% per annum and to receive 3-month LIBOR in return on a notional principal of \$100 million with payments being exchanged every 3 months. The swap has a remaining life of 14 months. The average of the bid and offer fixed rates currently being swapped for 3-month LIBOR is 12% per annum for all maturities. The 3-month LIBOR rate 1 month ago was 11.8% per annum. All rates are compounded quarterly. What is the value of the swap?
- 7.24. Suppose that the term structure of interest rates is flat in the United States and Australia. The USD interest rate is 7% per annum and the AUD rate is 9% per annum. The current value of the AUD is 0.62 USD. Under the terms of a swap agreement, a financial institution pays 8% per annum in AUD and receives 4% per annum in USD. The principals in the two currencies are \$12 million USD and 20 million AUD. Payments are exchanged every year, with one exchange having just taken place. The swap will last 2 more years. What is the value of the swap to the financial institution? Assume all interest rates are continuously compounded.
- 7.25. Company X is based in the United Kingdom and would like to borrow \$50 million at a fixed rate of interest for 5 years in US funds. Because the company is not well known in the United States, this has proved to be impossible. However, the company has been quoted 12% per annum on fixed-rate 5-year sterling funds. Company Y is based in the United States and would like to borrow the equivalent of \$50 million in sterling funds for 5 years at a fixed rate of interest. It has been unable to get a quote but has been offered US dollar funds at 10.5% per annum. Five-year government bonds currently yield 9.5% per annum in the United States and 10.5% in the United Kingdom. Suggest an appropriate currency swap that will net the financial intermediary 0.5% per annum.



CHAPTER

8

Securitization and the Credit Crisis of 2007

Derivatives such as forwards, futures, swaps, and options are concerned with transferring risk from one entity in the economy to another. The first seven chapters of this book have focused on forwards, futures, and swaps. Before moving on to discuss options, we consider another important way of transferring risk in the economy: securitization.

Securitization is of particular interest because of its role in the credit crisis (sometimes referred to as the “credit crunch”) that started in 2007. The crisis had its origins in financial products created from mortgages in the United States, but rapidly spread from the United States to other countries and from financial markets to the real economy. Some financial institutions failed. Many more had to be rescued by national governments. There can be no question that the first decade of the twenty-first century was disastrous for the financial sector.

In this chapter, we examine the nature of securitization and its role in the crisis. In the course of the chapter, we will learn about the US mortgage market, asset-backed securities, collateralized debt obligations, waterfalls, and the importance of incentives in financial markets.

8.1 SECURITIZATION

Traditionally, banks have funded their loans primarily from deposits. In the 1960s, US banks found that they could not keep pace with the demand for residential mortgages with this type of funding. This led to the development of the mortgage-backed security (MBS) market. Portfolios of mortgages were created and the cash flows (interest and principal payments) generated by the portfolios were packaged as securities and sold to investors. The US government created the Government National Mortgage Association (GNMA, also known as Ginnie Mae) in 1968. This organization guaranteed (for a fee) interest and principal payments on qualifying mortgages and created the securities that were sold to investors.

Thus, although banks originated the mortgages, they did not keep them on their balance sheets. Securitization allowed them to increase their lending faster than their


deposits were growing. GNMA's guarantee protected MBS investors against defaults by borrowers.¹

In the 1980s, the securitization techniques developed for the mortgage market were applied to asset classes such as automobile loans and credit card receivables in the United States. Securitization also became popular in other parts of the world. As the securitization market developed, investors became comfortable with situations where they did not have a guarantee against defaults by borrowers.

ABSs


A simple securitization arrangement of the type used during the 2000 to 2007 period is shown in Figure 8.1. This is known as an *asset-backed security* or ABS. A portfolio of income-producing assets such as loans is sold by the originating banks to a special purpose vehicle (SPV) and the cash flows from the assets are then allocated to tranches. Figure 8.1 is simplified in that there are three tranches. These are the senior tranche, the mezzanine tranche, and the equity tranche. The portfolio has a principal of \$100 million. This is divided as follows: \$80 million to the senior tranche, \$15 million to the mezzanine tranche, and \$5 million to the equity tranche. The senior tranche is promised a return of LIBOR plus 60 basis points, the mezzanine tranche is promised a return of LIBOR plus 250 basis points, and the equity tranche is promised a return of LIBOR plus 2,000 basis points.

Figure 8.1 An asset-backed security (simplified); bp = basis points (1bp = 0.01%).



¹ However, MBS investors do face uncertainty about mortgage prepayments. Prepayments tend to be greatest when interest rates are low and the reinvestment opportunities open to investors are not particularly attractive. In the early days of MBSs, many MBS investors realized lower returns than they expected because they did not take this into account.

Figure 8.2 The waterfall in an asset-backed security.



It sounds as though the equity tranche has the best deal, but this is not necessarily the case. The payments of interest and principal are not guaranteed. The equity tranche is more likely to lose part of its principal, and less likely to receive the promised interest payments on its outstanding principal, than the other tranches. Cash flows are allocated to tranches by specifying what is known as a waterfall. The general way a waterfall works is illustrated in Figure 8.2. A separate waterfall is applied to interest payments and the repayments of principal on the assets. Principal repayments are allocated to the senior tranche until its principal has been fully repaid. They are then allocated to the mezzanine tranche until its principal has been fully repaid. Only after this has happened do principal repayments go to the equity tranche. Interest payments are allocated to the senior tranche until the senior tranche has received its promised return on its outstanding principal. Assuming that this promised return can be made, interest payments are then allocated to the mezzanine tranche. If the promised return to the mezzanine tranche can be made and cash flows are left over, they are allocated to the equity tranche.

The extent to which the tranches get their principal back depends on losses on the underlying assets. The effect of the waterfall is roughly as follows. The first 5% of losses are borne by the equity tranche. If losses exceed 5%, the equity tranche loses all its principal and some losses are borne by the principal of the mezzanine tranche. If losses exceed 20%, the mezzanine tranche loses all its principal and some losses are borne by the principal of the senior tranche.

There are therefore two ways of looking at an ABS. One is with reference to the waterfall in Figure 8.2. Cash flows go first to the senior tranche, then to the mezzanine tranche, and then to the equity tranche. The other is in terms of losses. Losses of principal are first borne by the equity tranche, then by the mezzanine tranche, and then by the senior tranche. Rating agencies such as Moody's, S&P, and Fitch played a key role in securitization. The ABS in Figure 8.1 is designed so that the senior tranche is rated AAA. The mezzanine tranche is typically rated BBB. The equity tranche is typically unrated.

The description of ABSs that we have given so far is somewhat simplified. Typically, more than three tranches with a wide range of ratings are created. In the waterfall rules, as we have described them, the allocation of cash flows to tranches is sequential in that they always flow first to the most senior tranche, then to the next-most-senior tranche, and so on. In practice, the rules are somewhat more complicated than this and are described in a legal document that is several hundred pages long. Another complication is that there is often some over-collateralization where (a) the total principal of the tranches is less than the total principal of the underlying assets and (b) the weighted average return promised to the tranches is less than the weighted average return payable on the assets.

ABS CDOs

Finding investors to buy the senior AAA-rated tranches of ABSs was usually not difficult because the tranches promised returns which were very attractive when compared with the return on AAA-rated bonds. Equity tranches were typically retained by the originator of the assets or sold to a hedge fund.

Finding investors for mezzanine tranches was more difficult. This led to the creation of ABSs of ABSs. The way this was done is indicated in Figure 8.3. Many different mezzanine tranches, created in the way indicated in Figure 8.1, are put in a portfolio and the risks associated with the cash flows from the portfolio are tranching out in the same way as the risks associated with the cash flows from the assets are tranching out in Figure 8.1. The resulting structure is known as an *ABS CDO* or *Mezz ABS CDO*. In the example in Figure 8.3, the senior tranche of the ABS CDO accounts for 65% of the principal of the ABS mezzanine tranches, the mezzanine tranche of the ABS CDO accounts for 25% of the principal, and the equity tranche accounts for the remaining 10% of the principal. The structure is designed so that the senior tranche of the ABS CDO is rated AAA. This means that the total of the AAA-rated instruments created in the example that is considered here is about 90% (80% plus 65% of 15%) of the principal of the underlying portfolio. This seems high but, if the securitization were carried further with an ABS being created from the mezzanine tranches of ABS CDOs (and this did happen), the percentage would be pushed even higher.

In the example in Figure 8.3, the AAA-rated tranche of the ABS can expect to receive its promised return and get its principal back if losses on the underlying portfolio of

Figure 8.3 An ABS CDO (simplified).

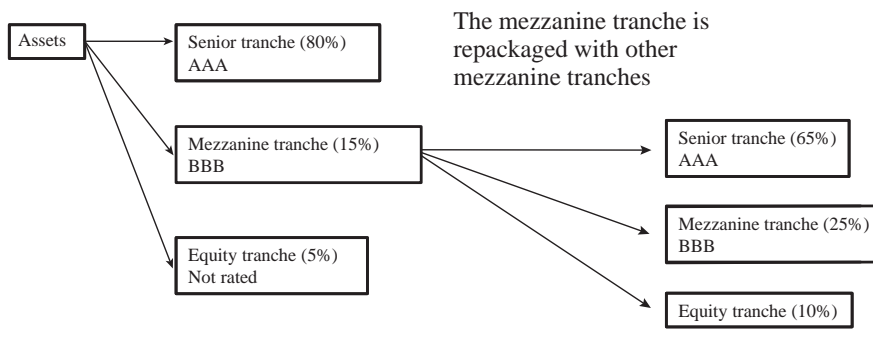


Table 8.1 Estimated losses to AAA-rated tranches of ABS CDO in Figure 8.3.

| <i>Losses on underlying assets</i> | <i>Losses to mezzanine tranche of ABS</i> | <i>Losses to equity tranche of ABS CDO</i> | <i>Losses to mezzanine tranche of ABS CDO</i> | <i>Losses to senior tranche of ABS CDO</i> |
|------------------------------------|---|--|---|--|
| 10% | 33.3% | 100.0% | 93.3% | 0.0% |
| 13% | 53.3% | 100.0% | 100.0% | 28.2% |
| 17% | 80.0% | 100.0% | 100.0% | 69.2% |
| 20% | 100.0% | 100.0% | 100.0% | 100.0% |

assets is less than 20% because all losses of principal would then be absorbed by the more junior tranches. The AAA-rated tranche of the ABS CDO in Figure 8.2 is more risky. It will receive the promised return and get its principal back if losses on the underlying assets are less than 10.25%. This is because a loss of 10.25% means that mezzanine tranches of ABSs have to absorb losses equal to 5.25% of the ABS principal. As these tranches have a total principal equal to 15% of the ABS principal, they lose 5.25/15 or 35% of their principal. The equity and mezzanine tranches of the ABS CDO are then wiped out, but the senior tranche just manages to survive intact.

The senior tranche of the ABS CDO suffers losses if losses on the underlying portfolios are more than 10.25%. Consider, for example, the situation where losses are 17% on the underlying portfolios. Of the 17%, 5% is borne by the equity tranche of the ABS and 12% by the mezzanine tranche of the ABS. Losses on the mezzanine tranches are therefore 12/15 or 80% of their principal. The first 35% is absorbed by the equity and mezzanine tranches of the ABS CDO. The senior tranche of the ABS CDO therefore loses 45/65 or 69.2% of its value. These and other results are summarized in Table 8.1.

8.2 THE US HOUSING MARKET


Figure 8.4 gives the S&P/Case–Shiller composite-10 index for house prices in the US between January 1987 and March 2010. This tracks house prices for the top ten metropolitan areas in the US. It shows that, in about the year 2000, house prices started to rise much faster than they had in the previous decade. The very low level of interest rates between 2002 and 2005 was an important contributory factor, but the bubble in house prices was largely fueled by mortgage-lending practices.

The 2000 to 2006 period was characterized by a huge increase in what is termed subprime mortgage lending. Subprime mortgages are mortgages that are considered to be significantly more risky than average. Before 2000, most mortgages classified as subprime were second mortgages. After 2000, this changed as financial institutions became more comfortable with the notion of a subprime first mortgage.

The Relaxation of Lending Standards

The relaxation of lending standards and the growth of subprime mortgages made house purchase possible for many families that had previously been considered to be not sufficiently creditworthy to qualify for a mortgage. These families increased the demand

Figure 8.4 The S&P/Case–Shiller Composite-10 index of U.S. real estate prices, 1987–2010.



for real estate and prices rose. To mortgage brokers and mortgage lenders, the combination of more lending and higher house prices was attractive. More lending meant bigger profits. Higher house prices meant that the lending was well covered by the underlying collateral. If the borrower defaulted, the resulting foreclosure would not lead to a loss.

Mortgage brokers and mortgage lenders naturally wanted to keep increasing their profits. Their problem was that, as house prices rose, it was more difficult for first-time buyers to afford a house. In order to continue to attract new entrants to the housing market, they had to find ways to relax their lending standards even more—and this is exactly what they did. The amount lent as a percentage of the house price increased. Adjustable-rate mortgages (ARMs) were developed where there was a low “teaser” rate of interest that would last for two or three years and be followed by a rate that was much higher.² A typical teaser rate was about 6% and the interest rate after the end of the teaser rate period was typically six-month LIBOR plus 6%.³ However, teaser rates as low as 1% or 2% have been reported. Lenders also became more cavalier in the way they reviewed mortgage applications. Indeed, the applicant’s income and other information reported on the application form were frequently not checked.

Subprime Mortgage Securitization

Subprime mortgages were frequently securitized in the way indicated in Figures 8.1 to 8.3. The investors in tranches created from subprime mortgages usually had no guarantees that interest and principal would be paid. Securitization played a part in the

² If real estate prices increased, lenders expected the borrowers to prepay and take out a new mortgage at the end of the teaser rate period. However, prepayment penalties, often zero on prime mortgages, were quite high on subprime mortgages.

³ A “2/28” ARM, for example, is an ARM where the rate is fixed for two years and then floats for the remaining 28 years.

crisis. The behavior of mortgage originators was influenced by their knowledge that mortgages would be securitized.⁴ When considering new mortgage applications, the question was not “Is this a credit we want to assume?” Instead it was “Is this a mortgage we can make money on by selling it to someone else?”

When mortgages were securitized, the only information received about the mortgages by the buyers of the products that were created from them was the loan-to-value ratio (i.e., the ratio of the size of the loan to the assessed value of the house) and the borrower’s FICO score.⁵ Other information on the mortgage application form was considered irrelevant and, as already mentioned, was often not even checked by lenders. The most important thing for the lender was whether the mortgage could be sold to others—and this depended largely on the loan-to-value ratio and the applicant’s FICO score.

It is interesting to note in passing that both the loan-to-value ratio and the FICO score were of doubtful quality. The property assessors who determined the value of a house at the time of a mortgage application sometimes succumbed to pressure from the lenders to come up with high values. Potential borrowers were sometimes counseled to take certain actions that would improve their FICO scores.⁶

Why was the government not regulating the behavior of mortgage lenders? The answer is that the US government had since the 1990s been trying to expand home ownership and had been applying pressure to mortgage lenders to increase loans to low- and moderate-income people. Some state legislators, such as those in Ohio and Georgia, were concerned about what was going on and wanted to curtail predatory lending.⁷ However, the courts decided that national standards should prevail.

A number of terms have been used to describe mortgage lending during the period leading up to the credit crunch. One is “liar loans” because individuals applying for a mortgage, knowing that no checks would be carried out, sometimes chose to lie on the application form. Another term used to describe some borrowers is “NINJA” (no income, no job, no assets).

The Bubble Bursts

One of the features of the US housing market is that mortgages are nonrecourse in many states. This means that, when there is a default, the lender is able to take possession of the house, but other assets of the borrower are off-limits. Consequently, the borrower has a free American-style put option. He or she can at any time sell the house to the lender for the principal outstanding on the mortgage. This feature encouraged speculative activity and was part of the cause of the house price bubble shown in Figure 8.4 that occurred between 2000 and 2006.

All bubbles burst eventually and this one was no exception. In 2007, many mortgage holders found that they could no longer afford their mortgages when the teaser rates ended. This led to foreclosures and large numbers of houses coming on the market,

⁴ See B.J. Keys, T. Mukherjee, A. Seru, and V. Vig, “Did Securitization Lead to Lax Screening? Evidence from Subprime Loans,” *Quarterly Journal of Economics*, 125, 1 (February 2010): 307–62.

⁵ FICO is a credit score developed by the Fair Isaac Corporation and is widely used in the US. It ranges from 300 to 850.

⁶ One such action might be to make regular payments on a new credit card for a few months.

⁷ Predatory lending describes the situation where a lender deceptively convinces borrowers to agree to unfair and abusive loan terms.

which in turn led to a decline in house prices. Other mortgage holders, who had borrowed 100%, or close to 100%, of the cost of a house found that they had negative equity. Market participants realized belatedly how costly the free put option could be. If the borrower had negative equity, the optimal decision was to exchange the house for the outstanding principal on the mortgage. The house was then sold by the lender, adding to the downward pressure on house prices.

It would be a mistake to assume that all mortgage defaulters were in the same position. Some were unable to meet mortgage payments and suffered greatly when they had to give up their homes. But many of the defaulters were speculators who bought multiple homes as rental properties and chose to exercise their put options. It was their tenants who suffered. There are also reports that some house owners (who were not speculators) were quite creative in extracting value from their put options. After handing the keys to their houses to the lender, they turned around and bought (sometimes at a bargain price) other houses that were in foreclosure. Imagine two people owning identical houses next to each other. Both have mortgages of \$250,000. Both houses are worth \$200,000 and in foreclosure can be expected to sell for \$170,000. What is the owners' optimal strategy? The answer is that each person should exercise the put option and buy the neighbor's house. (There were ways of doing this without getting a bad credit rating.)

The United States was not alone in having declining real estate prices. Prices declined in many other countries as well. Real estate prices in the United Kingdom were particularly badly affected.

The Losses

As foreclosures increased, the losses on mortgages also increased. It might be thought that a 35% reduction in house prices would lead to at most a 35% loss of principal on a defaulting mortgage. In fact, the losses were far greater than that. Houses in foreclosure were often in poor condition and sold for a small fraction of their values prior to the credit crisis. In 2008 and 2009, average losses as high as 75% were reported for mortgages on houses in foreclosure in some areas.

Investors in tranches that were formed from mortgages incurred big losses. The value of the ABS tranches created from subprime mortgages was monitored by a series of indices known as ABX. These indices indicated that the tranches originally rated BBB had lost about 80% of their value by the end of 2007 and about 97% of their value by mid-2009. The value of the ABS CDO tranches created from BBB tranches was monitored by a series of indices known as TABX. These indices indicated that the tranches originally rated AAA lost about 80% of their value by the end of 2007 and were essentially worthless by mid-2009.

Financial institutions such as UBS, Merrill Lynch, and Citigroup had big positions in some of the tranches and incurred huge losses, as did the insurance giant AIG, which provided protection against losses on ABS CDO tranches that had originally been rated AAA. Many financial institutions had to be rescued with government funds. There have been few worse years in financial history than 2008. Bear Stearns was taken over by J.P. Morgan Chase; Merrill Lynch was taken over by Bank of America; Goldman Sachs and Morgan Stanley, which had formerly been investment banks, became bank holding companies with both commercial and investment banking interests; and Lehman Brothers was allowed to fail.

The Credit Crisis

The losses on securities backed by residential mortgages led to a severe credit crisis. In 2006, banks were well capitalized, loans were relatively easy to obtain, and credit spreads were low. (The credit spread is the excess of the interest rate on a loan over the risk-free interest rate.) By 2008, the situation was totally different. The capital of banks had been badly eroded by their losses. They had become much more risk-averse and were reluctant to lend. Creditworthy individuals and corporations found borrowing difficult. Credit spreads had increased dramatically. The world experienced its worst recession in several generations. As discussed in Section 7.8, the LIBOR–OIS spread briefly reached 364 basis points in October 2008, indicating an extreme reluctance of banks to lend to each other. Another measure of this is the TED spread. This is the excess of the three-month LIBOR interest rate over the three-month Treasury interest. In normal market conditions, it is 30 to 50 basis points. It reached over 450 basis points in October 2008.

8.3 WHAT WENT WRONG?

“Irrational exuberance” is a phrase coined by Alan Greenspan, Chairman of the Federal Reserve Board, to describe the behavior of investors during the bull market of the 1990s. It can also be applied to the period leading up to the credit crisis. Mortgage lenders, the investors in tranches of ABSs and ABS CDOs that were created from residential mortgages, and the companies that sold protection on the tranches assumed that the good times would last for ever. They thought that US house prices would continue to increase. There might be declines in one or two areas, but the possibility of the widespread decline indicated by Figure 8.4 was a scenario not considered by most people.

Many factors contributed to the crisis that started in 2007. Mortgage originators used lax lending standards. Products were developed to enable mortgage originators to profitably transfer credit risk to investors. Rating agencies moved from their traditional business of rating bonds, where they had a great deal of experience, to rating structured products, which were relatively new and for which there were relatively little historical data. The products bought by investors were complex and in many instances investors and rating agencies had inaccurate or incomplete information about the quality of the underlying assets. Investors in the structured products that were created thought they had found a money machine and chose to rely on rating agencies rather than forming their own opinions about the underlying risks. The return earned by the products rated AAA was high compared with the returns on bonds rated AAA.

Structured products such as those in Figures 8.1 and 8.3 are highly dependent on the default correlation between the underlying assets. Default correlation measures the tendency for different borrowers to default at about the same time. If the default correlation between the underlying assets in Figure 8.1 is low, the AAA-rated tranches are extremely unlikely to experience losses. As this default correlation increases, they become more vulnerable. The tranches of ABS CDOs in Figure 8.3 are even more heavily dependent on default correlation.

If mortgages exhibit moderate default correlation (as they do in normal times), there is very little chance of a high overall default rate and the AAA-rated tranches of both

ABSs and ABS CDOs that are created from mortgages are fairly safe. However, as many investors found to their cost, default correlations tend to increase in stressed market conditions. This makes very high default rates possible.

There was a tendency to assume that a tranche with a particular rating could be equated to a bond with the that rating. The rating agencies published the criteria they used for rating tranches. S&P and Fitch rated a tranche so as to ensure that the probability of the tranche experiencing a loss was the same as the probability of a similarly rated bond experiencing a loss. Moody's rated a tranche so that the expected loss from the tranche was the the same as the expected loss from a similarly rated bond.⁸ The procedures used by rating agencies were therefore designed to ensure that one aspect of the loss distributions of tranches and bonds were matched. However, other aspects of the distributions were liable to be quite different.

The differences between tranches and bonds were accentuated by the fact tranches were often quite thin. The AAA tranches often accounted for about 80% of the principal as in Figure 8.1, but it was not unusual for there to be 15 to 20 other tranches. Each of these tranches would be 1% or 2% wide. Such thin tranches are likely to either incur no losses or be totally wiped out. The chance of investors recovering part of their principal (as bondholders usually do) is small. Consider, for example, a BBB tranche that is responsible for losses in the range 5% to 6%. If losses on the underlying portfolio are less than 5%, the tranche is safe. If losses are greater than 6%, the tranche is wiped out. Only in the case where losses are between 5% and 6% is a partial recovery made by investors.

The difference between a thin BBB-rated tranche and a BBB-rated bond was overlooked by many investors. The difference makes the tranches of ABS CDOs created from the BBB-rated tranches of ABSs much riskier than CDOs created in a similar way from BBB bonds. Losses on a portfolio of BBB bonds can reasonably be assumed to be unlikely to exceed 25% in stressed market conditions. Table 8.1 shows that 100% losses on a portfolio of BBB tranches can occur relatively easily—and this is even more true when the tranches are only 1% or 2% wide.

Regulatory Arbitrage

Many of the mortgages were originated by banks and it was banks that were the main investors in the tranches that were created from the mortgages. Why would banks choose to securitize mortgages and then buy the securitized products that were created? The answer concerns what is termed *regulatory arbitrage*. The regulatory capital banks were required to keep for the tranches created from a portfolio of mortgages was much less than the regulatory capital that would be required for the mortgages themselves.

Incentives

One of the lessons from the crisis is the importance of incentives. Economists use the term “agency costs” to describe the situation where incentives are such that the interests of two parties in a business relationship are not perfectly aligned. The process by which

⁸ For a discussion of the criteria used by rating agencies and the reasonableness of the ratings given the criteria, see J. Hull and A. White, “Ratings Arbitrage and Structured Products,” Working Paper, University of Toronto, and “The Risk of the Tranches Created from Mortgages,” *Financial Analysts Journal*, 66, 5 (2010): 54–67.

mortgages were originated, securitized, and sold to investors was unfortunately riddled with agency costs.

The incentive of the originators of mortgages was to make loans that would be acceptable to the creators of the ABS and ABS CDO tranches. The incentive of the individuals who valued houses on which the mortgages were written was to please the lender by providing as high a valuation as possible so that the loan-to-value ratio was as low as possible. (Pleasing the lender was likely to lead to more business from that lender.) The main concern of the creators of tranches was how the tranches would be rated. They wanted the volume of AAA-rated tranches that they created to be as high as possible and found ways of using the published criteria of rating agencies to achieve this. The rating agencies were paid by the issuers of the securities they rated and about half their income came from structured products.

Another source of agency costs concerns the incentives of the employees of financial institutions. Employee compensation falls into three categories: regular salary, the end-of-year bonus, and stock or stock options. Many employees at all levels of seniority in financial institutions, particularly traders, receive much of their compensation in the form of end-of-year bonuses. This form of compensation is focused on short-term performance. If an employee generates huge profits one year and is responsible for severe losses the next, the employee will receive a big bonus the first year and will not have to return it the following year. (The employee might lose his or her job as a result of the second year losses, but even that is not a disaster. Financial institutions seem to be surprisingly willing to recruit individuals with losses on their résumés.)

Imagine you are an employee of a financial institution in 2006 responsible for investing in ABS CDOs created from mortgages. Almost certainly you would have recognized that there was a bubble in the US housing market and would expect that bubble to burst sooner or later. However, it is possible that you would decide to continue with your ABS CDO investments. If the bubble did not burst until after the end of 2006, you would still get a nice bonus at the end of 2006.

8.4 THE AFTERMATH

Banks throughout the world are regulated by the Basel Committee on Banking Supervision and are subject to the legislation enacted by the governments of the countries in which they operate. One of the results of the credit crisis has been a “tsunami” of new regulation and new legislation.

The Basel Committee on Banking Supervision provides international standards which are applied by bank supervisors in countries throughout the world. The regulations produced by the committee prior to the credit crisis have become known as Basel I and Basel II, and are summarized in Business Snapshot 8.1. The regulations are mostly concerned with the amount of capital banks should be required to keep for the risks they are taking.⁹ At the end of 2009, the committee proposed what has been termed Basel III. This increases the amount of capital and the quality of the capital that banks are required to keep. It also requires banks to satisfy certain liquidity requirements. One of the lessons from the crisis is that the failures of financial institutions are frequently

⁹ For more details on the work of the Basel Committee and bank regulatory requirements, see J. Hull, *Risk Management and Financial Institutions*, 2nd edn., 2010.

Business Snapshot 8.1 Basel I, Basel II, and Basel III

As the activities of banks became more global in the 1980s, it became necessary for regulators in different countries to work together to determine an international regulatory framework. As a result, the Basel Committee on Banking Supervision was formed. In 1988, it published a set of rules for the capital banks were required to keep for credit risk. These capital requirements have become known as Basel I. They were modified to accommodate netting (which will be discussed in Chapter 23) in 1995. In 1996, a new capital requirement for market risk was published. This capital requirement was implemented in 1998. In 1999, significant changes were proposed for the calculation of the capital requirements for credit risk and a capital requirement for operational risk was introduced. These rules are referred to as Basel II and were finally implemented in 2007—just before the credit crisis. Following the crisis, a new set of rules known as Basel III were proposed.

caused by liquidity. This was discussed in Business Snapshot 4.3. Financial institutions often choose short-term sources of funding. When the market becomes concerned (rightly or wrongly) about the health of a financial institution, this source of funding is liable to dry up.

Prior to the crisis, many of the Basel Committee's regulations involved the calculation of value at risk (VaR). This is a measure of the size of the loss that could be incurred by a bank and will be discussed in Chapter 21. VaR will continue to figure prominently in the Basel Committee's regulations, but the committee has become more conscious of the need to estimate VaR using data on the movements in market variables during stressed market conditions rather than normal market conditions. It has also put more emphasis on *stress testing*. This is concerned with examining how the bank would perform in adverse future scenarios.

As has been mentioned in earlier chapters, many governments are introducing rules requiring clearing houses to be used for some over-the-counter derivatives. In the US, the Commodity Futures and Trading Commission has responsibility for determining which categories of OTC derivatives must be cleared and for regulating the clearing houses.

Some governments have introduced special taxes on banks and on the bonuses of bank employees to recoup the costs of the crisis. For example, in December 2009, the UK government announced a "super-tax" on bonuses of more than £25,000. In its budget of June 2010, it introduced a tax on the liabilities of banks and indicated it was considering a proposal of the International Monetary Fund for a "financial activities tax."

Legislation enacted in the US in 2010 limits the ability of federally insured banks to trade derivatives or engage in proprietary trading. This may result in banks spinning off some of these activities into separate companies.

SUMMARY

Securitization is a process used by banks to create securities from loans and other income-producing assets. The securities are sold to investors. This removes the loans from the banks' balance sheets and enables the banks to expand their lending faster than they would otherwise be able to. The first loans to be securitized were mortgages in the

1960s and 1970s in the US. Investors who bought the mortgage-backed securities were not exposed to the risk of borrowers defaulting because the loans were backed by the Government National Mortgage Association. Later automobile loans, corporate loans, credit card receivables, and subprime mortgages were securitized. In many cases, investors in the securities created from these instruments did not have a guarantee against defaults.

Securitization played a part in the credit crisis that started in 2007. Tranches were created from subprime mortgages and new tranches were then created from these tranches. The origins of the crisis can be found in the US housing market. The US government was keen to encourage home ownership. Interest rates were low. Mortgage brokers and mortgage lenders found it attractive to do more business by relaxing their lending standards. Securitization meant that the investors bearing the credit risk were not usually the same as the original lenders. Rating agencies gave AAA ratings to the senior tranches that were created. There was no shortage of buyers for these AAA-rated tranches because their yields were higher than the yields on other AAA-rated securities. Banks thought the “good times” would continue and, because compensation plans focused their attention on short-term profits, chose to ignore the housing bubble and its potential impact on some very complicated products they were trading.

House prices rose as both first-time buyers and speculators entered the market. Some mortgages had included a low “teaser rate” for two or three years. After the teaser rate ended, there was a significant increase in the interest rate for many borrowers. Unable to meet the higher interest rate they had no choice but to default. This led to foreclosures and an increase in the supply of houses to be sold. The price increases between 2000 and 2006 began to be reversed. Speculators and others who found that the amount owing on their mortgages was less than the value of their houses (i.e., they had negative equity) defaulted. This accentuated the price decline.

Banks are paying a price for the crisis. New legislation and regulation will reduce their profitability. For example, capital requirements are being increased, liquidity requirements are being introduced, OTC derivatives are being more carefully regulated, and new taxes have been introduced.

FURTHER READING

- Gorton, G. “The Subprime Panic,” *European Financial Management*, 15, 1 (2008): 10–46.
- Hull, J.C. “The Financial Crisis of 2007: Another Case of Irrational Exuberance,” in: *The Finance Crisis and Rescue: What Went Wrong? Why? What Lessons Can be Learned*. University of Toronto Press, 2008.
- Hull, J.C., and A. White. “The Risk of the Tranches Created from Mortgages,” *Financial Analysts Journal*, 66, 5 (2010): 54–67.
- Keys, B.J., T. Mukherjee, A. Seru, and V. Vig. “Did Securitization Lead to Lax Screening? Evidence from Subprime Loans,” *Quarterly Journal of Economics*, 125, 1 (February 2010): 307–62.
- Krinsman, A.N. “Subprime Mortgage Meltdown: How Did It Happen and How Will It End,” *Journal of Structured Finance*, Summer 2007, 13–19.
- Mian, A., and A. Sufi. “The Consequences of Mortgage Credit Expansion: Evidence from the U.S. Mortgage Default Crisis,” *Quarterly Journal of Economics*, 124, 4 (November 2009): 1449–96.
- Zimmerman, T. “The Great Subprime Meltdown,” *Journal of Structured Finance*, Fall 2007, 7–20.

Practice Questions (Answers in Solutions Manual)

- 8.1. What was the role of GNMA (Ginnie Mae) in the mortgage-backed securities market of the 1970s?
- 8.2. Explain what is meant by (a) an ABS and (b) an ABS CDO.
- 8.3. What is a mezzanine tranche?
- 8.4. What is the waterfall in a securitization?
- 8.5. What are the numbers in Table 8.1 for a loss rate of (a) 12% and (b) 15%?
- 8.6. What is a subprime mortgage?
- 8.7. Why do you think the increase in house prices during the 2000 to 2007 period is referred to as a bubble?
- 8.8. Why did mortgage lenders frequently not check on information provided by potential borrowers on mortgage application forms during the 2000 to 2007 period?
- 8.9. How were the risks in ABS CDOs misjudged by the market?
- 8.10. What is meant by the term “agency costs”? How did agency costs play a role in the credit crisis?
- 8.11. How is an ABS CDO created? What was the motivation to create ABS CDOs?
- 8.12. Explain the impact of an increase in default correlation on the risks of the senior tranche of an ABS. What is its impact on the risks of the equity tranche?
- 8.13. Explain why the AAA-rated tranche of an ABS CDO is more risky than the AAA-rated tranche of an ABS.
- 8.14. Explain why the end-of-year bonus is sometimes referred to as “short-term compensation.”

Further Questions

- 8.15. Suppose that the principal assigned to the senior, mezzanine, and equity tranches is 70%, 20%, and 10% for both the ABS and the ABS CDO in Figure 8.3. What difference does this make to Table 8.1?
- 8.16. Investigate what happens as the width of the mezzanine tranche of the ABS in Figure 8.3 is decreased with the reduction of mezzanine tranche principal being divided equally between the equity and senior tranches. In particular, what is the effect on Table 8.1?



9

CHAPTER

Mechanics of Options Markets

We introduced options in Chapter 1. This chapter explains how options markets are organized, what terminology is used, how the contracts are traded, how margin requirements are set, and so on. Later chapters will examine such topics as trading strategies involving options, the determination of option prices, and the ways in which portfolios of options can be hedged. This chapter is concerned primarily with stock options. It presents some introductory material on currency options, index options, and futures options. More details concerning these instruments can be found in Chapters 16 and 17.

Options are fundamentally different from forward and futures contracts. An option gives the holder of the option the right to do something, but the holder does not have to exercise this right. By contrast, in a forward or futures contract, the two parties have committed themselves to some action. It costs a trader nothing (except for the margin requirements) to enter into a forward or futures contract, whereas the purchase of an option requires an up-front payment.


When charts showing the gain or loss from options trading are produced, the usual practice is to ignore discounting, so that the profit is the final payoff minus the initial cost. This chapter follows this practice.

9.1 TYPES OF OPTIONS

As mentioned in Chapter 1, there are two basic types of options. A *call option* gives the holder of the option the right to buy an asset by a certain date for a certain price. A *put option* gives the holder the right to sell an asset by a certain date for a certain price. The date specified in the contract is known as the *expiration date* or the *maturity date*. The price specified in the contract is known as the *exercise price* or the *strike price*.

Options can be either American or European, a distinction that has nothing to do with geographical location. *American options* can be exercised at any time up to the expiration date, whereas *European options* can be exercised only on the expiration date itself. Most of the options that are traded on exchanges are American. However, European options are generally easier to analyze than American options, and some of the properties of an American option are frequently deduced from those of its European counterpart.

Figure 9.1 Profit from buying a European call option on one share of a stock. Option price = \$5; strike price = \$100.



Call Options


Consider the situation of an investor who buys a European call option with a strike price of \$100 to purchase 100 shares of a certain stock. Suppose that the current stock price is \$98, the expiration date of the option is in 4 months, and the price of an option to purchase one share is \$5. The initial investment is \$500. Because the option is European, the investor can exercise only on the expiration date. If the stock price on this date is less than \$100, the investor will clearly choose not to exercise. (There is no point in buying for \$100 a share that has a market value of less than \$100.) In these circumstances, the investor loses the whole of the initial investment of \$500. If the stock price is above \$100 on the expiration date, the option will be exercised. Suppose, for example, that the stock price is \$115. By exercising the option, the investor is able to buy 100 shares for \$100 per share. If the shares are sold immediately, the investor makes a gain of \$15 per share, or \$1,500, ignoring transactions costs. When the initial cost of the option is taken into account, the net profit to the investor is \$1,000.

Figure 9.1 shows how the investor's net profit or loss on an option to purchase one share varies with the final stock price in the example. It is important to realize that an investor sometimes exercises an option and makes a loss overall. Suppose that, in the example, the stock price is \$102 at the expiration of the option. The investor would exercise the option contract for a gain of $100 \times (\$102 - \$100) = \$200$ and realize a loss overall of \$300 when the initial cost of the option is taken into account. It is tempting to argue that the investor should not exercise the option in these circumstances. However, not exercising would lead to an overall loss of \$500, which is worse than the \$300 loss when the investor exercises. In general, call options should always be exercised at the expiration date if the stock price is above the strike price.

Put Options

Whereas the purchaser of a call option is hoping that the stock price will increase, the purchaser of a put option is hoping that it will decrease. Consider an investor who buys a European put option with a strike price of \$70 to sell 100 shares of a certain

Figure 9.2 Profit from buying a European put option on one share of a stock. Option price = \$7; strike price = \$70.



stock. Suppose that the current stock price is \$65, the expiration date of the option is in 3 months, and the price of an option to sell one share is \$7. The initial investment is \$700. Because the option is European, it will be exercised only if the stock price is below \$70 on the expiration date. Suppose that the stock price is \$55 on this date. The investor can buy 100 shares for \$55 per share and, under the terms of the put option, sell the same shares for \$70 to realize a gain of \$15 per share, or \$1,500. (Again, transactions costs are ignored.) When the \$700 initial cost of the option is taken into account, the investor's net profit is \$800. There is no guarantee that the investor will make a gain. If the final stock price is above \$70, the put option expires worthless, and the investor loses \$700. Figure 9.2 shows the way in which the investor's profit or loss on an option to sell one share varies with the terminal stock price in this example.


Early Exercise

As already mentioned, exchange-traded stock options are generally American rather than European. This means that the investor in the foregoing examples would not have to wait until the expiration date before exercising the option. We will see in later chapters that there are some circumstances when it is optimal to exercise American options before the expiration date.

9.2 OPTION POSITIONS

There are two sides to every option contract. On one side is the investor who has taken the long position (i.e., has bought the option). On the other side is the investor who has taken a short position (i.e., has sold or *written* the option). The writer of an option receives cash up front, but has potential liabilities later. The writer's profit or loss is the reverse of that for the purchaser of the option. Figures 9.3 and 9.4 show the variation of the profit or loss with the final stock price for writers of the options considered in Figures 9.1 and 9.2.

Figure 9.3 Profit from writing a European call option on one share of a stock. Option price = \$5; strike price = \$100.



There are four types of option positions:

1. A long position in a call option
2. A long position in a put option
3. A short position in a call option
4. A short position in a put option.

It is often useful to characterize a European option in terms of its payoff to the purchaser of the option. The initial cost of the option is then not included in the calculation. If K is the strike price and S_T is the final price of the underlying asset, the

Figure 9.4 Profit from writing a European put option on one share of a stock. Option price = \$7; strike price = \$70.





Figure 9.5 Payoffs from positions in European options: (a) long call; (b) short call; (c) long put; (d) short put. Strike price = K ; price of asset at maturity = S_T .



payoff from a long position in a European call option is

$$\max(S_T - K, 0)$$

This reflects the fact that the option will be exercised if $S_T > K$ and will not be exercised if $S_T \leq K$. The payoff to the holder of a short position in the European call option is

$$-\max(S_T - K, 0) = \min(K - S_T, 0)$$

The payoff to the holder of a long position in a European put option is

$$\max(K - S_T, 0)$$

and the payoff from a short position in a European put option is

$$-\max(K - S_T, 0) = \min(S_T - K, 0)$$

Figure 9.5 illustrates these payoffs.

9.3 UNDERLYING ASSETS

This section provides a first look at options on stocks, currencies, stock indices, and futures.

Stock Options

Most trading in stock options is on exchanges. In the United States, the exchanges include the Chicago Board Options Exchange (www.cboe.com), NASDAQ OMX (www.nasdaqtrader.com), which acquired the Philadelphia Stock Exchange in 2008, NYSE Euronext (www.euronext.com), which acquired the American Stock Exchange in 2008, the International Securities Exchange (www.iseoptions.com), and the Boston Options Exchange (www.bostonoptions.com). Options trade on more than 2,500 different stocks. One contract gives the holder the right to buy or sell 100 shares at the specified strike price. This contract size is convenient because the shares themselves are normally traded in lots of 100.

Foreign Currency Options

Most currency options trading is now in the over-the-counter market, but there is some exchange trading. Exchanges trading foreign currency options in the United States include NASDAQ OMX. This exchange offers European-style contracts on a variety of different currencies. One contract is to buy or sell 10,000 units of a foreign currency (1,000,000 units in the case of the Japanese yen) for US dollars. Foreign currency options contracts are discussed further in Chapter 16.

Index Options

Many different index options currently trade throughout the world in both the over-the-counter market and the exchange-traded market. The most popular exchange-traded contracts in the United States are those on the S&P 500 Index (SPX), the S&P 100 Index (OEX), the Nasdaq-100 Index (NDX), and the Dow Jones Industrial Index (DJX). All of these trade on the Chicago Board Options Exchange. Most of the contracts are European. An exception is the OEX contract on the S&P 100, which is American. One contract is usually to buy or sell 100 times the index at the specified strike price. Settlement is always in cash, rather than by delivering the portfolio underlying the index. Consider, for example, one call contract on an index with a strike price of 980. If it is exercised when the value of the index is 992, the writer of the contract pays the holder $(992 - 980) \times 100 = \$1,200$. Index options are discussed further in Chapter 16.

Futures Options

When an exchange trades a particular futures contract, it often also trades options on that contract. A futures option normally matures just before the delivery period in the futures contract. When a call option is exercised, the holder's gain equals the excess of the futures price over the strike price. When a put option is exercised, the holder's gain equals the excess of the strike price over the futures price. Futures options contracts are discussed further in Chapter 17.

9.4 SPECIFICATION OF STOCK OPTIONS

In the rest of this chapter, we will focus on stock options. As already mentioned, an exchange-traded stock option in the United States is an American-style option contract

to buy or sell 100 shares of the stock. Details of the contract (the expiration date, the strike price, what happens when dividends are declared, how large a position investors can hold, and so on) are specified by the exchange.

Expiration Dates

One of the items used to describe a stock option is the month in which the expiration date occurs. Thus, a January call trading on IBM is a call option on IBM with an expiration date in January. The precise expiration date is the Saturday immediately following the third Friday of the expiration month. The last day on which options trade is the third Friday of the expiration month. An investor with a long position in an option normally has until 4:30 p.m. Central Time on that Friday to instruct a broker to exercise the option. The broker then has until 10:59 p.m. the next day to complete the paperwork notifying the exchange that exercise is to take place.

Stock options in the United States are on a January, February, or March cycle. The January cycle consists of the months of January, April, July, and October. The February cycle consists of the months of February, May, August, and November. The March cycle consists of the months of March, June, September, and December. If the expiration date for the current month has not yet been reached, options trade with expiration dates in the current month, the following month, and the next two months in the cycle. If the expiration date of the current month has passed, options trade with expiration dates in the next month, the next-but-one month, and the next two months of the expiration cycle. For example, IBM is on a January cycle. At the beginning of January, options are traded with expiration dates in January, February, April, and July; at the end of January, they are traded with expiration dates in February, March, April, and July; at the beginning of May, they are traded with expiration dates in May, June, July, and October; and so on. When one option reaches expiration, trading in another is started. Longer-term options, known as LEAPS (long-term equity anticipation securities), also trade on about 800 stocks in the United States. These have expiration dates up to 39 months into the future. The expiration dates for LEAPS on stocks are always in January.

Strike Prices

The exchange normally chooses the strike prices at which options can be written so that they are spaced \$2.50, \$5, or \$10 apart. Typically the spacing is \$2.50 when the stock price is between \$5 and \$25, \$5 when the stock price is between \$25 and \$200, and \$10 for stock prices above \$200. As will be explained shortly, stock splits and stock dividends can lead to nonstandard strike prices.

When a new expiration date is introduced, the two or three strike prices closest to the current stock price are usually selected by the exchange. If the stock price moves outside the range defined by the highest and lowest strike price, trading is usually introduced in an option with a new strike price. To illustrate these rules, suppose that the stock price is \$84 when trading begins in the October options. Call and put options would probably first be offered with strike prices of \$80, \$85, and \$90. If the stock price rose above \$90, it is likely that a strike price of \$95 would be offered; if it fell below \$80, it is likely that a strike price of \$75 would be offered; and so on.

Terminology

For any given asset at any given time, many different option contracts may be trading. Consider a stock that has four expiration dates and five strike prices. If call and put options trade with every expiration date and every strike price, there are a total of 40 different contracts. All options of the same type (calls or puts) are referred to as an *option class*. For example, IBM calls are one class, whereas IBM puts are another class. An *option series* consists of all the options of a given class with the same expiration date and strike price. In other words, an option series refers to a particular contract that is traded. For example, IBM 70 October calls would constitute an option series.

Options are referred to as *in the money*, *at the money*, or *out of the money*. If S is the stock price and K is the strike price, a call option is in the money when $S > K$, at the money when $S = K$, and out of the money when $S < K$. A put option is in the money when $S < K$, at the money when $S = K$, and out of the money when $S > K$. Clearly, an option will be exercised only when it is in the money. In the absence of transactions costs, an in-the-money option will always be exercised on the expiration date if it has not been exercised previously.

The *intrinsic value* of an option is defined as the maximum of zero and the value the option would have if it were exercised immediately. For a call option, the intrinsic value is therefore $\max(S - K, 0)$. For a put option, it is $\max(K - S, 0)$. An in-the-money American option must be worth at least as much as its intrinsic value because the holder can realize the intrinsic value by exercising immediately. Often it is optimal for the holder of an in-the-money American option to wait rather than exercise immediately. The option is then said to have *time value*. The total value of an option can be thought of as the sum of its intrinsic value and its time value.

FLEX Options

The Chicago Board Options Exchange offers FLEX (short for flexible) options on equities and equity indices. These are options where the traders on the floor of the exchange agree to nonstandard terms. These nonstandard terms can involve a strike price or an expiration date that is different from what is usually offered by the exchange. It can also involve the option being European rather than American. FLEX options are an attempt by option exchanges to regain business from the over-the-counter markets. The exchange specifies a minimum size (e.g., 100 contracts) for FLEX option trades.

Dividends and Stock Splits

The early over-the-counter options were dividend protected. If a company declared a cash dividend, the strike price for options on the company's stock was reduced on the ex-dividend day by the amount of the dividend. Exchange-traded options are not usually adjusted for cash dividends. In other words, when a cash dividend occurs, there are no adjustments to the terms of the option contract. An exception is sometimes made for large cash dividends (see Business Snapshot 9.1).

Exchange-traded options are adjusted for stock splits. A stock split occurs when the existing shares are "split" into more shares. For example, in a 3-for-1 stock split, three new shares are issued to replace each existing share. Because a stock split does not change the assets or the earning ability of a company, we should not expect it to have

Business Snapshot 9.1 Gucci Group's Large Dividend

When there is a large cash dividend (typically one that is more than 10% of the stock price), a committee of the Options Clearing Corporation (OCC) at the Chicago Board Options Exchange can decide to adjust the terms of options traded on the exchange.

On May 28, 2003, Gucci Group NV (GUC) declared a cash dividend of 13.50 euros (approximately \$15.88) per common share and this was approved at the annual shareholders' meeting on July 16, 2003. The dividend was about 16% of the share price at the time it was declared. In this case, the OCC committee decided to adjust the terms of options. The result was that the holder of a call contract paid 100 times the strike price on exercise and received \$1,588 of cash in addition to 100 shares; the holder of a put contract received 100 times the strike price on exercise and delivered \$1,588 of cash in addition to 100 shares. These adjustments had the effect of reducing the strike price by \$15.88.

Adjustments for large dividends are not always made. For example, Deutsche Terminbörse chose not to adjust the terms of options traded on that exchange when Daimler-Benz surprised the market on March 10, 1998, with a dividend equal to about 12% of its stock price.

any effect on the wealth of the company's shareholders. All else being equal, the 3-for-1 stock split should cause the stock price to go down to one-third of its previous value. In general, an n -for- m stock split should cause the stock price to go down to m/n of its previous value. The terms of option contracts are adjusted to reflect expected changes in a stock price arising from a stock split. After an n -for- m stock split, the strike price is reduced to m/n of its previous value, and the number of shares covered by one contract is increased to n/m of its previous value. If the stock price declines in the way expected, the positions of both the writer and the purchaser of a contract remain unchanged.

Example 9.1

Consider a call option to buy 100 shares of a company for \$30 per share. Suppose the company makes a 2-for-1 stock split. The terms of the option contract are then changed so that it gives the holder the right to purchase 200 shares for \$15 per share.

Stock options are adjusted for stock dividends. A stock dividend involves a company issuing more shares to its existing shareholders. For example, a 20% stock dividend means that investors receive one new share for each five already owned. A stock dividend, like a stock split, has no effect on either the assets or the earning power of a company. The stock price can be expected to go down as a result of a stock dividend. The 20% stock dividend referred to is essentially the same as a 6-for-5 stock split. All else being equal, it should cause the stock price to decline to $5/6$ of its previous value. The terms of an option are adjusted to reflect the expected price decline arising from a stock dividend in the same way as they are for that arising from a stock split.

Example 9.2

Consider a put option to sell 100 shares of a company for \$15 per share. Suppose the company declares a 25% stock dividend. This is equivalent to a 5-for-4 stock split. The terms of the option contract are changed so that it gives the holder the right to sell 125 shares for \$12.

Adjustments are also made for rights issues. The basic procedure is to calculate the theoretical price of the rights and then to reduce the strike price by this amount.

Position Limits and Exercise Limits

The Chicago Board Options Exchange often specifies a *position limit* for option contracts. This defines the maximum number of option contracts that an investor can hold on one side of the market. For this purpose, long calls and short puts are considered to be on the same side of the market. Also considered to be on the same side are short calls and long puts. The *exercise limit* usually equals the position limit. It defines the maximum number of contracts that can be exercised by any individual (or group of individuals acting together) in any period of five consecutive business days. Options on the largest and most frequently traded stocks have position limits of 250,000 contracts. Smaller capitalization stocks have position limits of 200,000, 75,000, 50,000, or 25,000 contracts.

Position limits and exercise limits are designed to prevent the market from being unduly influenced by the activities of an individual investor or group of investors. However, whether the limits are really necessary is a controversial issue.

9.5 TRADING

Traditionally, exchanges have had to provide a large open area for individuals to meet and trade options. This has changed. Most derivatives exchanges are fully electronic, so traders do not have to physically meet. The International Securities Exchange (www.iseoptions.com) launched the first all-electronic options market for equities in the United States in May 2000. Over 95% of the orders at the Chicago Board Options Exchange are handled electronically. The remainder are mostly large or complex institutional orders that require the skills of traders.

Market Makers

Most options exchanges use market makers to facilitate trading. A market maker for a certain option is an individual who, when asked to do so, will quote both a bid and an offer price on the option. The bid is the price at which the market maker is prepared to buy, and the offer or asked is the price at which the market maker is prepared to sell. At the time the bid and offer prices are quoted, the market maker does not know whether the trader who asked for the quotes wants to buy or sell the option. The offer is always higher than the bid, and the amount by which the offer exceeds the bid is referred to as the *bid-offer* spread. The exchange sets upper limits for the bid-offer spread. For example, it might specify that the spread be no more than \$0.25 for options priced at less than \$0.50, \$0.50 for options priced between \$0.50 and \$10, \$0.75 for options priced between \$10 and \$20, and \$1 for options priced over \$20.

The existence of the market maker ensures that buy and sell orders can always be executed at some price without any delays. Market makers therefore add liquidity to the market. The market makers themselves make their profits from the bid-offer spread. They use methods such as those that will be discussed in Chapter 18 to hedge their risks.

Offsetting Orders

An investor who has purchased options can close out the position by issuing an offsetting order to sell the same number of options. Similarly, an investor who has written options can close out the position by issuing an offsetting order to buy the same number of options. (In this respect options markets are similar to futures markets.) If, when an option contract is traded, neither investor is closing an existing position, the open interest increases by one contract. If one investor is closing an existing position and the other is not, the open interest stays the same. If both investors are closing existing positions, the open interest goes down by one contract.

9.6 COMMISSIONS

The types of orders that can be placed with a broker for options trading are similar to those for futures trading (see Section 2.8). A market order is executed immediately, a limit order specifies the least favorable price at which the order can be executed, and so on.

For a retail investor, commissions vary significantly from broker to broker. Discount brokers generally charge lower commissions than full-service brokers. The actual amount charged is often calculated as a fixed cost plus a proportion of the dollar amount of the trade. Table 9.1 shows the sort of schedule that might be offered by a discount broker. Using this schedule, the purchase of eight contracts when the option price is \$3 would cost $\$20 + (0.02 \times \$2,400) = \$68$ in commissions.

If an option position is closed out by entering into an offsetting trade, the commission must be paid again. If the option is exercised, the commission is the same as it would be if the investor placed an order to buy or sell the underlying stock.

Consider an investor who buys one call contract with a strike price of \$50 when the stock price is \$49. We suppose the option price is \$4.50, so that the cost of the contract is \$450. Under the schedule in Table 9.1, the purchase or sale of one contract always costs \$30 (both the maximum and minimum commission is \$30 for the first contract). Suppose that the stock price rises and the option is exercised when the stock reaches \$60. Assuming that the investor pays 0.75% commission to exercise the option and a further 0.75% commission to sell the stock, there is an additional cost of

$$2 \times 0.0075 \times \$60 \times 100 = \$90$$

Table 9.1 A typical commission schedule for a discount broker.

| <i>Dollar amount of trade</i> | <i>Commission*</i> |
|-------------------------------|--------------------------------|
| < \$2,500 | \$20 + 2% of dollar amount |
| \$2,500 to \$10,000 | \$45 + 1% of dollar amount |
| > \$10,000 | \$120 + 0.25% of dollar amount |

* Maximum commission is \$30 per contract for the first five contracts plus \$20 per contract for each additional contract. Minimum commission is \$30 per contract for the first contract plus \$2 per contract for each additional contract.

The total commission paid is therefore \$120, and the net profit to the investor is

$$\$1,000 - \$450 - \$120 = \$430$$

Note that selling the option for \$10 instead of exercising it would save the investor \$60 in commissions. (The commission payable when an option is sold is only \$30 in our example.) As this example indicates, the commission system can push retail investors in the direction of selling options rather than exercising them.

A hidden cost in option trading (and in stock trading) is the market maker's bid–offer spread. Suppose that, in the example just considered, the bid price was \$4.00 and the offer price was \$4.50 at the time the option was purchased. We can reasonably assume that a “fair” price for the option is halfway between the bid and the offer price, or \$4.25. The cost to the buyer and to the seller of the market maker system is the difference between the fair price and the price paid. This is \$0.25 per option, or \$25 per contract.

9.7 MARGINS

When shares are purchased in the United States, an investor can borrow up to 50% of the price from the broker. This is known as *buying on margin*. If the share price declines so that the loan is substantially more than 50% of the stock's current value, there is a “margin call”, where the broker requests that cash be deposited by the investor. If the margin call is not met, the broker sells the stock.

When call and put options with maturities less than 9 months are purchased, the option price must be paid in full. Investors are not allowed to buy these options on margin because options already contain substantial leverage and buying on margin would raise this leverage to an unacceptable level. For options with maturities greater than 9 months investors can buy on margin, borrowing up to 25% of the option value.

A trader who writes options is required to maintain funds in a margin account. Both the trader's broker and the exchange want to be satisfied that the trader will not default if the option is exercised. The amount of margin required depends on the trader's position.

Writing Naked Options

A *naked option* is an option that is not combined with an offsetting position in the underlying stock. The initial and maintenance margin required by the CBOE for a written naked call option is the greater of the following two calculations:

1. A total of 100% of the proceeds of the sale plus 20% of the underlying share price less the amount, if any, by which the option is out of the money
2. A total of 100% of the option proceeds plus 10% of the underlying share price.

For a written naked put option, it is the greater of

1. A total of 100% of the proceeds of the sale plus 20% of the underlying share price less the amount, if any, by which the option is out of the money
2. A total of 100% of the option proceeds plus 10% of the exercise price.

The 20% in the preceding calculations is replaced by 15% for options on a broadly based stock index because a stock index is usually less volatile than the price of an individual stock.

Example 9.3

An investor writes four naked call option contracts on a stock. The option price is \$5, the strike price is \$40, and the stock price is \$38. Because the option is \$2 out of the money, the first calculation gives

$$400 \times (5 + 0.2 \times 38 - 2) = \$4,240$$

The second calculation gives

$$400 \times (5 + 0.1 \times 38) = \$3,520$$

The initial margin requirement is therefore \$4,240. Note that, if the option had been a put, it would be \$2 in the money and the margin requirement would be

$$400 \times (5 + 0.2 \times 38) = \$5,040$$

In both cases, the proceeds of the sale can be used to form part of the margin account.

A calculation similar to the initial margin calculation (but with the current market price replacing the proceeds of sale) is repeated every day. Funds can be withdrawn from the margin account when the calculation indicates that the margin required is less than the current balance in the margin account. When the calculation indicates that a greater margin is required, a margin call will be made.

Other Rules

In Chapter 11, we will examine option trading strategies such as covered calls, protective puts, spreads, combinations, straddles, and strangles. The CBOE has special rules for determining the margin requirements when these trading strategies are used. These are described in the *CBOE Margin Manual*, which is available on the CBOE website (www.cboe.com).

As an example of the rules, consider an investor who writes a covered call. This is a written call option when the shares that might have to be delivered are already owned. Covered calls are far less risky than naked calls, because the worst that can happen is that the investor is required to sell shares already owned at below their market value. No margin is required on the written option. However, the investor can borrow an amount equal to $0.5 \min(S, K)$, rather than the usual $0.5S$, on the stock position.

9.8 THE OPTIONS CLEARING CORPORATION

The Options Clearing Corporation (OCC) performs much the same function for options markets as the clearing house does for futures markets (see Chapter 2). It guarantees that options writers will fulfill their obligations under the terms of options contracts and keeps a record of all long and short positions. The OCC has a number of members, and all option trades must be cleared through a member. If a broker is not itself a member of an exchange's OCC, it must arrange to clear its trades with a member. Members are required to have a certain minimum amount of capital and to contribute to a special fund that can be used if any member defaults on an option obligation.

The funds used to purchase an option must be deposited with the OCC by the morning of the business day following the trade. The writer of the option maintains a margin account with a broker, as described earlier.¹ The broker maintains a margin account with the OCC member that clears its trades. The OCC member in turn maintains a margin account with the OCC.

Exercising an Option

When an investor notifies a broker to exercise an option, the broker in turn notifies the OCC member that clears its trades. This member then places an exercise order with the OCC. The OCC randomly selects a member with an outstanding short position in the same option. The member, using a procedure established in advance, selects a particular investor who has written the option. If the option is a call, this investor is required to sell stock at the strike price. If it is a put, the investor is required to buy stock at the strike price. The investor is said to be *assigned*. When an option is exercised, the open interest goes down by one.

At the expiration of the option, all in-the-money options should be exercised unless the transactions costs are so high as to wipe out the payoff from the option. Some brokers will automatically exercise options for a client at expiration when it is in their client's interest to do so. Many exchanges also have rules for exercising options that are in the money at expiration.

9.9 REGULATION

Options markets are regulated in a number of different ways. Both the exchange and Options Clearing Corporations have rules governing the behavior of traders. In addition, there are both federal and state regulatory authorities. In general, options markets have demonstrated a willingness to regulate themselves. There have been no major scandals or defaults by OCC members. Investors can have a high level of confidence in the way the market is run.

The Securities and Exchange Commission is responsible for regulating options markets in stocks, stock indices, currencies, and bonds at the federal level. The Commodity Futures Trading Commission is responsible for regulating markets for options on futures. The major options markets are in the states of Illinois and New York. These states actively enforce their own laws on unacceptable trading practices.

9.10 TAXATION

Determining the tax implications of option trading strategies can be tricky, and an investor who is in doubt about this should consult a tax specialist. In the United States, the general rule is that (unless the taxpayer is a professional trader) gains and losses from the trading of stock options are taxed as capital gains or losses. The way that

¹ The margin requirements described in the previous section are the minimum requirements specified by the OCC. A broker may require higher margins from its clients. However, it cannot require lower margins. Some brokers do not allow their retail clients to write uncovered options at all.

capital gains and losses are taxed in the United States was discussed in Section 2.10. For both the holder and the writer of a stock option, a gain or loss is recognized when (a) the option expires unexercised or (b) the option position is closed out. If the option is exercised, the gain or loss from the option is rolled into the position taken in the stock and recognized when the stock position is closed out. For example, when a call option is exercised, the party with a long position is deemed to have purchased the stock at the strike price plus the call price. This is then used as a basis for calculating this party's gain or loss when the stock is eventually sold. Similarly, the party with the short call position is deemed to have sold the stock at the strike price plus the call price. When a put option is exercised, the seller of the option is deemed to have bought the stock for the strike price less the original put price and the purchaser of the option is deemed to have sold the stock for the strike price less the original put price.

Wash Sale Rule

One tax consideration in option trading in the United States is the wash sale rule. To understand this rule, imagine an investor who buys a stock when the price is \$60 and plans to keep it for the long term. If the stock price drops to \$40, the investor might be tempted to sell the stock and then immediately repurchase it, so that the \$20 loss is realized for tax purposes. To prevent this sort of thing, the tax authorities have ruled that when the repurchase is within 30 days of the sale (i.e., between 30 days before the sale and 30 days after the sale), any loss on the sale is not deductible. The disallowance also applies where, within the 61-day period, the taxpayer enters into an option or similar contract to acquire the stock. Thus, selling a stock at a loss and buying a call option within a 30-day period will lead to the loss being disallowed. The wash sale rule does not apply if the taxpayer is a dealer in stocks or securities and the loss is sustained in the ordinary course of business.

Constructive Sales

Prior to 1997, if a United States taxpayer shorted a security while holding a long position in a substantially identical security, no gain or loss was recognized until the short position was closed out. This means that short positions could be used to defer recognition of a gain for tax purposes. The situation was changed by the Tax Relief Act of 1997. An appreciated property is now treated as "constructively sold" when the owner does one of the following:

1. Enters into a short sale of the same or substantially identical property
2. Enters into a futures or forward contract to deliver the same or substantially identical property
3. Enters into one or more positions that eliminate substantially all of the loss and opportunity for gain.

It should be noted that transactions reducing only the risk of loss or only the opportunity for gain should not result in constructive sales. Therefore an investor holding a long position in a stock can buy in-the-money put options on the stock without triggering a constructive sale.

Tax practitioners sometimes use options to minimize tax costs or maximize tax benefits (see Business Snapshot 9.2). Tax authorities in many jurisdictions have

Business Snapshot 9.2 Tax Planning Using Options

As a simple example of a possible tax planning strategy using options, suppose that Country A has a tax regime where the tax is low on interest and dividends and high on capital gains, while Country B has a tax regime where tax is high on interest and dividends and low on capital gains. It is advantageous for a company to receive the income from a security in Country A and the capital gain, if there is one, in Country B. The company would like to keep capital losses in Country A, where they can be used to offset capital gains on other items. All of this can be accomplished by arranging for a subsidiary company in Country A to have legal ownership of the security and for a subsidiary company in Country B to buy a call option on the security from the company in Country A, with the strike price of the option equal to the current value of the security. During the life of the option, income from the security is earned in Country A. If the security price rises sharply, the option will be exercised and the capital gain will be realized in Country B. If it falls sharply, the option will not be exercised and the capital loss will be realized in Country A.

proposed legislation designed to combat the use of derivatives for tax purposes. Before entering into any tax-motivated transaction, a corporate treasurer or private individual should explore in detail how the structure could be unwound in the event of legislative change and how costly this process could be.

9.11 WARRANTS, EMPLOYEE STOCK OPTIONS, AND CONVERTIBLES

Warrants are options issued by a financial institution or nonfinancial corporation. For example, a financial institution might issue put warrants on one million ounces of gold and then proceed to create a market for the warrants. To exercise the warrant, the holder would contact the financial institution. A common use of warrants by a nonfinancial corporation is at the time of a bond issue. The corporation issues call warrants on its own stock and then attaches them to the bond issue to make it more attractive to investors.

Employee stock options are call options issued to employees by their company to motivate them to act in the best interests of the company's shareholders (see Chapter 15). They are usually at the money at the time of issue. They are now a cost on the income statement of the company in most countries, making them a less attractive form of compensation than they used to be.

Convertible bonds, often referred to as *convertibles*, are bonds issued by a company that can be converted into equity at certain times using a predetermined exchange ratio. They are therefore bonds with an embedded call option on the company's stock.

One feature of warrants, employee stock options, and convertibles is that a predetermined number of options are issued. By contrast, the number of options on a particular stock that trade on the CBOE or another exchange is not predetermined. As people take positions in a particular option series, the number of options outstanding increases; as people close out positions, it declines. Warrants issued by a company on its own stock, employee stock options, and convertibles are different from exchange-traded options in another important way. When these instruments are exercised, the company issues more shares of its own stock and sells them to the option holder for the strike price. The

exercise of the instruments therefore leads to an increase in the number of shares of the company's stock that are outstanding. By contrast, when an exchange-traded call option is exercised, the party with the short position buys in the market shares that have already been issued and sells them to the party with the long position for the strike price. The company whose stock underlies the option is not involved in any way.

9.12 OVER-THE-COUNTER OPTIONS MARKETS

Most of this chapter has focused on exchange-traded options markets. The over-the-counter market for options has become increasingly important since the early 1980s and is now larger than the exchange-traded market. As explained in Chapter 1, in the over-the-counter market, financial institutions, corporate treasurers, and fund managers trade over the phone. There is a wide range of assets underlying the options. Over-the-counter options on foreign exchange and interest rates are particularly popular. The chief potential disadvantage of the over-the-counter market is that the option writer may default. This means that the purchaser is subject to some credit risk. In an attempt to overcome this disadvantage, market participants usually require counterparties to post collateral. This was discussed in Section 2.5.

The instruments traded in the over-the-counter market are often structured by financial institutions to meet the precise needs of their clients. Sometimes this involves choosing exercise dates, strike prices, and contract sizes that are different from those offered by an exchange. In other cases the structure of the option is different from standard calls and puts. The option is then referred to as an *exotic option*. Chapter 25 describes a number of different types of exotic options.

SUMMARY

There are two types of options: calls and puts. A call option gives the holder the right to buy the underlying asset for a certain price by a certain date. A put option gives the holder the right to sell the underlying asset by a certain date for a certain price. There are four possible positions in options markets: a long position in a call, a short position in a call, a long position in a put, and a short position in a put. Taking a short position in an option is known as writing it. Options are currently traded on stocks, stock indices, foreign currencies, futures contracts, and other assets.

An exchange must specify the terms of the option contracts it trades. In particular, it must specify the size of the contract, the precise expiration time, and the strike price. In the United States one stock option contract gives the holder the right to buy or sell 100 shares. The expiration of a stock option contract is 10:59 p.m. Central Time on the Saturday immediately following the third Friday of the expiration month. Options with several different expiration months trade at any given time. Strike prices are at $\$2\frac{1}{2}$, \$5, or \$10 intervals, depending on the stock price. The strike price is generally fairly close to the stock price when trading in an option begins.

The terms of a stock option are not normally adjusted for cash dividends. However, they are adjusted for stock dividends, stock splits, and rights issues. The aim of the adjustment is to keep the positions of both the writer and the buyer of a contract unchanged.

Most option exchanges use market makers. A market maker is an individual who is prepared to quote both a bid price (at which he or she is prepared to buy) and an offer price (at which he or she is prepared to sell). Market makers improve the liquidity of the market and ensure that there is never any delay in executing market orders. They themselves make a profit from the difference between their bid and offer prices (known as their bid–offer spread). The exchange has rules specifying upper limits for the bid–offer spread.

Writers of options have potential liabilities and are required to maintain margins with their brokers. If it is not a member of the Options Clearing Corporation, the broker will maintain a margin account with a firm that is a member. This firm will in turn maintain a margin account with the Options Clearing Corporation. The Options Clearing Corporation is responsible for keeping a record of all outstanding contracts, handling exercise orders, and so on.

Not all options are traded on exchanges. Many options are traded by phone in the over-the-counter market. An advantage of over-the-counter options is that they can be tailored by a financial institution to meet the particular needs of a corporate treasurer or fund manager.

FURTHER READING

Arzac, E. R. “PERCs, DECs, and Other Mandatory Convertibles,” *Journal of Applied Corporate Finance*, 10, 1 (1997): 54–63.

Chicago Board Options Exchange, *Characteristics and Risks of Standardized Options*. Available online at www.optionsclearing.com/about/publications/character-risks.jsp. First published 1994; last updated 2009.

McMillan, L. G. *McMillan on Options*, 2nd edn. Hoboken, NJ: Wiley, 2004.

Practice Questions (Answers in Solutions Manual)

- 9.1. An investor buys a European put on a share for \$3. The stock price is \$42 and the strike price is \$40. Under what circumstances does the investor make a profit? Under what circumstances will the option be exercised? Draw a diagram showing the variation of the investor’s profit with the stock price at the maturity of the option.
- 9.2. An investor sells a European call on a share for \$4. The stock price is \$47 and the strike price is \$50. Under what circumstances does the investor make a profit? Under what circumstances will the option be exercised? Draw a diagram showing the variation of the investor’s profit with the stock price at the maturity of the option.
- 9.3. An investor sells a European call option with strike price of K and maturity T and buys a put with the same strike price and maturity. Describe the investor’s position.
- 9.4. Explain why margins are required when clients write options but not when they buy options.
- 9.5. A stock option is on a February, May, August, and November cycle. What options trade on (a) April 1 and (b) May 30?
- 9.6. A company declares a 2-for-1 stock split. Explain how the terms change for a call option with a strike price of \$60.

- 9.7. “Employee stock options issued by a company are different from regular exchange-traded call options on the company’s stock because they can affect the capital structure of the company.” Explain this statement.
- 9.8. A corporate treasurer is designing a hedging program involving foreign currency options. What are the pros and cons of using (a) NASDAQ OMX and (b) the over-the-counter market for trading?
- 9.9. Suppose that a European call option to buy a share for \$100.00 costs \$5.00 and is held until maturity. Under what circumstances will the holder of the option make a profit? Under what circumstances will the option be exercised? Draw a diagram illustrating how the profit from a long position in the option depends on the stock price at maturity of the option.
- 9.10. Suppose that a European put option to sell a share for \$60 costs \$8 and is held until maturity. Under what circumstances will the seller of the option (the party with the short position) make a profit? Under what circumstances will the option be exercised? Draw a diagram illustrating how the profit from a short position in the option depends on the stock price at maturity of the option.
- 9.11. Describe the terminal value of the following portfolio: a newly entered-into long forward contract on an asset and a long position in a European put option on the asset with the same maturity as the forward contract and a strike price that is equal to the forward price of the asset at the time the portfolio is set up. Show that the European put option has the same value as a European call option with the same strike price and maturity.
- 9.12. A trader buys a call option with a strike price of \$45 and a put option with a strike price of \$40. Both options have the same maturity. The call costs \$3 and the put costs \$4. Draw a diagram showing the variation of the trader’s profit with the asset price.
- 9.13. Explain why an American option is always worth at least as much as a European option on the same asset with the same strike price and exercise date.
- 9.14. Explain why an American option is always worth at least as much as its intrinsic value.
- 9.15. Explain carefully the difference between writing a put option and buying a call option.
- 9.16. The treasurer of a corporation is trying to choose between options and forward contracts to hedge the corporation’s foreign exchange risk. Discuss the advantages and disadvantages of each.
- 9.17. Consider an exchange-traded call option contract to buy 500 shares with a strike price of \$40 and maturity in 4 months. Explain how the terms of the option contract change when there is: (a) a 10% stock dividend; (b) a 10% cash dividend; and (c) a 4-for-1 stock split.
- 9.18. “If most of the call options on a stock are in the money, it is likely that the stock price has risen rapidly in the last few months.” Discuss this statement.
- 9.19. What is the effect of an unexpected cash dividend on (a) a call option price and (b) a put option price?
- 9.20. Options on General Motors stock are on a March, June, September, and December cycle. What options trade on (a) March 1, (b) June 30, and (c) August 5?
- 9.21. Explain why the market maker’s bid–offer spread represents a real cost to options investors.
- 9.22. A United States investor writes five naked call option contracts. The option price is \$3.50, the strike price is \$60.00, and the stock price is \$57.00. What is the initial margin requirement?

Further Questions

- 9.23. The price of a stock is \$40. The price of a 1-year European put option on the stock with a strike price of \$30 is quoted as \$7 and the price of a 1-year European call option on the stock with a strike price of \$50 is quoted as \$5. Suppose that an investor buys 100 shares, shorts 100 call options, and buys 100 put options. Draw a diagram illustrating how the investor's profit or loss varies with the stock price over the next year. How does your answer change if the investor buys 100 shares, shorts 200 call options, and buys 200 put options?
- 9.24. "If a company does not do better than its competitors but the stock market goes up, executives do very well from their stock options. This makes no sense." Discuss this viewpoint. Can you think of alternatives to the usual employee stock option plan that take the viewpoint into account.
- 9.25. Use DerivaGem to calculate the value of an American put option on a non-dividend-paying stock when the stock price is \$30, the strike price is \$32, the risk-free rate is 5%, the volatility is 30%, and the time to maturity is 1.5 years. (Choose "Binomial American" for the "option type" and 50 time steps.)
- What is the option's intrinsic value?
 - What is the option's time value?
 - What would a time value of zero indicate? What is the value of an option with zero time value?
 - Using a trial and error approach, calculate how low the stock price would have to be for the time value of the option to be zero.
- 9.26. On July 20, 2004, Microsoft surprised the market by announcing a \$3 dividend. The ex-dividend date was November 17, 2004, and the payment date was December 2, 2004. Its stock price at the time was about \$28. It also changed the terms of its employee stock options so that each exercise price was adjusted downward to

$$\text{Predividend exercise price} \times \frac{\text{Closing price} - \$3.00}{\text{Closing price}}$$

The number of shares covered by each stock option outstanding was adjusted upward to

$$\text{Number of shares predividend} \times \frac{\text{Closing price}}{\text{Closing price} - \$3.00}$$

"Closing Price" means the official NASDAQ closing price of a share of Microsoft common stock on the last trading day before the ex-dividend date. Evaluate this adjustment. Compare it with the system used by exchanges to adjust for extraordinary dividends (see Business Snapshot 9.1).



10

CHAPTER

Properties of Stock Options

In this chapter, we look at the factors affecting stock option prices. We use a number of different arbitrage arguments to explore the relationships between European option prices, American option prices, and the underlying stock price. The most important of these relationships is put–call parity, which is a relationship between the price of a European call option, the price of a European put option, and the underlying stock price.

The chapter examines whether American options should be exercised early. It shows that it is never optimal to exercise an American call option on a non-dividend-paying stock prior to the option’s expiration, but that under some circumstances the early exercise of an American put option on such a stock is optimal. When there are dividends, it can be optimal to exercise either calls or puts early.

10.1 FACTORS AFFECTING OPTION PRICES

There are six factors affecting the price of a stock option:

1. The current stock price, S_0
2. The strike price, K
3. The time to expiration, T
4. The volatility of the stock price, σ
5. The risk-free interest rate, r
6. The dividends that are expected to be paid.

In this section, we consider what happens to option prices when there is a change to one of these factors, with all the other factors remaining fixed. The results are summarized in Table 10.1.

Figures 10.1 and 10.2 show how European call and put prices depend on the first five factors in the situation where $S_0 = 50$, $K = 50$, $r = 5\%$ per annum, $\sigma = 30\%$ per annum, $T = 1$ year, and there are no dividends. In this case the call price is 7.116 and the put price is 4.677.

Table 10.1 Summary of the effect on the price of a stock option of increasing one variable while keeping all others fixed.*

| <i>Variable</i> | <i>European call</i> | <i>European put</i> | <i>American call</i> | <i>American put</i> |
|----------------------------|--------------------------|-------------------------|--------------------------|-------------------------|
| Current stock price | + | − | + | − |
| Strike price | − | + | − | + |
| Time to expiration | ? | ? | + | + |
| Volatility | + | + | + | + |
| Risk-free rate | + | − | + | − |
| Amount of future dividends | − | + | − | + |

* + indicates that an increase in the variable causes the option price to increase;
 − indicates that an increase in the variable causes the option price to decrease;
 ? indicates that the relationship is uncertain.

Stock Price and Strike Price

If a call option is exercised at some future time, the payoff will be the amount by which the stock price exceeds the strike price. Call options therefore become more valuable as the stock price increases and less valuable as the strike price increases. For a put option, the payoff on exercise is the amount by which the strike price exceeds the stock price. Put options therefore behave in the opposite way from call options: they become less valuable as the stock price increases and more valuable as the strike price increases. Figure 10.1a–d illustrate the way in which put and call prices depend on the stock price and strike price.

Time to Expiration

Now consider the effect of the expiration date. Both put and call American options become more valuable (or at least do not decrease in value) as the time to expiration increases. Consider two American options that differ only as far as the expiration date is concerned. The owner of the long-life option has all the exercise opportunities open to the owner of the short-life option—and more. The long-life option must therefore always be worth at least as much as the short-life option.

Although European put and call options usually become more valuable as the time to expiration increases (see Figure 10.1e, f), this is not always the case. Consider two European call options on a stock: one with an expiration date in 1 month, the other with an expiration date in 2 months. Suppose that a very large dividend is expected in 6 weeks. The dividend will cause the stock price to decline, so that the short-life option could be worth more than the long-life option.¹

Volatility

The precise way in which volatility is defined is discussed in Chapter 14. Roughly speaking, the *volatility* of a stock price is a measure of how uncertain we are about

¹ We assume that, when the life of the option is changed, the dividends on the stock and their timing remain unchanged.

future stock price movements. As volatility increases, the chance that the stock will do very well or very poorly increases. For the owner of a stock, these two outcomes tend to offset each other. However, this is not so for the owner of a call or put. The owner of a call benefits from price increases but has limited downside risk in the event of price decreases because the most the owner can lose is the price of the option. Similarly, the owner of a put benefits from price decreases, but has limited downside risk in the event of price increases. The values of both calls and puts therefore increase as volatility increases (see Figure 10.2a, b).

Figure 10.1 Effect of changes in stock price, strike price, and expiration date on option prices when $S_0 = 50$, $K = 50$, $r = 5\%$, $\sigma = 30\%$, and $T = 1$.


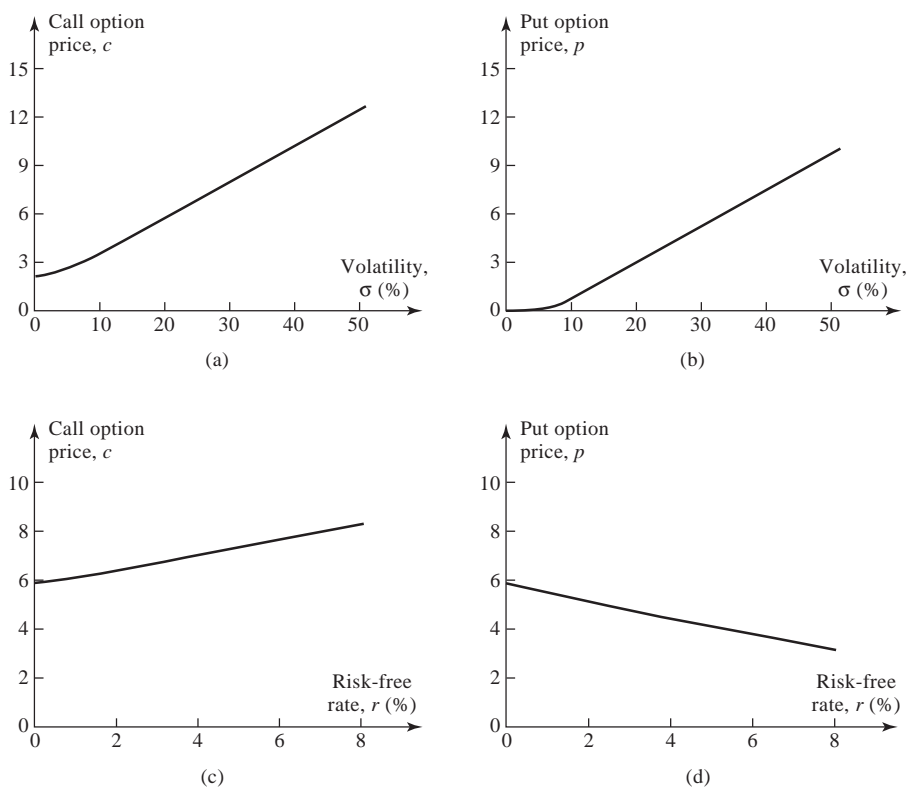


Figure 10.2 Effect of changes in volatility and risk-free interest rate on option prices when $S_0 = 50$, $K = 50$, $r = 5\%$, $\sigma = 30\%$, and $T = 1$.



Risk-Free Interest Rate

The risk-free interest rate affects the price of an option in a less clear-cut way. As interest rates in the economy increase, the expected return required by investors from the stock tends to increase. In addition, the present value of any future cash flow received by the holder of the option decreases. The combined impact of these two effects is to increase the value of call options and decrease the value of put options (see Figure 10.2c, d).

It is important to emphasize that we are assuming that interest rates change while all other variables stay the same. In particular we are assuming in Table 10.1 that interest rates change while the stock price remains the same. In practice, when interest rates rise (fall), stock prices tend to fall (rise). The combined effect of an interest rate increase and the accompanying stock price decrease can be to decrease the value of a call option and increase the value of a put option. Similarly, the combined effect of an interest rate decrease and the accompanying stock price increase can be to increase the value of a call option and decrease the value of a put option.

Amount of Future Dividends

Dividends have the effect of reducing the stock price on the ex-dividend date. This is bad news for the value of call options and good news for the value of put options.

Consider a dividend whose ex-dividend date is during the life of an option. The value of the option is negatively related to the size of the dividend if the option is a call and positively related to the size of the dividend if the option is a put.

10.2 ASSUMPTIONS AND NOTATION

In this chapter, we will make assumptions similar to those made when deriving forward and futures prices in Chapter 5. We assume that there are some market participants, such as large investment banks, for which the following statements are true:

1. There are no transactions costs.
2. All trading profits (net of trading losses) are subject to the same tax rate.
3. Borrowing and lending are possible at the risk-free interest rate.

We assume that these market participants are prepared to take advantage of arbitrage opportunities as they arise. As discussed in Chapters 1 and 5, this means that any available arbitrage opportunities disappear very quickly. For the purposes of our analysis, it is therefore reasonable to assume that there are no arbitrage opportunities.

We will use the following notation:

S_0 : Current stock price

K : Strike price of option

T : Time to expiration of option

S_T : Stock price on the expiration date

r : Continuously compounded risk-free rate of interest for an investment maturing in time T

C : Value of American call option to buy one share

P : Value of American put option to sell one share

c : Value of European call option to buy one share

p : Value of European put option to sell one share

It should be noted that r is the nominal rate of interest, not the real rate of interest. We can assume that $r > 0$. Otherwise, a risk-free investment would provide no advantages over cash. (Indeed, if $r < 0$, cash would be preferable to a risk-free investment.)

10.3 UPPER AND LOWER BOUNDS FOR OPTION PRICES

In this section we derive upper and lower bounds for option prices. These bounds do not depend on any particular assumptions about the factors mentioned in Section 10.1 (except $r > 0$). If an option price is above the upper bound or below the lower bound, then there are profitable opportunities for arbitrageurs.

Upper Bounds

An American or European call option gives the holder the right to buy one share of a stock for a certain price. No matter what happens, the option can never be worth more

than the stock. Hence, the stock price is an upper bound to the option price:

$$c \leq S_0 \quad \text{and} \quad C \leq S_0 \quad (10.1)$$

If these relationships were not true, an arbitrageur could easily make a riskless profit by buying the stock and selling the call option.

An American put option gives the holder the right to sell one share of a stock for K . No matter how low the stock price becomes, the option can never be worth more than K . Hence,

$$P \leq K \quad (10.2)$$

For European options, we know that at maturity the option cannot be worth more than K . It follows that it cannot be worth more than the present value of K today:

$$p \leq Ke^{-rT} \quad (10.3)$$

If this were not true, an arbitrageur could make a riskless profit by writing the option and investing the proceeds of the sale at the risk-free interest rate.

Lower Bound for Calls on Non-Dividend-Paying Stocks

A lower bound for the price of a European call option on a non-dividend-paying stock is

$$S_0 - Ke^{-rT}$$

We first look at a numerical example and then consider a more formal argument.

Suppose that $S_0 = \$20$, $K = \$18$, $r = 10\%$ per annum, and $T = 1$ year. In this case,

$$S_0 - Ke^{-rT} = 20 - 18e^{-0.1} = 3.71$$

or \$3.71. Consider the situation where the European call price is \$3.00, which is less than the theoretical minimum of \$3.71. An arbitrageur can short the stock and buy the call to provide a cash inflow of $\$20.00 - \$3.00 = \$17.00$. If invested for 1 year at 10% per annum, the \$17.00 grows to $17e^{0.1} = \$18.79$. At the end of the year, the option expires. If the stock price is greater than \$18.00, the arbitrageur exercises the option for \$18.00, closes out the short position, and makes a profit of

$$\$18.79 - \$18.00 = \$0.79$$

If the stock price is less than \$18.00, the stock is bought in the market and the short position is closed out. The arbitrageur then makes an even greater profit. For example, if the stock price is \$17.00, the arbitrageur's profit is

$$\$18.79 - \$17.00 = \$1.79$$

For a more formal argument, we consider the following two portfolios:

Portfolio A: one European call option plus a zero-coupon bond that provides a payoff of K at time T

Portfolio B: one share of the stock.

In portfolio A, the zero-coupon bond will be worth K at time T . If $S_T > K$, the call option is exercised at maturity and portfolio A is worth S_T . If $S_T < K$, the call option expires worthless and the portfolio is worth K . Hence, at time T , portfolio A is worth

$$\max(S_T, K)$$

Portfolio B is worth S_T at time T . Hence, portfolio A is always worth as much as, and can be worth more than, portfolio B at the option's maturity. It follows that in the absence of arbitrage opportunities this must also be true today. The zero-coupon bond is worth Ke^{-rT} today. Hence,

$$c + Ke^{-rT} \geq S_0$$

or

$$c \geq S_0 - Ke^{-rT}$$

Because the worst that can happen to a call option is that it expires worthless, its value cannot be negative. This means that $c \geq 0$ and therefore

$$c \geq \max(S_0 - Ke^{-rT}, 0) \quad (10.4)$$

Example 10.1

Consider a European call option on a non-dividend-paying stock when the stock price is \$51, the strike price is \$50, the time to maturity is 6 months, and the risk-free interest rate is 12% per annum. In this case, $S_0 = 51$, $K = 50$, $T = 0.5$, and $r = 0.12$. From equation (10.4), a lower bound for the option price is $S_0 - Ke^{-rT}$, or

$$51 - 50e^{-0.12 \times 0.5} = \$3.91$$

Lower Bound for European Puts on Non-Dividend-Paying Stocks

For a European put option on a non-dividend-paying stock, a lower bound for the price is

$$Ke^{-rT} - S_0$$

Again, we first consider a numerical example and then look at a more formal argument.

Suppose that $S_0 = \$37$, $K = \$40$, $r = 5\%$ per annum, and $T = 0.5$ years. In this case,

$$Ke^{-rT} - S_0 = 40e^{-0.05 \times 0.5} - 37 = \$2.01$$

Consider the situation where the European put price is \$1.00, which is less than the theoretical minimum of \$2.01. An arbitrageur can borrow \$38.00 for 6 months to buy both the put and the stock. At the end of the 6 months, the arbitrageur will be required to repay $38e^{0.05 \times 0.5} = \38.96 . If the stock price is below \$40.00, the arbitrageur exercises the option to sell the stock for \$40.00, repays the loan, and makes a profit of

$$\$40.00 - \$38.96 = \$1.04$$

If the stock price is greater than \$40.00, the arbitrageur discards the option, sells the stock, and repays the loan for an even greater profit. For example, if the stock price is \$42.00, the arbitrageur's profit is

$$\$42.00 - \$38.96 = \$3.04$$

For a more formal argument, we consider the following two portfolios:

Portfolio C: one European put option plus one share

Portfolio D: a zero-coupon bond paying off K at time T .

If $S_T < K$, then the option in portfolio C is exercised at option maturity and the portfolio becomes worth K . If $S_T > K$, then the put option expires worthless and the portfolio is worth S_T at this time. Hence, portfolio C is worth

$$\max(S_T, K)$$

in time T . Portfolio D is worth K in time T . Hence, portfolio C is always worth as much as, and can sometimes be worth more than, portfolio D in time T . It follows that in the absence of arbitrage opportunities portfolio C must be worth at least as much as portfolio D today. Hence,

$$p + S_0 \geq Ke^{-rT}$$

or

$$p \geq Ke^{-rT} - S_0$$

Because the worst that can happen to a put option is that it expires worthless, its value cannot be negative. This means that

$$p \geq \max(Ke^{-rT} - S_0, 0) \quad (10.5)$$

Example 10.2

Consider a European put option on a non-dividend-paying stock when the stock price is \$38, the strike price is \$40, the time to maturity is 3 months, and the risk-free rate of interest is 10% per annum. In this case $S_0 = 38$, $K = 40$, $T = 0.25$, and $r = 0.10$. From equation (10.5), a lower bound for the option price is $Ke^{-rT} - S_0$, or

$$40e^{-0.1 \times 0.25} - 38 = \$1.01$$

10.4 PUT-CALL PARITY

We now derive an important relationship between the prices of European put and call options that have the same strike price and time to maturity. Consider the following two portfolios that were used in the previous section:

Portfolio A: one European call option plus a zero-coupon bond that provides a payoff of K at time T

Portfolio C: one European put option plus one share of the stock.

We continue to assume that the stock pays no dividends. The call and put options have the same strike price K and the same time to maturity T .

As discussed in the previous section, the zero-coupon bond in portfolio A will be worth K at time T . If the stock price S_T at time T proves to be above K , then the call option in portfolio A will be exercised. This means that portfolio A is worth $(S_T - K) + K = S_T$ at time T in these circumstances. If S_T proves to be less than K , then the call option in portfolio A will expire worthless and the portfolio will be worth K at time T .

Table 10.2 Values of Portfolio A and Portfolio C at time T .

| | | $S_T > K$ | $S_T < K$ |
|-------------|------------------|-----------|-----------|
| Portfolio A | Call option | $S_T - K$ | 0 |
| | Zero-coupon bond | K | K |
| | <i>Total</i> | S_T | K |
| Portfolio C | Put Option | 0 | $K - S_T$ |
| | Share | S_T | S_T |
| | <i>Total</i> | S_T | K |

In portfolio C, the share will be worth S_T at time T . If S_T proves to be below K , then the put option in portfolio C will be exercised. This means that portfolio C is worth $(K - S_T) + S_T = K$ at time T in these circumstances. If S_T proves to be greater than K , then the put option in portfolio C will expire worthless and the portfolio will be worth S_T at time T .

The situation is summarized in Table 10.2. If $S_T > K$, both portfolios are worth S_T at time T ; if $S_T < K$, both portfolios are worth K at time T . In other words, both are worth

$$\max(S_T, K)$$

when the options expire at time T . Because they are European, the options cannot be exercised prior to time T . Since the portfolios have identical values at time T , they must have identical values today. If this were not the case, an arbitrageur could buy the less expensive portfolio and sell the more expensive one. Because the portfolios are guaranteed to cancel each other out at time T , this trading strategy would lock in an arbitrage profit equal to the difference in the values of the two portfolios.

The components of portfolio A are worth c and Ke^{-rT} today, and the components of portfolio C are worth p and S_0 today. Hence,

$$c + Ke^{-rT} = p + S_0 \quad (10.6)$$

This relationship is known as *put-call parity*. It shows that the value of a European call with a certain exercise price and exercise date can be deduced from the value of a European put with the same exercise price and exercise date, and vice versa.

To illustrate the arbitrage opportunities when equation (10.6) does not hold, suppose that the stock price is \$31, the exercise price is \$30, the risk-free interest rate is 10% per annum, the price of a three-month European call option is \$3, and the price of a 3-month European put option is \$2.25. In this case,

$$\begin{aligned} c + Ke^{-rT} &= 3 + 30e^{-0.1 \times 3/12} = \$32.26 \\ p + S_0 &= 2.25 + 31 = \$33.25 \end{aligned}$$

Portfolio C is overpriced relative to portfolio A. An arbitrageur can buy the securities in portfolio A and short the securities in portfolio C. The strategy involves buying the call and shorting both the put and the stock, generating a positive cash flow of

$$-3 + 2.25 + 31 = \$30.25$$

up front. When invested at the risk-free interest rate, this amount grows to

$$30.25e^{0.1 \times 0.25} = \$31.02$$

in three months. If the stock price at expiration of the option is greater than \$30, the call will be exercised. If it is less than \$30, the put will be exercised. In either case, the arbitrageur ends up buying one share for \$30. This share can be used to close out the short position. The net profit is therefore

$$\$31.02 - \$30.00 = \$1.02$$

For an alternative situation, suppose that the call price is \$3 and the put price is \$1. In this case,

$$c + Ke^{-rT} = 3 + 30e^{-0.1 \times 3/12} = \$32.26$$

$$p + S_0 = 1 + 31 = \$32.00$$

Portfolio A is overpriced relative to portfolio C. An arbitrageur can short the securities in portfolio A and buy the securities in portfolio C to lock in a profit. The strategy involves shorting the call and buying both the put and the stock with an initial investment of

$$\$31 + \$1 - \$3 = \$29$$

When the investment is financed at the risk-free interest rate, a repayment of $29e^{0.1 \times 0.25} = \29.73 is required at the end of the three months. As in the previous case, either the call or the put will be exercised. The short call and long put option position therefore leads to the stock being sold for \$30.00. The net profit is therefore

$$\$30.00 - \$29.73 = \$0.27$$

These examples are illustrated in Table 10.3. Business Snapshot 10.1 shows how options

Table 10.3 Arbitrage opportunities when put–call parity does not hold.

Stock price = \$31; interest rate = 10%; call price = \$3. Both put and call have strike price of \$30 and three months to maturity.

| <i>Three-month put price = \$2.25</i> | <i>Three-month put price = \$1</i> |
|--|--|
| <i>Action now:</i> | <i>Action now:</i> |
| Buy call for \$3 | Borrow \$29 for 3 months |
| Short put to realize \$2.25 | Short call to realize \$3 |
| Short the stock to realize \$31 | Buy put for \$1 |
| Invest \$30.25 for 3 months | Buy the stock for \$31 |
| <i>Action in 3 months if $S_T > 30$:</i> | <i>Action in 3 months if $S_T > 30$:</i> |
| Receive \$31.02 from investment | Call exercised: sell stock for \$30 |
| Exercise call to buy stock for \$30 | Use \$29.73 to repay loan |
| Net profit = \$1.02 | Net profit = \$0.27 |
| <i>Action in 3 months if $S_T < 30$:</i> | <i>Action in 3 months if $S_T < 30$:</i> |
| Receive \$31.02 from investment | Exercise put to sell stock for \$30 |
| Put exercised: buy stock for \$30 | Use \$29.73 to repay loan |
| Net profit = \$1.02 | Net profit = \$0.27 |

Business Snapshot 10.1 Put–Call Parity and Capital Structure

The pioneers of option pricing were Fischer Black, Myron Scholes, and Robert Merton. In the early 1970s they showed that options can be used to characterize the capital structure of a company. Today this analysis is widely used by financial institutions to assess a company's credit risk.

To illustrate the analysis, consider a company that has assets that are financed with zero-coupon bonds and equity. Suppose that the bonds mature in five years at which time a principal payment of K is required. The company pays no dividends. If the assets are worth more than K in five years, the equity holders choose to repay the bond holders. If the assets are worth less than K , the equity holders choose to declare bankruptcy and the bond holders end up owning the company.

The value of the equity in five years is therefore $\max(A_T - K, 0)$, where A_T is the value of the company's assets at that time. This shows that the equity holders have a five-year European call option on the assets of the company with a strike price of K . What about the bondholders? They get $\min(A_T, K)$ in five years. This is the same as $K - \max(K - A_T, 0)$. This shows that today the bonds are worth the present value of K minus the value of a five-year European put option on the assets with a strike price of K .

To summarize, if c and p are the values, respectively, of the call and put options on the company's assets, then

$$\text{Value of equity} = c$$

$$\text{Value of debt} = PV(K) - p$$

Denote the value of the assets of the company today by A_0 . The value of the assets must equal the total value of the instruments used to finance the assets. This means that it must equal the sum of the value of the equity and the value of the debt, so that

$$A_0 = c + [PV(K) - p]$$

Rearranging this equation, we have

$$c + PV(K) = p + A_0$$

This is the put–call parity result in equation (10.6) for call and put options on the assets of the company.

and put–call parity can help us understand the positions of the debt holders and equity holders in a company.

American Options

Put–call parity holds only for European options. However, it is possible to derive some results for American option prices. It can be shown (see Problem 10.18) that, when there are no dividends,

$$S_0 - K \leq C - P \leq S_0 - Ke^{-rT} \quad (10.7)$$

Example 10.3

An American call option on a non-dividend-paying stock with strike price \$20.00 and maturity in 5 months is worth \$1.50. Suppose that the current stock price is

\$19.00 and the risk-free interest rate is 10% per annum. From equation (10.7), we have

$$19 - 20 \leq C - P \leq 19 - 20e^{-0.1 \times 5/12}$$

or

$$1 \geq P - C \geq 0.18$$

showing that $P - C$ lies between \$1.00 and \$0.18. With C at \$1.50, P must lie between \$1.68 and \$2.50. In other words, upper and lower bounds for the price of an American put with the same strike price and expiration date as the American call are \$2.50 and \$1.68.

10.5 CALLS ON A NON-DIVIDEND-PAYING STOCK

In this section, we first show that it is never optimal to exercise an American call option on a non-dividend-paying stock before the expiration date.

To illustrate the general nature of the argument, consider an American call option on a non-dividend-paying stock with one month to expiration when the stock price is \$70 and the strike price is \$40. The option is deep in the money, and the investor who owns the option might well be tempted to exercise it immediately. However, if the investor plans to hold the stock obtained by exercising the option for more than one month, this is not the best strategy. A better course of action is to keep the option and exercise it at the end of the month. The \$40 strike price is then paid out one month later than it would be if the option were exercised immediately, so that interest is earned on the \$40 for one month. Because the stock pays no dividends, no income from the stock is sacrificed. A further advantage of waiting rather than exercising immediately is that there is some chance (however remote) that the stock price will fall below \$40 in one month. In this case the investor will not exercise in one month and will be glad that the decision to exercise early was not taken!

This argument shows that there are no advantages to exercising early if the investor plans to keep the stock for the remaining life of the option (one month, in this case). What if the investor thinks the stock is currently overpriced and is wondering whether to exercise the option and sell the stock? In this case, the investor is better off selling the option than exercising it.² The option will be bought by another investor who does want to hold the stock. Such investors must exist. Otherwise the current stock price would not be \$70. The price obtained for the option will be greater than its intrinsic value of \$30, for the reasons mentioned earlier.

For a more formal argument, we can use equation (10.4):

$$c \geq S_0 - Ke^{-rT}$$

Because the owner of an American call has all the exercise opportunities open to the owner of the corresponding European call, we must have $C \geq c$. Hence,

$$C \geq S_0 - Ke^{-rT}$$

Given $r > 0$, it follows that $C > S_0 - K$ when $T > 0$. This means that C is always greater

² As an alternative strategy, the investor can keep the option and short the stock to lock in a better profit than \$10.

than the option's intrinsic value prior to maturity. If it were optimal to exercise at a particular time prior to maturity, C would equal the option's intrinsic value at that time. It follows that it can never be optimal to exercise early.

To summarize, there are two reasons an American call on a non-dividend-paying stock should not be exercised early. One relates to the insurance that it provides. A call option, when held instead of the stock itself, in effect insures the holder against the stock price falling below the strike price. Once the option has been exercised and the strike price has been exchanged for the stock price, this insurance vanishes. The other reason concerns the time value of money. From the perspective of the option holder, the later the strike price is paid out the better.

Bounds

Because American call options are never exercised early when there are no dividends, they are equivalent to European call options, so that $C = c$. From equations (10.1) and (10.4), it follows that upper and lower bounds are given by

$$\max(S_0 - Ke^{-rT}, 0) \leq c, C \leq S_0$$

These bounds are illustrated in Figure 10.3.

The general way in which the call price varies with the stock price, S_0 , is shown in Figure 10.4. As r or T or the stock price volatility increases, the line relating the call price to the stock price moves in the direction indicated by the arrows.

10.6 PUTS ON A NON-DIVIDEND-PAYING STOCK

It can be optimal to exercise an American put option on a non-dividend-paying stock early. Indeed, at any given time during its life, a put option should always be exercised early if it is sufficiently deep in the money.

To illustrate, consider an extreme situation. Suppose that the strike price is \$10 and the stock price is virtually zero. By exercising immediately, an investor makes an immediate gain of \$10. If the investor waits, the gain from exercise might be less than \$10, but it cannot be more than \$10, because negative stock prices are impossible.

Figure 10.3 Bounds for European and American call options when there are no dividends.





Figure 10.4 Variation of price of an American or European call option on a non-dividend-paying stock with the stock price. Curve moves in the direction of the arrows when there is an increase in the interest rate, time to maturity, or stock price volatility.



Furthermore, receiving \$10 now is preferable to receiving \$10 in the future. It follows that the option should be exercised immediately.

Like a call option, a put option can be viewed as providing insurance. A put option, when held in conjunction with the stock, insures the holder against the stock price falling below a certain level. However, a put option is different from a call option in that it may be optimal for an investor to forgo this insurance and exercise early in order to realize the strike price immediately. In general, the early exercise of a put option becomes more attractive as S_0 decreases, as r increases, and as the volatility decreases.

Bounds

From equations (10.3) and (10.5), upper and lower bounds for a European put option when there are no dividends are given by

$$\max(Ke^{-rT} - S_0, 0) \leq p \leq Ke^{-rT}$$

For an American put option on a non-dividend-paying stock, the condition

$$P \geq \max(K - S_0, 0)$$


must apply because the option can be exercised at any time. This is a stronger condition than the one for a European put option in equation (10.5). Using the result in equation (10.2), bounds for an American put option on a non-dividend-paying stock are

$$\max(K - S_0, 0) \leq P \leq K$$

Figure 10.5 illustrates the bounds.

Figure 10.6 shows the general way in which the price of an American put option varies with S_0 . As we argued earlier, provided that $r > 0$, it is always optimal to exercise an American put immediately when the stock price is sufficiently low. When early exercise is optimal, the value of the option is $K - S_0$. The curve representing the value

Figure 10.5 Bounds for European and American put options when there are no dividends.



of the put therefore merges into the put's intrinsic value, $K - S_0$, for a sufficiently small value of S_0 . In Figure 10.6, this value of S_0 is shown as point A. The line relating the put price to the stock price moves in the direction indicated by the arrows when r decreases, when the volatility increases, and when T increases.

Because there are some circumstances when it is desirable to exercise an American put option early, it follows that an American put option is always worth more than the corresponding European put option. Furthermore, because an American put is sometimes worth its intrinsic value (see Figure 10.6), it follows that a European put option must sometimes be worth less than its intrinsic value. This means that the curve representing the relationship between the put price and the stock price for a European option must be below the corresponding curve for an American option.

Figure 10.7 shows the variation of the European put price with the stock price. Note that point B in Figure 10.7, at which the price of the option is equal to its intrinsic value, must represent a higher value of the stock price than point A in Figure 10.6 because the curve in Figure 10.7 is below that in Figure 10.6. Point E in Figure 10.7 is where $S_0 = 0$ and the European put price is Ke^{-rT} .

Figure 10.6 Variation of price of an American put option with stock price. Curve moves in the direction of the arrows when the time to maturity or stock price volatility increases or when the interest rate decreases.





Figure 10.7 Variation of price of a European put option with the stock price.



10.7 EFFECT OF DIVIDENDS

The results produced so far in this chapter have assumed that we are dealing with options on a non-dividend-paying stock. In this section, we examine the impact of dividends. We assume that the dividends that will be paid during the life of the option are known. Most exchange-traded stock options have a life of less than one year, so this assumption is not too unreasonable in many situations. We will use D to denote the present value of the dividends during the life of the option. In the calculation of D , a dividend is assumed to occur at the time of its ex-dividend date.

Lower Bound for Calls and Puts

We can redefine portfolios A and B as follows:

- Portfolio A:* one European call option plus an amount of cash equal to $D + Ke^{-rT}$
- Portfolio B:* one share

A similar argument to the one used to derive equation (10.4) shows that

$$c \geq \max(S_0 - D - Ke^{-rT}, 0) \tag{10.8}$$

We can also redefine portfolios C and D as follows:

- Portfolio C:* one European put option plus one share
- Portfolio D:* an amount of cash equal to $D + Ke^{-rT}$

A similar argument to the one used to derive equation (10.5) shows that

$$p \geq \max(D + Ke^{-rT} - S_0, 0) \tag{10.9}$$

Early Exercise

When dividends are expected, we can no longer assert that an American call option will not be exercised early. Sometimes it is optimal to exercise an American call immediately prior to an ex-dividend date. It is never optimal to exercise a call at other times. This point is discussed further in Section 14.12.

Put–Call Parity

Comparing the value at option maturity of the redefined portfolios A and C shows that, with dividends, the put–call parity result in equation (10.6) becomes

$$c + D + Ke^{-rT} = p + S_0 \quad (10.10)$$

Dividends cause equation (10.7) to be modified (see Problem 10.19) to

$$S_0 - D - K \leq C - P \leq S_0 - Ke^{-rT} \quad (10.11)$$

SUMMARY

There are six factors affecting the value of a stock option: the current stock price, the strike price, the expiration date, the stock price volatility, the risk-free interest rate, and the dividends expected during the life of the option. The value of a call generally increases as the current stock price, the time to expiration, the volatility, and the risk-free interest rate increase. The value of a call decreases as the strike price and expected dividends increase. The value of a put generally increases as the strike price, the time to expiration, the volatility, and the expected dividends increase. The value of a put decreases as the current stock price and the risk-free interest rate increase.

It is possible to reach some conclusions about the value of stock options without making any assumptions about the volatility of stock prices. For example, the price of a call option on a stock must always be worth less than the price of the stock itself. Similarly, the price of a put option on a stock must always be worth less than the option's strike price.

A European call option on a non-dividend-paying stock must be worth more than

$$\max(S_0 - Ke^{-rT}, 0)$$

where S_0 is the stock price, K is the strike price, r is the risk-free interest rate, and T is the time to expiration. A European put option on a non-dividend-paying stock must be worth more than

$$\max(Ke^{-rT} - S_0, 0)$$

When dividends with present value D will be paid, the lower bound for a European call option becomes

$$\max(S_0 - D - Ke^{-rT}, 0)$$

and the lower bound for a European put option becomes

$$\max(Ke^{-rT} + D - S_0, 0)$$

Put–call parity is a relationship between the price, c , of a European call option on a stock and the price, p , of a European put option on a stock. For a non-dividend-paying stock, it is

$$c + Ke^{-rT} = p + S_0$$

For a dividend-paying stock, the put–call parity relationship is

$$c + D + Ke^{-rT} = p + S_0$$

Put–call parity does not hold for American options. However, it is possible to use arbitrage arguments to obtain upper and lower bounds for the difference between the price of an American call and the price of an American put.

In Chapter 14, we will carry the analyses in this chapter further by making specific assumptions about the probabilistic behavior of stock prices. The analysis will enable us to derive exact pricing formulas for European stock options. In Chapters 12 and 20, we will see how numerical procedures can be used to price American options.

FURTHER READING

Black, F., and M. Scholes. “The Pricing of Options and Corporate Liabilities,” *Journal of Political Economy*, 81 (May/June 1973): 637–59.

Broadie, M., and J. Detemple. “American Option Valuation: New Bounds, Approximations, and a Comparison of Existing Methods,” *Review of Financial Studies*, 9, 4 (1996): 1211–50.

Merton, R.C.. “On the Pricing of Corporate Debt: The Risk Structure of Interest Rates,” *Journal of Finance*, 29, 2 (1974): 449–70.

Merton, R. C. “Theory of Rational Option Pricing,” *Bell Journal of Economics and Management Science*, 4 (Spring 1973): 141–83.

Merton, R. C. “The Relationship between Put and Call Prices: Comment,” *Journal of Finance*, 28 (March 1973): 183–84.

Stoll, H. R. “The Relationship between Put and Call Option Prices,” *Journal of Finance*, 24 (December 1969): 801–24.

Practice Questions (Answers in Solutions Manual)

- 10.1. List the six factors that affect stock option prices.
- 10.2. What is a lower bound for the price of a 4-month call option on a non-dividend-paying stock when the stock price is \$28, the strike price is \$25, and the risk-free interest rate is 8% per annum?
- 10.3. What is a lower bound for the price of a 1-month European put option on a non-dividend-paying stock when the stock price is \$12, the strike price is \$15, and the risk-free interest rate is 6% per annum?
- 10.4. Give two reasons why the early exercise of an American call option on a non-dividend-paying stock is not optimal. The first reason should involve the time value of money. The second should apply even if interest rates are zero.
- 10.5. “The early exercise of an American put is a trade-off between the time value of money and the insurance value of a put.” Explain this statement.

- 10.6. Why is an American call option on a dividend-paying stock always worth at least as much as its intrinsic value. Is the same true of a European call option? Explain your answer.
- 10.7. The price of a non-dividend-paying stock is \$19 and the price of a 3-month European call option on the stock with a strike price of \$20 is \$1. The risk-free rate is 4% per annum. What is the price of a 3-month European put option with a strike price of \$20?
- 10.8. Explain why the arguments leading to put–call parity for European options cannot be used to give a similar result for American options.
- 10.9. What is a lower bound for the price of a 6-month call option on a non-dividend-paying stock when the stock price is \$80, the strike price is \$75, and the risk-free interest rate is 10% per annum?
- 10.10. What is a lower bound for the price of a 2-month European put option on a non-dividend-paying stock when the stock price is \$58, the strike price is \$65, and the risk-free interest rate is 5% per annum?
- 10.11. A 4-month European call option on a dividend-paying stock is currently selling for \$5. The stock price is \$64, the strike price is \$60, and a dividend of \$0.80 is expected in 1 month. The risk-free interest rate is 12% per annum for all maturities. What opportunities are there for an arbitrageur?
- 10.12. A 1-month European put option on a non-dividend-paying stock is currently selling for \$2.50. The stock price is \$47, the strike price is \$50, and the risk-free interest rate is 6% per annum. What opportunities are there for an arbitrageur?
- 10.13. Give an intuitive explanation of why the early exercise of an American put becomes more attractive as the risk-free rate increases and volatility decreases.
- 10.14. The price of a European call that expires in 6 months and has a strike price of \$30 is \$2. The underlying stock price is \$29, and a dividend of \$0.50 is expected in 2 months and again in 5 months. The term structure is flat, with all risk-free interest rates being 10%. What is the price of a European put option that expires in 6 months and has a strike price of \$30?
- 10.15. Explain the arbitrage opportunities in Problem 10.14 if the European put price is \$3.
- 10.16. The price of an American call on a non-dividend-paying stock is \$4. The stock price is \$31, the strike price is \$30, and the expiration date is in 3 months. The risk-free interest rate is 8%. Derive upper and lower bounds for the price of an American put on the same stock with the same strike price and expiration date.
- 10.17. Explain carefully the arbitrage opportunities in Problem 10.16 if the American put price is greater than the calculated upper bound.
- 10.18. Prove the result in equation (10.7). (*Hint*: For the first part of the relationship, consider (a) a portfolio consisting of a European call plus an amount of cash equal to K , and (b) a portfolio consisting of an American put option plus one share.)
- 10.19. Prove the result in equation (10.11). (*Hint*: For the first part of the relationship, consider (a) a portfolio consisting of a European call plus an amount of cash equal to $D + K$, and (b) a portfolio consisting of an American put option plus one share.)
- 10.20. Consider a 5-year call option on a non-dividend-paying stock granted to employees. The option can be exercised at any time after the end of the first year. Unlike a regular exchange-traded call option, the employee stock option cannot be sold. What is the likely impact of this restriction on the early-exercise decision?

10.21. Use the software DerivaGem to verify that Figures 10.1 and 10.2 are correct.

Further Questions

10.22. A European call option and put option on a stock both have a strike price of \$20 and an expiration date in 3 months. Both sell for \$3. The risk-free interest rate is 10% per annum, the current stock price is \$19, and a \$1 dividend is expected in 1 month. Identify the arbitrage opportunity open to a trader.

10.23. Suppose that c_1 , c_2 , and c_3 are the prices of European call options with strike prices K_1 , K_2 , and K_3 , respectively, where $K_3 > K_2 > K_1$ and $K_3 - K_2 = K_2 - K_1$. All options have the same maturity. Show that

$$c_2 \leq 0.5(c_1 + c_3)$$

(Hint: Consider a portfolio that is long one option with strike price K_1 , long one option with strike price K_3 , and short two options with strike price K_2 .)

10.24. What is the result corresponding to that in Problem 10.23 for European put options?

10.25. Suppose that you are the manager and sole owner of a highly leveraged company. All the debt will mature in 1 year. If at that time the value of the company is greater than the face value of the debt, you will pay off the debt. If the value of the company is less than the face value of the debt, you will declare bankruptcy and the debt holders will own the company.

(a) Express your position as an option on the value of the company.

(b) Express the position of the debt holders in terms of options on the value of the company.

(c) What can you do to increase the value of your position?

10.26. Consider an option on a stock when the stock price is \$41, the strike price is \$40, the risk-free rate is 6%, the volatility is 35%, and the time to maturity is 1 year. Assume that a dividend of \$0.50 is expected after 6 months.

(a) Use DerivaGem to value the option assuming it is a European call.

(b) Use DerivaGem to value the option assuming it is a European put.

(c) Verify that put–call parity holds.

(d) Explore using DerivaGem what happens to the price of the options as the time to maturity becomes very large. For this purpose, assume there are no dividends. Explain the results you get.

10.27. Consider a put option on a non-dividend-paying stock when the stock price is \$40, the strike price is \$42, the risk-free interest rate is 2%, the volatility is 25% per annum, and the time to maturity is three months. Use DerivaGem to determine the following:

(a) The price of the option if it is European (use Analytic: European)

(b) The price of the option if it is American (use Binomial: American with 100 tree steps)

(c) Point B in Figure 10.7.

11

CHAPTER

Trading Strategies Involving Options



We discussed the profit pattern from an investment in a single option in Chapter 9. In this chapter, we look at what can be achieved when an option is traded in conjunction with other assets. In particular, we examine the properties of portfolios consisting of positions in (a) an option and a zero-coupon bond, (b) an option and the asset underlying the option, and (c) two or more options on the same underlying asset.

Further trading strategies involving options are considered in later chapters. For example, Chapter 16 shows how stock indices can be used to manage the risks in a stock portfolio and explains how range forward contracts can be used to hedge a foreign exchange exposure; Chapter 18 covers the way in which Greek letters are used to manage the risks when derivatives are traded; Chapter 25 covers exotic options and what is known as static options replication.

11.1 PRINCIPAL-PROTECTED NOTES

Options are often used to create what are termed *principal-protected notes* for the retail market. These are products that appeal to conservative investors. The return earned by the investor depends on the performance of a stock, a stock index, or other risky asset, but the initial principal amount invested is not at risk. An example will illustrate how a simple principal-protected note can be created.

Example 11.1

Suppose that the 3-year interest rate is 6% with continuous compounding. This means that $1,000e^{-0.06 \times 3} = \835.27 will grow to \$1,000 in 3 years. The difference between \$1,000 and \$835.27 is \$164.73. Suppose that a stock portfolio is worth \$1,000 and provides a dividend yield of 1.5% per annum. Suppose further that a 3-year at-the-money European call option on the stock portfolio can be purchased for less than \$164.73. (From DerivaGem, it can be verified that this will be the case if the volatility of the value of the portfolio is less than about 15%.) A bank can offer clients a \$1,000 investment opportunity consisting of:

1. A 3-year zero-coupon bond with a principal of \$1,000
2. A 3-year at-the-money European call option on the stock portfolio.

If the value of the portfolio increases the investor gets whatever \$1,000 invested in the portfolio would have grown to. (This is because the zero-coupon bond pays off \$1,000 and this equals the strike price of the option.) If the value of the portfolio goes down, the option has no value, but payoff from the zero-coupon bond ensures that the investor receives the original \$1,000 principal invested.

The attraction of a principal-protected note is that an investor is able to take a risky position without risking any principal. The worst that can happen is that the investor loses the chance to earn interest, or other income such as dividends, on the initial investment for the life of the note.

There are many variations on the product we have described. An investor who thinks that the price of an asset will decline can buy a principal-protected note consisting of a zero-coupon bond plus a put option. The investor's payoff in 3 years is then \$1,000 plus the payoff (if any) from the put option.

Is a principal-protected note a good deal from the retail investor's perspective? A bank will always build in a profit for itself when it creates a principal-protected note. This means that, in Example 11.1, the zero-coupon bond plus the call option will always cost the bank less than \$1,000. In addition, investors are taking the risk that the bank will not be in a position to make the payoff on the principal-protected note at maturity. (Some retail investors lost money on principal-protected notes created by Lehman Brothers when it failed in 2008.) In some situations, therefore, an investor will be better off if he or she buys the underlying option in the usual way and invests the remaining principal in a risk-free investment. However, this is not always the case. The investor is likely to face wider bid-offer spreads on the option than the bank and is likely to earn lower interest rates than the bank. It is therefore possible that the bank can add value for the investor while making a profit itself.

Now let us look at the principal-protected notes from the perspective of the bank. The economic viability of the structure in Example 11.1 depends critically on the level of interest rates and the volatility of the portfolio. If the interest rate is 3% instead of 6%, the bank has only $1,000 - 1,000e^{-0.03 \times 3} = \86.07 with which to buy the call option. If interest rates are 6%, but the volatility is 25% instead of 15%, the price of the option would be about \$221. In either of these circumstances, the product described in Example 11.1 cannot be profitably created by the bank. However, there are a number of ways the bank can still create a viable 3-year product. For example, the strike price of the option can be increased so that the value of the portfolio has to rise by, say, 15% before the investor makes a gain; the investor's return could be capped; the return of the investor could depend on the average price of the asset instead of the final price; a knockout barrier could be specified. The derivatives involved in some of these alternatives will be discussed later in the book. (Capping the option corresponds to the creation of a bull spread for the investor and will be discussed later in this chapter.)

One way in which a bank can sometimes create a profitable principal-protected note when interest rates are low or volatilities are high is by increasing its life. Consider the situation in Example 11.1 when (a) the interest rate is 3% rather than 6% and (b) the stock portfolio has a volatility of 15% and provides a dividend yield of 1.5%. DerivaGem shows that a 3-year at-the-money European option costs about \$119. This is more than the funds available to purchase it ($1,000 - 1,000e^{-0.03 \times 3} = \86.07). A 10-year at-the-money option costs about \$217. This is less than the funds available to purchase it ($1,000 - 1,000e^{-0.03 \times 10} = \259.18), making the structure profitable. When

the life is increased to 20 years, the option cost is about \$281, which is much less than the funds available to purchase it ($1,000 - 1,000e^{-0.03 \times 20} = \451.19), so that the structure is even more profitable.

A critical variable for the bank in our example is the dividend yield. The higher it is, the more profitable the product is for the bank. If the dividend yield were zero, the principal-protected note in Example 11.1 cannot be profitable for the bank no matter how long it lasts. (This follows from equation (10.4).)

11.2 TRADING AN OPTION AND THE UNDERLYING ASSET

For convenience, we will assume that the asset underlying the options considered in the rest of the chapter is a stock. (Similar trading strategies can be developed for other underlying assets. We will also follow the usual practice of calculating the profit from a trading strategy as the final payoff minus the initial cost without discounting.)

There are a number of different trading strategies involving a single option on a stock and the stock itself. The profits from these are illustrated in Figure 11.1. In this figure and in other figures throughout this chapter, the dashed line shows the relationship between profit and the stock price for the individual securities constituting the portfolio, whereas the solid line shows the relationship between profit and the stock price for the whole portfolio.

In Figure 11.1a, the portfolio consists of a long position in a stock plus a short position in a European call option. This is known as *writing a covered call*. The long stock position “covers” or protects the investor from the payoff on the short call that becomes necessary if there is a sharp rise in the stock price. In Figure 11.1b, a short position in a stock is combined with a long position in a call option. This is the reverse of writing a covered call. In Figure 11.1c, the investment strategy involves buying a European put option on a stock and the stock itself. The approach is referred to as a *protective put* strategy. In Figure 11.1d, a short position in a put option is combined with a short position in the stock. This is the reverse of a protective put.


The profit patterns in Figures 11.1a, b, c, d have the same general shape as the profit patterns discussed in Chapter 9 for short put, long put, long call, and short call, respectively. Put–call parity provides a way of understanding why this is so. From Chapter 10, the put–call parity relationship is

$$p + S_0 = c + Ke^{-rT} + D \quad (11.1)$$

where p is the price of a European put, S_0 is the stock price, c is the price of a European call, K is the strike price of both call and put, r is the risk-free interest rate, T is the time to maturity of both call and put, and D is the present value of the dividends anticipated during the life of the options.

Equation (11.1) shows that a long position in a European put combined with a long position in the stock is equivalent to a long European call position plus a certain amount ($= Ke^{-rT} + D$) of cash. This explains why the profit pattern in Figure 11.1c is similar to the profit pattern from a long call position. The position in Figure 11.1d is the reverse of that in Figure 11.1c and therefore leads to a profit pattern similar to that from a short call position.

Figure 11.1 Profit patterns (a) long position in a stock combined with short position in a call; (b) short position in a stock combined with long position in a call; (c) long position in a put combined with long position in a stock; (d) short position in a put combined with short position in a stock.



Equation (11.1) can be rearranged to become

$$S_0 - c = Ke^{-rT} + D - p$$

This shows that a long position in a stock combined with a short position in a European call is equivalent to a short European put position plus a certain amount ($= Ke^{-rT} + D$) of cash. This equality explains why the profit pattern in Figure 11.1a is similar to the profit pattern from a short put position. The position in Figure 11.1b is the reverse of that in Figure 11.1a and therefore leads to a profit pattern similar to that from a long put position.

11.3 SPREADS

A spread trading strategy involves taking a position in two or more options of the same type (i.e., two or more calls or two or more puts).

Bull Spreads

One of the most popular types of spreads is a *bull spread*. This can be created by buying a European call option on a stock with a certain strike price and selling a European call option on the same stock with a higher strike price. Both options have the same expiration date. The strategy is illustrated in Figure 11.2. The profits from the two option positions taken separately are shown by the dashed lines. The profit from the whole strategy is the sum of the profits given by the dashed lines and is indicated by the solid line. Because a call price always decreases as the strike price increases, the value of the option sold is always less than the value of the option bought. A bull spread, when created from calls, therefore requires an initial investment.

Suppose that K_1 is the strike price of the call option bought, K_2 is the strike price of the call option sold, and S_T is the stock price on the expiration date of the options. Table 11.1 shows the total payoff that will be realized from a bull spread in different circumstances. If the stock price does well and is greater than the higher strike price, the payoff is the difference between the two strike prices, or $K_2 - K_1$. If the stock price on the expiration date lies between the two strike prices, the payoff is $S_T - K_1$. If the stock price on the expiration date is below the lower strike price, the payoff is zero. The profit in Figure 11.2 is calculated by subtracting the initial investment from the payoff.

A bull spread strategy limits the investor's upside as well as downside risk. The strategy can be described by saying that the investor has a call option with a strike price equal to K_1 and has chosen to give up some upside potential by selling a call option with strike price K_2 ($K_2 > K_1$). In return for giving up the upside potential, the investor gets the price of the option with strike price K_2 . Three types of bull spreads can be distinguished:

1. Both calls are initially out of the money.
2. One call is initially in the money; the other call is initially out of the money.
3. Both calls are initially in the money.

Figure 11.2 Profit from bull spread created using call options.




Table 11.1 Payoff from a bull spread created using calls.

| <i>Stock price range</i> | <i>Payoff from long call option</i> | <i>Payoff from short call option</i> | <i>Total payoff</i> |
|--------------------------|-------------------------------------|--------------------------------------|---------------------|
| $S_T \leq K_1$ | 0 | 0 | 0 |
| $K_1 < S_T < K_2$ | $S_T - K_1$ | 0 | $S_T - K_1$ |
| $S_T \geq K_2$ | $S_T - K_1$ | $-(S_T - K_2)$ | $K_2 - K_1$ |


The most aggressive bull spreads are those of type 1. They cost very little to set up and have a small probability of giving a relatively high payoff ($= K_2 - K_1$). As we move from type 1 to type 2 and from type 2 to type 3, the spreads become more conservative.

Example 11.2

An investor buys for \$3 a 3-month European call with a strike price of \$30 and sells for \$1 a 3-month European call with a strike price of \$35. The payoff from this bull spread strategy is \$5 if the stock price is above \$35, and zero if it is below \$30. If the stock price is between \$30 and \$35, the payoff is the amount by which the stock price exceeds \$30. The cost of the strategy is $\$3 - \$1 = \$2$. So the profit is:

| <i>Stock price range</i> | <i>Profit</i> |
|--------------------------|---------------|
| $S_T \leq 30$ | -2 |
| $30 < S_T < 35$ | $S_T - 32$ |
| $S_T \geq 35$ | 3 |

Bull spreads can also be created by buying a European put with a low strike price and selling a European put with a high strike price, as illustrated in Figure 11.3. Unlike bull spreads created from calls, those created from puts involve a positive up-front cash flow to the investor (ignoring margin requirements) and a payoff that is either negative or zero.

Figure 11.3 Profit from bull spread created using put options.

Bear Spreads

An investor who enters into a bull spread is hoping that the stock price will increase. By contrast, an investor who enters into a *bear spread* is hoping that the stock price will decline. Bear spreads can be created by buying a European put with one strike price and selling a European put with another strike price. The strike price of the option purchased is greater than the strike price of the option sold. (This is in contrast to a bull spread, where the strike price of the option purchased is always less than the strike price of the option sold.) In Figure 11.4, the profit from the spread is shown by the solid line. A bear spread created from puts involves an initial cash outflow because the price of the put sold is less than the price of the put purchased. In essence, the investor has bought a put with a certain strike price and chosen to give up some of the profit potential by selling a put with a lower strike price. In return for the profit given up, the investor gets the price of the option sold.

Assume that the strike prices are K_1 and K_2 , with $K_1 < K_2$. Table 11.2 shows the payoff that will be realized from a bear spread in different circumstances. If the stock price is greater than K_2 , the payoff is zero. If the stock price is less than K_1 , the payoff is $K_2 - K_1$. If the stock price is between K_1 and K_2 , the payoff is $K_2 - S_T$. The profit is calculated by subtracting the initial cost from the payoff.

Example 11.3

An investor buys for \$3 a 3-month European put with a strike price of \$35 and sells for \$1 a 3-month European put with a strike price of \$30. The payoff from this bear spread strategy is zero if the stock price is above \$35, and \$5 if it is below \$30. If the stock price is between \$30 and \$35, the payoff is $35 - S_T$. The options cost $\$3 - \$1 = \$2$ up front. So the profit is:

| Stock price range | Profit |
|-------------------|------------|
| $S_T \leq 30$ | +3 |
| $30 < S_T < 35$ | $33 - S_T$ |
| $S_T \geq 35$ | -2 |

Figure 11.4 Profit from bear spread created using put options.




Table 11.2 Payoff from a bear spread created with put options.

| Stock price range | Payoff from long put option | Payoff from short put option | Total payoff |
|-------------------|-----------------------------|------------------------------|--------------|
| $S_T \leq K_1$ | $K_2 - S_T$ | $-(K_1 - S_T)$ | $K_2 - K_1$ |
| $K_1 < S_T < K_2$ | $K_2 - S_T$ | 0 | $K_2 - S_T$ |
| $S_T \geq K_2$ | 0 | 0 | 0 |

Like bull spreads, bear spreads limit both the upside profit potential and the downside risk. Bear spreads can be created using calls instead of puts. The investor buys a call with a high strike price and sells a call with a low strike price, as illustrated in Figure 11.5. Bear spreads created with calls involve an initial cash inflow (ignoring margin requirements).

Box Spreads

A box spread is a combination of a bull call spread with strike prices K_1 and K_2 and a bear put spread with the same two strike prices. As shown in Table 11.3, the payoff from a box spread is always $K_2 - K_1$. The value of a box spread is therefore always the present value of this payoff or $(K_2 - K_1)e^{-rT}$. If it has a different value there is an arbitrage opportunity. If the market price of the box spread is too low, it is profitable to buy the box. This involves buying a call with strike price K_1 , buying a put with strike price K_2 , selling a call with strike price K_2 , and selling a put with strike price K_1 . If the market price of the box spread is too high, it is profitable to sell the box. This involves buying a call with strike price K_2 , buying a put with strike price K_1 , selling a call with strike price K_1 , and selling a put with strike price K_2 .

It is important to realize that a box-spread arbitrage only works with European options. Many of the options that trade on exchanges are American. As shown in Business Snapshot 11.1, inexperienced traders who treat American options as European are liable to lose money.


Figure 11.5 Profit from bear spread created using call options.

Table 11.3 Payoff from a box spread.

| Stock price range | Payoff from bull call spread | Payoff from bear put spread | Total payoff |
|-------------------|------------------------------|-----------------------------|--------------|
| $S_T \leq K_1$ | 0 | $K_2 - K_1$ | $K_2 - K_1$ |
| $K_1 < S_T < K_2$ | $S_T - K_1$ | $K_2 - S_T$ | $K_2 - K_1$ |
| $S_T \geq K_2$ | $K_2 - K_1$ | 0 | $K_2 - K_1$ |

Butterfly Spreads

A *butterfly spread* involves positions in options with three different strike prices. It can be created by buying a European call option with a relatively low strike price K_1 , buying a European call option with a relatively high strike price K_3 , and selling two European call options with a strike price K_2 that is halfway between K_1 and K_3 . Generally, K_2 is close to the current stock price. The pattern of profits from the strategy is shown in Figure 11.6. A butterfly spread leads to a profit if the stock price stays close to K_2 , but gives rise to a small loss if there is a significant stock price move in either direction. It is therefore an appropriate strategy for an investor who feels that large stock price moves are unlikely. The strategy requires a small investment initially. The payoff from a butterfly spread is shown in Table 11.4.

Suppose that a certain stock is currently worth \$61. Consider an investor who feels that a significant price move in the next 6 months is unlikely. Suppose that the market prices of 6-month European calls are as follows:

| Strike price (\$) | Call price (\$) |
|-------------------|-----------------|
| 55 | 10 |
| 60 | 7 |
| 65 | 5 |


Figure 11.6 Profit from butterfly spread using call options.

Table 11.4 Payoff from a butterfly spread.

| Stock price range | Payoff from first long call | Payoff from second long call | Payoff from short calls | Total payoff* |
|----------------------|-----------------------------|------------------------------|-------------------------|---------------|
| $S_T \leq K_1$ | 0 | 0 | 0 | 0 |
| $K_1 < S_T \leq K_2$ | $S_T - K_1$ | 0 | 0 | $S_T - K_1$ |
| $K_2 < S_T < K_3$ | $S_T - K_1$ | 0 | $-2(S_T - K_2)$ | $K_3 - S_T$ |
| $S_T \geq K_3$ | $S_T - K_1$ | $S_T - K_3$ | $-2(S_T - K_2)$ | 0 |


* These payoffs are calculated using the relationship $K_2 = 0.5(K_1 + K_3)$.

The investor could create a butterfly spread by buying one call with a \$55 strike price, buying one call with a \$65 strike price, and selling two calls with a \$60 strike price. It costs $\$10 + \$5 - (2 \times \$7) = \1 to create the spread. If the stock price in 6 months is greater than \$65 or less than \$55, the total payoff is zero, and the investor incurs a net loss of \$1. If the stock price is between \$56 and \$64, a profit is made. The maximum profit, \$4, occurs when the stock price in 6 months is \$60.

Butterfly spreads can be created using put options. The investor buys two European puts, one with a low strike price and one with a high strike price, and sells two European puts with an intermediate strike price, as illustrated in Figure 11.7. The butterfly spread in the example considered above would be created by buying one put with a strike price of \$55, another with a strike price of \$65, and selling two puts with a strike price of \$60. The use of put options results in exactly the same spread as the use of call options. Put-call parity can be used to show that the initial investment is the same in both cases.

A butterfly spread can be sold or shorted by following the reverse strategy. Options are sold with strike prices of K_1 and K_3 , and two options with the middle strike price K_2 are purchased. This strategy produces a modest profit if there is a significant movement in the stock price.

Figure 11.7 Profit from butterfly spread using put options.



Business Snapshot 11.1 Losing Money with Box Spreads

Suppose that a stock has a price of \$50 and a volatility of 30%. No dividends are expected and the risk-free rate is 8%. A trader offers you the chance to sell on the CBOE a 2-month box spread where the strike prices are \$55 and \$60 for \$5.10. Should you do the trade?

The trade certainly sounds attractive. In this case $K_1 = 55$, $K_2 = 60$, and the payoff is certain to be \$5 in 2 months. By selling the box spread for \$5.10 and investing the funds for 2 months you would have more than enough funds to meet the \$5 payoff in 2 months. The theoretical value of the box spread today is $5 \times e^{-0.08 \times 2/12} = \4.93 .

Unfortunately there is a snag. CBOE stock options are American and the \$5 payoff from the box spread is calculated on the assumption that the options comprising the box are European. Option prices for this example (calculated using DerivaGem) are shown in the table below. A bull call spread where the strike prices are \$55 and \$60 costs $0.96 - 0.26 = \$0.70$. (This is the same for both European and American options because, as we saw in Chapter 10, the price of a European call is the same as the price of an American call when there are no dividends.) A bear put spread with the same strike prices costs $9.46 - 5.23 = \$4.23$ if the options are European and $10.00 - 5.44 = \$4.56$ if they are American. The combined value of both spreads if they are created with European options is $0.70 + 4.23 = \$4.93$. This is the theoretical box spread price calculated above. The combined value of buying both spreads if they are American is $0.70 + 4.56 = \$5.26$. Selling a box spread created with American options for \$5.10 would not be a good trade. You would realize this almost immediately as the trade involves selling a \$60 strike put and this would be exercised against you almost as soon as you sold it!


| <i>Option type</i> | <i>Strike price</i> | <i>European option price</i> | <i>American option price</i> |
|--------------------|---------------------|------------------------------|------------------------------|
| Call | 60 | 0.26 | 0.26 |
| Call | 55 | 0.96 | 0.96 |
| Put | 60 | 9.46 | 10.00 |
| Put | 55 | 5.23 | 5.44 |

Calendar Spreads

Up to now we have assumed that the options used to create a spread all expire at the same time. We now move on to *calendar spreads* in which the options have the same strike price and different expiration dates.

A calendar spread can be created by selling a European call option with a certain strike price and buying a longer-maturity European call option with the same strike price. The longer the maturity of an option, the more expensive it usually is. A calendar spread therefore usually requires an initial investment. Profit diagrams for calendar spreads are usually produced so that they show the profit when the short-maturity option expires on the assumption that the long-maturity option is closed out at that time. The profit pattern for a calendar spread produced from call options is shown in Figure 11.8. The pattern is similar to the profit from the butterfly spread in Figure 11.6. The investor

Figure 11.8 Profit from calendar spread created using two call options, calculated at the time when the short-maturity call option expires.



makes a profit if the stock price at the expiration of the short-maturity option is close to the strike price of the short-maturity option. However, a loss is incurred when the stock price is significantly above or significantly below this strike price.

To understand the profit pattern from a calendar spread, first consider what happens if the stock price is very low when the short-maturity option expires. The short-maturity option is worthless and the value of the long-maturity option is close to zero. The investor therefore incurs a loss that is close to the cost of setting up the spread initially. Consider next what happens if the stock price, S_T , is very high when the short-maturity option expires. The short-maturity option costs the investor $S_T - K$, and the long-maturity option is worth close to $S_T - K$, where K is the strike price of the options. Again, the investor makes a net loss that is close to the cost of setting up the spread initially. If S_T is close to K , the short-maturity option costs the investor either a small amount or nothing at all. However, the long-maturity option is still quite valuable. In this case a significant net profit is made.

In a *neutral calendar spread*, a strike price close to the current stock price is chosen. A *bullish calendar spread* involves a higher strike price, whereas a *bearish calendar spread* involves a lower strike price.


Calendar spreads can be created with put options as well as call options. The investor buys a long-maturity put option and sells a short-maturity put option. As shown in Figure 11.9, the profit pattern is similar to that obtained from using calls.

A *reverse calendar spread* is the opposite to that in Figures 11.8 and 11.9. The investor buys a short-maturity option and sells a long-maturity option. A small profit arises if the stock price at the expiration of the short-maturity option is well above or well below the strike price of the short-maturity option. However, a significant loss results if it is close to the strike price.

Diagonal Spreads

Bull, bear, and calendar spreads can all be created from a long position in one call and a short position in another call. In the case of bull and bear spreads, the calls have

Figure 11.9 Profit from calendar spread created using two put options, calculated at the time when the short-maturity put option expires.



different strike prices and the same expiration date. In the case of calendar spreads, the calls have the same strike price and different expiration dates.

In a *diagonal spread* both the expiration date and the strike price of the calls are different. This increases the range of profit patterns that are possible.

11.4 COMBINATIONS

A *combination* is an option trading strategy that involves taking a position in both calls and puts on the same stock. We will consider straddles, strips, straps, and strangles.

Straddle

One popular combination is a *straddle*, which involves buying a European call and put with the same strike price and expiration date. The profit pattern is shown in Figure 11.10. The strike price is denoted by K . If the stock price is close to this strike price at expiration

Figure 11.10 Profit from a straddle.




Table 11.5 Payoff from a straddle.

| <i>Range of stock price</i> | <i>Payoff from call</i> | <i>Payoff from put</i> | <i>Total payoff</i> |
|-----------------------------|-------------------------|------------------------|---------------------|
| $S_T \leq K$ | 0 | $K - S_T$ | $K - S_T$ |
| $S_T > K$ | $S_T - K$ | 0 | $S_T - K$ |

of the options, the straddle leads to a loss. However, if there is a sufficiently large move in either direction, a significant profit will result. The payoff from a straddle is calculated in Table 11.5.

A straddle is appropriate when an investor is expecting a large move in a stock price but does not know in which direction the move will be. Consider an investor who feels that the price of a certain stock, currently valued at \$69 by the market, will move significantly in the next 3 months. The investor could create a straddle by buying both a put and a call with a strike price of \$70 and an expiration date in 3 months. Suppose that the call costs \$4 and the put costs \$3. If the stock price stays at \$69, it is easy to see that the strategy costs the investor \$6. (An up-front investment of \$7 is required, the call expires worthless, and the put expires worth \$1.) If the stock price moves to \$70, a loss of \$7 is experienced. (This is the worst that can happen.) However, if the stock price jumps up to \$90, a profit of \$13 is made; if the stock moves down to \$55, a profit of \$8 is made; and so on. As discussed in Business Snapshot 11.2 an investor should carefully consider whether the jump that he or she anticipates is already reflected in option prices before putting on a straddle trade.

The straddle in Figure 11.10 is sometimes referred to as a *bottom straddle* or *straddle purchase*. A *top straddle* or *straddle write* is the reverse position. It is created by selling a call and a put with the same exercise price and expiration date. It is a highly risky strategy. If the stock price on the expiration date is close to the strike price, a significant profit results. However, the loss arising from a large move is unlimited.

Strips and Straps

A *strip* consists of a long position in one European call and two European puts with the same strike price and expiration date. A *strap* consists of a long position in two European calls and one European put with the same strike price and expiration date. The profit patterns from strips and straps are shown in Figure 11.11. In a strip the investor is betting that there will be a big stock price move and considers a decrease in the stock price to be more likely than an increase. In a strap the investor is also betting that there will be a big stock price move. However, in this case, an increase in the stock price is considered to be more likely than a decrease.

Strangles

In a *strangle*, sometimes called a *bottom vertical combination*, an investor buys a European put and a European call with the same expiration date and different strike prices. The profit pattern that is obtained is shown in Figure 11.12. The call strike price, K_2 , is higher than the put strike price, K_1 . The payoff function for a strangle is calculated in Table 11.6.

Business Snapshot 11.2 How to Make Money from Trading Straddles

Suppose that a big move is expected in a company's stock price because there is a takeover bid for the company or the outcome of a major lawsuit involving the company is about to be announced. Should you trade a straddle?

A straddle seems a natural trading strategy in this case. However, if your view of the company's situation is much the same as that of other market participants, this view will be reflected in the prices of options. Options on the stock will be significantly more expensive than options on a similar stock for which no jump is expected. The \wedge -shaped profit pattern from the straddle in Figure 11.10 will have moved downward, so that a bigger move in the stock price is necessary for you to make a profit.

For a straddle to be an effective strategy, you must believe that there are likely to be big movements in the stock price and these beliefs must be different from those of most other investors. Market prices incorporate the beliefs of market participants. To make money from any investment strategy, you must take a view that is different from most of the rest of the market—and you must be right!

A strangle is a similar strategy to a straddle. The investor is betting that there will be a large price move, but is uncertain whether it will be an increase or a decrease. Comparing Figures 11.12 and 11.10, we see that the stock price has to move farther in a strangle than in a straddle for the investor to make a profit. However, the downside risk if the stock price ends up at a central value is less with a strangle.

The profit pattern obtained with a strangle depends on how close together the strike prices are. The farther they are apart, the less the downside risk and the farther the stock price has to move for a profit to be realized.

The sale of a strangle is sometimes referred to as a *top vertical combination*. It can be appropriate for an investor who feels that large stock price moves are unlikely. However, as with sale of a straddle, it is a risky strategy involving unlimited potential loss to the investor.

Figure 11.11 Profit from a strip and a strap.





Figure 11.12 Profit from a strangle.

11.5 OTHER PAYOFFS

This chapter has demonstrated just a few of the ways in which options can be used to produce an interesting relationship between profit and stock price. If European options expiring at time T were available with every single possible strike price, any payoff function at time T could in theory be obtained. The easiest illustration of this involves butterfly spreads. Recall that a butterfly spread is created by buying options with strike prices K_1 and K_3 and selling two options with strike price K_2 , where $K_1 < K_2 < K_3$ and $K_3 - K_2 = K_2 - K_1$. Figure 11.13 shows the payoff from a butterfly spread. The pattern could be described as a spike. As K_1 and K_3 move closer together, the spike becomes smaller. Through the judicious combination of a large number of very small spikes, any payoff function can be approximated.

SUMMARY


Principal-protected notes can be created from a zero-coupon bond and a European call option. They are attractive to some investors because the issuer of the product guarantees that the purchaser will receive his or her principal back regardless of the performance of the asset underlying the option.

A number of common trading strategies involve a single option and the underlying stock. For example, writing a covered call involves buying the stock and selling a call option on the stock; a protective put involves buying a put option and buying the stock. The former is similar to selling a put option; the latter is similar to buying a call option.

Table 11.6 Payoff from a strangle.

| <i>Range of stock price</i> | <i>Payoff from call</i> | <i>Payoff from put</i> | <i>Total payoff</i> |
|-----------------------------|-------------------------|------------------------|---------------------|
| $S_T \leq K_1$ | 0 | $K_1 - S_T$ | $K_1 - S_T$ |
| $K_1 < S_T < K_2$ | 0 | 0 | 0 |
| $S_T \geq K_2$ | $S_T - K_2$ | 0 | $S_T - K_2$ |

Figure 11.13 “Spike payoff” from a butterfly spread that can be used as a building block to create other payoffs.



Spreads involve either taking a position in two or more calls or taking a position in two or more puts. A bull spread can be created by buying a call (put) with a low strike price and selling a call (put) with a high strike price. A bear spread can be created by buying a put (call) with a high strike price and selling a put (call) with a low strike price. A butterfly spread involves buying calls (puts) with a low and high strike price and selling two calls (puts) with some intermediate strike price. A calendar spread involves selling a call (put) with a short time to expiration and buying a call (put) with a longer time to expiration. A diagonal spread involves a long position in one option and a short position in another option such that both the strike price and the expiration date are different.

Combinations involve taking a position in both calls and puts on the same stock. A straddle combination involves taking a long position in a call and a long position in a put with the same strike price and expiration date. A strip consists of a long position in one call and two puts with the same strike price and expiration date. A strap consists of a long position in two calls and one put with the same strike price and expiration date. A strangle consists of a long position in a call and a put with different strike prices and the same expiration date. There are many other ways in which options can be used to produce interesting payoffs. It is not surprising that option trading has steadily increased in popularity and continues to fascinate investors.

FURTHER READING

- Bharadwaj, A. and J. B. Wiggins. “Box Spread and Put–Call Parity Tests for the S&P Index LEAPS Markets,” *Journal of Derivatives*, 8, 4 (Summer 2001): 62–71.
- Chaput, J. S., and L. H. Ederington, “Option Spread and Combination Trading,” *Journal of Derivatives*, 10, 4 (Summer 2003): 70–88.
- McMillan, L. G. *McMillan on Options*, 2nd edn. Hoboken, NJ: Wiley, 2004.
- Rendleman, R. J. “Covered Call Writing from an Expected Utility Perspective,” *Journal of Derivatives*, 8, 3 (Spring 2001): 63–75.
- Ronn, A. G. and E. I. Ronn. “The Box–Spread Arbitrage Conditions,” *Review of Financial Studies*, 2, 1 (1989): 91–108.

Practice Questions (Answers in Solutions Manual)

- 11.1. What is meant by a protective put? What position in call options is equivalent to a protective put?

- 11.2. Explain two ways in which a bear spread can be created.
- 11.3. When is it appropriate for an investor to purchase a butterfly spread?
- 11.4. Call options on a stock are available with strike prices of \$15, $\$17\frac{1}{2}$, and \$20, and expiration dates in 3 months. Their prices are \$4, \$2, and $\$2\frac{1}{2}$, respectively. Explain how the options can be used to create a butterfly spread. Construct a table showing how profit varies with stock price for the butterfly spread.
- 11.5. What trading strategy creates a reverse calendar spread?
- 11.6. What is the difference between a strangle and a straddle?
- 11.7. A call option with a strike price of \$50 costs \$2. A put option with a strike price of \$45 costs \$3. Explain how a strangle can be created from these two options. What is the pattern of profits from the strangle?
- 11.8. Use put–call parity to relate the initial investment for a bull spread created using calls to the initial investment for a bull spread created using puts.
- 11.9. Explain how an aggressive bear spread can be created using put options.
- 11.10. Suppose that put options on a stock with strike prices \$30 and \$35 cost \$4 and \$7, respectively. How can the options be used to create (a) a bull spread and (b) a bear spread? Construct a table that shows the profit and payoff for both spreads.
- 11.11. Use put–call parity to show that the cost of a butterfly spread created from European puts is identical to the cost of a butterfly spread created from European calls.
- 11.12. A call with a strike price of \$60 costs \$6. A put with the same strike price and expiration date costs \$4. Construct a table that shows the profit from a straddle. For what range of stock prices would the straddle lead to a loss?
- 11.13. Construct a table showing the payoff from a bull spread when puts with strike prices K_1 and K_2 , with $K_2 > K_1$, are used.
- 11.14. An investor believes that there will be a big jump in a stock price, but is uncertain as to the direction. Identify six different strategies the investor can follow and explain the differences among them.
- 11.15. How can a forward contract on a stock with a particular delivery price and delivery date be created from options?
- 11.16. “A box spread comprises four options. Two can be combined to create a long forward position and two can be combined to create a short forward position.” Explain this statement.
- 11.17. What is the result if the strike price of the put is higher than the strike price of the call in a strangle?
- 11.18. One Australian dollar is currently worth \$0.64. A 1-year butterfly spread is set up using European call options with strike prices of \$0.60, \$0.65, and \$0.70. The risk-free interest rates in the United States and Australia are 5% and 4% respectively, and the volatility of the exchange rate is 15%. Use the DerivaGem software to calculate the cost of setting up the butterfly spread position. Show that the cost is the same if European put options are used instead of European call options.
- 11.19. An index provides a dividend yield of 1% and has a volatility of 20%. The risk-free interest rate is 4%. How long does a principal-protected note, created as in Example 11.1, have to last for it to be profitable for the bank issuing it? Use DerivaGem.

Further Questions

- 11.20. Three put options on a stock have the same expiration date and strike prices of \$55, \$60, and \$65. The market prices are \$3, \$5, and \$8, respectively. Explain how a butterfly spread can be created. Construct a table showing the profit from the strategy. For what range of stock prices would the butterfly spread lead to a loss?
- 11.21. A diagonal spread is created by buying a call with strike price K_2 and exercise date T_2 and selling a call with strike price K_1 and exercise date T_1 , where $T_2 > T_1$. Draw a diagram showing the profit when (a) $K_2 > K_1$ and (b) $K_2 < K_1$.
- 11.22. Draw a diagram showing the variation of an investor's profit and loss with the terminal stock price for a portfolio consisting of:
- (a) One share and a short position in one call option
 - (b) Two shares and a short position in one call option
 - (c) One share and a short position in two call options
 - (d) One share and a short position in four call options.
- In each case, assume that the call option has an exercise price equal to the current stock price.
- 11.23. Suppose that the price of a non-dividend-paying stock is \$32, its volatility is 30%, and the risk-free rate for all maturities is 5% per annum. Use DerivaGem to calculate the cost of setting up the following positions:
- (a) A bull spread using European call options with strike prices of \$25 and \$30 and a maturity of 6 months
 - (b) A bear spread using European put options with strike prices of \$25 and \$30 and a maturity of 6 months
 - (c) A butterfly spread using European call options with strike prices of \$25, \$30, and \$35 and a maturity of 1 year
 - (d) A butterfly spread using European put options with strike prices of \$25, \$30, and \$35 and a maturity of 1 year
 - (e) A straddle using options with a strike price of \$30 and a 6-month maturity
 - (f) A strangle using options with strike prices of \$25 and \$35 and a 6-month maturity.
- In each case provide a table showing the relationship between profit and final stock price. Ignore the impact of discounting.
- 11.24. What trading position is created from a long strangle and a short straddle when both have the same time to maturity? Assume that the strike price in the straddle is halfway between the two strike prices of the strangle.
- 11.25. Describe the trading position created in which a call option is bought with strike price K_2 and a put option is sold with strike price K_1 when both have the same time to maturity and $K_2 > K_1$. What does the position become when $K_1 = K_2$?
- 11.26. A bank decides to create a five-year principal-protected note on a non-dividend-paying stock by offering investors a zero-coupon bond plus a bull spread created from calls. The risk-free rate is 4% and the stock price volatility is 25%. The low-strike-price option in the bull spread is at the money. What is the maximum ratio of the high strike price to the low strike price in the bull spread. Use DerivaGem.

12

CHAPTER

Binomial Trees



A useful and very popular technique for pricing an option involves constructing a *binomial tree*. This is a diagram representing different possible paths that might be followed by the stock price over the life of an option. The underlying assumption is that the stock price follows a *random walk*. In each time step, it has a certain probability of moving up by a certain percentage amount and a certain probability of moving down by a certain percentage amount. In the limit, as the time step becomes smaller, this model is the same as the Black–Scholes–Merton model we will be discussing in Chapter 14. Indeed, in the appendix to this chapter, we show that the European option price given by the binomial tree converges to the Black–Scholes–Merton price as the time step becomes smaller.


The material in this chapter is important for a number of reasons. First, it explains the nature of the no-arbitrage arguments that are used for valuing options. Second, it explains the binomial tree numerical procedure that is widely used for valuing American options and other derivatives. Third, it introduces a very important principle known as risk-neutral valuation.

The general approach adopted here is similar to that in an important paper published by Cox, Ross, and Rubinstein in 1979. More details on numerical procedures using binomial trees are given in Chapter 20.

12.1 A ONE-STEP BINOMIAL MODEL AND A NO-ARBITRAGE ARGUMENT

We start by considering a very simple situation. A stock price is currently \$20, and it is known that at the end of 3 months it will be either \$22 or \$18. We are interested in valuing a European call option to buy the stock for \$21 in 3 months. This option will have one of two values at the end of the 3 months. If the stock price turns out to be \$22, the value of the option will be \$1; if the stock price turns out to be \$18, the value of the option will be zero. The situation is illustrated in Figure 12.1.

It turns out that a relatively simple argument can be used to price the option in this example. The only assumption needed is that arbitrage opportunities do not exist. We set up a portfolio of the stock and the option in such a way that there is no uncertainty

Figure 12.1 Stock price movements for numerical example in Section 12.1.

about the value of the portfolio at the end of the 3 months. We then argue that, because the portfolio has no risk, the return it earns must equal the risk-free interest rate. This enables us to work out the cost of setting up the portfolio and therefore the option's price. Because there are two securities (the stock and the stock option) and only two possible outcomes, it is always possible to set up the riskless portfolio.

Consider a portfolio consisting of a long position in Δ shares of the stock and a short position in one call option (Δ is the capital Greek letter "delta"). We calculate the value of Δ that makes the portfolio riskless. If the stock price moves up from \$20 to \$22, the value of the shares is 22Δ and the value of the option is 1, so that the total value of the portfolio is $22\Delta - 1$. If the stock price moves down from \$20 to \$18, the value of the shares is 18Δ and the value of the option is zero, so that the total value of the portfolio is 18Δ . The portfolio is riskless if the value of Δ is chosen so that the final value of the portfolio is the same for both alternatives. This means that

$$22\Delta - 1 = 18\Delta$$

or

$$\Delta = 0.25$$

A riskless portfolio is therefore

Long: 0.25 shares

Short: 1 option.

If the stock price moves up to \$22, the value of the portfolio is

$$22 \times 0.25 - 1 = 4.5$$

If the stock price moves down to \$18, the value of the portfolio is

$$18 \times 0.25 = 4.5$$

Regardless of whether the stock price moves up or down, the value of the portfolio is always 4.5 at the end of the life of the option. This shows that Δ is the number of shares necessary to hedge a short position in one option. It is one of the "Greek letters" that will be discussed later in this chapter and in Chapter 18.

Riskless portfolios must, in the absence of arbitrage opportunities, earn the risk-free rate of interest. Suppose that, in this case, the risk-free rate is 12% per annum. It

follows that the value of the portfolio today must be the present value of 4.5, or

$$4.5e^{-0.12 \times 3/12} = 4.367$$

The value of the stock price today is known to be \$20. Suppose the option price is denoted by f . The value of the portfolio today is

$$20 \times 0.25 - f = 5 - f$$

It follows that

$$5 - f = 4.367$$

or

$$f = 0.633$$

This shows that, in the absence of arbitrage opportunities, the current value of the option must be 0.633. If the value of the option were more than 0.633, the portfolio would cost less than 4.367 to set up and would earn more than the risk-free rate. If the value of the option were less than 0.633, shorting the portfolio would provide a way of borrowing money at less than the risk-free rate.


A Generalization

We can generalize the no-arbitrage argument just presented by considering a stock whose price is S_0 and an option on the stock (or any derivative dependent on the stock) whose current price is f . We suppose that the option lasts for time T and that during the life of the option the stock price can either move up from S_0 to a new level, S_0u , where $u > 1$, or down from S_0 to a new level, S_0d , where $d < 1$. The percentage increase in the stock price when there is an up movement is $u - 1$; the percentage decrease when there is a down movement is $1 - d$. If the stock price moves up to S_0u , we suppose that the payoff from the option is f_u ; if the stock price moves down to S_0d , we suppose the payoff from the option is f_d . The situation is illustrated in Figure 12.2.

As before, we imagine a portfolio consisting of a long position in Δ shares and a short position in one option. We calculate the value of Δ that makes the portfolio riskless. If there is an up movement in the stock price, the value of the portfolio at the end of the life of the option is

$$S_0u\Delta - f_u$$

Figure 12.2 Stock and option prices in a general one-step tree.



If there is a down movement in the stock price, the value becomes

$$S_0d\Delta - f_d$$

The two are equal when

$$S_0u\Delta - f_u = S_0d\Delta - f_d$$

or

$$\Delta = \frac{f_u - f_d}{S_0u - S_0d} \quad (12.1)$$

In this case, the portfolio is riskless and, for there to be no arbitrage opportunities, it must earn the risk-free interest rate. Equation (12.1) shows that Δ is the ratio of the change in the option price to the change in the stock price as we move between the nodes at time T .

If we denote the risk-free interest rate by r , the present value of the portfolio is

$$(S_0u\Delta - f_u)e^{-rT}$$

The cost of setting up the portfolio is

$$S_0\Delta - f$$

It follows that

$$S_0\Delta - f = (S_0u\Delta - f_u)e^{-rT}$$

or

$$f = S_0\Delta(1 - ue^{-rT}) + f_ue^{-rT}$$

Substituting from equation (12.1) for Δ , we obtain

$$f = S_0 \left(\frac{f_u - f_d}{S_0u - S_0d} \right) (1 - ue^{-rT}) + f_ue^{-rT}$$

or

$$f = \frac{f_u(1 - de^{-rT}) + f_d(ue^{-rT} - 1)}{u - d}$$

or

$$f = e^{-rT} [pf_u + (1 - p)f_d] \quad (12.2)$$

where

$$p = \frac{e^{rT} - d}{u - d} \quad (12.3)$$

Equations (12.2) and (12.3) enable an option to be priced when stock price movements are given by a one-step binomial tree. The only assumption needed for the equation is that there are no arbitrage opportunities in the market.

In the numerical example considered previously (see Figure 12.1), $u = 1.1$, $d = 0.9$, $r = 0.12$, $T = 0.25$, $f_u = 1$, and $f_d = 0$. From equation (12.3), we have

$$p = \frac{e^{0.12 \times 3/12} - 0.9}{1.1 - 0.9} = 0.6523$$

and, from equation (12.2), we have

$$f = e^{-0.12 \times 0.25} (0.6523 \times 1 + 0.3477 \times 0) = 0.633$$

The result agrees with the answer obtained earlier in this section.

Irrelevance of the Stock's Expected Return

The option pricing formula in equation (12.2) does not involve the probabilities of the stock price moving up or down. For example, we get the same option price when the probability of an upward movement is 0.5 as we do when it is 0.9. This is surprising and seems counterintuitive. It is natural to assume that, as the probability of an upward movement in the stock price increases, the value of a call option on the stock increases and the value of a put option on the stock decreases. This is not the case.

The key reason is that we are not valuing the option in absolute terms. We are calculating its value in terms of the price of the underlying stock. The probabilities of future up or down movements are already incorporated into the stock price: we do not need to take them into account again when valuing the option in terms of the stock price.

12.2 RISK-NEUTRAL VALUATION

We are now in a position to introduce a very important principle in the pricing of derivatives known as *risk-neutral valuation*. This states that, when valuing a derivative, we can make the assumption that investors are *risk-neutral*. This assumption means investors do not increase the expected return they require from an investment to compensate for increased risk. A world where investors are risk-neutral is referred to as a *risk-neutral world*. The world we live in is, of course, not a risk-neutral world. The higher the risks investors take, the higher the expected returns they require. However, it turns out that assuming a risk-neutral world gives us the right option price for the world we live in, as well as for a risk-neutral world. Almost miraculously, it finesses the problem that we know hardly anything about the risk aversion of the buyers and sellers of options.

Risk-neutral valuation seems a surprising result when it is first encountered. Options are risky investments. Should not a person's risk preferences affect how they are priced? The answer is that, when we are pricing an option in terms of the price of the underlying stock, risk preferences are unimportant. As investors become more risk-averse, stock prices decline, but the formulas relating option prices to stock prices remain the same.

A risk-neutral world has two features that simplify the pricing of derivatives:

1. The expected return on a stock (or any other investment) is the risk-free rate.
2. The discount rate used for the expected payoff on an option (or any other instrument) is the risk-free rate.

Returning to equation (12.2), the parameter p should be interpreted as the probability of an up movement in a risk-neutral world, so that $1 - p$ is the probability of a down movement in this world. The expression

$$pf_u + (1 - p)f_d$$

is the expected future payoff from the option in a risk-neutral world and equation (12.2) states that the value of the option today is its expected future payoff in a risk-neutral world discounted at the risk-free rate. This is an application of risk-neutral valuation.

To prove the validity of our interpretation of p , we note that, when p is the probability of an up movement, the expected stock price $E(S_T)$ at time T is given by

$$E(S_T) = pS_0u + (1 - p)S_0d$$

or

$$E(S_T) = pS_0(u - d) + S_0d$$

Substituting from equation (12.3) for p gives

$$E(S_T) = S_0e^{rT} \quad (12.4)$$

This shows that the stock price grows, on average, at the risk-free rate when p is the probability of an up movement. In other words, the stock price behaves exactly as we would expect it to behave in a risk-neutral world when p is the probability of an up movement.

Risk-neutral valuation is a very important general result in the pricing of derivatives. It states that, when we assume the world is risk-neutral, we get the right price for a derivative in all worlds, not just in a risk-neutral one. We have shown that risk-neutral valuation is correct when a simple binomial model is assumed for how the price of the stock evolves. It can be shown that the result is true regardless of the assumptions we make about the evolution of the stock price.

To apply risk-neutral valuation to the pricing of a derivative, we first calculate what the probabilities of different outcomes would be if the world were risk-neutral. We then calculate the expected payoff from the derivative and discount that expected payoff at the risk-free rate of interest.

The One-Step Binomial Example Revisited

We now return to the example in Figure 12.1 and illustrate that risk-neutral valuation gives the same answer as no-arbitrage arguments. In Figure 12.1, the stock price is currently \$20 and will move either up to \$22 or down to \$18 at the end of 3 months. The option considered is a European call option with a strike price of \$21 and an expiration date in 3 months. The risk-free interest rate is 12% per annum.

We define p as the probability of an upward movement in the stock price in a risk-neutral world. We can calculate p from equation (12.3). Alternatively, we can argue that the expected return on the stock in a risk-neutral world must be the risk-free rate of 12%. This means that p must satisfy

$$22p + 18(1 - p) = 20e^{0.12 \times 3/12}$$

or

$$4p = 20e^{0.12 \times 3/12} - 18$$

That is, p must be 0.6523.

At the end of the 3 months, the call option has a 0.6523 probability of being worth 1 and a 0.3477 probability of being worth zero. Its expected value is therefore

$$0.6523 \times 1 + 0.3477 \times 0 = 0.6523$$

In a risk-neutral world this should be discounted at the risk-free rate. The value of the option today is therefore

$$0.6523e^{-0.12 \times 3/12}$$

or \$0.633. This is the same as the value obtained earlier, demonstrating that no-arbitrage arguments and risk-neutral valuation give the same answer.

Real World vs. Risk-Neutral World

It should be emphasized that p is the probability of an up movement in a risk-neutral world. In general, this is not the same as the probability of an up movement in the real world. In our example $p = 0.6523$. When the probability of an up movement is 0.6523, the expected return on both the stock and the option is the risk-free rate of 12%. Suppose that, in the real world, the expected return on the stock is 16% and p^* is the probability of an up movement in this world. It follows that

$$22p^* + 18(1 - p^*) = 20e^{0.16 \times 3/12}$$

so that $p^* = 0.7041$.

The expected payoff from the option in the real world is then given by

$$p^* \times 1 + (1 - p^*) \times 0$$

or 0.7041. Unfortunately, it is not easy to know the correct discount rate to apply to the expected payoff in the real world. The return the market requires on the stock is 16% and this is the discount rate that would be used for the expected cash flows from an investment in the stock. A position in a call option is riskier than a position in the stock. As a result the discount rate to be applied to the payoff from a call option is greater than 16%, but we do not know how much greater than 16% it should be.¹ Using risk-neutral valuation solves this problem because we know that in a risk-neutral world the expected return on all assets (and therefore the discount rate to use for all expected payoffs) is the risk-free rate.


12.3 TWO-STEP BINOMIAL TREES

We can extend the analysis to a two-step binomial tree such as that shown in Figure 12.3. Here the stock price starts at \$20 and in each of two time steps may go up by 10% or down by 10%. Each time step is 3 months long and the risk-free interest rate is 12% per annum. We consider a 6-month option with a strike price of \$21.

The objective of the analysis is to calculate the option price at the initial node of the tree. This can be done by repeatedly applying the principles established earlier in the chapter. Figure 12.4 shows the same tree as Figure 12.3, but with both the stock price and the option price at each node. (The stock price is the upper number and the option price is the lower number.) The option prices at the final nodes of the tree are easily calculated. They are the payoffs from the option. At node D the stock price is 24.2 and the option price is $24.2 - 21 = 3.2$; at nodes E and F the option is out of the money and its value is zero.

At node C the option price is zero, because node C leads to either node E or node F and at both of those nodes the option price is zero. We calculate the option price at node B by focusing our attention on the part of the tree shown in Figure 12.5. Using the

¹ Since we know the correct value of the option is 0.633, we can deduce that the correct real-world discount rate is 42.58%. This is because $0.633 = 0.7041e^{-0.4258 \times 3/12}$.

Figure 12.3 Stock prices in a two-step tree.

notation introduced earlier in the chapter, $u = 1.1$, $d = 0.9$, $r = 0.12$, and $T = 0.25$, so that $p = 0.6523$, and equation (12.2) gives the value of the option at node B as

$$e^{-0.12 \times 3/12} (0.6523 \times 3.2 + 0.3477 \times 0) = 2.0257$$

It remains for us to calculate the option price at the initial node A. We do so by focusing on the first step of the tree. We know that the value of the option at node B is 2.0257 and



Figure 12.4 Stock and option prices in a two-step tree. The upper number at each node is the stock price and the lower number is the option price.

Figure 12.5 Evaluation of option price at node B of Figure 12.4.

that at node C it is zero. Equation (12.2) therefore gives the value at node A as

$$e^{-0.12 \times 3/12} (0.6523 \times 2.0257 + 0.3477 \times 0) = 1.2823$$

The value of the option is \$1.2823.

Note that this example was constructed so that u and d (the proportional up and down movements) were the same at each node of the tree and so that the time steps were of the same length. As a result, the risk-neutral probability, p , as calculated by equation (12.3) is the same at each node.

A Generalization

We can generalize the case of two time steps by considering the situation in Figure 12.6. The stock price is initially S_0 . During each time step, it either moves up to u times its initial value or moves down to d times its initial value. The notation for the value of the option is shown on the tree. (For example, after two up movements the value of the option is f_{uu} .) We suppose that the risk-free interest rate is r and the length of the time step is Δt years.

Because the length of a time step is now Δt rather than T , equations (12.2) and (12.3) become

$$f = e^{-r\Delta t} [pf_u + (1-p)f_d] \quad (12.5)$$

$$p = \frac{e^{r\Delta t} - d}{u - d} \quad (12.6)$$

Repeated application of equation (12.5) gives

$$f_u = e^{-r\Delta t} [pf_{uu} + (1-p)f_{ud}] \quad (12.7)$$


$$f_d = e^{-r\Delta t} [pf_{ud} + (1-p)f_{dd}] \quad (12.8)$$

$$f = e^{-r\Delta t} [pf_u + (1-p)f_d] \quad (12.9)$$

Substituting from equations (12.7) and (12.8) into (12.9), we get

$$f = e^{-2r\Delta t} [p^2 f_{uu} + 2p(1-p)f_{ud} + (1-p)^2 f_{dd}] \quad (12.10)$$

This is consistent with the principle of risk-neutral valuation mentioned earlier. The

Figure 12.6 Stock and option prices in general two-step tree.

variables p^2 , $2p(1-p)$, and $(1-p)^2$ are the probabilities that the upper, middle, and lower final nodes will be reached. The option price is equal to its expected payoff in a risk-neutral world discounted at the risk-free interest rate.

As we add more steps to the binomial tree, the risk-neutral valuation principle continues to hold. The option price is always equal to its expected payoff in a risk-neutral world discounted at the risk-free interest rate.

12.4 A PUT EXAMPLE

The procedures described in this chapter can be used to price puts as well as calls. Consider a 2-year European put with a strike price of \$52 on a stock whose current price is \$50. We suppose that there are two time steps of 1 year, and in each time step the stock price either moves up by 20% or moves down by 20%. We also suppose that the risk-free interest rate is 5%.

The tree is shown in Figure 12.7. In this case $u = 1.2$, $d = 0.8$, $\Delta t = 1$, and $r = 0.05$. From equation (12.6) the value of the risk-neutral probability, p , is given by


$$p = \frac{e^{0.05 \times 1} - 0.8}{1.2 - 0.8} = 0.6282$$

The possible final stock prices are: \$72, \$48, and \$32. In this case, $f_{uu} = 0$, $f_{ud} = 4$, and $f_{dd} = 20$. From equation (12.10),

$$f = e^{-2 \times 0.05 \times 1} (0.6282^2 \times 0 + 2 \times 0.6282 \times 0.3718 \times 4 + 0.3718^2 \times 20) = 4.1923$$

The value of the put is \$4.1923. This result can also be obtained using equation (12.5)

Figure 12.7 Using a two-step tree to value a European put option. At each node, the upper number is the stock price and the lower number is the option price.



and working back through the tree one step at a time. Figure 12.7 shows the intermediate option prices that are calculated.

12.5 AMERICAN OPTIONS


Up to now all the options we have considered have been European. We now move on to consider how American options can be valued using a binomial tree such as that in Figure 12.4 or 12.7. The procedure is to work back through the tree from the end to the beginning, testing at each node to see whether early exercise is optimal. The value of the option at the final nodes is the same as for the European option. At earlier nodes the value of the option is the greater of

1. The value given by equation (12.5)
2. The payoff from early exercise.

Figure 12.8 shows how Figure 12.7 is affected if the option under consideration is American rather than European. The stock prices and their probabilities are unchanged. The values for the option at the final nodes are also unchanged. At node B, equation (12.5) gives the value of the option as 1.4147, whereas the payoff from early exercise is negative ($= -8$). Clearly early exercise is not optimal at node B, and the value of the option at this node is 1.4147. At node C, equation (12.5) gives the value of the option as 9.4636, whereas the payoff from early exercise is 12. In this case, early exercise is optimal and the value of the option at the node is 12. At the initial node A, the value given by equation (12.5) is

$$e^{-0.05 \times 1} (0.6282 \times 1.4147 + 0.3718 \times 12.0) = 5.0894$$

Figure 12.8 Using a two-step tree to value an American put option. At each node, the upper number is the stock price and the lower number is the option price.



and the payoff from early exercise is 2. In this case early exercise is not optimal. The value of the option is therefore \$5.0894.

12.6 DELTA

At this stage, it is appropriate to introduce *delta*, an important parameter (sometimes referred to as a “Greek letter” or simply a “Greek”) in the pricing and hedging of options.

The delta (Δ) of a stock option is the ratio of the change in the price of the stock option to the change in the price of the underlying stock. It is the number of units of the stock we should hold for each option shorted in order to create a riskless portfolio. It is the same as the Δ introduced earlier in this chapter. The construction of a riskless portfolio is sometimes referred to as *delta hedging*. The delta of a call option is positive, whereas the delta of a put option is negative.

From Figure 12.1, we can calculate the value of the delta of the call option being considered as

$$\frac{1 - 0}{22 - 18} = 0.25$$

This is because when the stock price changes from \$18 to \$22, the option price changes from \$0 to \$1. (This is also the value of Δ calculated in Section 12.1.)

In Figure 12.4 the delta corresponding to stock price movements over the first time step is

$$\frac{2.0257 - 0}{22 - 18} = 0.5064$$

The delta for stock price movements over the second time step is

$$\frac{3.2 - 0}{24.2 - 19.8} = 0.7273$$

if there is an upward movement over the first time step, and

$$\frac{0 - 0}{19.8 - 16.2} = 0$$

if there is a downward movement over the first time step.

From Figure 12.7, delta is

$$\frac{1.4147 - 9.4636}{60 - 40} = -0.4024$$

at the end of the first time step, and either

$$\frac{0 - 4}{72 - 48} = -0.1667 \quad \text{or} \quad \frac{4 - 20}{48 - 32} = -1.0000$$

at the end of the second time step.

The two-step examples show that delta changes over time. (In Figure 12.4, delta changes from 0.5064 to either 0.7273 or 0; and, in Figure 12.7, it changes from -0.4024 to either -0.1667 or -1.0000 .) Thus, in order to maintain a riskless hedge using an option and the underlying stock, we need to adjust our holdings in the stock periodically. We will return to this feature of options in Chapter 18.

12.7 MATCHING VOLATILITY WITH u AND d

In practice, when constructing a binomial tree to represent the movements in a stock price, we choose the parameters u and d to match the volatility of the stock price. A question that arises is whether we should match volatility in the real world or the risk-neutral world. As we will now show, this does not matter. For small Δt and particular values of u and d , the volatility being assumed is the same in both the real world and the risk-neutral world.

Figure 12.9a shows stock price movements over one step of a binomial tree in the real world and Figure 12.9b shows these movements in a risk-neutral world. The step is of length Δt . The stock price starts at S_0 and moves either up to S_0u or down to S_0d . These are the only two possible outcomes in both the real world and the risk-neutral world. The probability of an up movement in the real world is denoted by p^* and, consistent with our earlier notation, in the risk-neutral world this probability is p .


The expected stock price at the end of the first time step in the real world is $S_0e^{\mu\Delta t}$, where μ is the expected return. On the tree the expected stock price at this time is

$$p^*S_0u + (1 - p^*)S_0d$$

In order to match the expected return on the stock with the tree's parameters, we must therefore have

$$p^*S_0u + (1 - p^*)S_0d = S_0e^{\mu\Delta t}$$

Figure 12.9 Change in stock price in time Δt in (a) the real world and (b) the risk-neutral world.



or

$$p^* = \frac{e^{\mu\Delta t} - d}{u - d} \quad (12.11)$$

As we will explain in Chapter 14, the volatility σ of a stock price is defined so that $\sigma\sqrt{\Delta t}$ is the standard deviation of the return on the stock price in a short period of time of length Δt . Equivalently, the variance of the return is $\sigma^2\Delta t$. On the tree in Figure 12.9a, the variance of the stock price return is²

$$p^*u^2 + (1 - p^*)d^2 - [p^*u + (1 - p^*)d]^2$$

In order to match the stock price volatility with the tree's parameters, we must therefore have

$$p^*u^2 + (1 - p^*)d^2 - [p^*u + (1 - p^*)d]^2 = \sigma^2\Delta t \quad (12.12)$$

Substituting from equation (12.11) into equation (12.12) gives

$$e^{\mu\Delta t}(u + d) - ud - e^{2\mu\Delta t} = \sigma^2\Delta t$$

When terms in Δt^2 and higher powers of Δt are ignored, one solution to this equation is³

$$u = e^{\sigma\sqrt{\Delta t}}$$

$$d = e^{-\sigma\sqrt{\Delta t}}$$

These are the values of u and d proposed by Cox, Ross, and Rubinstein (1979) for matching volatility.

² The return is either $u - 1$ or $d - 1$. Subtracting 1 from a variable makes no difference to its variance. The variance of the return is therefore the variance of a variable that has probability p^* of being u and probability $1 - p^*$ of being d . The variance of a variable X equals $E(X^2) - [E(X)]^2$, where E denotes expected value.

³ We are here using the series expansion

$$e^x = 1 + x + \frac{x^2}{2!} + \frac{x^3}{3!} + \dots$$

In Figure 12.9b, the expected stock price at the end of the time step is $S_0 e^{r\Delta t}$, as shown in equation (12.4). The variance of the stock price return in the risk-neutral world is

$$pu^2 + (1-p)d^2 - [pu + (1-p)d]^2 = [e^{r\Delta t}(u+d) - ud - e^{2r\Delta t}]$$

Substituting $u = e^{\sigma\sqrt{\Delta t}}$ and $d = e^{-\sigma\sqrt{\Delta t}}$, we find this equals $\sigma^2 \Delta t$ when terms in Δt^2 and higher powers of Δt are ignored.

This analysis shows that when we move from the real world to the risk-neutral world the expected return on the stock changes, but its volatility remains the same (at least in the limit as Δt tends to zero). This is an illustration of an important general result known as *Girsanov's theorem*. When we move from a world with one set of risk preferences to a world with another set of risk preferences, the expected growth rates in variables change, but their volatilities remain the same. We will examine the impact of risk preferences on the behavior of market variables in more detail in Chapter 27. Moving from one set of risk preferences to another is sometimes referred to as *changing the measure*. The real-world measure is sometimes referred to as the *P-measure*, while the risk-neutral world measure is referred to as the *Q-measure*.⁴

12.8 THE BINOMIAL TREE FORMULAS

The analysis in the previous section shows that, when the length of the time step on a binomial tree is Δt , we should match volatility by setting

$$u = e^{\sigma\sqrt{\Delta t}} \quad (12.13)$$

and

$$d = e^{-\sigma\sqrt{\Delta t}} \quad (12.14)$$

Also, from equation (12.6),

$$p = \frac{a-d}{u-d} \quad (12.15)$$

where

$$a = e^{r\Delta t} \quad (12.16)$$

Equations (12.13) to (12.16) define the tree.

Consider again the American put option in Figure 12.8, where the stock price is \$50, the strike price is \$52, the risk-free rate is 5%, the life of the option is 2 years, and there are two time steps. In this case, $\Delta t = 1$. Suppose that the volatility σ is 30%. Then, from equations (12.13) to (12.16), we have

$$u = e^{0.3 \times 1} = 1.3499, \quad d = \frac{1}{1.3499} = 0.7408, \quad a = e^{0.05 \times 1} = 1.0513$$


and

$$p = \frac{1.053 - 0.7408}{1.3499 - 0.7408} = 0.5097$$

The tree is shown in Figure 12.10. The value of the put option is 7.43. (This is different

⁴ With the notation we have been using, p is the probability under the Q-measure, while p^* is the probability under the P-measure.

Figure 12.10 Two-step tree to value a 2-year American put option when the stock price is 50, strike price is 52, risk-free rate is 5%, and volatility is 30%.



from the value obtained in Figure 12.8 by assuming $u = 1.2$ and $d = 0.8$.) Note that the option is exercised at the end of the first time step if the lower node is reached.

12.9 INCREASING THE NUMBER OF STEPS

The binomial model presented above is unrealistically simple. Clearly, an analyst can expect to obtain only a very rough approximation to an option price by assuming that stock price movements during the life of the option consist of one or two binomial steps.

When binomial trees are used in practice, the life of the option is typically divided into 30 or more time steps. In each time step there is a binomial stock price movement. With 30 time steps there are 31 terminal stock prices and 2^{30} , or about 1 billion, possible stock price paths are implicitly considered.

The equations defining the tree are equations (12.13) to (12.16), regardless of the number of time steps. Suppose, for example, that there are five steps instead of two in the example we considered in Figure 12.10. The parameters would be $\Delta t = 2/5 = 0.4$, $r = 0.05$, and $\sigma = 0.3$. These values give $u = e^{0.3 \times \sqrt{0.4}} = 1.2089$, $d = 1/1.2089 = 0.8272$, $a = e^{0.05 \times 0.4} = 1.0202$, and $p = (1.0202 - 0.8272)/(1.2089 - 0.8272) = 0.5056$.

As the number of time steps is increased (so that Δt becomes smaller), the binomial tree model makes the same assumptions about stock price behavior as the Black–Scholes–Merton model, which will be presented in Chapter 14. When the binomial tree is used to price a European option, the price converges to the Black–Scholes–Merton price, as expected, as the number of time steps is increased. This is proved in the appendix to this chapter.

12.10 USING DerivaGem

The software accompanying this book, DerivaGem, is a useful tool for becoming comfortable with binomial trees. After loading the software in the way described at the end of this book, go to the Equity_FX_Index_Futures_Options worksheet. Choose Equity as the Underlying Type and select Binomial American as the Option Type. Enter the stock price, volatility, risk-free rate, time to expiration, exercise price, and tree steps, as 50, 30%, 5%, 2, 52, and 2, respectively. Click on the *Put* button and then on *Calculate*. The price of the option is shown as 7.428 in the box labeled Price. Now click on *Display Tree* and you will see the equivalent of Figure 12.10. (The red numbers in the software indicate the nodes where the option is exercised.)

Return to the Equity_FX_Index_Futures_Options worksheet and change the number of time steps to 5. Hit *Enter* and click on *Calculate*. You will find that the value of the option changes to 7.671. By clicking on *Display Tree* the five-step tree is displayed, together with the values of u , d , a , and p calculated above.

DerivaGem can display trees that have up to 10 steps, but the calculations can be done for up to 500 steps. In our example, 500 steps gives the option price (to two decimal places) as 7.47. This is an accurate answer. By changing the Option Type to Binomial European, we can use the tree to value a European option. Using 500 time steps, the value of a European option with the same parameters as the American option is 6.76. (By changing the Option Type to Black–Scholes European, we can display the value of the option using the Black–Scholes–Merton formula that will be presented in Chapter 14. This is also 6.76.)

By changing the Underlying Type, we can consider options on assets other than stocks. These will now be discussed.

12.11 OPTIONS ON OTHER ASSETS

We introduced options on indices, currencies, and futures contracts in Chapter 9 and will cover them in more detail in Chapters 16 and 17. It turns out that we can construct and use binomial trees for these options in exactly the same way as for options on stocks except that the equations for p change. As in the case of options on stocks, equation (12.2) applies so that the value at a node (before the possibility of early exercise is considered) is p times the value if there is an up movement plus $1 - p$ times the value if there is a down movement, discounted at the risk-free rate.

Options on Stocks Paying a Continuous Dividend Yield

Consider a stock paying a known dividend yield at rate q . The total return from dividends and capital gains in a risk-neutral world is r . The dividends provide a return of q . Capital gains must therefore provide a return of $r - q$. If the stock starts at S_0 , its expected value after one time step of length Δt must be $S_0 e^{(r-q)\Delta t}$. This means that

$$pS_0u + (1 - p)S_0d = S_0e^{(r-q)\Delta t}$$

so that

$$p = \frac{e^{(r-q)\Delta t} - d}{u - d}$$

As in the case of options on non-dividend-paying stocks, we match volatility by setting $u = e^{\sigma\sqrt{\Delta t}}$ and $d = 1/u$. This means that we can use equations (12.13) to (12.16), except that we set $a = e^{(r-q)\Delta t}$ instead of $a = e^{r\Delta t}$.

Options on Stock Indices

When calculating a futures price for a stock index in Chapter 5 we assumed that the stocks underlying the index provided a dividend yield at rate q . We make a similar assumption here. The valuation of an option on a stock index is therefore very similar to the valuation of an option on a stock paying a known dividend yield.

Example 12.1

A stock index is currently 810 and has a volatility of 20% and a dividend yield of 2%. The risk-free rate is 5%. Figure 12.11 shows the output from DerivaGem for valuing a European 6-month call option with a strike price of 800 using a two-step tree. In this case,


$$\begin{aligned}\Delta t &= 0.25, & u &= e^{0.20 \times \sqrt{0.25}} = 1.1052, \\ d &= 1/u = 0.9048, & a &= e^{(0.05 - 0.02) \times 0.25} = 1.0075 \\ p &= (1.0075 - 0.9048) / (1.1052 - 0.9048) = 0.5126\end{aligned}$$

The value of the option is 53.39.

Figure 12.11 Two-step tree to value a European 6-month call option on an index when the index level is 810, strike price is 800, risk-free rate is 5%, volatility is 20%, and dividend yield is 2% (DerivaGem output).

At each node:
Upper value = Underlying Asset Price
Lower value = Option Price
Shading indicates where option is exercised

Strike price = 800
Discount factor per step = 0.9876
Time step, $dt = 0.2500$ years, 91.25 days
Growth factor per step, $a = 1.0075$
Probability of up move, $p = 0.5126$
Up step size, $u = 1.1052$
Down step size, $d = 0.9048$



Node Time:
0.0000 0.2500 0.5000

Options on Currencies

As pointed out in Section 5.10, a foreign currency can be regarded as an asset providing a yield at the foreign risk-free rate of interest, r_f . By analogy with the stock index case we can construct a tree for options on a currency by using equations (12.13) to (12.16) and setting $a = e^{(r-r_f)\Delta t}$.

Example 12.2

The Australian dollar is currently worth 0.6100 US dollars and this exchange rate has a volatility of 12%. The Australian risk-free rate is 7% and the US risk-free rate is 5%. Figure 12.12 shows the output from DerivaGem for valuing a 3-month American call option with a strike price of 0.6000 using a three-step tree. In this case,


$$\begin{aligned} \Delta t &= 0.08333, & u &= e^{0.12 \times \sqrt{0.08333}} = 1.0352 \\ d &= 1/u = 0.9660, & a &= e^{(0.05 - 0.07) \times 0.08333} = 0.9983 \\ p &= (0.9983 - 0.9660) / (1.0352 - 0.9660) = 0.4673 \end{aligned}$$

The value of the option is 0.019.

Figure 12.12 Three-step tree to value an American 3-month call option on a currency when the value of the currency is 0.6100, strike price is 0.6000, risk-free rate is 5%, volatility is 12%, and foreign risk-free rate is 7% (DerivaGem output).

At each node:
 Upper value = Underlying Asset Price
 Lower value = Option Price
 Shading indicates where option is exercised

Strike price = 0.6
 Discount factor per step = 0.9958
 Time step, dt = 0.0833 years, 30.42 days
 Growth factor per step, a = 0.9983
 Probability of up move, p = 0.4673
 Up step size, u = 1.0352
 Down step size, d = 0.9660



Options on Futures

It costs nothing to take a long or a short position in a futures contract. It follows that in a risk-neutral world a futures price should have an expected growth rate of zero. (We discuss this point in more detail in Section 17.7.) As above, we define p as the probability of an up movement in the futures price, u as the percentage up movement, and d as the percentage down movement. If F_0 is the initial futures price, the expected futures price at the end of one time step of length Δt should also be F_0 . This means that

$$pF_0u + (1 - p)F_0d = F_0$$

so that

$$p = \frac{1 - d}{u - d}$$

and we can use equations (12.13) to (12.16) with $a = 1$.

Example 12.3

A futures price is currently 31 and has a volatility of 30%. The risk-free rate is 5%. Figure 12.13 shows the output from DerivaGem for valuing a 9-month American put option with a strike price of 30 using a three-step tree. In this case,

$$\begin{aligned}\Delta t &= 0.25, & u &= e^{0.3\sqrt{0.25}} = 1.1618 \\ d &= 1/u = 1/1.1618 = 0.8607, & a &= 1, \\ p &= (1 - 0.8607)/(1.1618 - 0.8607) = 0.4626\end{aligned}$$

The value of the option is 2.84.

SUMMARY

This chapter has provided a first look at the valuation of options on stocks and other assets using trees. In the simple situation where movements in the price of a stock during the life of an option are governed by a one-step binomial tree, it is possible to set up a riskless portfolio consisting of a position in the stock option and a position in the stock. In a world with no arbitrage opportunities, riskless portfolios must earn the risk-free interest. This enables the stock option to be priced in terms of the stock. It is interesting to note that no assumptions are required about the probabilities of up and down movements in the stock price at each node of the tree.

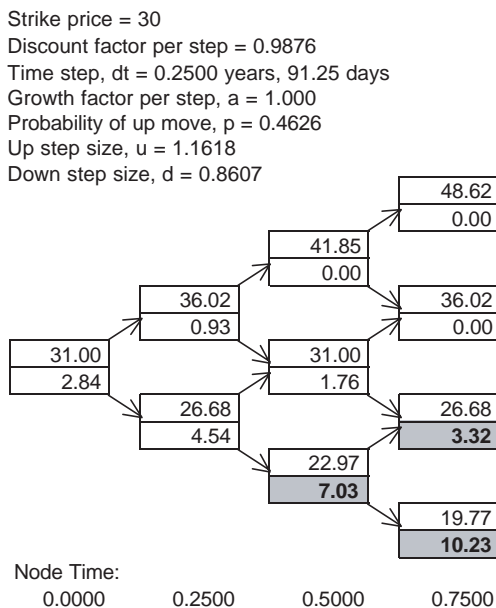
When stock price movements are governed by a multistep binomial tree, we can treat each binomial step separately and work back from the end of the life of the option to the beginning to obtain the current value of the option. Again only no-arbitrage arguments are used, and no assumptions are required about the probabilities of up and down movements in the stock price at each node.

A very important principle states that we can assume the world is risk-neutral when valuing an option. This chapter has shown, through both numerical examples and algebra, that no-arbitrage arguments and risk-neutral valuation are equivalent and lead to the same option prices.

The delta of a stock option, Δ , considers the effect of a small change in the underlying stock price on the change in the option price. It is the ratio of the change in the option

Figure 12.13 Three-step tree to value an American 9-month put option on a futures contract when the futures price is 31, strike price is 30, risk-free rate is 5%, and volatility is 30% (DerivaGem output).

At each node:
 Upper value = Underlying Asset Price
 Lower value = Option Price
 Shading indicates where option is exercised



price to the change in the stock price. For a riskless position, an investor should buy Δ shares for each option sold. An inspection of a typical binomial tree shows that delta changes during the life of an option. This means that to hedge a particular option position, we must change our holding in the underlying stock periodically.

Constructing binomial trees for valuing options on stock indices, currencies, and futures contracts is very similar to doing so for valuing options on stocks. In Chapter 20, we will return to binomial trees and provide more details on how they are used in practice.

FURTHER READING

Coval, J.E. and T. Shumway. "Expected Option Returns," *Journal of Finance*, 56, 3 (2001): 983-1009.

Cox, J. C., S. A. Ross, and M. Rubinstein. "Option Pricing: A Simplified Approach," *Journal of Financial Economics* 7 (October 1979): 229-64.

Rendleman, R., and B. Bartter. "Two State Option Pricing," *Journal of Finance* 34 (1979): 1092-1110.

Shreve, S.E. *Stochastic Calculus for Finance I: The Binomial Asset Pricing Model*. New York: Springer, 2004.

Practice Questions (Answers in Solutions Manual)

- 12.1. A stock price is currently \$40. It is known that at the end of 1 month it will be either \$42 or \$38. The risk-free interest rate is 8% per annum with continuous compounding. What is the value of a 1-month European call option with a strike price of \$39?
- 12.2. Explain the no-arbitrage and risk-neutral valuation approaches to valuing a European option using a one-step binomial tree.
- 12.3. What is meant by the “delta” of a stock option?
- 12.4. A stock price is currently \$50. It is known that at the end of 6 months it will be either \$45 or \$55. The risk-free interest rate is 10% per annum with continuous compounding. What is the value of a 6-month European put option with a strike price of \$50?
- 12.5. A stock price is currently \$100. Over each of the next two 6-month periods it is expected to go up by 10% or down by 10%. The risk-free interest rate is 8% per annum with continuous compounding. What is the value of a 1-year European call option with a strike price of \$100?
- 12.6. For the situation considered in Problem 12.5, what is the value of a 1-year European put option with a strike price of \$100? Verify that the European call and European put prices satisfy put–call parity.
- 12.7. What are the formulas for u and d in terms of volatility?
- 12.8. Consider the situation in which stock price movements during the life of a European option are governed by a two-step binomial tree. Explain why it is not possible to set up a position in the stock and the option that remains riskless for the whole of the life of the option.
- 12.9. A stock price is currently \$50. It is known that at the end of 2 months it will be either \$53 or \$48. The risk-free interest rate is 10% per annum with continuous compounding. What is the value of a 2-month European call option with a strike price of \$49? Use no-arbitrage arguments.
- 12.10. A stock price is currently \$80. It is known that at the end of 4 months it will be either \$75 or \$85. The risk-free interest rate is 5% per annum with continuous compounding. What is the value of a 4-month European put option with a strike price of \$80? Use no-arbitrage arguments.
- 12.11. A stock price is currently \$40. It is known that at the end of 3 months it will be either \$45 or \$35. The risk-free rate of interest with quarterly compounding is 8% per annum. Calculate the value of a 3-month European put option on the stock with an exercise price of \$40. Verify that no-arbitrage arguments and risk-neutral valuation arguments give the same answers.
- 12.12. A stock price is currently \$50. Over each of the next two 3-month periods it is expected to go up by 6% or down by 5%. The risk-free interest rate is 5% per annum with continuous compounding. What is the value of a 6-month European call option with a strike price of \$51?
- 12.13. For the situation considered in Problem 12.12, what is the value of a 6-month European put option with a strike price of \$51? Verify that the European call and European put prices satisfy put–call parity. If the put option were American, would it ever be optimal to exercise it early at any of the nodes on the tree?

- 12.14. A stock price is currently \$25. It is known that at the end of 2 months it will be either \$23 or \$27. The risk-free interest rate is 10% per annum with continuous compounding. Suppose S_T is the stock price at the end of 2 months. What is the value of a derivative that pays off S_T^2 at this time?
- 12.15. Calculate u , d , and p when a binomial tree is constructed to value an option on a foreign currency. The tree step size is 1 month, the domestic interest rate is 5% per annum, the foreign interest rate is 8% per annum, and the volatility is 12% per annum.

Further Questions

- 12.16. A stock price is currently \$50. It is known that at the end of 6 months it will be either \$60 or \$42. The risk-free rate of interest with continuous compounding is 12% per annum. Calculate the value of a 6-month European call option on the stock with an exercise price of \$48. Verify that no-arbitrage arguments and risk-neutral valuation arguments give the same answers.
- 12.17. A stock price is currently \$40. Over each of the next two 3-month periods it is expected to go up by 10% or down by 10%. The risk-free interest rate is 12% per annum with continuous compounding.
- What is the value of a 6-month European put option with a strike price of \$42?
 - What is the value of a 6-month American put option with a strike price of \$42?
- 12.18. Using a “trial-and-error” approach, estimate how high the strike price has to be in Problem 11.17 for it to be optimal to exercise the option immediately.
- 12.19. A stock price is currently \$30. During each 2-month period for the next 4 months it will increase by 8% or reduce by 10%. The risk-free interest rate is 5%. Use a two-step tree to calculate the value of a derivative that pays off $[\max(30 - S_T, 0)]^2$, where S_T is the stock price in 4 months. If the derivative is American-style, should it be exercised early?
- 12.20. Consider a European call option on a non-dividend-paying stock where the stock price is \$40, the strike price is \$40, the risk-free rate is 4% per annum, the volatility is 30% per annum, and the time to maturity is 6 months.
- Calculate u , d , and p for a two-step tree.
 - Value the option using a two-step tree.
 - Verify that DerivaGem gives the same answer.
 - Use DerivaGem to value the option with 5, 50, 100, and 500 time steps.
- 12.21. Repeat Problem 12.20 for an American put option on a futures contract. The strike price and the futures price are \$50, the risk-free rate is 10%, the time to maturity is 6 months, and the volatility is 40% per annum.
- 12.22. Footnote 1 shows that the correct discount rate to use for the real-world expected payoff in the case of the call option considered in Figure 12.1 is 42.6%. Show that if the option is a put rather than a call the discount rate is -52.5% . Explain why the two real-world discount rates are so different.

APPENDIX

DERIVATION OF THE BLACK-SCHOLES-MERTON OPTION-PRICING FORMULA FROM A BINOMIAL TREE

One way of deriving the famous Black-Scholes-Merton result for valuing a European option on a non-dividend-paying stock is by allowing the number of time steps in a binomial tree to approach infinity.

Suppose that a tree with n time steps is used to value a European call option with strike price K and life T . Each step is of length T/n . If there have been j upward movements and $n - j$ downward movements on the tree, the final stock price is $S_0 u^j d^{n-j}$, where u is the proportional up movement, d is the proportional down movement, and S_0 is the initial stock price. The payoff from a European call option is then

$$\max(S_0 u^j d^{n-j} - K, 0)$$

From the properties of the binomial distribution, the probability of exactly j upward and $n - j$ downward movements is given by

$$\frac{n!}{(n-j)! j!} p^j (1-p)^{n-j}$$

It follows that the expected payoff from the call option is

$$\sum_{j=0}^n \frac{n!}{(n-j)! j!} p^j (1-p)^{n-j} \max(S_0 u^j d^{n-j} - K, 0)$$

As the tree represents movements in a risk-neutral world, we can discount this at the risk-free rate r to obtain the option price:

$$c = e^{-rT} \sum_{j=0}^n \frac{n!}{(n-j)! j!} p^j (1-p)^{n-j} \max(S_0 u^j d^{n-j} - K, 0) \quad (12A.1)$$

The terms in equation (12A.1) are nonzero when the final stock price is greater than the strike price, that is, when

$$S_0 u^j d^{n-j} > K$$

or

$$\ln(S_0/K) > -j \ln(u) - (n-j) \ln(d)$$

Since $u = e^{\sigma\sqrt{T/n}}$ and $d = e^{-\sigma\sqrt{T/n}}$, this condition becomes

$$\ln(S_0/K) > n\sigma\sqrt{T/n} - 2j\sigma\sqrt{T/n}$$

or

$$j > \frac{n}{2} - \frac{\ln(S_0/K)}{2\sigma\sqrt{T/n}}$$

Equation (12A.1) can therefore be written

$$c = e^{-rT} \sum_{j > \alpha} \frac{n!}{(n-j)! j!} p^j (1-p)^{n-j} (S_0 u^j d^{n-j} - K)$$

where

$$\alpha = \frac{n}{2} - \frac{\ln(S_0/K)}{2\sigma\sqrt{T/n}}$$

For convenience, we define

$$U_1 = \sum_{j>\alpha} \frac{n!}{(n-j)! j!} p^j (1-p)^{n-j} u^j d^{n-j} \quad (12A.2)$$

and

$$U_2 = \sum_{j>\alpha} \frac{n!}{(n-j)! j!} p^j (1-p)^{n-j} \quad (12A.3)$$

so that

$$c = e^{-rT} (S_0 U_1 - K U_2) \quad (12A.4)$$

Consider first U_2 . As is well known, the binomial distribution approaches a normal distribution as the number of trials approaches infinity. Specifically, when there are n trials and p is the probability of success, the probability distribution of the number of successes is approximately normal with mean np and standard deviation $\sqrt{np(1-p)}$. The variable U_2 in equation (12A.3) is the probability of the number of successes being more than α . From the properties of the normal distribution, it follows that, for large n ,

$$U_2 = N\left(\frac{np - \alpha}{\sqrt{np(1-p)}}\right) \quad (12A.5)$$

where N is the cumulative normal distribution function. Substituting for α , we obtain

$$U_2 = N\left(\frac{\ln(S_0/K)}{2\sigma\sqrt{T}\sqrt{p(1-p)}} + \frac{\sqrt{n}(p - \frac{1}{2})}{\sqrt{p(1-p)}}\right) \quad (12A.6)$$

From equations (12.13) to (12.15), we have

$$p = \frac{e^{rT/n} - e^{-\sigma\sqrt{T/n}}}{e^{\sigma\sqrt{T/n}} - e^{-\sigma\sqrt{T/n}}}$$

By expanding the exponential functions in a series, we see that, as n tends to infinity, $p(1-p)$ tends to $\frac{1}{4}$ and $\sqrt{n}(p - \frac{1}{2})$ tends to

$$\frac{(r - \sigma^2/2)\sqrt{T}}{2\sigma}$$

so that in the limit, as n tends to infinity, equation (12A.6) becomes

$$U_2 = N\left(\frac{\ln(S_0/K) + (r - \sigma^2/2)T}{\sigma\sqrt{T}}\right) \quad (12A.7)$$

We now move on to evaluate U_1 . From equation (12A.2), we have

$$U_1 = \sum_{j>\alpha} \frac{n!}{(n-j)! j!} (pu)^j [(1-p)d]^{n-j} \quad (12A.8)$$

Define

$$p^* = \frac{pu}{pu + (1-p)d} \quad (12A.9)$$

It then follows that

$$1 - p^* = \frac{(1-p)d}{pu + (1-p)d}$$

and we can write equation (12A.8) as

$$U_1 = [pu + (1-p)d]^n \sum_{j>\alpha} \frac{n!}{(n-j)! j!} (p^*)^j (1-p^*)^{n-j}$$

Since the expected return in the risk-neutral world is the risk-free rate r , it follows that $pu + (1-p)d = e^{rT/n}$ and

$$U_1 = e^{rT} \sum_{j>\alpha} \frac{n!}{(n-j)! j!} (p^*)^j (1-p^*)^{n-j}$$

This shows that U_1 involves a binomial distribution where the probability of an up movement is p^* rather than p . Approximating the binomial distribution with a normal distribution, we obtain, similarly to equation (12A.5),

$$U_1 = e^{rT} N\left(\frac{np^* - \alpha}{\sqrt{np^*(1-p^*)}}\right)$$

and substituting for α gives, as with equation (12A.6),

$$U_2 = e^{rT} N\left(\frac{\ln(S_0/K)}{2\sigma\sqrt{T}\sqrt{p^*(1-p^*)}} + \frac{\sqrt{n}(p^* - \frac{1}{2})}{\sqrt{p^*(1-p^*)}}\right)$$

Substituting for u and d in equation (12A.9) gives

$$p^* = \left(\frac{e^{rT/n} - e^{-\sigma\sqrt{T/n}}}{e^{\sigma\sqrt{T/n}} - e^{-\sigma\sqrt{T/n}}}\right) \left(\frac{e^{\sigma\sqrt{T/n}}}{e^{rT/n}}\right)$$

By expanding the exponential functions in a series we see that, as n tends to infinity, $p^*(1-p^*)$ tends to $\frac{1}{4}$ and $\sqrt{n}(p^* - \frac{1}{2})$ tends to

$$\frac{(r + \sigma^2/2)\sqrt{T}}{2\sigma}$$

with the result that

$$U_1 = e^{rT} N\left(\frac{\ln(S_0/K) + (r + \sigma^2/2)T}{\sigma\sqrt{T}}\right) \quad (12A.10)$$

From equations (12A.4), (12A.7), and (12A.10), we have

$$c = S_0 N(d_1) - Ke^{-rT} N(d_2)$$

where

$$d_1 = \frac{\ln(S_0/K) + (r + \sigma^2/2)T}{\sigma\sqrt{T}}$$

and

$$d_2 = \frac{\ln(S_0/K) + (r - \sigma^2/2)T}{\sigma\sqrt{T}} = d_1 - \sigma\sqrt{T}$$

This is the Black–Scholes–Merton formula for the valuation of a European call option. It will be discussed in Chapter 14. An alternative derivation is given in the appendix to that chapter.

CHAPTER 13

Wiener Processes and Itô's Lemma



Any variable whose value changes over time in an uncertain way is said to follow a *stochastic process*. Stochastic processes can be classified as *discrete time* or *continuous time*. A discrete-time stochastic process is one where the value of the variable can change only at certain fixed points in time, whereas a continuous-time stochastic process is one where changes can take place at any time. Stochastic processes can also be classified as *continuous variable* or *discrete variable*. In a continuous-variable process, the underlying variable can take any value within a certain range, whereas in a discrete-variable process, only certain discrete values are possible.

This chapter develops a continuous-variable, continuous-time stochastic process for stock prices. Learning about this process is the first step to understanding the pricing of options and other more complicated derivatives. It should be noted that, in practice, we do not observe stock prices following continuous-variable, continuous-time processes. Stock prices are restricted to discrete values (e.g., multiples of a cent) and changes can be observed only when the exchange is open for trading. Nevertheless, the continuous-variable, continuous-time process proves to be a useful model for many purposes.

Many people feel that continuous-time stochastic processes are so complicated that they should be left entirely to “rocket scientists.” This is not so. The biggest hurdle to understanding these processes is the notation. Here we present a step-by-step approach aimed at getting the reader over this hurdle. We also explain an important result known as *Itô's lemma* that is central to the pricing of derivatives.

13.1 THE MARKOV PROPERTY

A *Markov process* is a particular type of stochastic process where only the current value of a variable is relevant for predicting the future. The past history of the variable and the way that the present has emerged from the past are irrelevant.

Stock prices are usually assumed to follow a Markov process. Suppose that the price of IBM stock is \$100 now. If the stock price follows a Markov process, our predictions for the future should be unaffected by the price one week ago, one month

ago, or one year ago. The only relevant piece of information is that the price is now \$100.¹ Predictions for the future are uncertain and must be expressed in terms of probability distributions. The Markov property implies that the probability distribution of the price at any particular future time is not dependent on the particular path followed by the price in the past.

The Markov property of stock prices is consistent with the weak form of market efficiency. This states that the present price of a stock impounds all the information contained in a record of past prices. If the weak form of market efficiency were not true, technical analysts could make above-average returns by interpreting charts of the past history of stock prices. There is very little evidence that they are in fact able to do this.

It is competition in the marketplace that tends to ensure that weak-form market efficiency holds. There are many investors watching the stock market closely. Trying to make a profit from it leads to a situation where a stock price, at any given time, reflects the information in past prices. Suppose that it was discovered that a particular pattern in stock prices always gave a 65% chance of subsequent steep price rises. Investors would attempt to buy a stock as soon as the pattern was observed, and demand for the stock would immediately rise. This would lead to an immediate rise in its price and the observed effect would be eliminated, as would any profitable trading opportunities.

13.2 CONTINUOUS-TIME STOCHASTIC PROCESSES

Consider a variable that follows a Markov stochastic process. Suppose that its current value is 10 and that the change in its value during a year is $\phi(0, 1)$, where $\phi(m, v)$ denotes a probability distribution that is normally distributed with mean m and variance v .² What is the probability distribution of the change in the value of the variable during 2 years?

The change in 2 years is the sum of two normal distributions, each of which has a mean of zero and variance of 1.0. Because the variable is Markov, the two probability distributions are independent. When we add two independent normal distributions, the result is a normal distribution where the mean is the sum of the means and the variance is the sum of the variances. The mean of the change during 2 years in the variable we are considering is, therefore, zero and the variance of this change is 2.0. Hence, the change in the variable over 2 years has the distribution $\phi(0, 2)$. The standard deviation of the distribution is $\sqrt{2}$.

Consider next the change in the variable during 6 months. The variance of the change in the value of the variable during 1 year equals the variance of the change during the first 6 months plus the variance of the change during the second 6 months. We assume these are the same. It follows that the variance of the change during a 6-month period must be 0.5. Equivalently, the standard deviation of the change is $\sqrt{0.5}$. The probability distribution for the change in the value of the variable during 6 months is $\phi(0, 0.5)$.

¹ Statistical properties of the stock price history of IBM may be useful in determining the characteristics of the stochastic process followed by the stock price (e.g., its volatility). The point being made here is that the particular path followed by the stock in the past is irrelevant.

² Variance is the square of standard deviation. The variance of a 1-year change in the value of the variable we are considering is therefore 1.0.

A similar argument shows that the probability distribution for the change in the value of the variable during 3 months is $\phi(0, 0.25)$. More generally, the change during any time period of length T is $\phi(0, T)$. In particular, the change during a very short time period of length Δt is $\phi(0, \Delta t)$.

Note that, when Markov processes are considered, the variances of the changes in successive time periods are additive. The standard deviations of the changes in successive time periods are not additive. The variance of the change in the variable in our example is 1.0 per year, so that the variance of the change in 2 years is 2.0 and the variance of the change in 3 years is 3.0. The standard deviations of the changes in 2 and 3 years are $\sqrt{2}$ and $\sqrt{3}$, respectively. Strictly speaking, we should not refer to the standard deviation of the variable as 1.0 per year. The results explain why uncertainty is sometimes referred to as being proportional to the square root of time.

Wiener Processes

The process followed by the variable we have been considering is known as a *Wiener process*. It is a particular type of Markov stochastic process with a mean change of zero and a variance rate of 1.0 per year. It has been used in physics to describe the motion of a particle that is subject to a large number of small molecular shocks and is sometimes referred to as *Brownian motion*.

Expressed formally, a variable z follows a Wiener process if it has the following two properties:

- P 1. The change Δz during a small period of time Δt is

$$\Delta z = \epsilon \sqrt{\Delta t} \quad (13.1)$$

where ϵ has a standardized normal distribution $\phi(0, 1)$.

- P 2. The values of Δz for any two different short intervals of time, Δt , are independent.

It follows from the first property that Δz itself has a normal distribution with

$$\begin{aligned} \text{mean of } \Delta z &= 0 \\ \text{standard deviation of } \Delta z &= \sqrt{\Delta t} \\ \text{variance of } \Delta z &= \Delta t \end{aligned}$$

The second property implies that z follows a Markov process.

Consider the change in the value of z during a relatively long period of time, T . This can be denoted by $z(T) - z(0)$. It can be regarded as the sum of the changes in z in N small time intervals of length Δt , where


$$N = \frac{T}{\Delta t}$$

Thus,


$$z(T) - z(0) = \sum_{i=1}^N \epsilon_i \sqrt{\Delta t} \quad (13.2)$$

where the ϵ_i ($i = 1, 2, \dots, N$) are distributed $\phi(0, 1)$. We know from the second property of Wiener processes that the ϵ_i are independent of each other. It follows


Figure 13.1 How a Wiener process is obtained when $\Delta t \rightarrow 0$ in equation (12.1).



Relatively large value of Δt



Smaller value of Δt



The true process obtained as $\Delta t \rightarrow 0$

from equation (13.2) that $z(T) - z(0)$ is normally distributed, with

$$\begin{aligned} \text{mean of } [z(T) - z(0)] &= 0 \\ \text{variance of } [z(T) - z(0)] &= N \Delta t = T \\ \text{standard deviation of } [z(T) - z(0)] &= \sqrt{T} \end{aligned}$$

This is consistent with the discussion earlier in this section.

Example 13.1

Suppose that the value, z , of a variable that follows a Wiener process is initially 25 and that time is measured in years. At the end of 1 year, the value of the variable is normally distributed with a mean of 25 and a standard deviation of 1.0. At the end of 5 years, it is normally distributed with a mean of 25 and a standard deviation of $\sqrt{5}$, or 2.236. Our uncertainty about the value of the variable at a certain time in the future, as measured by its standard deviation, increases as the square root of how far we are looking ahead.

In ordinary calculus, it is usual to proceed from small changes to the limit as the small changes become closer to zero. Thus, $dx = a dt$ is the notation used to indicate that $\Delta x = a \Delta t$ in the limit as $\Delta t \rightarrow 0$. We use similar notational conventions in stochastic calculus. So, when we refer to dz as a Wiener process, we mean that it has the properties for Δz given above in the limit as $\Delta t \rightarrow 0$.

Figure 13.1 illustrates what happens to the path followed by z as the limit $\Delta t \rightarrow 0$ is approached. Note that the path is quite “jagged.” This is because the standard deviation of the movement in z in time Δt equals $\sqrt{\Delta t}$ and, when Δt is small, $\sqrt{\Delta t}$ is much bigger than Δt . Two intriguing properties of Wiener processes, related to this $\sqrt{\Delta t}$ property, are as follows:

1. The expected length of the path followed by z in any time interval is infinite.
2. The expected number of times z equals any particular value in any time interval is infinite.

Generalized Wiener Process

The mean change per unit time for a stochastic process is known as the *drift rate* and the variance per unit time is known as the *variance rate*. The basic Wiener process, dz , that has been developed so far has a drift rate of zero and a variance rate of 1.0. The drift rate of zero means that the expected value of z at any future time is equal to its current value. The variance rate of 1.0 means that the variance of the change in z in a time interval of length T equals T . A *generalized Wiener process* for a variable x can be defined in terms of dx as

$$dx = a dt + b dz \tag{13.3}$$

where a and b are constants.

To understand equation (13.3), it is useful to consider the two components on the right-hand side separately. The $a dt$ term implies that x has an expected drift rate of a per unit of time. Without the $b dz$ term, the equation is $dx = a dt$, which implies that $dx/dt = a$. Integrating with respect to time, we get

$$x = x_0 + at$$

where x_0 is the value of x at time 0. In a period of time of length T , the variable x increases by an amount aT . The $b dz$ term on the right-hand side of equation (13.3) can be regarded as adding noise or variability to the path followed by x . The amount of this noise or variability is b times a Wiener process. A Wiener process has a variance rate per unit time of 1.0. It follows that b times a Wiener process has a variance rate per unit time of b^2 . In a small time interval Δt , the change Δx in the value of x is given by equations (13.1) and (13.3) as

$$\Delta x = a \Delta t + b\epsilon\sqrt{\Delta t}$$

where, as before, ϵ has a standard normal distribution. Thus Δx has a normal distribution with


$$\begin{aligned} \text{mean of } \Delta x &= a \Delta t \\ \text{standard deviation of } \Delta x &= b\sqrt{\Delta t} \\ \text{variance of } \Delta x &= b^2 \Delta t \end{aligned}$$

Similar arguments to those given for a Wiener process show that the change in the value of x in any time interval T is normally distributed with

$$\begin{aligned} \text{mean of change in } x &= aT \\ \text{standard deviation of change in } x &= b\sqrt{T} \\ \text{variance of change in } x &= b^2T \end{aligned}$$

To summarize, the generalized Wiener process given in equation (13.3) has an expected drift rate (i.e., average drift per unit of time) of a and a variance rate (i.e., variance per unit of time) of b^2 . It is illustrated in Figure 13.2.

Figure 13.2 Generalized Wiener process with $a = 0.3$ and $b = 1.5$.



Example 13.2

Consider the situation where the cash position of a company, measured in thousands of dollars, follows a generalized Wiener process with a drift of 20 per year and a variance rate of 900 per year. Initially, the cash position is 50. At the end of 1 year the cash position will have a normal distribution with a mean of 70 and a standard deviation of $\sqrt{900}$, or 30. At the end of 6 months it will have a normal distribution with a mean of 60 and a standard deviation of $30\sqrt{0.5} = 21.21$. Our uncertainty about the cash position at some time in the future, as measured by its standard deviation, increases as the square root of how far ahead we are looking. Note that the cash position can become negative. (We can interpret this as a situation where the company is borrowing funds.)

Itô Process

A further type of stochastic process, known as an *Itô process*, can be defined. This is a generalized Wiener process in which the parameters a and b are functions of the value of the underlying variable x and time t . An Itô process can be written algebraically as

$$dx = a(x, t) dt + b(x, t) dz \quad (13.4)$$

Both the expected drift rate and variance rate of an Itô process are liable to change over time. In a small time interval between t and $t + \Delta t$, the variable changes from x to $x + \Delta x$, where

$$\Delta x = a(x, t)\Delta t + b(x, t)\epsilon\sqrt{\Delta t}$$

This relationship involves a small approximation. It assumes that the drift and variance rate of x remain constant, equal to their values at time t , during the time interval between t and $t + \Delta t$.

Note that the process in equation (13.4) is Markov because the change in x at time t depends only on the value of x at time t , not on its history. A non-Markov process could be defined by letting a and b in equation (13.4) depend on values of x prior to time t .

13.3 THE PROCESS FOR A STOCK PRICE

In this section we discuss the stochastic process usually assumed for the price of a non-dividend-paying stock.

It is tempting to suggest that a stock price follows a generalized Wiener process; that is, that it has a constant expected drift rate and a constant variance rate. However, this model fails to capture a key aspect of stock prices. This is that the expected percentage return required by investors from a stock is independent of the stock's price. If investors require a 14% per annum expected return when the stock price is \$10, then, *ceteris paribus*, they will also require a 14% per annum expected return when it is \$50.

Clearly, the assumption of constant expected drift rate is inappropriate and needs to be replaced by the assumption that the expected return (i.e., expected drift divided by the stock price) is constant. If S is the stock price at time t , then the expected drift rate in S should be assumed to be μS for some constant parameter μ . This means that in a short interval of time, Δt , the expected increase in S is $\mu S \Delta t$. The parameter μ is the expected rate of return on the stock, expressed in decimal form.

If the coefficient of dz is zero, so that there is no uncertainty, then this model implies that

$$\Delta S = \mu S \Delta t$$

In the limit, as $\Delta t \rightarrow 0$,

$$dS = \mu S dt$$

or

$$\frac{dS}{S} = \mu dt$$

Integrating between time 0 and time T , we get

$$S_T = S_0 e^{\mu T} \quad (13.5)$$

where S_0 and S_T are the stock price at time 0 and time T . Equation (13.5) shows that, when there is no uncertainty, the stock price grows at a continuously compounded rate of μ per unit of time.

In practice, of course, there is uncertainty. A reasonable assumption is that the variability of the percentage return in a short period of time, Δt , is the same regardless of the stock price. In other words, an investor is just as uncertain of the percentage return when the stock price is \$50 as when it is \$10. This suggests that the standard deviation of the change in a short period of time Δt should be proportional to the stock price and leads to the model

$$dS = \mu S dt + \sigma S dz$$

or

$$\frac{dS}{S} = \mu dt + \sigma dz \quad (13.6)$$

Equation (13.6) is the most widely used model of stock price behavior. The variable μ is the stock's expected rate of return. The variable σ is the volatility of the stock price. The variable σ^2 is referred to as its variance rate. The model in equation (13.6) represents the stock price process in the real world. In a risk-neutral world, μ equals the risk-free rate r .

Discrete-Time Model

The model of stock price behavior we have developed is known as *geometric Brownian motion*. The discrete-time version of the model is

$$\frac{\Delta S}{S} = \mu \Delta t + \sigma \epsilon \sqrt{\Delta t} \quad (13.7)$$

or

$$\Delta S = \mu S \Delta t + \sigma S \epsilon \sqrt{\Delta t} \quad (13.8)$$

The variable ΔS is the change in the stock price S in a small time interval Δt , and as before ϵ has a standard normal distribution (i.e., a normal distribution with a mean of zero and standard deviation of 1.0). The parameter μ is the expected rate of return per unit of time from the stock. The parameter σ is the volatility of the stock price. In this chapter we will assume these parameters are constant.

The left-hand side of equation (13.7) is the return provided by the stock in a short period of time, Δt . The term $\mu \Delta t$ is the expected value of this return, and the

term $\sigma\epsilon\sqrt{\Delta t}$ is the stochastic component of the return. The variance of the stochastic component (and, therefore, of the whole return) is $\sigma^2\Delta t$. This is consistent with the definition of the volatility σ given in Section 12.7; that is, σ is such that $\sigma\sqrt{\Delta t}$ is the standard deviation of the return in a short time period Δt .

Equation (13.7) shows that $\Delta S/S$ is normally distributed with mean $\mu \Delta t$ and standard deviation $\sigma\sqrt{\Delta t}$. In other words,

$$\frac{\Delta S}{S} \sim \phi(\mu \Delta t, \sigma^2 \Delta t) \quad (13.9)$$

Example 13.3

Consider a stock that pays no dividends, has a volatility of 30% per annum, and provides an expected return of 15% per annum with continuous compounding. In this case, $\mu = 0.15$ and $\sigma = 0.30$. The process for the stock price is

$$\frac{dS}{S} = 0.15 dt + 0.30 dz$$

If S is the stock price at a particular time and ΔS is the increase in the stock price in the next small interval of time,

$$\frac{\Delta S}{S} = 0.15\Delta t + 0.30\epsilon\sqrt{\Delta t}$$

where ϵ has a standard normal distribution. Consider a time interval of 1 week, or 0.0192 year, so that $\Delta t = 0.0192$. Then

$$\frac{\Delta S}{S} = 0.15 \times 0.0192 + 0.30 \times \sqrt{0.0192} \epsilon$$

or

$$\Delta S = 0.00288S + 0.0416S\epsilon$$

Monte Carlo Simulation

A Monte Carlo simulation of a stochastic process is a procedure for sampling random outcomes for the process. We will use it as a way of developing some understanding of the nature of the stock price process in equation (13.6).

Consider the situation in Example 13.3 where the expected return from a stock is 15% per annum and the volatility is 30% per annum. The stock price change over 1 week was shown to be

$$\Delta S = 0.00288S + 0.0416S\epsilon \quad (13.10)$$

A path for the stock price over 10 weeks can be simulated by sampling repeatedly for ϵ from $\phi(0, 1)$ and substituting into equation (13.10). The expression `=RAND()` in Excel produces a random sample between 0 and 1. The inverse cumulative normal distribution is `NORMSINV`. The instruction to produce a random sample from a standard normal distribution in Excel is therefore `=NORMSINV(RAND())`. Table 13.1 shows one path for a stock price that was sampled in this way. The initial stock price is assumed to be \$100. For the first period, ϵ is sampled as 0.52. From equation (13.10), the change during the first time period is

$$\Delta S = 0.00288 \times 100 + 0.0416 \times 100 \times 0.52 = 2.45$$

Therefore, at the beginning of the second time period, the stock price is \$102.45. The

Table 13.1 Simulation of stock price when $\mu = 0.15$ and $\sigma = 0.30$ during 1-week periods.

| <i>Stock price at start of period</i> | <i>Random sample for ϵ</i> | <i>Change in stock price during period</i> |
|---|--|--|
| 100.00 | 0.52 | 2.45 |
| 102.45 | 1.44 | 6.43 |
| 108.88 | -0.86 | -3.58 |
| 105.30 | 1.46 | 6.70 |
| 112.00 | -0.69 | -2.89 |
| 109.11 | -0.74 | -3.04 |
| 106.06 | 0.21 | 1.23 |
| 107.30 | -1.10 | -4.60 |
| 102.69 | 0.73 | 3.41 |
| 106.11 | 1.16 | 5.43 |
| 111.54 | 2.56 | 12.20 |

value of ϵ sampled for the next period is 1.44. From equation (13.10), the change during the second time period is

$$\Delta S = 0.00288 \times 102.45 + 0.0416 \times 102.45 \times 1.44 = 6.43$$

So, at the beginning of the next period, the stock price is \$108.88, and so on.³ Note that, because the process we are simulating is Markov, the samples for ϵ should be independent of each other.

Table 13.1 assumes that stock prices are measured to the nearest cent. It is important to realize that the table shows only one possible pattern of stock price movements. Different random samples would lead to different price movements. Any small time interval Δt can be used in the simulation. In the limit as $\Delta t \rightarrow 0$, a perfect description of the stochastic process is obtained. The final stock price of 111.54 in Table 13.1 can be regarded as a random sample from the distribution of stock prices at the end of 10 weeks. By repeatedly simulating movements in the stock price, a complete probability distribution of the stock price at the end of this time is obtained. Monte Carlo simulation is discussed in more detail in Chapter 20.

13.4 THE PARAMETERS

The process for a stock price developed in this chapter involves two parameters, μ and σ . The parameter μ is the expected return (annualized) earned by an investor in a short period of time. Most investors require higher expected returns to induce them to take higher risks. It follows that the value of μ should depend on the risk of the return from the stock.⁴ It should also depend on the level of interest rates in the economy. The higher the level of interest rates, the higher the expected return required on any given stock.

³ In practice, it is more efficient to sample $\ln S$ rather than S , as will be discussed in Section 20.6.

⁴ More precisely, μ depends on that part of the risk that cannot be diversified away by the investor.

Fortunately, we do not have to concern ourselves with the determinants of μ in any detail because the value of a derivative dependent on a stock is, in general, independent of μ . The parameter σ , the stock price volatility, is, by contrast, critically important to the determination of the value of many derivatives. We will discuss procedures for estimating σ in Chapter 14. Typical values of σ for a stock are in the range 0.15 to 0.60 (i.e., 15% to 60%).

The standard deviation of the proportional change in the stock price in a small interval of time Δt is $\sigma\sqrt{\Delta t}$. As a rough approximation, the standard deviation of the proportional change in the stock price over a relatively long period of time T is $\sigma\sqrt{T}$. This means that, as an approximation, volatility can be interpreted as the standard deviation of the change in the stock price in 1 year. In Chapter 14, we will show that the volatility of a stock price is exactly equal to the standard deviation of the continuously compounded return provided by the stock in 1 year.

13.5 CORRELATED PROCESSES

So far we have considered how the stochastic process for a single variable can be represented. We now extend the analysis to the situation where there are two or more variables following correlated stochastic processes. Suppose that the processes followed by two variables x_1 and x_2 are

$$dx_1 = a_1 dt + b_1 dz_1 \quad \text{and} \quad dx_2 = a_2 dt + b_2 dz_2$$

where dz_1 and dz_2 are Wiener processes.

As has been explained, the discrete-time versions of these processes are

$$\Delta x_1 = a_1 \Delta t + b_1 \epsilon_1 \sqrt{\Delta t} \quad \text{and} \quad \Delta x_2 = a_2 \Delta t + b_2 \epsilon_2 \sqrt{\Delta t}$$

where ϵ_1 and ϵ_2 are samples from a standard normal distribution $\phi(0, 1)$.

The variables x_1 and x_2 can be simulated in the way described in Section 13.3. If they are uncorrelated with each other, the random samples ϵ_1 and ϵ_2 that are used to obtain movements in a particular period of time Δt should be independent of each other.

If x_1 and x_2 have a nonzero correlation ρ , then the ϵ_1 and ϵ_2 that are used to obtain movements in a particular period of time should be sampled from a bivariate normal distribution. Each variable in the bivariate normal distribution has a standard normal distribution and the correlation between the variables is ρ . In this situation, we would refer to the Wiener processes dz_1 and dz_2 as having a correlation ρ .

Obtaining samples for uncorrelated standard normal variables in cells in Excel involves putting the instruction “=NORMSINV(RAND))” in each of the cells. To sample standard normal variables ϵ_1 and ϵ_2 with correlation ρ , we can set

$$\epsilon_1 = u \quad \text{and} \quad \epsilon_2 = \rho u + \sqrt{1 - \rho^2} v$$

where u and v are sampled as uncorrelated variables with standard normal distributions.

Note that, in the processes we have assumed for x_1 and x_2 , the parameters a_1 , a_2 , b_1 , and b_2 can be functions of x_1 , x_2 , and t . In particular, a_1 and b_1 can be functions of x_2 as well as x_1 and t ; and a_2 and b_2 can be functions of x_1 as well as x_2 and t .

The results here can be generalized. When there are three different variables following correlated stochastic processes, we have to sample three different ϵ 's. These have a trivariate normal distribution. When there are n correlated variables, we have n different ϵ 's and these must be sampled from an appropriate multivariate normal distribution. The way this is done is discussed in Chapter 20.

13.6 ITÔ'S LEMMA

The price of a stock option is a function of the underlying stock's price and time. More generally, we can say that the price of any derivative is a function of the stochastic variables underlying the derivative and time. A serious student of derivatives must, therefore, acquire some understanding of the behavior of functions of stochastic variables. An important result in this area was discovered by the mathematician K. Itô in 1951,⁵ and is known as *Itô's lemma*.

Suppose that the value of a variable x follows the Itô process

$$dx = a(x, t) dt + b(x, t) dz \tag{13.11}$$

where dz is a Wiener process and a and b are functions of x and t . The variable x has a drift rate of a and a variance rate of b^2 . Itô's lemma shows that a function G of x and t follows the process

$$dG = \left(\frac{\partial G}{\partial x} a + \frac{\partial G}{\partial t} + \frac{1}{2} \frac{\partial^2 G}{\partial x^2} b^2 \right) dt + \frac{\partial G}{\partial x} b dz \tag{13.12}$$

where the dz is the same Wiener process as in equation (13.11). Thus, G also follows an Itô process, with a drift rate of

$$\frac{\partial G}{\partial x} a + \frac{\partial G}{\partial t} + \frac{1}{2} \frac{\partial^2 G}{\partial x^2} b^2$$

and a variance rate of

$$\left(\frac{\partial G}{\partial x} \right)^2 b^2$$

A completely rigorous proof of Itô's lemma is beyond the scope of this book. In the appendix to this chapter, we show that the lemma can be viewed as an extension of well-known results in differential calculus.

Earlier, we argued that

$$dS = \mu S dt + \sigma S dz \tag{13.13}$$

with μ and σ constant, is a reasonable model of stock price movements. From Itô's lemma, it follows that the process followed by a function G of S and t is

$$dG = \left(\frac{\partial G}{\partial S} \mu S + \frac{\partial G}{\partial t} + \frac{1}{2} \frac{\partial^2 G}{\partial S^2} \sigma^2 S^2 \right) dt + \frac{\partial G}{\partial S} \sigma S dz \tag{13.14}$$

Note that both S and G are affected by the same underlying source of uncertainty, dz . This proves to be very important in the derivation of the Black–Scholes–Merton results.

⁵ See K. Itô, "On Stochastic Differential Equations," *Memoirs of the American Mathematical Society*, 4 (1951): 1–51.

Application to Forward Contracts

To illustrate Itô's lemma, consider a forward contract on a non-dividend-paying stock. Assume that the risk-free rate of interest is constant and equal to r for all maturities. From equation (5.1),

$$F_0 = S_0 e^{rT}$$

where F_0 is the forward price at time zero, S_0 is the spot price at time zero, and T is the time to maturity of the forward contract.

We are interested in what happens to the forward price as time passes. We define F as the forward price at a general time t , and S as the stock price at time t , with $t < T$. The relationship between F and S is given by

$$F = S e^{r(T-t)} \quad (13.15)$$

Assuming that the process for S is given by equation (13.13), we can use Itô's lemma to determine the process for F . From equation (13.15),

$$\frac{\partial F}{\partial S} = e^{r(T-t)}, \quad \frac{\partial^2 F}{\partial S^2} = 0, \quad \frac{\partial F}{\partial t} = -r S e^{r(T-t)}$$

From equation (13.14), the process for F is given by

$$dF = [e^{r(T-t)} \mu S - r S e^{r(T-t)}] dt + e^{r(T-t)} \sigma S dz$$

Substituting F for $S e^{r(T-t)}$ gives

$$dF = (\mu - r)F dt + \sigma F dz \quad (13.16)$$

Like S , the forward price F follows geometric Brownian motion. It has an expected growth rate of $\mu - r$ rather than μ . The growth rate in F is the excess return of S over the risk-free rate.

13.7 THE LOGNORMAL PROPERTY

We now use Itô's lemma to derive the process followed by $\ln S$ when S follows the process in equation (13.13). We define

$$G = \ln S$$

Since

$$\frac{\partial G}{\partial S} = \frac{1}{S}, \quad \frac{\partial^2 G}{\partial S^2} = -\frac{1}{S^2}, \quad \frac{\partial G}{\partial t} = 0$$

it follows from equation (13.14) that the process followed by G is

$$dG = \left(\mu - \frac{\sigma^2}{2} \right) dt + \sigma dz \quad (13.17)$$

Since μ and σ are constant, this equation indicates that $G = \ln S$ follows a generalized Wiener process. It has constant drift rate $\mu - \sigma^2/2$ and constant variance rate σ^2 . The

change in $\ln S$ between time 0 and some future time T is therefore normally distributed, with mean $(\mu - \sigma^2/2)T$ and variance $\sigma^2 T$. This means that

$$\ln S_T - \ln S_0 \sim \phi \left[\left(\mu - \frac{\sigma^2}{2} \right) T, \sigma^2 T \right] \quad (13.18)$$

or

$$\ln S_T \sim \phi \left[\ln S_0 + \left(\mu - \frac{\sigma^2}{2} \right) T, \sigma^2 T \right] \quad (13.19)$$

where S_T is the stock price at a future time T , S_0 is the stock price at time 0, and as before $\phi(m, v)$ denotes a normal distribution with mean m and variance v .

Equation (13.19) shows that $\ln S_T$ is normally distributed. A variable has a lognormal distribution if the natural logarithm of the variable is normally distributed. The model of stock price behavior we have developed in this chapter therefore implies that a stock's price at time T , given its price today, is lognormally distributed. The standard deviation of the logarithm of the stock price is $\sigma\sqrt{T}$. It is proportional to the square root of how far ahead we are looking.

SUMMARY

Stochastic processes describe the probabilistic evolution of the value of a variable through time. A Markov process is one where only the present value of the variable is relevant for predicting the future. The past history of the variable and the way in which the present has emerged from the past is irrelevant.

A Wiener process dz is a process describing the evolution of a normally distributed variable. The drift of the process is zero and the variance rate is 1.0 per unit time. This means that, if the value of the variable is x_0 at time 0, then at time T it is normally distributed with mean x_0 and standard deviation \sqrt{T} .

A generalized Wiener process describes the evolution of a normally distributed variable with a drift of a per unit time and a variance rate of b^2 per unit time, where a and b are constants. This means that if, as before, the value of the variable is x_0 at time 0, it is normally distributed with a mean of $x_0 + aT$ and a standard deviation of $b\sqrt{T}$ at time T .

An Itô process is a process where the drift and variance rate of x can be a function of both x itself and time. The change in x in a very short period of time is, to a good approximation, normally distributed, but its change over longer periods of time is liable to be nonnormal.

One way of gaining an intuitive understanding of a stochastic process for a variable is to simulate the behavior of the variable. This involves dividing a time interval into many small time steps and randomly sampling possible paths for the variable. The future probability distribution for the variable can then be calculated. Monte Carlo simulation is discussed further in Chapter 20.

Itô's lemma is a way of calculating the stochastic process followed by a function of a variable from the stochastic process followed by the variable itself. As we shall see in Chapter 14, Itô's lemma plays a very important part in the pricing of derivatives. A key point is that the Wiener process dz underlying the stochastic process for the variable is exactly the same as the Wiener process underlying the stochastic process for the function of the variable. Both are subject to the same underlying source of uncertainty.

The stochastic process usually assumed for a stock price is geometric Brownian motion. Under this process the return to the holder of the stock in a small period of time is normally distributed and the returns in two nonoverlapping periods are independent. The value of the stock price at a future time has a lognormal distribution. The Black–Scholes–Merton model, which we cover in the next chapter, is based on the geometric Brownian motion assumption.

FURTHER READING

On Efficient Markets and the Markov Property of Stock Prices

Brealey, R. A. *An Introduction to Risk and Return from Common Stock*, 2nd edn. Cambridge, MA: MIT Press, 1986.

Cootner, P. H. (ed.) *The Random Character of Stock Market Prices*. Cambridge, MA: MIT Press, 1964.

On Stochastic Processes

Cox, D. R., and H. D. Miller. *The Theory of Stochastic Processes*. London: Chapman & Hall, 1977.

Feller, W. *Introduction to Probability Theory and Its Applications*. New York: Wiley, 1968.

Karlin, S., and H. M. Taylor. *A First Course in Stochastic Processes*, 2nd edn. New York: Academic Press, 1975.

Neftci, S. *Introduction to Mathematics of Financial Derivatives*, 2nd edn. New York: Academic Press, 2000.

Shreve, S. E. *Stochastic Calculus for Finance II: Continuous-Time Models*. New York: Springer, 2004.

Practice Questions (Answers in Solutions Manual)

- 13.1. What would it mean to assert that the temperature at a certain place follows a Markov process? Do you think that temperatures do, in fact, follow a Markov process?
- 13.2. Can a trading rule based on the past history of a stock's price ever produce returns that are consistently above average? Discuss.
- 13.3. A company's cash position, measured in millions of dollars, follows a generalized Wiener process with a drift rate of 0.5 per quarter and a variance rate of 4.0 per quarter. How high does the company's initial cash position have to be for the company to have a less than 5% chance of a negative cash position by the end of 1 year?
- 13.4. Variables X_1 and X_2 follow generalized Wiener processes, with drift rates μ_1 and μ_2 and variances σ_1^2 and σ_2^2 . What process does $X_1 + X_2$ follow if:
 - (a) The changes in X_1 and X_2 in any short interval of time are uncorrelated?
 - (b) There is a correlation ρ between the changes in X_1 and X_2 in any short time interval?
- 13.5. Consider a variable S that follows the process

$$dS = \mu dt + \sigma dz$$

For the first three years, $\mu = 2$ and $\sigma = 3$; for the next three years, $\mu = 3$ and $\sigma = 4$. If the initial value of the variable is 5, what is the probability distribution of the value of the variable at the end of year 6?

- 13.6. Suppose that G is a function of a stock price S and time. Suppose that σ_S and σ_G are the volatilities of S and G . Show that, when the expected return of S increases by $\lambda\sigma_S$, the growth rate of G increases by $\lambda\sigma_G$, where λ is a constant.
- 13.7. Stock A and stock B both follow geometric Brownian motion. Changes in any short interval of time are uncorrelated with each other. Does the value of a portfolio consisting of one of stock A and one of stock B follow geometric Brownian motion? Explain your answer.
- 13.8. The process for the stock price in equation (13.8) is

$$\Delta S = \mu S \Delta t + \sigma S \epsilon \sqrt{\Delta t}$$

where μ and σ are constant. Explain carefully the difference between this model and each of the following:

$$\Delta S = \mu \Delta t + \sigma \epsilon \sqrt{\Delta t}$$

$$\Delta S = \mu S \Delta t + \sigma \epsilon \sqrt{\Delta t}$$

$$\Delta S = \mu \Delta t + \sigma S \epsilon \sqrt{\Delta t}$$

Why is the model in equation (13.8) a more appropriate model of stock price behavior than any of these three alternatives?

- 13.9. It has been suggested that the short-term interest rate r follows the stochastic process

$$dr = a(b - r)dt + rc dz$$

where a , b , c are positive constants and dz is a Wiener process. Describe the nature of this process.

- 13.10. Suppose that a stock price S follows geometric Brownian motion with expected return μ and volatility σ :

$$dS = \mu S dt + \sigma S dz$$

What is the process followed by the variable S^n ? Show that S^n also follows geometric Brownian motion.

- 13.11. Suppose that x is the yield to maturity with continuous compounding on a zero-coupon bond that pays off \$1 at time T . Assume that x follows the process

$$dx = a(x_0 - x)dt + sx dz$$

where a , x_0 , and s are positive constants and dz is a Wiener process. What is the process followed by the bond price?

- 13.12. A stock whose price is \$30 has an expected return of 9% and a volatility of 20%. In Excel, simulate the stock price path over 5 years using monthly time steps and random samples from a normal distribution. Chart the simulated stock price path. By hitting F9, observe how the path changes as the random samples change.

Further Questions

- 13.13. Suppose that a stock price has an expected return of 16% per annum and a volatility of 30% per annum. When the stock price at the end of a certain day is \$50, calculate the following:
- The expected stock price at the end of the next day
 - The standard deviation of the stock price at the end of the next day
 - The 95% confidence limits for the stock price at the end of the next day.

- 13.14. A company's cash position, measured in millions of dollars, follows a generalized Wiener process with a drift rate of 0.1 per month and a variance rate of 0.16 per month. The initial cash position is 2.0.
- What are the probability distributions of the cash position after 1 month, 6 months, and 1 year?
 - What are the probabilities of a negative cash position at the end of 6 months and 1 year?
 - At what time in the future is the probability of a negative cash position greatest?
- 13.15. Suppose that x is the yield on a perpetual government bond that pays interest at the rate of \$1 per annum. Assume that x is expressed with continuous compounding, that interest is paid continuously on the bond, and that x follows the process

$$dx = a(x_0 - x) dt + sx dz$$

where a , x_0 , and s are positive constants, and dz is a Wiener process. What is the process followed by the bond price? What is the expected instantaneous return (including interest and capital gains) to the holder of the bond?

- 13.16. If S follows the geometric Brownian motion process in equation (13.6), what is the process followed by
- $y = 2S$
 - $y = S^2$
 - $y = e^S$
 - $y = e^{r(T-t)}/S$.

In each case express the coefficients of dt and dz in terms of y rather than S .

- 13.17. A stock price is currently 50. Its expected return and volatility are 12% and 30%, respectively. What is the probability that the stock price will be greater than 80 in 2 years? (*Hint*: $S_T > 80$ when $\ln S_T > \ln 80$.)
- 13.18. Stock A, whose price is \$30, has an expected return of 11% and a volatility of 25%. Stock B, whose price is \$40, has an expected return of 15% and a volatility of 30%. The processes driving the returns are correlated with correlation parameter ρ . In Excel, simulate the two stock price paths over 3 months using daily time steps and random samples from normal distributions. Chart the results and by hitting F9 observe how the paths change as the random samples change. Consider values for ρ equal to 0.25, 0.75, and 0.95.

APPENDIX DERIVATION OF ITÔ'S LEMMA

In this appendix, we show how Itô's lemma can be regarded as a natural extension of other, simpler results. Consider a continuous and differentiable function G of a variable x . If Δx is a small change in x and ΔG is the resulting small change in G , a well-known result from ordinary calculus is

$$\Delta G \approx \frac{dG}{dx} \Delta x \tag{13A.1}$$

In other words, ΔG is approximately equal to the rate of change of G with respect to x multiplied by Δx . The error involves terms of order Δx^2 . If more precision is required, a Taylor series expansion of ΔG can be used:

$$\Delta G = \frac{dG}{dx} \Delta x + \frac{1}{2} \frac{d^2G}{dx^2} \Delta x^2 + \frac{1}{6} \frac{d^3G}{dx^3} \Delta x^3 + \dots$$

For a continuous and differentiable function G of two variables x and y , the result analogous to equation (13A.1) is

$$\Delta G \approx \frac{\partial G}{\partial x} \Delta x + \frac{\partial G}{\partial y} \Delta y \tag{13A.2}$$

and the Taylor series expansion of ΔG is

$$\Delta G = \frac{\partial G}{\partial x} \Delta x + \frac{\partial G}{\partial y} \Delta y + \frac{1}{2} \frac{\partial^2 G}{\partial x^2} \Delta x^2 + \frac{\partial^2 G}{\partial x \partial y} \Delta x \Delta y + \frac{1}{2} \frac{\partial^2 G}{\partial y^2} \Delta y^2 + \dots \tag{13A.3}$$

In the limit, as Δx and Δy tend to zero, equation (13A.3) becomes

$$dG = \frac{\partial G}{\partial x} dx + \frac{\partial G}{\partial y} dy \tag{13A.4}$$

We now extend equation (13A.4) to cover functions of variables following Itô processes. Suppose that a variable x follows the Itô process

$$dx = a(x, t) dt + b(x, t) dz \tag{13A.5}$$

and that G is some function of x and of time t . By analogy with equation (13A.3), we can write

$$\Delta G = \frac{\partial G}{\partial x} \Delta x + \frac{\partial G}{\partial t} \Delta t + \frac{1}{2} \frac{\partial^2 G}{\partial x^2} \Delta x^2 + \frac{\partial^2 G}{\partial x \partial t} \Delta x \Delta t + \frac{1}{2} \frac{\partial^2 G}{\partial t^2} \Delta t^2 + \dots \tag{13A.6}$$

Equation (13A.5) can be discretized to

$$\Delta x = a(x, t) \Delta t + b(x, t) \epsilon \sqrt{\Delta t}$$

or, if arguments are dropped,

$$\Delta x = a \Delta t + b \epsilon \sqrt{\Delta t} \tag{13A.7}$$

This equation reveals an important difference between the situation in equation (13A.6) and the situation in equation (13A.3). When limiting arguments were used to move from equation (13A.3) to equation (13A.4), terms in Δx^2 were ignored because they were second-order terms. From equation (13A.7), we have

$$\Delta x^2 = b^2 \epsilon^2 \Delta t + \text{terms of higher order in } \Delta t \quad (13A.8)$$

This shows that the term involving Δx^2 in equation (13A.6) has a component that is of order Δt and cannot be ignored.

The variance of a standardized normal distribution is 1.0. This means that

$$E(\epsilon^2) - [E(\epsilon)]^2 = 1$$

where E denotes expected value. Since $E(\epsilon) = 0$, it follows that $E(\epsilon^2) = 1$. The expected value of $\epsilon^2 \Delta t$, therefore, is Δt . The variance of $\epsilon^2 \Delta t$ is, from the properties of the standard normal distribution, $2\Delta t^2$. We know that the variance of the change in a stochastic variable in time Δt is proportional to Δt , not Δt^2 . The variance of $\epsilon^2 \Delta t$ is therefore too small for it to have a stochastic component. As a result, we can treat $\epsilon^2 \Delta t$ as nonstochastic and equal to its expected value, Δt , as Δt tends to zero. It follows from equation (13A.8) that Δx^2 becomes nonstochastic and equal to $b^2 \Delta t$ as Δt tends to zero. Taking limits as Δx and Δt tend to zero in equation (13A.6), and using this last result, we obtain

$$dG = \frac{\partial G}{\partial x} dx + \frac{\partial G}{\partial t} dt + \frac{1}{2} \frac{\partial^2 G}{\partial x^2} b^2 dt \quad (13A.9)$$

This is Itô's lemma. If we substitute for dx from equation (13A.5), equation (13A.9) becomes

$$dG = \left(\frac{\partial G}{\partial x} a + \frac{\partial G}{\partial t} + \frac{1}{2} \frac{\partial^2 G}{\partial x^2} b^2 \right) dt + \frac{\partial G}{\partial x} b dz.$$

Technical Note 29 at www.rotman.utoronto.ca/~hull/TechnicalNotes provides proofs of extensions to Itô's lemma. When G is a function of variables x_1, x_2, \dots, x_n and

$$dx_i = a_i dt + b_i dz_i$$

we have

$$dG = \left(\sum_{i=1}^n \frac{\partial G}{\partial x_i} a_i + \frac{\partial G}{\partial t} + \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n \frac{\partial^2 G}{\partial x_i \partial x_j} b_i b_j \rho_{ij} \right) dt + \sum_{i=1}^n \frac{\partial G}{\partial x_i} b_i dz_i \quad (13A.10)$$

Also, when G is a function of a variable x with several sources of uncertainty so that

$$dx = a dt + \sum_{i=1}^m b_i dz_i$$

we have

$$dG = \left(\sum_{i=1}^m \frac{\partial G}{\partial x} a + \frac{\partial G}{\partial t} + \frac{1}{2} \frac{\partial^2 G}{\partial x^2} \sum_{i=1}^m \sum_{j=1}^m b_i b_j \rho_{ij} \right) dt + \frac{\partial G}{\partial x} \sum_{i=1}^m b_i dz_i \quad (13A.11)$$

In these equations, ρ_{ij} is the correlation between dz_i and dz_j (see Section 13.5).

14

CHAPTER

The Black–Scholes–Merton Model



In the early 1970s, Fischer Black, Myron Scholes, and Robert Merton achieved a major breakthrough in the pricing of European stock options.¹ This was the development of what has become known as the Black–Scholes–Merton (or Black–Scholes) model. The model has had a huge influence on the way that traders price and hedge derivatives. In 1997, the importance of the model was recognized when Robert Merton and Myron Scholes were awarded the Nobel prize for economics. Sadly, Fischer Black died in 1995; otherwise he too would undoubtedly have been one of the recipients of this prize.

How did Black, Scholes, and Merton make their breakthrough? Previous researchers had made the similar assumptions and had correctly calculated the expected payoff from a European option. However, as explained in Section 12.2, it is difficult to know the correct discount rate to use for this payoff. Black and Scholes used the capital asset pricing model (see the appendix to Chapter 3) to determine a relationship between the market's required return on the option to the required return on the stock. This was not easy because the relationship depends on both the stock price and time. Merton's approach was different from that of Black and Scholes. It involved setting up a riskless portfolio consisting of the option and the underlying stock and arguing that the return on the portfolio over a short period of time must be the risk-free return. This is similar to what we did in Section 12.1—but more complicated because the portfolio changes continuously through time. Merton's approach was more general than that of Black and Scholes because it did not rely on the assumptions of the capital asset pricing model.

This chapter covers Merton's approach to deriving the Black–Scholes–Merton model. It explains how volatility can be either estimated from historical data or implied from option prices using the model. It shows how the risk-neutral valuation argument introduced in Chapter 12 can be used. It also shows how the Black–Scholes–Merton model can be extended to deal with European call and put options on dividend-paying stocks and presents some results on the pricing of American call options on dividend-paying stocks.

¹ See F. Black and M. Scholes, "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, 81 (May/June 1973): 637–59; R.C. Merton, "Theory of Rational Option Pricing," *Bell Journal of Economics and Management Science*, 4 (Spring 1973): 141–83.

14.1 LOGNORMAL PROPERTY OF STOCK PRICES

The model of stock price behavior used by Black, Scholes, and Merton is the model we developed in Chapter 13. It assumes that percentage changes in the stock price in a short period of time are normally distributed. Define

- μ : Expected return on stock per year
- σ : Volatility of the stock price per year.

The mean of the return in time Δt is $\mu \Delta t$ and the standard deviation of the return is $\sigma\sqrt{\Delta t}$, so that

$$\frac{\Delta S}{S} \sim \phi(\mu \Delta t, \sigma^2 \Delta t) \quad (14.1)$$

where ΔS is the change in the stock price S in time Δt , and $\phi(m, v)$ denotes a normal distribution with mean m and variance v . (This is equation (13.9).)

As shown in Section 13.7, the model implies that

$$\ln S_T - \ln S_0 \sim \phi\left[\left(\mu - \frac{\sigma^2}{2}\right)T, \sigma^2 T\right]$$

so that

$$\ln \frac{S_T}{S_0} \sim \phi\left[\left(\mu - \frac{\sigma^2}{2}\right)T, \sigma^2 T\right] \quad (14.2)$$

and

$$\ln S_T \sim \phi\left[\ln S_0 + \left(\mu - \frac{\sigma^2}{2}\right)T, \sigma^2 T\right] \quad (14.3)$$

where S_T is the stock price at a future time T and S_0 is the stock price at time 0. Equation (14.3) shows that $\ln S_T$ is normally distributed, so that S_T has a lognormal distribution. The mean of $\ln S_T$ is $\ln S_0 + (\mu - \sigma^2/2)T$ and the standard deviation is $\sigma\sqrt{T}$.

Example 14.1

Consider a stock with an initial price of \$40, an expected return of 16% per annum, and a volatility of 20% per annum. From equation (14.3), the probability distribution of the stock price S_T in 6 months' time is given by

$$\ln S_T \sim \phi[\ln 40 + (0.16 - 0.2^2/2) \times 0.5, 0.2^2 \times 0.5]$$

$$\ln S_T \sim \phi(3.759, 0.02)$$

There is a 95% probability that a normally distributed variable has a value within 1.96 standard deviations of its mean. In this case, the standard deviation is $\sqrt{0.02} = 0.141$. Hence, with 95% confidence,

$$3.759 - 1.96 \times 0.141 < \ln S_T < 3.759 + 1.96 \times 0.141$$


This can be written

$$e^{3.759-1.96 \times 0.141} < S_T < e^{3.759+1.96 \times 0.141}$$

or

$$32.55 < S_T < 56.56$$

Thus, there is a 95% probability that the stock price in 6 months will lie between 32.55 and 56.56.

Figure 14.1 Lognormal distribution.

A variable that has a lognormal distribution can take any value between zero and infinity. Figure 14.1 illustrates the shape of a lognormal distribution. Unlike the normal distribution, it is skewed so that the mean, median, and mode are all different. From equation (14.3) and the properties of the lognormal distribution, it can be shown that the expected value $E(S_T)$ of S_T is given by

$$E(S_T) = S_0 e^{\mu T} \quad (14.4)$$

This fits in with the definition of μ as the expected rate of return. The variance $\text{var}(S_T)$ of S_T , can be shown to be given by²

$$\text{var}(S_T) = S_0^2 e^{2\mu T} (e^{\sigma^2 T} - 1) \quad (14.5)$$

Example 14.2

Consider a stock where the current price is \$20, the expected return is 20% per annum, and the volatility is 40% per annum. The expected stock price, $E(S_T)$, and the variance of the stock price, $\text{var}(S_T)$, in 1 year are given by

$$E(S_T) = 20e^{0.2 \times 1} = 24.43 \quad \text{and} \quad \text{var}(S_T) = 400e^{2 \times 0.2 \times 1} (e^{0.4^2 \times 1} - 1) = 103.54$$

The standard deviation of the stock price in 1 year is $\sqrt{103.54}$, or 10.18.

14.2 THE DISTRIBUTION OF THE RATE OF RETURN

The lognormal property of stock prices can be used to provide information on the probability distribution of the continuously compounded rate of return earned on a stock between times 0 and T . If we define the continuously compounded rate of return

² See Technical Note 2 at www.rotman.utoronto.ca/~hull/TechnicalNotes for a proof of the results in equations (14.4) and (14.5). For a more extensive discussion of the properties of the lognormal distribution, see J. Aitchison and J. A. C. Brown, *The Lognormal Distribution*. Cambridge University Press, 1966.

per annum realized between times 0 and T as x , then

$$S_T = S_0 e^{xT}$$

so that

$$x = \frac{1}{T} \ln \frac{S_T}{S_0} \quad (14.6)$$

From equation (14.2), it follows that

$$x \sim \phi\left(\mu - \frac{\sigma^2}{2}, \frac{\sigma^2}{T}\right) \quad (14.7)$$

Thus, the continuously compounded rate of return per annum is normally distributed with mean $\mu - \sigma^2/2$ and standard deviation σ/\sqrt{T} . As T increases, the standard deviation of x declines. To understand the reason for this, consider two cases: $T = 1$ and $T = 20$. We are more certain about the average return per year over 20 years than we are about the return in any one year.

Example 14.3

Consider a stock with an expected return of 17% per annum and a volatility of 20% per annum. The probability distribution for the average rate of return (continuously compounded) realized over 3 years is normal, with mean

$$0.17 - \frac{0.2^2}{2} = 0.15$$

or 15% per annum, and standard deviation

$$\sqrt{\frac{0.2^2}{3}} = 0.1155$$

or 11.55% per annum. Because there is a 95% chance that a normally distributed variable will lie within 1.96 standard deviations of its mean, we can be 95% confident that the average return realized over 3 years will be between $15 - 1.96 \times 11.55 = -7.6\%$ and $15 + 1.96 \times 11.55 = +37.6\%$ per annum.

14.3 THE EXPECTED RETURN

The expected return, μ , required by investors from a stock depends on the riskiness of the stock. The higher the risk, the higher the expected return. It also depends on the level of interest rates in the economy. The higher the level of interest rates, the higher the expected return required on any given stock. Fortunately, we do not have to concern ourselves with the determinants of μ in any detail. It turns out that the value of a stock option, when expressed in terms of the value of the underlying stock, does not depend on μ at all. Nevertheless, there is one aspect of the expected return from a stock that frequently causes confusion and needs to be explained.

Equation (14.1) shows that $\mu \Delta t$ is the expected percentage change in the stock price in a very short period of time, Δt . It is natural to assume from this that μ is the expected continuously compounded return on the stock. However, this is not the case. The continuously compounded return, x , actually realized over a period of time of length T

is given by equation (14.6) as

$$x = \frac{1}{T} \ln \frac{S_T}{S_0}$$

and, as indicated in equation (14.7), the expected value $E(x)$ of x is $\mu - \sigma^2/2$.

The reason why the expected continuously compounded return is different from μ is subtle, but important. Suppose we consider a very large number of very short periods of time of length Δt . Define S_i as the stock price at the end of the i th interval and ΔS_i as $S_{i+1} - S_i$. Under the assumptions we are making for stock price behavior, the average of the returns on the stock in each interval is close to μ . In other words, $\mu \Delta t$ is close to the arithmetic mean of the $\Delta S_i/S_i$. However, the expected return over the whole period covered by the data, expressed with a compounding interval of Δt , is close to $\mu - \sigma^2/2$, not μ .³ Business Snapshot 14.1 provides a numerical example concerning the mutual fund industry to illustrate why this is so. For a mathematical explanation of what is going on, we start with equation (14.4):

$$E(S_T) = S_0 e^{\mu T}$$

Taking logarithms, we get

$$\ln[E(S_T)] = \ln(S_0) + \mu T$$

It is now tempting to set $\ln[E(S_T)] = E[\ln(S_T)]$, so that $E[\ln(S_T)] - \ln(S_0) = \mu T$, or $E[\ln(S_T/S_0)] = \mu T$, which leads to $E(x) = \mu$. However, we cannot do this because \ln is a nonlinear function. In fact, $\ln[E(S_T)] > E[\ln(S_T)]$, so that $E[\ln(S_T/S_0)] < \mu T$, which leads to $E(x) < \mu$. (As pointed out above, $E(x) = \mu - \sigma^2/2$.)

14.4 VOLATILITY

The volatility, σ , of a stock is a measure of our uncertainty about the returns provided by the stock. Stocks typically have a volatility between 15% and 60%.

From equation (14.7), the volatility of a stock price can be defined as the standard deviation of the return provided by the stock in 1 year when the return is expressed using continuous compounding.

When Δt is small, equation (14.1) shows that $\sigma^2 \Delta t$ is approximately equal to the variance of the percentage change in the stock price in time Δt . This means that $\sigma \sqrt{\Delta t}$ is approximately equal to the standard deviation of the percentage change in the stock price in time Δt . Suppose that $\sigma = 0.3$, or 30%, per annum and the current stock price is \$50. The standard deviation of the percentage change in the stock price in 1 week is approximately

$$30 \times \sqrt{\frac{1}{52}} = 4.16\%$$

A 1-standard-deviation move in the stock price in 1 week is therefore $50 \times 0.0416 = 2.08$.

Uncertainty about a future stock price, as measured by its standard deviation, increases—at least approximately—with the square root of how far ahead we are looking. For example, the standard deviation of the stock price in 4 weeks is approximately twice the standard deviation in 1 week.

³ The arguments in this section show that the term “expected return” is ambiguous. It can refer either to μ or to $\mu - \sigma^2/2$. Unless otherwise stated, it will be used to refer to μ throughout this book.

Business Snapshot 14.1 Mutual Fund Returns Can Be Misleading

The difference between μ and $\mu - \sigma^2/2$ is closely related to an issue in the reporting of mutual fund returns. Suppose that the following is a sequence of returns per annum reported by a mutual fund manager over the last five years (measured using annual compounding): 15%, 20%, 30%, -20%, 25%.

The arithmetic mean of the returns, calculated by taking the sum of the returns and dividing by 5, is 14%. However, an investor would actually earn less than 14% per annum by leaving the money invested in the fund for 5 years. The dollar value of \$100 at the end of the 5 years would be

$$100 \times 1.15 \times 1.20 \times 1.30 \times 0.80 \times 1.25 = \$179.40$$

By contrast, a 14% return with annual compounding would give

$$100 \times 1.14^5 = \$192.54$$

The return that gives \$179.40 at the end of five years is 12.4%. This is because

$$100 \times (1.124)^5 = 179.40$$

What average return should the fund manager report? It is tempting for the manager to make a statement such as: “The average of the returns per year that we have realized in the last 5 years is 14%.” Although true, this is misleading. It is much less misleading to say: “The average return realized by someone who invested with us for the last 5 years is 12.4% per year.” In some jurisdictions, regulations require fund managers to report returns the second way.

This phenomenon is an example of a result that is well known by mathematicians. The geometric mean of a set of numbers (not all the same) is always less than the arithmetic mean. In our example, the return multipliers each year are 1.15, 1.20, 1.30, 0.80, and 1.25. The arithmetic mean of these numbers is 1.140, but the geometric mean is only 1.124.

Estimating Volatility from Historical Data

To estimate the volatility of a stock price empirically, the stock price is usually observed at fixed intervals of time (e.g., every day, week, or month). Define:

$n + 1$: Number of observations

S_i : Stock price at end of i th interval, with $i = 0, 1, \dots, n$

τ : Length of time interval in years

and let

$$u_i = \ln\left(\frac{S_i}{S_{i-1}}\right) \quad \text{for } i = 1, 2, \dots, n$$

The usual estimate, s , of the standard deviation of the u_i is given by

$$s = \sqrt{\frac{1}{n-1} \sum_{i=1}^n (u_i - \bar{u})^2}$$

or

$$s = \sqrt{\frac{1}{n-1} \sum_{i=1}^n u_i^2 - \frac{1}{n(n-1)} \left(\sum_{i=1}^n u_i \right)^2}$$

where \bar{u} is the mean of the u_i .⁴

From equation (14.2), the standard deviation of the u_i is $\sigma\sqrt{\tau}$. The variable s is therefore an estimate of $\sigma\sqrt{\tau}$. It follows that σ itself can be estimated as $\hat{\sigma}$, where

$$\hat{\sigma} = \frac{s}{\sqrt{\tau}}$$

The standard error of this estimate can be shown to be approximately $\hat{\sigma}/\sqrt{2n}$.

Choosing an appropriate value for n is not easy. More data generally lead to more accuracy, but σ does change over time and data that are too old may not be relevant for predicting the future volatility. A compromise that seems to work reasonably well is to use closing prices from daily data over the most recent 90 to 180 days. Alternatively, as a rule of thumb, n can be set equal to the number of days to which the volatility is to be applied. Thus, if the volatility estimate is to be used to value a 2-year option, daily data for the last 2 years are used. More sophisticated approaches to estimating volatility involving GARCH models are discussed in Chapter 22.

Example 14.4

Table 14.1 shows a possible sequence of stock prices during 21 consecutive trading days. In this case, $n = 20$, so that

$$\sum_{i=1}^n u_i = 0.09531 \quad \text{and} \quad \sum_{i=1}^n u_i^2 = 0.00326$$

and the estimate of the standard deviation of the daily return is

$$\sqrt{\frac{0.00326}{19} - \frac{0.09531^2}{20 \times 19}} = 0.01216$$

or 1.216%. Assuming that there are 252 trading days per year, $\tau = 1/252$ and the data give an estimate for the volatility per annum of $0.01216\sqrt{252} = 0.193$, or 19.3%. The standard error of this estimate is

$$\frac{0.193}{\sqrt{2 \times 20}} = 0.031$$

or 3.1% per annum.

The foregoing analysis assumes that the stock pays no dividends, but it can be adapted to accommodate dividend-paying stocks. The return, u_i , during a time interval that includes an ex-dividend day is given by

$$u_i = \ln \frac{S_i + D}{S_{i-1}}$$

where D is the amount of the dividend. The return in other time intervals is still

$$u_i = \ln \frac{S_i}{S_{i-1}}$$

⁴ The mean \bar{u} is often assumed to be zero when estimates of historical volatilities are made.

Table 14.1 Computation of volatility.

| Day <i>i</i> | Closing stock price (dollars), S_i | Price relative S_i/S_{i-1} | Daily return $u_i = \ln(S_i/S_{i-1})$ |
|-----------------|---|---------------------------------|--|
| 0 | 20.00 | | |
| 1 | 20.10 | 1.00500 | 0.00499 |
| 2 | 19.90 | 0.99005 | -0.01000 |
| 3 | 20.00 | 1.00503 | 0.00501 |
| 4 | 20.50 | 1.02500 | 0.02469 |
| 5 | 20.25 | 0.98780 | -0.01227 |
| 6 | 20.90 | 1.03210 | 0.03159 |
| 7 | 20.90 | 1.00000 | 0.00000 |
| 8 | 20.90 | 1.00000 | 0.00000 |
| 9 | 20.75 | 0.99282 | -0.00720 |
| 10 | 20.75 | 1.00000 | 0.00000 |
| 11 | 21.00 | 1.01205 | 0.01198 |
| 12 | 21.10 | 1.00476 | 0.00475 |
| 13 | 20.90 | 0.99052 | -0.00952 |
| 14 | 20.90 | 1.00000 | 0.00000 |
| 15 | 21.25 | 1.01675 | 0.01661 |
| 16 | 21.40 | 1.00706 | 0.00703 |
| 17 | 21.40 | 1.00000 | 0.00000 |
| 18 | 21.25 | 0.99299 | -0.00703 |
| 19 | 21.75 | 1.02353 | 0.02326 |
| 20 | 22.00 | 1.01149 | 0.01143 |

However, as tax factors play a part in determining returns around an ex-dividend date, it is probably best to discard altogether data for intervals that include an ex-dividend date.

Trading Days vs. Calendar Days

An important issue is whether time should be measured in calendar days or trading days when volatility parameters are being estimated and used. As shown in Business Snapshot 14.2, research shows that volatility is much higher when the exchange is open for trading than when it is closed. As a result, practitioners tend to ignore days when the exchange is closed when estimating volatility from historical data and when calculating the life of an option. The volatility per annum is calculated from the volatility per trading day using the formula

$$\text{Volatility per annum} = \text{Volatility per trading day} \times \sqrt{\frac{\text{Number of trading days per annum}}{\text{Number of trading days per annum}}}$$

This is what we did in Example 14.4 when calculating volatility from the data in Table 14.1. The number of trading days in a year is usually assumed to be 252 for stocks.

Business Snapshot 14.2 What Causes Volatility?

It is natural to assume that the volatility of a stock is caused by new information reaching the market. This new information causes people to revise their opinions about the value of the stock. The price of the stock changes and volatility results. This view of what causes volatility is not supported by research. With several years of daily stock price data, researchers can calculate:

1. The variance of stock price returns between the close of trading on one day and the close of trading on the next day when there are no intervening nontrading days
2. The variance of the stock price returns between the close of trading on Friday and the close of trading on Monday

The second of these is the variance of returns over a 3-day period. The first is a variance over a 1-day period. We might reasonably expect the second variance to be three times as great as the first variance. Fama (1965), French (1980), and French and Roll (1986) show that this is not the case. These three research studies estimate the second variance to be, respectively, 22%, 19%, and 10.7% higher than the first variance.

At this stage one might be tempted to argue that these results are explained by more news reaching the market when the market is open for trading. But research by Roll (1984) does not support this explanation. Roll looked at the prices of orange juice futures. By far the most important news for orange juice futures prices is news about the weather and this is equally likely to arrive at any time. When Roll did a similar analysis to that just described for stocks, he found that the second (Friday-to-Monday) variance for orange juice futures is only 1.54 times the first variance.

The only reasonable conclusion from all this is that volatility is to a large extent caused by trading itself. (Traders usually have no difficulty accepting this conclusion!)

The life of an option is also usually measured using trading days rather than calendar days. It is calculated as T years, where


$$T = \frac{\text{Number of trading days until option maturity}}{252}$$

14.5 THE IDEA UNDERLYING THE BLACK–SCHOLES–MERTON DIFFERENTIAL EQUATION

The Black–Scholes–Merton differential equation is an equation that must be satisfied by the price of any derivative dependent on a non-dividend-paying stock. The equation is derived in the next section. Here we consider the nature of the arguments we will use.

These are similar to the no-arbitrage arguments we used to value stock options in Chapter 12 for the situation where stock price movements are binomial. They involve setting up a riskless portfolio consisting of a position in the derivative and a position in the stock. In the absence of arbitrage opportunities, the return from the portfolio must be the risk-free interest rate, r . This leads to the Black-Scholes-Merton differential equation.

Figure 14.2 Relationship between call price and stock price. Current stock price is S_0 .



The reason a riskless portfolio can be set up is that the stock price and the derivative price are both affected by the same underlying source of uncertainty: stock price movements. In any short period of time, the price of the derivative is perfectly correlated with the price of the underlying stock. When an appropriate portfolio of the stock and the derivative is established, the gain or loss from the stock position always offsets the gain or loss from the derivative position so that the overall value of the portfolio at the end of the short period of time is known with certainty.

Suppose, for example, that at a particular point in time the relationship between a small change ΔS in the stock price and the resultant small change Δc in the price of a European call option is given by

$$\Delta c = 0.4 \Delta S$$

This means that the slope of the line representing the relationship between c and S is 0.4, as indicated in Figure 14.2. The riskless portfolio would consist of:

1. A long position in 0.4 shares
2. A short position in one call option.

Suppose, for example, that the stock price increases by 10 cents. The option price will increase by 4 cents and the $0.4 \times 10 = 4$ cents gain on the shares is equal to the 4 cents loss on the short option position.

There is one important difference between the Black–Scholes–Merton analysis and our analysis using a binomial model in Chapter 12. In Black–Scholes–Merton, the position in the stock and the derivative is riskless for only a very short period of time. (Theoretically, it remains riskless only for an instantaneously short period of time.) To remain riskless, it must be adjusted, or *rebalanced*, frequently.⁵ For example, the relationship between Δc and ΔS in our example might change from $\Delta c = 0.4 \Delta S$ today to $\Delta c = 0.5 \Delta S$ in 2 weeks. This would mean that, in order to maintain the riskless position, an extra 0.1 share would have to be purchased for each call option sold. It is nevertheless true that the return from the riskless portfolio in any very short period of time must be the risk-free interest rate. This is the key element in the Black–Scholes–Merton analysis and leads to their pricing formulas.

⁵ We discuss the rebalancing of portfolios in more detail in Chapter 18.

Assumptions

The assumptions we use to derive the Black–Scholes–Merton differential equation are as follows:

1. The stock price follows the process developed in Chapter 13 with μ and σ constant.
2. The short selling of securities with full use of proceeds is permitted.
3. There are no transactions costs or taxes. All securities are perfectly divisible.
4. There are no dividends during the life of the derivative.
5. There are no riskless arbitrage opportunities.
6. Security trading is continuous.
7. The risk-free rate of interest, r , is constant and the same for all maturities.

As we discuss in later chapters, some of these assumptions can be relaxed. For example, σ and r can be known functions of t . We can even allow interest rates to be stochastic provided that the stock price distribution at maturity of the option is still lognormal.

14.6 DERIVATION OF THE BLACK–SCHOLES–MERTON DIFFERENTIAL EQUATION

In this section, the notation is different from elsewhere in the book. We consider a derivative's price at a general time t (not at time zero). If T is the maturity date, the time to maturity is $T - t$.

The stock price process we are assuming is the one we developed in Section 13.3:

$$dS = \mu S dt + \sigma S dz \quad (14.8)$$

Suppose that f is the price of a call option or other derivative contingent on S . The variable f must be some function of S and t . Hence, from equation (13.14),

$$df = \left(\frac{\partial f}{\partial S} \mu S + \frac{\partial f}{\partial t} + \frac{1}{2} \frac{\partial^2 f}{\partial S^2} \sigma^2 S^2 \right) dt + \frac{\partial f}{\partial S} \sigma S dz \quad (14.9)$$

The discrete versions of equations (14.8) and (14.9) are

$$\Delta S = \mu S \Delta t + \sigma S \Delta z \quad (14.10)$$

and

$$\Delta f = \left(\frac{\partial f}{\partial S} \mu S + \frac{\partial f}{\partial t} + \frac{1}{2} \frac{\partial^2 f}{\partial S^2} \sigma^2 S^2 \right) \Delta t + \frac{\partial f}{\partial S} \sigma S \Delta z \quad (14.11)$$

where Δf and ΔS are the changes in f and S in a small time interval Δt . Recall from the discussion of Itô's lemma in Section 13.6 that the Wiener processes underlying f and S are the same. In other words, the Δz ($= \epsilon \sqrt{\Delta t}$) in equations (14.10) and (14.11) are the same. It follows that a portfolio of the stock and the derivative can be constructed so that the Wiener process is eliminated. The portfolio is

- 1: derivative
- + $\partial f / \partial S$: shares.

The holder of this portfolio is short one derivative and long an amount $\partial f/\partial S$ of shares. Define Π as the value of the portfolio. By definition

$$\Pi = -f + \frac{\partial f}{\partial S} S \quad (14.12)$$

The change $\Delta\Pi$ in the value of the portfolio in the time interval Δt is given by

$$\Delta\Pi = -\Delta f + \frac{\partial f}{\partial S} \Delta S \quad (14.13)$$

Substituting equations (14.10) and (14.11) into equation (14.13) yields

$$\Delta\Pi = \left(-\frac{\partial f}{\partial t} - \frac{1}{2} \frac{\partial^2 f}{\partial S^2} \sigma^2 S^2 \right) \Delta t \quad (14.14)$$

Because this equation does not involve Δz , the portfolio must be riskless during time Δt . The assumptions listed in the preceding section imply that the portfolio must instantaneously earn the same rate of return as other short-term risk-free securities. If it earned more than this return, arbitrageurs could make a riskless profit by borrowing money to buy the portfolio; if it earned less, they could make a riskless profit by shorting the portfolio and buying risk-free securities. It follows that

$$\Delta\Pi = r\Pi \Delta t \quad (14.15)$$

where r is the risk-free interest rate. Substituting from equations (14.12) and (14.14) into (14.15), we obtain

$$\left(\frac{\partial f}{\partial t} + \frac{1}{2} \frac{\partial^2 f}{\partial S^2} \sigma^2 S^2 \right) \Delta t = r \left(f - \frac{\partial f}{\partial S} S \right) \Delta t$$

so that

$$\frac{\partial f}{\partial t} + rS \frac{\partial f}{\partial S} + \frac{1}{2} \sigma^2 S^2 \frac{\partial^2 f}{\partial S^2} = rf \quad (14.16)$$

Equation (14.16) is the Black–Scholes–Merton differential equation. It has many solutions, corresponding to all the different derivatives that can be defined with S as the underlying variable. The particular derivative that is obtained when the equation is solved depends on the *boundary conditions* that are used. These specify the values of the derivative at the boundaries of possible values of S and t . In the case of a European call option, the key boundary condition is

$$f = \max(S - K, 0) \quad \text{when } t = T$$

In the case of a European put option, it is

$$f = \max(K - S, 0) \quad \text{when } t = T$$

One point that should be emphasized about the portfolio used in the derivation of equation (14.16) is that it is not permanently riskless. It is riskless only for an infinitesimally short period of time. As S and t change, $\partial f/\partial S$ also changes. To keep the portfolio riskless, it is therefore necessary to frequently change the relative proportions of the derivative and the stock in the portfolio.

Example 14.5

A forward contract on a non-dividend-paying stock is a derivative dependent on the stock. As such, it should satisfy equation (14.16). From equation (5.5), we know that the value of the forward contract, f , at a general time t is given in terms of the stock price S at this time by

$$f = S - Ke^{-r(T-t)}$$

where K is the delivery price. This means that

$$\frac{\partial f}{\partial t} = -rKe^{-r(T-t)}, \quad \frac{\partial f}{\partial S} = 1, \quad \frac{\partial^2 f}{\partial S^2} = 0$$

When these are substituted into the left-hand side of equation (14.16), we obtain

$$-rKe^{-r(T-t)} + rS$$

This equals rf , showing that equation (14.16) is indeed satisfied.

The Prices of Tradeable Derivatives

Any function $f(S, t)$ that is a solution of the differential equation (14.16) is the theoretical price of a derivative that could be traded. If a derivative with that price existed, it would not create any arbitrage opportunities. Conversely, if a function $f(S, t)$ does not satisfy the differential equation (14.16), it cannot be the price of a derivative without creating arbitrage opportunities for traders.

To illustrate this point, consider first the function e^S . This does not satisfy the differential equation (14.16). It is therefore not a candidate for being the price of a derivative dependent on the stock price. If an instrument whose price was always e^S existed, there would be an arbitrage opportunity. As a second example, consider the function

$$\frac{e^{(\sigma^2 - 2r)(T-t)}}{S}$$

This does satisfy the differential equation, and so is, in theory, the price of a tradeable security. (It is the price of a derivative that pays off $1/S_T$ at time T .) For other examples of tradeable derivatives, see Problems 14.11, 14.12, 14.23, and 14.28.

14.7 RISK-NEUTRAL VALUATION

We introduced risk-neutral valuation in connection with the binomial model in Chapter 12. It is without doubt the single most important tool for the analysis of derivatives. It arises from one key property of the Black–Scholes–Merton differential equation (14.16). This property is that the equation does not involve any variables that are affected by the risk preferences of investors. The variables that do appear in the equation are the current stock price, time, stock price volatility, and the risk-free rate of interest. All are independent of risk preferences.

The Black–Scholes–Merton differential equation would not be independent of risk preferences if it involved the expected return, μ , on the stock. This is because the value of μ does depend on risk preferences. The higher the level of risk aversion by investors,

the higher μ will be for any given stock. It is fortunate that μ happens to drop out in the derivation of the differential equation.

Because the Black–Scholes–Merton differential equation is independent of risk preferences, an ingenious argument can be used. If risk preferences do not enter the equation, they cannot affect its solution. Any set of risk preferences can, therefore, be used when evaluating f . In particular, the very simple assumption that all investors are risk neutral can be made.

In a world where investors are risk neutral, the expected return on all investment assets is the risk-free rate of interest, r . The reason is that risk-neutral investors do not require a premium to induce them to take risks. It is also true that the present value of any cash flow in a risk-neutral world can be obtained by discounting its expected value at the risk-free rate. The assumption that the world is risk neutral does, therefore, considerably simplify the analysis of derivatives.

Consider a derivative that provides a payoff at one particular time. It can be valued using risk-neutral valuation by using the following procedure:

1. Assume that the expected return from the underlying asset is the risk-free interest rate, r (i.e., assume $\mu = r$).
2. Calculate the expected payoff from the derivative.
3. Discount the expected payoff at the risk-free interest rate.

It is important to appreciate that risk-neutral valuation (or the assumption that all investors are risk neutral) is merely an artificial device for obtaining solutions to the Black–Scholes–Merton differential equation. The solutions that are obtained are valid in all worlds, not just those where investors are risk neutral. When we move from a risk-neutral world to a risk-averse world, two things happen. The expected growth rate in the stock price changes and the discount rate that must be used for any payoffs from the derivative changes. It happens that these two changes always offset each other exactly.

Application to Forward Contracts on a Stock

We valued forward contracts on a non-dividend-paying stock in Section 5.7. In Example 14.5, we verified that the pricing formula satisfies the Black–Scholes–Merton differential equation. In this section we derive the pricing formula from risk-neutral valuation. We make the assumption that interest rates are constant and equal to r . This is somewhat more restrictive than the assumption in Chapter 5.

Consider a long forward contract that matures at time T with delivery price, K . As indicated in Figure 1.2, the value of the contract at maturity is

$$S_T - K$$

where S_T is the stock price at time T . From the risk-neutral valuation argument, the value of the forward contract at time 0 is its expected value at time T in a risk-neutral world discounted at the risk-free rate of interest. Denoting the value of the forward contract at time zero by f , this means that

$$f = e^{-rT} \hat{E}(S_T - K)$$

where \hat{E} denotes the expected value in a risk-neutral world. Since K is a constant, this

equation becomes

$$f = e^{-rT} \hat{E}(S_T) - Ke^{-rT} \quad (14.17)$$

The expected return μ on the stock becomes r in a risk-neutral world. Hence, from equation (14.4), we have

$$\hat{E}(S_T) = S_0 e^{rT} \quad (14.18)$$

Substituting equation (14.18) into equation (14.17) gives

$$f = S_0 - Ke^{-rT} \quad (14.19)$$

This is in agreement with equation (5.5).

14.8 BLACK–SCHOLES–MERTON PRICING FORMULAS

The most famous solutions to the differential equation (14.16) are the Black–Scholes–Merton formulas for the prices of European call and put options. These formulas are:

$$c = S_0 N(d_1) - Ke^{-rT} N(d_2) \quad (14.20)$$

and

$$p = Ke^{-rT} N(-d_2) - S_0 N(-d_1) \quad (14.21)$$


where

$$d_1 = \frac{\ln(S_0/K) + (r + \sigma^2/2)T}{\sigma\sqrt{T}}$$

$$d_2 = \frac{\ln(S_0/K) + (r - \sigma^2/2)T}{\sigma\sqrt{T}} = d_1 - \sigma\sqrt{T}$$

The function $N(x)$ is the cumulative probability distribution function for a standardized normal distribution. In other words, it is the probability that a variable with a standard normal distribution, $\phi(0, 1)$, will be less than x . It is illustrated in Figure 14.3. The

Figure 14.3 Shaded area represents $N(x)$.



remaining variables should be familiar. The variables c and p are the European call and European put price, S_0 is the stock price at time zero, K is the strike price, r is the continuously compounded risk-free rate, σ is the stock price volatility, and T is the time to maturity of the option.

One way of deriving the Black–Scholes–Merton formulas is by solving the differential equation (14.16) subject to the boundary condition mentioned in Section 14.6.⁶ (See Problem 14.17 to prove that the call price in equation (14.20) satisfies the differential equation.) Another approach is to use risk-neutral valuation. Consider a European call option. The expected value of the option at maturity in a risk-neutral world is

$$\hat{E}[\max(S_T - K, 0)]$$

where, as before, \hat{E} denotes the expected value in a risk-neutral world. From the risk-neutral valuation argument, the European call option price c is this expected value discounted at the risk-free rate of interest, that is,

$$c = e^{-rT} \hat{E}[\max(S_T - K, 0)] \quad (14.22)$$

The appendix at the end of this chapter shows that this equation leads to the result in equation (14.20).

To provide an interpretation of the terms in equation (14.20), we note that it can be written

$$c = e^{-rT} [S_0 N(d_1) e^{rT} - KN(d_2)] \quad (14.23)$$

The expression $N(d_2)$ is the probability that the option will be exercised in a risk-neutral world, so that $KN(d_2)$ is the strike price times the probability that the strike price will be paid. The expression $S_0 N(d_1) e^{rT}$ is the expected value in a risk-neutral world of a variable that is equal to S_T if $S_T > K$ and to zero otherwise.

Since it is never optimal to exercise early an American call option on a non-dividend-paying stock (see Section 10.5), equation (14.20) is the value of an American call option on a non-dividend-paying stock. Unfortunately, no exact analytic formula for the value of an American put option on a non-dividend-paying stock has been produced. Numerical procedures for calculating American put values are discussed in Chapter 20.

When the Black–Scholes–Merton formula is used in practice the interest rate r is set equal to the zero-coupon risk-free interest rate for a maturity T . As we show in later chapters, this is theoretically correct when r is a known function of time. It is also theoretically correct when the interest rate is stochastic provided that the stock price at time T is lognormal and the volatility parameter is chosen appropriately. As mentioned earlier, time is normally measured as the number of trading days left in the life of the option divided by the number of trading days in 1 year.

⁶ The differential equation gives the call and put prices at a general time t . For example, the call price that satisfies the differential equation is $c = SN(d_1) - Ke^{-r(T-t)}N(d_2)$, where

$$d_1 = \frac{\ln(S/K) + (r + \sigma^2/2)(T - t)}{\sigma\sqrt{T - t}}$$

and $d_2 = d_1 - \sigma\sqrt{T - t}$.

Properties of the Black–Scholes–Merton Formulas

We now show that the Black–Scholes–Merton formulas have the right general properties by considering what happens when some of the parameters take extreme values.

When the stock price, S_0 , becomes very large, a call option is almost certain to be exercised. It then becomes very similar to a forward contract with delivery price K . From equation (5.5), we expect the call price to be

$$S_0 - Ke^{-rT}$$

This is, in fact, the call price given by equation (14.20) because, when S_0 becomes very large, both d_1 and d_2 become very large, and $N(d_1)$ and $N(d_2)$ become close to 1.0. When the stock price becomes very large, the price of a European put option, p , approaches zero. This is consistent with equation (14.21) because $N(-d_1)$ and $N(-d_2)$ are both close to zero in this case.

Consider next what happens when the volatility σ approaches zero. Because the stock is virtually riskless, its price will grow at rate r to S_0e^{rT} at time T and the payoff from a call option is

$$\max(S_0e^{rT} - K, 0)$$

Discounting at rate r , the value of the call today is

$$e^{-rT} \max(S_0e^{rT} - K, 0) = \max(S_0 - Ke^{-rT}, 0)$$

To show that this is consistent with equation (14.20), consider first the case where $S_0 > Ke^{-rT}$. This implies that $\ln(S_0/K) + rT > 0$. As σ tends to zero, d_1 and d_2 tend to $+\infty$, so that $N(d_1)$ and $N(d_2)$ tend to 1.0 and equation (14.20) becomes

$$c = S_0 - Ke^{-rT}$$

When $S_0 < Ke^{-rT}$, it follows that $\ln(S_0/K) + rT < 0$. As σ tends to zero, d_1 and d_2 tend to $-\infty$, so that $N(d_1)$ and $N(d_2)$ tend to zero and equation (14.20) gives a call price of zero. The call price is therefore $\max(S_0 - Ke^{-rT}, 0)$ as σ tends to zero. Similarly, it can be shown that the put price is always $\max(Ke^{-rT} - S_0, 0)$ as σ tends to zero.

14.9 CUMULATIVE NORMAL DISTRIBUTION FUNCTION

When implementing equations (14.20) and (14.21), it is necessary to evaluate the cumulative normal distribution function $N(x)$. Tables for $N(x)$ are provided at the end of this book. The NORMSDIST function in Excel also provides a convenient way of calculating $N(x)$.

Example 14.6

The stock price 6 months from the expiration of an option is \$42, the exercise price of the option is \$40, the risk-free interest rate is 10% per annum, and the volatility

is 20% per annum. This means that $S_0 = 42$, $K = 40$, $r = 0.1$, $\sigma = 0.2$, $T = 0.5$,

$$d_1 = \frac{\ln(42/40) + (0.1 + 0.2^2/2) \times 0.5}{0.2\sqrt{0.5}} = 0.7693$$

$$d_2 = \frac{\ln(42/40) + (0.1 - 0.2^2/2) \times 0.5}{0.2\sqrt{0.5}} = 0.6278$$

and

$$Ke^{-rT} = 40e^{-0.05} = 38.049$$

Hence, if the option is a European call, its value c is given by

$$c = 42N(0.7693) - 38.049N(0.6278)$$

If the option is a European put, its value p is given by

$$p = 38.049N(-0.6278) - 42N(-0.7693)$$

Using the NORMSDIST function in Excel gives

$$N(0.7693) = 0.7791, \quad N(-0.7693) = 0.2209$$

$$N(0.6278) = 0.7349, \quad N(-0.6278) = 0.2651$$

so that

$$c = 4.76, \quad p = 0.81$$

Ignoring the time value of money, the stock price has to rise by \$2.76 for the purchaser of the call to break even. Similarly, the stock price has to fall by \$2.81 for the purchaser of the put to break even.

14.10 WARRANTS AND EMPLOYEE STOCK OPTIONS

The exercise of a regular call option on a company has no effect on the number of the company's shares outstanding. If the writer of the option does not own the company's shares, he or she must buy them in the market in the usual way and then sell them to the option holder for the strike price. As explained in Chapter 9, warrants and employee stock options are different from regular call options in that exercise leads to the company issuing more shares and then selling them to the option holder for the strike price. As the strike price is less than the market price, this dilutes the interest of the existing shareholders.

How should potential dilution affect the way we value outstanding warrants and employee stock options? The answer is that it should not! Assuming markets are efficient the stock price will reflect potential dilution from all outstanding warrants and employee stock options. This is explained in Business Snapshot 14.3.⁷

Consider next the situation a company is in when it is contemplating a new issue of warrants (or employee stock options). We suppose that the company is interested in

⁷ Analysts sometimes assume that the sum of the values of the warrants and the equity (rather than just the value of the equity) is lognormal. The result is a Black-Scholes type of equation for the value of the warrant in terms of the value of the warrant. See Technical Note 3 at www.rotman.utoronto.ca/~hull/TechnicalNotes for an explanation of this model.

Business Snapshot 14.3 Warrants, Employee Stock Options, and Dilution

Consider a company with 100,000 shares each worth \$50. It surprises the market with an announcement that it is granting 100,000 stock options to its employees with a strike price of \$50. If the market sees little benefit to the shareholders from the employee stock options in the form of reduced salaries and more highly motivated managers, the stock price will decline immediately after the announcement of the employee stock options. If the stock price declines to \$45, the dilution cost to the current shareholders is \$5 per share or \$500,000 in total.

Suppose that the company does well so that by the end of three years the share price is \$100. Suppose further that all the options are exercised at this point. The payoff to the employees is \$50 per option. It is tempting to argue that there will be further dilution in that 100,000 shares worth \$100 per share are now merged with 100,000 shares for which only \$50 is paid, so that (a) the share price reduces to \$75 and (b) the payoff to the option holders is only \$25 per option. However, this argument is flawed. The exercise of the options is anticipated by the market and already reflected in the share price. The payoff from each option exercised is \$50.

This example illustrates the general point that when markets are efficient the impact of dilution from executive stock options or warrants is reflected in the stock price as soon as they are announced and does not need to be taken into account again when the options are valued.

calculating the cost of the issue assuming that there are no compensating benefits. We assume that the company has N shares worth S_0 each and the number of new options contemplated is M , with each option giving the holder the right to buy one share for K . The value of the company today is NS_0 . This value does not change as a result of the warrant issue. Suppose that without the warrant issue the share price will be S_T at the warrant's maturity. This means that (with or without the warrant issue) the total value of the equity and the warrants at time T will NS_T . If the warrants are exercised, there is a cash inflow from the strike price increasing this to $NS_T + MK$. This value is distributed among $N + M$ shares, so that the share price immediately after exercise becomes

$$\frac{NS_T + MK}{N + M}$$

Therefore the payoff to an option holder if the option is exercised is

$$\frac{NS_T + MK}{N + M} - K$$

or

$$\frac{N}{N + M}(S_T - K)$$

This shows that the value of each option is the value of

$$\frac{N}{N + M}$$

regular call options on the company's stock. Therefore the total cost of the options is M times this. Since we are assuming that there are no benefits to the company from the

warrant issue, the total value of the company's equity will decline by the total cost of the options as soon as the decision to issue the warrants becomes generally known. This means that the reduction in the stock price is

$$\frac{M}{N + M}$$

times the value of a regular call option with strike price K and maturity T .

Example 14.7

A company with 1 million shares worth \$40 each is considering issuing 200,000 warrants each giving the holder the right to buy one share with a strike price of \$60 in 5 years. It wants to know the cost of this. The interest rate is 3% per annum, and the volatility is 30% per annum. The company pays no dividends. From equation (14.20), the value of a 5-year European call option on the stock is \$7.04. In this case, $N = 1,000,000$ and $M = 200,000$, so that the value of each warrant is

$$\frac{1,000,000}{1,000,000 + 200,000} \times 7.04 = 5.87$$

or \$5.87. The total cost of the warrant issue is $200,000 \times 5.87 = \$1.17$ million. Assuming the market perceives no benefits from the warrant issue, we expect the stock price to decline by \$1.17 to \$38.83.

14.11 IMPLIED VOLATILITIES

The one parameter in the Black–Scholes–Merton pricing formulas that cannot be directly observed is the volatility of the stock price. In Section 14.4, we discussed how this can be estimated from a history of the stock price. In practice, traders usually work with what are known as *implied volatilities*. These are the volatilities implied by option prices observed in the market.⁸

To illustrate how implied volatilities are calculated, suppose that the value of a European call option on a non-dividend-paying stock is 1.875 when $S_0 = 21$, $K = 20$, $r = 0.1$, and $T = 0.25$. The implied volatility is the value of σ that, when substituted into equation (14.20), gives $c = 1.875$. Unfortunately, it is not possible to invert equation (14.20) so that σ is expressed as a function of S_0 , K , r , T , and c . However, an iterative search procedure can be used to find the implied σ . For example, we can start by trying $\sigma = 0.20$. This gives a value of c equal to 1.76, which is too low. Because c is an increasing function of σ , a higher value of σ is required. We can next try a value of 0.30 for σ . This gives a value of c equal to 2.10, which is too high and means that σ must lie between 0.20 and 0.30. Next, a value of 0.25 can be tried for σ . This also proves to be too high, showing that σ lies between 0.20 and 0.25. Proceeding in this way, we can halve the range for σ at each iteration and the correct value of σ can be calculated to any required accuracy.⁹ In this example, the implied volatility is 0.235, or 23.5%, per

⁸ Implied volatilities for European and American options can be calculated using DerivaGem.

⁹ This method is presented for illustration. Other more powerful methods, such as the Newton–Raphson method, are often used in practice (see footnote 3 of Chapter 4). DerivaGem can be used to calculate implied volatilities.

annum. A similar procedure can be used in conjunction with binomial trees to find implied volatilities for American options.

Implied volatilities are used to monitor the market's opinion about the volatility of a particular stock. Whereas historical volatilities (see Section 14.4) are backward looking, implied volatilities are forward looking. Traders often quote the implied volatility of an option rather than its price. This is convenient because the implied volatility tends to be less variable than the option price. As will be explained in Chapter 19, the implied volatilities of actively traded options are used by traders to estimate appropriate implied volatilities for other options.

The VIX Index


The CBOE publishes indices of implied volatility. The most popular index, the SPX VIX, is an index of the implied volatility of 30-day options on the S&P 500 calculated from a wide range of calls and puts.¹⁰ Information on the way the index is calculated is in Section 25.15. Trading in futures on the VIX started in 2004 and trading in options on the VIX started in 2006. A trade involving futures or options on the S&P 500 is a bet on both the future level of the S&P 500 and the volatility of the S&P 500. By contrast, a futures or options contract on the VIX is a bet only on volatility. One contract is on 1,000 times the index.

Example 14.8

Suppose that a trader buys an April futures contract on the VIX when the futures price is 18.5 (corresponding to a 30-day S&P 500 volatility of 18.5%) and closes out the contract when the futures price is 19.3 (corresponding to an S&P 500 volatility of 19.3%). The trader makes a gain of \$800.

Figure 14.4 shows the VIX index between January 2004 and July 2010. Between 2004 and mid-2007 it tended to stay between 10 and 20. It reached 30 during the second half

Figure 14.4 The VIX index, January 2004 to July 2010.



¹⁰ Similarly, the VXN is an index of the volatility of the NASDAQ 100 index and the VXD is an index of the volatility of the Dow Jones Industrial Average.

of 2007 and a record 80 in October and November 2008 after Lehman's bankruptcy. By early 2010, it had declined to a more normal levels, but in May 2010 it spiked at over 45 because of the European sovereign debt crisis.

14.12 DIVIDENDS

Up to now, we have assumed that the stock on which the option is written pays no dividends. In this section, we modify the Black–Scholes–Merton model to take account of dividends. We assume that the amount and timing of the dividends during the life of an option can be predicted with certainty. When options last for relatively short periods of time, this assumption is not too unreasonable. (For long-life options it is usual to assume that the dividend yield rather the cash dividend payments are known. Options can then be valued as will be described in the Chapter 16.) The date on which the dividend is paid should be assumed to be the ex-dividend date. On this date the stock price declines by the amount of the dividend.¹¹

European Options

European options can be analyzed by assuming that the stock price is the sum of two components: a riskless component that corresponds to the known dividends during the life of the option and a risky component. The riskless component, at any given time, is the present value of all the dividends during the life of the option discounted from the ex-dividend dates to the present at the risk-free rate. By the time the option matures, the dividends will have been paid and the riskless component will no longer exist. The Black–Scholes–Merton formula is therefore correct if S_0 is equal to the risky component of the stock price and σ is the volatility of the process followed by the risky component.¹²

Operationally, this means that the Black–Scholes–Merton formulas can be used provided that the stock price is reduced by the present value of all the dividends during the life of the option, the discounting being done from the ex-dividend dates at the risk-free rate. As already mentioned, a dividend is counted as being during the life of the option only if its ex-dividend date occurs during the life of the option.

Example 14.9

Consider a European call option on a stock when there are ex-dividend dates in two months and five months. The dividend on each ex-dividend date is expected to be \$0.50. The current share price is \$40, the exercise price is \$40, the stock price

¹¹ For tax reasons the stock price may go down by somewhat less than the cash amount of the dividend. To take account of this phenomenon, we need to interpret the word 'dividend' in the context of option pricing as the reduction in the stock price on the ex-dividend date caused by the dividend. Thus, if a dividend of \$1 per share is anticipated and the share price normally goes down by 80% of the dividend on the ex-dividend date, the dividend should be assumed to be \$0.80 for the purposes of the analysis.

¹² In theory, this is not quite the same as the volatility of the stochastic process followed by the whole stock price. The volatility of the risky component is approximately equal to the volatility of the whole stock price multiplied by $S_0/(S_0 - D)$, where D is the present value of the dividends. However, an adjustment is only necessary when volatilities are estimated using historical data. An implied volatility is calculated after the present value of dividends have been subtracted from the stock price and is the volatility of the risky component.

volatility is 30% per annum, the risk-free rate of interest is 9% per annum, and the time to maturity is six months. The present value of the dividends is

$$0.5e^{-0.09 \times 2/12} + 0.5e^{-0.09 \times 5/12} = 0.9742$$

The option price can therefore be calculated from the Black–Scholes–Merton formula, with $S_0 = 40 - 0.9742 = 39.0258$, $K = 40$, $r = 0.09$, $\sigma = 0.3$, and $T = 0.5$:

$$d_1 = \frac{\ln(39.0258/40) + (0.09 + 0.3^2/2) \times 0.5}{0.3\sqrt{0.5}} = 0.2020$$

$$d_2 = \frac{\ln(39.0258/40) + (0.09 - 0.3^2/2) \times 0.5}{0.3\sqrt{0.5}} = -0.0102$$

Using the NORMSDIST function in Excel gives

$$N(d_1) = 0.5800, \quad N(d_2) = 0.4959$$

and, from equation (14.20), the call price is

$$39.0258 \times 0.5800 - 40e^{-0.09 \times 0.5} \times 0.4959 = 3.67$$

or \$3.67.

American Options

Consider next American call options. Chapter 10 showed that in the absence of dividends American options should never be exercised early. An extension to the argument shows that, when there are dividends, it can only be optimal to exercise at a time immediately before the stock goes ex-dividend. We assume that n ex-dividend dates are anticipated and that they are at times t_1, t_2, \dots, t_n , with $t_1 < t_2 < \dots < t_n$. The dividends corresponding to these times will be denoted by D_1, D_2, \dots, D_n , respectively.

We start by considering the possibility of early exercise just prior to the final ex-dividend date (i.e., at time t_n). If the option is exercised at time t_n , the investor receives

$$S(t_n) - K$$

where $S(t)$ denotes the stock price at time t . If the option is not exercised, the stock price drops to $S(t_n) - D_n$. As shown by equation (10.4), the value of the option is then greater than

$$S(t_n) - D_n - Ke^{-r(T-t_n)}$$

It follows that, if

$$S(t_n) - D_n - Ke^{-r(T-t_n)} \geq S(t_n) - K$$

that is,

$$D_n \leq K[1 - e^{-r(T-t_n)}] \tag{14.24}$$

it cannot be optimal to exercise at time t_n . On the other hand, if

$$D_n > K[1 - e^{-r(T-t_n)}] \tag{14.25}$$

for any reasonable assumption about the stochastic process followed by the stock price, it can be shown that it is always optimal to exercise at time t_n for a sufficiently high

value of $S(t_n)$. The inequality in (14.25) will tend to be satisfied when the final ex-dividend date is fairly close to the maturity of the option (i.e., $T - t_n$ is small) and the dividend is large.

Consider next time t_{n-1} , the penultimate ex-dividend date. If the option is exercised immediately prior to time t_{n-1} , the investor receives $S(t_{n-1}) - K$. If the option is not exercised at time t_{n-1} , the stock price drops to $S(t_{n-1}) - D_{n-1}$ and the earliest subsequent time at which exercise could take place is t_n . Hence, from equation (10.4), a lower bound to the option price if it is not exercised at time t_{n-1} is

$$S(t_{n-1}) - D_{n-1} - Ke^{-r(t_n - t_{n-1})}$$

It follows that if

$$S(t_{n-1}) - D_{n-1} - Ke^{-r(t_n - t_{n-1})} \geq S(t_{n-1}) - K$$

or

$$D_{n-1} \leq K[1 - e^{-r(t_n - t_{n-1})}]$$

it is not optimal to exercise immediately prior to time t_{n-1} . Similarly, for any $i < n$, if

$$D_i \leq K[1 - e^{-r(t_{i+1} - t_i)}] \quad (14.26)$$

it is not optimal to exercise immediately prior to time t_i .

The inequality in (14.26) is approximately equivalent to

$$D_i \leq Kr(t_{i+1} - t_i)$$

Assuming that K is fairly close to the current stock price, this inequality is satisfied when the dividend yield on the stock is less than the risk-free rate of interest. This is often the case.

We can conclude from this analysis that, in many circumstances, the most likely time for the early exercise of an American call is immediately before the final ex-dividend date, t_n . Furthermore, if inequality (14.26) holds for $i = 1, 2, \dots, n - 1$ and inequality (14.24) holds, we can be certain that early exercise is never optimal.

Black's Approximation

Black suggests an approximate procedure for taking account of early exercise in call options.¹³ This involves calculating, as described earlier in this section, the prices of European options that mature at times T and t_n , and then setting the American price equal to the greater of the two. This approximation seems to work well in most cases.¹⁴

Example 14.10

Consider the situation in Example 14.9, but suppose that the option is American rather than European. In this case $D_1 = D_2 = 0.5$, $S_0 = 40$, $K = 40$, $r = 0.09$,

¹³ See F. Black, "Fact and Fantasy in the Use of Options," *Financial Analysts Journal*, 31 (July/August 1975): 36–41, 61–72.

¹⁴ For an exact formula, suggested by Roll, Geske, and Whaley, for valuing American calls when there is only one ex-dividend date, see Technical Note 4 at www.rotman.utoronto.ca/~hull/TechnicalNotes. This involves the cumulative bivariate normal distribution function. A procedure for calculating this function is given in Technical Note 5 in the same place.

$t_1 = 2/12$, and $t_2 = 5/12$. Since

$$K[1 - e^{-r(t_2-t_1)}] = 40(1 - e^{-0.09 \times 0.25}) = 0.89$$

is greater than 0.5, it follows (see inequality (14.26)) that the option should never be exercised immediately before the first ex-dividend date. In addition, since

$$K[1 - e^{-r(T-t_2)}] = 40(1 - e^{-0.09 \times 0.0833}) = 0.30$$

is less than 0.5, it follows (see inequality (14.25)) that, when it is sufficiently deep in the money, the option should be exercised immediately before the second ex-dividend date.

We now use Black's approximation to value the option. The present value of the first dividend is

$$0.5e^{-0.1667 \times 0.09} = 0.4926$$

so that the value of the option, on the assumption that it expires just before the final ex-dividend date, can be calculated using the Black–Scholes–Merton formula with $S_0 = 40 - 0.4926 = 39.5074$, $K = 40$, $r = 0.09$, $\sigma = 0.30$, and $T = 0.4167$. It is \$3.52. Black's approximation involves taking the greater of this and the value of the option when it can only be exercised at the end of 6 months. From Example 14.9, we know that the latter is \$3.67. Black's approximation, therefore, gives the value of the American call as \$3.67.

The option can be valued using a binomial tree, as will be described in Section 20.3. As shown by DerivaGem, this approach with 500 time steps gives \$3.72 as the value. (Note that DerivaGem requires dividends to be input in chronological order in the table; the time to a dividend is in the first column and the amount of the dividend is in the second column.) There are two reasons for differences between the Binomial Model (BM) and Black's approximation (BA). The first concerns the timing of the early exercise decision; the second concerns the way volatility is applied. The timing of the early exercise decision tends to make BM greater than BA. In BA, the assumption is that the holder has to decide today whether the option will be exercised after 5 months or after 6 months; BM allows the decision on early exercise at the 5-month point to depend on the stock price at that time. The way in which volatility is applied tends to make BA greater than BM. In BA, when we assume exercise takes place after 5 months, the volatility is applied to the stock price less the present value of the first dividend; when we assume exercise takes place after 6 months, the volatility is applied to the stock price less the present value of both dividends.

SUMMARY

We started this chapter by examining the properties of the process for stock prices introduced in Chapter 13. The process implies that the price of a stock at some future time, given its price today, is lognormal. It also implies that the continuously compounded return from the stock in a period of time is normally distributed. Our uncertainty about future stock prices increases as we look further ahead. The standard deviation of the logarithm of the stock price is proportional to the square root of how far ahead we are looking.

To estimate the volatility σ of a stock price empirically, the stock price is observed at fixed intervals of time (e.g., every day, every week, or every month). For each time period, the natural logarithm of the ratio of the stock price at the end of the time period to the stock price at the beginning of the time period is calculated. The volatility is estimated as the standard deviation of these numbers divided by the square root of the length of the time period in years. Usually, days when the exchanges are closed are ignored in measuring time for the purposes of volatility calculations.

The differential equation for the price of any derivative dependent on a stock can be obtained by creating a riskless portfolio of the option and the stock. Because the derivative and the stock price both depend on the same underlying source of uncertainty, this can always be done. The portfolio that is created remains riskless for only a very short period of time. However, the return on a riskless portfolio must always be the risk-free interest rate if there are to be no arbitrage opportunities.

The expected return on the stock does not enter into the Black–Scholes–Merton differential equation. This leads to a useful result known as risk-neutral valuation. This result states that when valuing a derivative dependent on a stock price, we can assume that the world is risk neutral. This means that we can assume that the expected return from the stock is the risk-free interest rate, and then discount expected payoffs at the risk-free interest rate. The Black–Scholes–Merton equations for European call and put options can be derived by either solving their differential equation or by using risk-neutral valuation.

An implied volatility is the volatility that, when used in conjunction with the Black–Scholes–Merton option pricing formula, gives the market price of the option. Traders monitor implied volatilities. They often quote the implied volatility of an option rather than its price. They have developed procedures for using the volatilities implied by the prices of actively traded options to estimate volatilities for other options.

The Black–Scholes–Merton results can be extended to cover European call and put options on dividend-paying stocks. The procedure is to use the Black–Scholes–Merton formula with the stock price reduced by the present value of the dividends anticipated during the life of the option, and the volatility equal to the volatility of the stock price net of the present value of these dividends.

In theory, it can be optimal to exercise American call options immediately before any ex-dividend date. In practice, it is often only necessary to consider the final ex-dividend date. Fischer Black has suggested an approximation. This involves setting the American call option price equal to the greater of two European call option prices. The first European call option expires at the same time as the American call option; the second expires immediately prior to the final ex-dividend date.

FURTHER READING

On the Distribution of Stock Price Changes

Blattberg, R., and N. Gonedes, “A Comparison of the Stable and Student Distributions as Statistical Models for Stock Prices,” *Journal of Business*, 47 (April 1974): 244–80.

Fama, E. F., “The Behavior of Stock Market Prices,” *Journal of Business*, 38 (January 1965): 34–105.

Kon, S. J., “Models of Stock Returns—A Comparison,” *Journal of Finance*, 39 (March 1984): 147–65.

Richardson, M., and T. Smith, “A Test for Multivariate Normality in Stock Returns,” *Journal of Business*, 66 (1993): 295–321.

On the Black–Scholes–Merton Analysis

Black, F. “Fact and Fantasy in the Use of Options and Corporate Liabilities,” *Financial Analysts Journal*, 31 (July/August 1975): 36–41, 61–72.

Black, F. “How We Came Up with the Option Pricing Formula,” *Journal of Portfolio Management*, 15, 2 (1989): 4–8.

Black, F., and M. Scholes, “The Pricing of Options and Corporate Liabilities,” *Journal of Political Economy*, 81 (May/June 1973): 637–59.

Merton, R. C., “Theory of Rational Option Pricing,” *Bell Journal of Economics and Management Science*, 4 (Spring 1973): 141–83.

On Risk-Neutral Valuation

Cox, J. C., and S. A. Ross, “The Valuation of Options for Alternative Stochastic Processes,” *Journal of Financial Economics*, 3 (1976): 145–66.

Smith, C. W., “Option Pricing: A Review,” *Journal of Financial Economics*, 3 (1976): 3–54.

On the Causes of Volatility

Fama, E. F. “The Behavior of Stock Market Prices,” *Journal of Business*, 38 (January 1965): 34–105.

French, K. R. “Stock Returns and the Weekend Effect,” *Journal of Financial Economics*, 8 (March 1980): 55–69.

French, K. R., and R. Roll “Stock Return Variances: The Arrival of Information and the Reaction of Traders,” *Journal of Financial Economics*, 17 (September 1986): 5–26.

Roll R. “Orange Juice and Weather,” *American Economic Review*, 74, 5 (December 1984): 861–80.

Practice Questions (Answers in Solutions Manual)

- 14.1. What does the Black–Scholes–Merton stock option pricing model assume about the probability distribution of the stock price in one year? What does it assume about the continuously compounded rate of return on the stock during the year?
- 14.2. The volatility of a stock price is 30% per annum. What is the standard deviation of the percentage price change in one trading day?
- 14.3. Explain the principle of risk-neutral valuation.
- 14.4. Calculate the price of a 3-month European put option on a non-dividend-paying stock with a strike price of \$50 when the current stock price is \$50, the risk-free interest rate is 10% per annum, and the volatility is 30% per annum.
- 14.5. What difference does it make to your calculations in Problem 14.4 if a dividend of \$1.50 is expected in 2 months?
- 14.6. What is *implied volatility*? How can it be calculated?
- 14.7. A stock price is currently \$40. Assume that the expected return from the stock is 15% and that its volatility is 25%. What is the probability distribution for the rate of return (with continuous compounding) earned over a 2-year period?
- 14.8. A stock price follows geometric Brownian motion with an expected return of 16% and a volatility of 35%. The current price is \$38.
 - (a) What is the probability that a European call option on the stock with an exercise price of \$40 and a maturity date in 6 months will be exercised?

- (b) What is the probability that a European put option on the stock with the same exercise price and maturity will be exercised?
- 14.9. Using the notation in this chapter, prove that a 95% confidence interval for S_T is between

$$S_0 e^{(\mu - \sigma^2/2)T - 1.96\sigma\sqrt{T}} \quad \text{and} \quad S_0 e^{(\mu - \sigma^2/2)T + 1.96\sigma\sqrt{T}}$$

- 14.10. A portfolio manager announces that the average of the returns realized in each year of the last 10 years is 20% per annum. In what respect is this statement misleading?
- 14.11. Assume that a non-dividend-paying stock has an expected return of μ and a volatility of σ . An innovative financial institution has just announced that it will trade a security that pays off a dollar amount equal to $\ln S_T$ at time T , where S_T denotes the value of the stock price at time T .

- (a) Use risk-neutral valuation to calculate the price of the security at time t in terms of the stock price, S , at time t .
- (b) Confirm that your price satisfies the differential equation (14.16).
- 14.12. Consider a derivative that pays off S_T^n at time T , where S_T is the stock price at that time. When the stock price follows geometric Brownian motion, it can be shown that its price at time t ($t \leq T$) has the form

$$h(t, T)S^n$$

where S is the stock price at time t and h is a function only of t and T .

- (a) By substituting into the Black–Scholes–Merton partial differential equation, derive an ordinary differential equation satisfied by $h(t, T)$.
- (b) What is the boundary condition for the differential equation for $h(t, T)$?
- (c) Show that

$$h(t, T) = e^{[0.5\sigma^2n(n-1) + r(n-1)](T-t)}$$

where r is the risk-free interest rate and σ is the stock price volatility.

- 14.13. What is the price of a European call option on a non-dividend-paying stock when the stock price is \$52, the strike price is \$50, the risk-free interest rate is 12% per annum, the volatility is 30% per annum, and the time to maturity is 3 months?
- 14.14. What is the price of a European put option on a non-dividend-paying stock when the stock price is \$69, the strike price is \$70, the risk-free interest rate is 5% per annum, the volatility is 35% per annum, and the time to maturity is 6 months?
- 14.15. Consider an American call option on a stock. The stock price is \$70, the time to maturity is 8 months, the risk-free rate of interest is 10% per annum, the exercise price is \$65, and the volatility is 32%. A dividend of \$1 is expected after 3 months and again after 6 months. Show that it can never be optimal to exercise the option on either of the two dividend dates. Use DerivaGem to calculate the price of the option.
- 14.16. A call option on a non-dividend-paying stock has a market price of $2\frac{1}{2}$. The stock price is \$15, the exercise price is \$13, the time to maturity is 3 months, and the risk-free interest rate is 5% per annum. What is the implied volatility?
- 14.17. With the notation used in this chapter:

- (a) What is $N'(x)$?
- (b) Show that $SN'(d_1) = Ke^{-r(T-t)}N'(d_2)$, where S is the stock price at time t and

$$d_1 = \frac{\ln(S/K) + (r + \sigma^2/2)(T-t)}{\sigma\sqrt{T-t}}, \quad d_2 = \frac{\ln(S/K) + (r - \sigma^2/2)(T-t)}{\sigma\sqrt{T-t}}$$

- (c) Calculate $\partial d_1/\partial S$ and $\partial d_2/\partial S$.
 (d) Show that when

$$c = SN(d_1) - Ke^{-r(T-t)}N(d_2)$$

it follows that

$$\frac{\partial c}{\partial t} = -rKe^{-r(T-t)}N(d_2) - SN'(d_1)\frac{\sigma}{2\sqrt{T-t}}$$

where c is the price of a call option on a non-dividend-paying stock.

- (e) Show that $\partial c/\partial S = N(d_1)$.
 (f) Show that c satisfies the Black–Scholes–Merton differential equation.
 (g) Show that c satisfies the boundary condition for a European call option, i.e., that $c = \max(S - K, 0)$ as $t \rightarrow T$.
- 14.18. Show that the Black–Scholes–Merton formulas for call and put options satisfy put–call parity.
- 14.19. A stock price is currently \$50 and the risk-free interest rate is 5%. Use the DerivaGem software to translate the following table of European call options on the stock into a table of implied volatilities, assuming no dividends. Are the option prices consistent with the assumptions underlying Black–Scholes–Merton?

| Strike price (\$) | Maturity (months) | | |
|-------------------|-------------------|-----|------|
| | 3 | 6 | 12 |
| 45 | 7.0 | 8.3 | 10.5 |
| 50 | 3.7 | 5.2 | 7.5 |
| 55 | 1.6 | 2.9 | 5.1 |

- 14.20. Explain carefully why Black’s approach to evaluating an American call option on a dividend-paying stock may give an approximate answer even when only one dividend is anticipated. Does the answer given by Black’s approach understate or overstate the true option value? Explain your answer.
- 14.21. Consider an American call option on a stock. The stock price is \$50, the time to maturity is 15 months, the risk-free rate of interest is 8% per annum, the exercise price is \$55, and the volatility is 25%. Dividends of \$1.50 are expected in 4 months and 10 months. Show that it can never be optimal to exercise the option on either of the two dividend dates. Calculate the price of the option.
- 14.22. Show that the probability that a European call option will be exercised in a risk-neutral world is, with the notation introduced in this chapter, $N(d_2)$. What is an expression for the value of a derivative that pays off \$100 if the price of a stock at time T is greater than K ?
- 14.23. Show that S^{-2r/σ^2} could be the price of a traded derivative security.
- 14.24. A company has an issue of executive stock options outstanding. Should dilution be taken into account when the options are valued? Explain your answer.
- 14.25. A company’s stock price is \$50 and 10 million shares are outstanding. The company is considering giving its employees 3 million at-the-money 5-year call options. Option exercises will be handled by issuing more shares. The stock price volatility is 25%, the 5-year risk-free rate is 5%, and the company does not pay dividends. Estimate the cost to the company of the employee stock option issue.

Further Questions

- 14.26. A stock price is currently \$50. Assume that the expected return from the stock is 18% and its volatility is 30%. What is the probability distribution for the stock price in 2 years? Calculate the mean and standard deviation of the distribution. Determine the 95% confidence interval.
- 14.27. Suppose that observations on a stock price (in dollars) at the end of each of 15 consecutive weeks are as follows:
30.2, 32.0, 31.1, 30.1, 30.2, 30.3, 30.6, 33.0, 32.9, 33.0, 33.5, 33.5, 33.7, 33.5, 33.2
Estimate the stock price volatility. What is the standard error of your estimate?
- 14.28. A financial institution plans to offer a security that pays off a dollar amount equal to S_T^2 at time T , where S_T is the price at time T of a stock that pays no dividends.
(a) Use risk-neutral valuation to calculate the price of the security at time t in terms of the stock price S at time t . (*Hint*: The expected value of S_T^2 can be calculated from the mean and variance of S_T given in Section 14.1.)
(b) Confirm that your price satisfies the differential equation (14.16).
- 14.29. Consider an option on a non-dividend-paying stock when the stock price is \$30, the exercise price is \$29, the risk-free interest rate is 5%, the volatility is 25% per annum, and the time to maturity is 4 months.
(a) What is the price of the option if it is a European call?
(b) What is the price of the option if it is an American call?
(c) What is the price of the option if it is a European put?
(d) Verify that put–call parity holds.
- 14.30. Assume that the stock in Problem 14.29 is due to go ex-dividend in $1\frac{1}{2}$ months. The expected dividend is 50 cents.
(a) What is the price of the option if it is a European call?
(b) What is the price of the option if it is a European put?
(c) If the option is an American call, are there any circumstances under which it will be exercised early?
- 14.31. Consider an American call option when the stock price is \$18, the exercise price is \$20, the time to maturity is 6 months, the volatility is 30% per annum, and the risk-free interest rate is 10% per annum. Two equal dividends are expected during the life of the option with ex-dividend dates at the end of 2 months and 5 months. Assume the dividends are 40 cents. Use Black's approximation and the DerivaGem software to value the option. How high can the dividends be without the American option being worth more than the corresponding European option?

APPENDIX

PROOF OF THE BLACK–SCHOLES–MERTON FORMULA USING RISK-NEUTRAL VALUATION

We will prove the Black–Scholes result by first proving another key result that will also be useful in future chapters.

Key Result

If V is lognormally distributed and the standard deviation of $\ln V$ is w , then

$$E[\max(V - K, 0)] = E(V)N(d_1) - KN(d_2) \quad (14A.1)$$

where

$$d_1 = \frac{\ln[E(V)/K] + w^2/2}{w}$$

$$d_2 = \frac{\ln[E(V)/K] - w^2/2}{w}$$

and E denotes the expected value.

Proof of Key Result

Define $g(V)$ as the probability density function of V . It follows that

$$E[\max(V - K, 0)] = \int_K^\infty (V - K)g(V) dV \quad (14A.2)$$

The variable $\ln V$ is normally distributed with standard deviation w . From the properties of the lognormal distribution, the mean of $\ln V$ is m , where¹⁵

$$m = \ln[E(V)] - w^2/2 \quad (14A.3)$$

Define a new variable

$$Q = \frac{\ln V - m}{w} \quad (14A.4)$$

This variable is normally distributed with a mean of zero and a standard deviation of 1.0. Denote the density function for Q by $h(Q)$ so that

$$h(Q) = \frac{1}{\sqrt{2\pi}} e^{-Q^2/2}$$

Using equation (14A.4) to convert the expression on the right-hand side of equation (14A.2) from an integral over V to an integral over Q , we get

$$E[\max(V - K, 0)] = \int_{(\ln K - m)/w}^\infty (e^{Qw+m} - K)h(Q) dQ$$

or

$$E[\max(V - K, 0)] = \int_{(\ln K - m)/w}^\infty e^{Qw+m} h(Q) dQ - K \int_{(\ln K - m)/w}^\infty h(Q) dQ \quad (14A.5)$$

¹⁵ For a proof of this, see Technical Note 2 at www.rotman.utoronto.ca/~hull/TechnicalNotes.

Now

$$\begin{aligned}
 e^{Qw+m}h(Q) &= \frac{1}{\sqrt{2\pi}}e^{(-Q^2+2Qw+2m)/2} \\
 &= \frac{1}{\sqrt{2\pi}}e^{[-(Q-w)^2+2m+w^2]/2} \\
 &= \frac{e^{m+w^2/2}}{\sqrt{2\pi}}e^{[-(Q-w)^2]/2} \\
 &= e^{m+w^2/2}h(Q-w)
 \end{aligned}$$

This means that equation (14A.5) becomes

$$E[\max(V - K, 0)] = e^{m+w^2/2} \int_{(\ln K - m)/w}^{\infty} h(Q - w)dQ - K \int_{(\ln K - m)/w}^{\infty} h(Q)dQ \quad (14A.6)$$

If we define $N(x)$ as the probability that a variable with a mean of zero and a standard deviation of 1.0 is less than x , the first integral in equation (14A.6) is

$$1 - N[(\ln K - m)/w - w]$$

or

$$N[(-\ln K + m)/w + w]$$

Substituting for m from equation (14A.3) leads to

$$N\left(\frac{\ln[E(V)/K] + w^2/2}{w}\right) = N(d_1)$$

Similarly the second integral in equation (14A.6) is $N(d_2)$. Equation (14A.6), therefore, becomes

$$E[\max(V - K, 0)] = e^{m+w^2/2}N(d_1) - KN(d_2)$$

Substituting for m from equation (14A.3) gives the key result.

The Black–Scholes–Merton Result

We now consider a call option on a non-dividend-paying stock maturing at time T . The strike price is K , the risk-free rate is r , the current stock price is S_0 , and the volatility is σ . As shown in equation (14.22), the call price c is given by

$$c = e^{-rT} \hat{E}[\max(S_T - K, 0)] \quad (14A.7)$$

where S_T is the stock price at time T and \hat{E} denotes the expectation in a risk-neutral world. Under the stochastic process assumed by Black–Scholes–Merton, S_T is log-normal. Also, from equations (14.3) and (14.4), $\hat{E}(S_T) = S_0e^{rT}$ and the standard deviation of $\ln S_T$ is $\sigma\sqrt{T}$.

From the key result just proved, equation (14A.7) implies

$$c = e^{-rT}[S_0e^{rT}N(d_1) - KN(d_2)]$$

or

$$c = S_0N(d_1) - Ke^{-rT}N(d_2)$$

where

$$d_1 = \frac{\ln[\hat{E}(S_T)/K] + \sigma^2 T/2}{\sigma\sqrt{T}} = \frac{\ln(S_0/K) + (r + \sigma^2/2)T}{\sigma\sqrt{T}}$$

and

$$d_2 = \frac{\ln[\hat{E}(S_T)/K] - \sigma^2 T/2}{\sigma\sqrt{T}} = \frac{\ln(S_0/K) + (r - \sigma^2/2)T}{\sigma\sqrt{T}}$$

This is the Black–Scholes–Merton result.

15

CHAPTER

Employee Stock Options



Employee stock options are call options on a company's stock granted by the company to its employees. The options give the employees a stake in the fortunes of the company. If the company does well so that the company's stock price moves above the strike price, employees gain by exercising the options and then selling the stock they acquire at the market price.

Employee stock options have become very popular in the last 20 years. Many companies, particularly technology companies, feel that the only way they can attract and keep the best employees is to offer them very attractive stock option packages. Some companies grant options only to senior management; others grant them to people at all levels in the organization. Microsoft was one of the first companies to use employee stock options. All Microsoft employees were granted options and, as the company's stock price rose, it is estimated that over 10,000 of them became millionaires. In 2003 Microsoft announced that it would discontinue the use of options and award shares of Microsoft to employees instead. But many other companies throughout the world continue to be enthusiastic users of employee stock options.

Employee stock options are popular with start-up companies. Often these companies do not have the resources to pay key employees as much as they could earn with an established company and they solve this problem by supplementing the salaries of the employees with stock options. If the company does well and shares are sold to the public in an IPO, the options are likely to prove to be very valuable. Some newly formed companies have even granted options to students who worked for just a few months during their summer break—and in some cases this has led to windfalls of hundreds of thousands of dollars for the students!

This chapter explains how stock option plans work and how their popularity has been influenced by their accounting treatment. It discusses whether employee stock options help to align the interests of shareholders with those of top executives running a company. It also describes how these options are valued and looks at backdating scandals.

15.1 CONTRACTUAL ARRANGEMENTS

Employee stock options often last as long as 10 to 15 years. Very often the strike price is set equal to the stock price on the grant date so that the option is initially at the money.

The following are usually features employee stock option plans:

1. There is a vesting period during which the options cannot be exercised. This vesting period can be as long as four years.
2. When employees leave their jobs (voluntarily or involuntarily) during the vesting period, they forfeit their options.
3. When employees leave (voluntarily or involuntarily) after the vesting period, they forfeit options that are out of the money and they have to exercise vested options that are in the money almost immediately.
4. Employees are not permitted to sell the options.
5. When an employee exercises options, the company issues new shares and sells them to the employee for the strike price.

The Early Exercise Decision

The fourth feature of employee stock option plans noted above has important implications. If employees, for whatever reason, want to realize a cash benefit from options that have vested, they must exercise the options and sell the underlying shares. They cannot sell the options to someone else. This leads to a tendency for employee stock options to be exercised earlier than similar regular exchange-traded or over-the-counter call options.

Consider a call option on a stock paying no dividends. In Section 10.5 we showed that, if it is a regular call option, it should never be exercised early. The holder of the option will always do better by selling the option rather than exercising it before the end of its life. However, the arguments we used in Section 10.5 are not applicable to employee stock options because they cannot be sold. The only way employees can realize a cash benefit from the options (or diversify their holdings) is by exercising the options and selling the stock. It is therefore not unusual for an employee stock option to be exercised well before it would be optimal to exercise the option if it were a regular exchange-traded or over-the-counter option.

Should an employee ever exercise his or her options before maturity and then keep the stock rather than selling it? Assume that the option's strike price is constant during the life of the option and the option can be exercised at any time. To answer the question we consider two options: the employee stock option and an otherwise identical regular option that can be sold in the market. We refer to the first option as option A and the second as option B. If the stock pays no dividends, we know that option B should never be exercised early. It follows that it is not optimal to exercise option A and keep the stock. If the employee wants to maintain a stake in his or her company, a better strategy is to keep the option. This delays paying the strike price and maintains the insurance value of the option, as described in Section 10.5. Only when it is optimal to exercise option B can it be a rational strategy for an employee to exercise option A before maturity and keep the stock.¹ As discussed in Section 14.12, it is optimal to exercise option B only when a relatively high dividend is imminent.

In practice the early exercise behavior of employees varies widely from company to company. In some companies, there is a culture of not exercising early; in others, employees tend to exercise options and sell the stock soon after the end of the vesting period, even if the options are only slightly in the money.

¹ The only exception to this could be when an executive wants to own the stock for its voting rights.

15.2 DO OPTIONS ALIGN THE INTERESTS OF SHAREHOLDERS AND MANAGERS?

For investors to have confidence in capital markets, it is important that the interests of shareholders and managers are reasonably well aligned. This means that managers should be motivated to make decisions that are in the best interests of shareholders. Managers are the agents of the shareholders and, as mentioned in Chapter 8, economists use the term *agency costs* to describe the losses experienced when the interests of agents and principals are not aligned.

Do employee stock options help align the interests of employees and shareholders? The answer to this question is not straightforward. There can be little doubt that they serve a useful purpose for a start-up company. The options are an excellent way for the main shareholders, who are usually also senior executives, to motivate employees to work long hours. If the company is successful and there is an IPO, the employees will do very well; but if the company is unsuccessful, the options will be worthless.

It is the options granted to the senior executives of publicly traded companies that are most controversial. It has been estimated that employee stock options account for about 50% of the remuneration of top executives in the United States. Executive stock options are sometimes referred to as an executive's "pay for performance." If the company's stock price goes up, so that shareholders make gains, the executive is rewarded. However, this overlooks the asymmetric payoffs of options. If the company does badly then the shareholders lose money, but all that happens to the executives is that they fail to make a gain. Unlike the shareholders, they do not experience a loss.² A better type of pay for performance involves the simpler strategy of giving stock to executives. The gains and losses of the executives then mirror those of other shareholders.

What temptations do stock options create for a senior executive? Suppose an executive plans to exercise a large number of stock options in three months and sell the stock. He or she might be tempted to time announcements of good news—or even move earnings from one quarter to another—so that the stock price increases just before the options are exercised. Alternatively, if at-the-money options are due to be granted to the executive in three months, the executive might be tempted to take actions that reduce the stock price just before the grant date. The type of behavior we are talking about here is of course totally unacceptable—and may well be illegal. But the backdating scandals, which are discussed later in this chapter, show that the way some executives have handled issues related to stock options leaves much to be desired.

Even when there is no impropriety of the type we have just mentioned, executive stock options are liable to have the effect of motivating executives to focus on short-term profits at the expense of longer-term performance. In some cases they might even take risks they would not otherwise take (and risks that are not in the interests of the shareholders) because of the asymmetric payoffs of options. Managers of large funds worry that, because stock options are such a huge component of an executive's compensation, they are liable to be a big source of distraction. Senior management may spend too much time thinking about all the different aspects of their compensation and not enough time running the company!

² When options have moved out of the money, companies have sometimes replaced them with new at-the-money options. This practice known as "repricing" leads to the executive's gains and losses being even less closely tied to those of the shareholders.

A manager's inside knowledge and ability to affect outcomes and announcements is always liable to interact with his or her trading in a way that is to the disadvantage of other shareholders. One radical suggestion for mitigating this problem is to require executives to give notice to the market—perhaps one week's notice—of an intention to buy or sell their company's stock.³ (Once the notice of an intention to trade had been given, it would be binding on the executive.) This allows the market to form its own conclusions about why the executive is trading. As a result, the price may increase before the executive buys and decrease before the executive sells.

15.3 ACCOUNTING ISSUES

An employee stock option represents a cost to the company and a benefit to the employee just like any other form of compensation. This point, which for many is self-evident, is actually quite controversial. Many corporate executives appear to believe that an option has no value unless it is in the money. As a result, they argue that an at-the-money option issued by the company is not a cost to the company. The reality is that, if options are valuable to employees, they must represent a cost to the company's shareholders—and therefore to the company. There is no free lunch. The cost to the company of the options arises from the fact that the company has agreed that, if its stock does well, it will sell shares to employees at a price less than that which would apply in the open market.

Prior to 1995 the cost charged to the income statement of a company when it issued stock options was the intrinsic value. Most options were at the money when they were first issued, so that this cost was zero. In 1995, accounting standard FAS 123 was issued. Many people expected it to require the expensing of options at their fair value. However, as a result of intense lobbying, the 1995 version of FAS 123 only encouraged companies to expense the fair value of the options they granted on the income statement. It did not require them to do so. If fair value was not expensed on the income statement, it had to be reported in a footnote to the company's accounts.

Accounting standards have now changed to require the expensing of stock options at their fair value on the income statement. In February 2004 the International Accounting Standards Board issued IAS 2 requiring companies to start expensing stock options in 2005. In December 2004 FAS 123 was revised to require the expensing of employee stock options in the United States starting in 2005.

The effect of the new accounting standards is to require options to be valued on the grant date and the amount to be recorded as an expense in the income statement for the year in which the grant is made. Valuation at a time later than the grant date is not required. It can be argued that options should be revalued at financial year ends (or every quarter) until they are exercised or reach the end of their lives.⁴ This would treat them in the same way as other derivative transactions entered into by the company. If the option became more valuable from one year to the next, there would then be an

³ This would apply to the exercise of options because, if an executive wants to exercise options and sell the stock that is acquired, then he or she would have to give notice of intention to sell.

⁴ See J. Hull and A. White, "Accounting for Employee Stock Options: A Practical Approach to Handling the Valuation Issues," *Journal of Derivatives Accounting*, 1, 1 (2004): 3–9.

additional amount to be expensed. However, if it declined in value, there would be a positive impact on income.

This approach would have a number of advantages. The cumulative charge to the company would reflect the actual cost of the options (either zero if the options are not exercised or the option payoff if they are exercised). Although the charge in any year would depend on the option pricing model used, the cumulative charge over the life of the option would not.⁵ Arguably there would be much less incentive for the company to engage in the backdating practices described later in the chapter. The disadvantage usually cited for accounting in this way is that it is undesirable because it introduces volatility into the income statement.⁶

Nontraditional Option Plans

It is easy to understand why pre-2005 employee stock options tended to be at the money on the grant date and have strike prices that did not change during the life of the option. Any departure from this standard arrangement was likely to require the options to be expensed. Now that accounting rules have changed so that all options are expensed at fair value, many companies are considering alternatives to the standard arrangement.

One argument against the standard arrangement is that employees do well when the stock market goes up, even if their own company's stock price does less well than the market. One way of overcoming this problem is to tie the strike price of the options to the performance of the S&P 500. Suppose that on the option grant date the stock price is \$30 and the S&P 500 is 1,500. The strike price would initially be set at \$30. If the S&P 500 increased by 10% to 1,650, then the strike price would also increase by 10% to \$33. If the S&P 500 moved down by 15% to 1,275, then the strike price would also move down by 15% to \$25.50. The effect of this is that the company's stock price performance has to beat the performance of the S&P 500 to become in the money. As an alternative to using the S&P 500 as the reference index, the company could use an index of the prices of stocks in the same industrial sector as the company.

In another variation on the standard arrangement, the strike price increases through time in a predetermined way such that the shares of the stock have to provide a certain minimum return per year for the options to be in the money. In some cases, profit targets are specified and the options vest only if the profit targets are met.⁷

15.4 VALUATION

Accounting standards give companies quite a bit of latitude in choosing a method for valuing employee stock options. In this section we review some of the alternatives.

⁵ Interestingly, if an option is settled in cash rather than by the company issuing new shares, it is subject to the accounting treatment proposed here. (However, there is no economic difference between an option that is settled in cash and one that is settled by selling new shares to the employee.)

⁶ In fact the income statement is likely be less volatile if stock options are revalued. When the company does well, income is reduced by revaluing the executive stock options. When the company does badly, it is increased.

⁷ This type of option is difficult to value because the payoff depends on reported accounting numbers as well as the stock price. Valuations usually assume that targets will be met.

The “Quick and Dirty” Approach

A frequently used approach is based on what is known as the option’s *expected life*. This is the average time for which employees hold the option before it is exercised or expires. The expected life can be approximately estimated from historical data on the early exercise behavior of employees and reflects the vesting period, the impact of employees leaving the company, and the tendency we mentioned in Section 15.1 for employee stock options to be exercised earlier than regular options. The Black–Scholes–Merton model is used with the life of the option, T , set equal to the expected life. The volatility is usually estimated from several years of historical data as described in Section 14.4.

It should be emphasized that using the Black–Scholes–Merton formula in this way has no theoretical validity. There is no reason why the value of a European stock option with the time to maturity, T , set equal to the expected life should be approximately the same as the value of the American-style employee stock option that we are interested in. However, the results given by the model are not totally unreasonable. Companies, when reporting their employee stock option expense, will frequently mention the volatility and expected life used in their Black–Scholes–Merton computations.

Example 15.1

A company grants 1,000,000 options to its executives on November 1, 2011. The stock price on that date is \$30 and the strike price of the options is also \$30. The options last for 10 years and vest after three years. The company has issued similar at-the-money options for the last 10 years. The average time to exercise or expiry of these options is 4.5 years. The company therefore decides to use an “expected life” of 4.5 years. It estimates the long-term volatility of the stock price, using 5 years of historical data, to be 25%. The present value of dividends during the next 4.5 years is estimated to be \$4. The 4.5-year zero-coupon risk-free interest rate is 5%. The option is therefore valued using the Black–Scholes–Merton model (adjusted for dividends in the way described in Section 14.12) with $S_0 = 30 - 4 = 26$, $K = 30$, $r = 5\%$, $\sigma = 25\%$, and $T = 4.5$. The Black–Scholes–Merton formula gives the value of one option as \$6.31. Hence, the income statement expense is $1,000,000 \times 6.31$, or \$6,310,000.

Binomial Tree Approach

A more sophisticated approach to valuing employee stock options involves building a binomial tree as outlined in Chapter 12 and adjusting the rules used when rolling back through the tree to reflect (a) whether the option has vested, (b) the probability of the employee leaving the company, and (c) the probability of the employee choosing to exercise the option. The terms of the option define whether the option has vested at different nodes of the tree. Historical data on turnover rates for employees can be used to estimate the probability of the option being either prematurely exercised or forfeited at a node because the employee leaves the company. The probability of an employee choosing to exercise the option at different nodes of the tree is more difficult to quantify. Clearly this probability increases as the ratio of the stock price to the strike price increases and as the time to the option’s maturity declines. If enough historical data is available, the probability of exercise as a function of these two variables can be estimated—at least approximately.

Example 15.2

Suppose a company grants stock options that last 8 years and vest after 3 years. The stock price and strike price are both \$40. The stock price volatility is 30%, the risk-free rate is 5%, and the company pays no dividends. Figure 15.1 shows how a four-step tree could be used to value the option. (This is for illustration; in practice more time steps would be used.) In this case, $\sigma = 0.3$, $\Delta t = 2$, and $r = 0.05$, so that, with the notation of Chapter 12, $a = e^{0.05 \times 2} = 1.1052$, $u = e^{0.3\sqrt{2}} = 1.5285$, $d = 1/u = 0.6543$, and $p = (a - d)/(u - d) = 0.5158$. The probability on the “up branches” is 0.5158 and the probability on the “down branches” is 0.4842. There are three nodes where early exercise could be desirable: D, G, and H. (The option has not vested at node B and is not in the money at the other nodes prior to maturity.) We assume that the probabilities that the holder will choose to exercise at nodes D, G, and H (conditional on no earlier exercise) have been estimated as 40%, 80%, and 30%, respectively. We suppose that the probability of an employee leaving the company during each time step is 5%. (This corresponds to an employee turnover rate of approximately 2.5% per year.) For the purposes of the calculation, it is assumed that employees always leave at the end of a time period. If an employee leaves the company before an option has vested or when the option is out of the money, the option is forfeited. In other cases the option must be exercised immediately.

Figure 15.1 Valuation of employee stock option in Example 15.2.

At each node:

Upper value = Underlying asset price

Lower value = Option price

Values in bold are a result of early exercise.

Strike price = 40

Discount factor per step = 0.9048

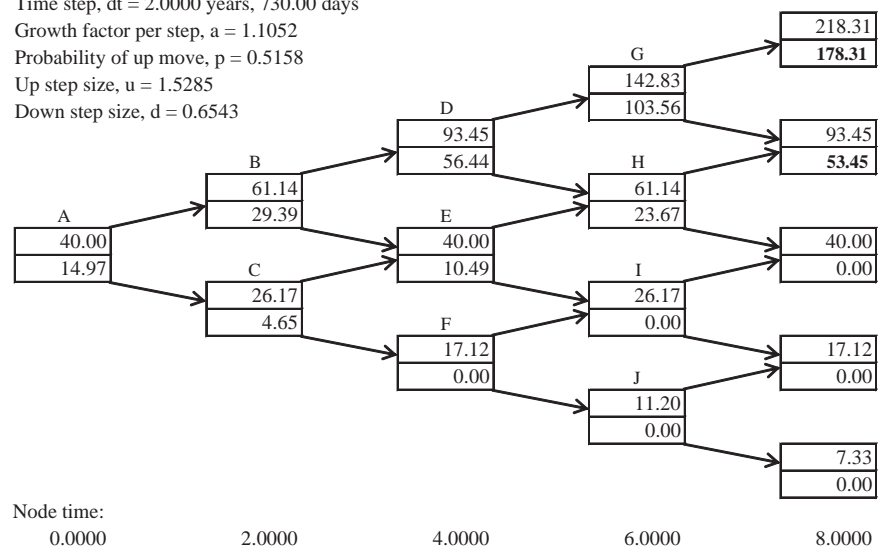
Time step, $dt = 2.0000$ years, 730.00 days

Growth factor per step, $a = 1.1052$

Probability of up move, $p = 0.5158$

Up step size, $u = 1.5285$

Down step size, $d = 0.6543$



The value of the option at the final nodes is its intrinsic value. Consider the nodes at time 6 years. Nodes I and J are easy. Since these nodes are certain to lead to nodes where the option is worth nothing, the value of the option is zero at these nodes. At node H there is a 30% chance that the employee will choose to exercise the option. In cases where the employee does not choose to exercise, there is a 5% chance that the employee leaves the company and has to exercise. The total probability of exercise is therefore $0.3 + 0.7 \times 0.05 = 0.335$. If the option is exercised, its value is $61.14 - 40 = 21.14$. If it is not exercised, its value is

$$e^{-0.05 \times 2}(0.5158 \times 53.45 + 0.4842 \times 0) = 24.95$$

The value of the option at node H is therefore

$$0.335 \times 21.14 + 0.665 \times 24.95 = 23.67$$

The value at node G is similarly

$$0.81 \times 102.83 + 0.19 \times 106.64 = 103.56$$

We now move on to the nodes at time 4 years. At node F the option is clearly worth zero. At node E there is a 5% chance that the employee will forfeit the option because he or she leaves the company and a 95% chance that the option will be retained. In the latter case the option is worth

$$e^{-0.05 \times 2}(0.5158 \times 23.67 + 0.4842 \times 0) = 11.05$$

The option is therefore worth $0.95 \times 11.05 = 10.49$. At node D there is a 0.43 probability that the option will be exercised and a 0.57 chance that it will be retained. The value of the option is 56.44.

Consider next the initial node and the nodes at time 2 years. The option has not vested at these nodes. There is a 5% chance that the option will be forfeited and a 95% chance that it will be retained for a further 2 years. This leads to the valuations shown in Figure 15.1. The valuation of the option at the initial node is 14.97. (This compares with a valuation of 17.98 for a regular option using the same tree.)

The Exercise Multiple Approach

Hull and White suggest a simple model where an employee exercises as soon as the option has vested and the ratio of the stock price to the strike price is above a certain level.⁸ They refer to the ratio of stock price to strike price that triggers exercise as the “exercise multiple”. The option can be valued using a binomial or trinomial tree. As outlined in Section 26.6, it is important to construct a binomial or trinomial tree where nodes lie on the stock prices that will lead to exercise. For example, if the strike price is \$30 and the assumption is that employees exercise when the ratio of the stock price to the strike price is 1.5, the tree should be constructed so that there are nodes at a stock price level of \$45. The tree calculations are similar to those for Example 15.2 and take account of the probability of an employee leaving the company.⁹ To estimate the exercise multiple, it is necessary to calculate from historical data the average ratio of

⁸ See J. Hull and A. White, “How to value employee stock options,” *Financial Analysts Journal*, 60, 1 (January/February 2004): 3–9.

⁹ Software implementing this approach is on www.rotman.utoronto.ca/~hull.

stock price to strike price at the time of exercise. (Exercises at maturity and those arising from the termination of the employee's job are not included in the calculation of the average.) This may be easier to estimate from historical data than the expected life because the latter is quite heavily dependent on the particular path that has been followed by the stock's price.

A Market-Based Approach

One way of valuing an employee stock option is to see what the market would pay for it. Cisco was the first to try this in 2006. It proposed selling options with the exact terms of its employee stock options to institutional investors. This approach was rejected by the SEC on the grounds that the range of investors bidding for the options was not wide enough.

Zions Bancorp has suggested an alternative approach. It proposed that securities providing payoffs mirroring those actually realized by its employees be sold. Suppose that the strike price for a particular grant to employees is \$40 and it turns out that 1% of employees exercise after exactly 5 years when the stock price is \$60, 2% exercise after exactly 6 years when the stock price is \$65, and so on. Then 1% of the securities owned by an investor will provide a \$20 payoff after 5 years, 2% will provide a payoff of \$25 after 6 years, and so on.

Zions Bancorp tested the idea using its own stock option grant to its employees. It sold the securities using a Dutch auction process. In this individuals or companies can submit a bid indicating the price they are prepared to pay and the number of options they are prepared to buy. The clearing price is the highest bid such that the aggregate number of options sought at that price or a higher price equals or exceeds the number of options for sale. Buyers who have bid more than the clearing price get their orders filled at the clearing price and the buyer who bid the clearing price gets the remainder. Zions Bancorp announced that it had received SEC approval for its market-based approach in October 2007.

Dilution

The fact that a company issues new stock when an employee stock option is exercised leads to some dilution for existing stock holders because new shares are being sold to employees at below the current stock price. It is natural to assume that this dilution takes place at the time the option is exercised. However, this is not the case. As explained in Section 14.10, stock prices are diluted when the market first hears about a stock option grant. The possible exercise of options is anticipated and immediately reflected in the stock price. This point is emphasized by the example in Business Snapshot 14.3.

The stock price immediately after a grant is announced to the public reflects any dilution. Provided that this stock price is used in the valuation of the option, it is not necessary to adjust the option price for dilution. In many instances the market expects a company to make regular stock option grants and so the market price of the stock anticipates dilution even before the announcement is made.

If a company is contemplating a stock option grant that will surprise the market, the cost can be calculated as described in Example 14.7. This cost can be compared with benefits such as lower regular employee remuneration and less employee turnover.

15.5 BACKDATING SCANDALS

No discussion of employee stock options would be complete without mentioning backdating scandals. Backdating is the practice of marking a document with a date that precedes the current date.

Suppose that a company decides to grant at-the-money options to its employees on April 30 when the stock price is \$50. If the stock price was \$42 on April 3, it is tempting to behave as if those the options were granted on April 3 and use a strike price of \$42. This is legal provided that the company reports the options as \$8 in the money on the date when the decision to grant the options is made, April 30. But it is illegal for the company to report the options as at-the-money and granted on April 3. The value on April 3 of an option with a strike price of \$42 is much less than its value on April 30. Shareholders are misled about the true cost of the decision to grant options if the company reports the options as granted on April 3.

How prevalent is backdating? To answer this question, researchers have investigated whether a company's stock price has, on average, a tendency to be low at the time of the grant date that the company reports. Early research by Yermack shows that stock prices tend to increase after reported grant dates.¹⁰ Lie extended Yermack's work, showing that stock prices also tended to decrease before reported grant dates.¹¹ Furthermore he showed that the pre- and post-grant stock price patterns had become more pronounced over time. His results are summarized in Figure 15.2, which shows average abnormal returns around the grant date for the 1993–94, 1995–98, and 1999–2002 periods. (Abnormal returns are the returns after adjustments for returns on the market portfolio and the beta of the stock.) Standard statistical tests show that it is almost impossible for the patterns shown in Figure 15.2 to be observed by chance. This led both academics and regulators to conclude in 2002 that backdating had become a common practice. In August 2002 the SEC required option grants by public companies to be reported within two business days. Heron and Lie showed that this led to a dramatic reduction in the abnormal returns around the grant dates—particularly for those companies that complied with this requirement.¹² It might be argued that the patterns in Figure 15.2 are explained by managers simply choosing grant dates after bad news or before good news, but the Heron and Lie study provides compelling evidence that this is not the case.


Estimates of the number of companies that illegally backdated stock option grants in the United States vary widely. Tens and maybe hundreds of companies seem to have engaged in the practice. Many companies seem to have adopted the view that it was acceptable to backdate up to one month. Some CEOs resigned when their backdating practices came to light. In August 2007, Gregory Reyes of Brocade Communications Systems, Inc., became the first CEO to be tried for backdating stock option grants. Allegedly, Mr. Reyes said to a human resources employee: “It is not illegal if you do not get caught.” In June 2010, he was sentenced to 18 months in prison and fined \$15 million.

¹⁰ See D. Yermack, “Good timing: CEO stock option awards and company news announcements,” *Journal of Finance*, 52 (1997), 449–476.

¹¹ See E. Lie, “On the timing of CEO stock option awards,” *Management Science*, 51, 5 (May 2005), 802–12.

¹² See R. Heron and E. Lie, “Does backdating explain the stock price pattern around executive stock option grants,” *Journal of Financial Economics*, 83, 2 (February 2007), 271–95.

Figure 15.2 Erik Lie's results providing evidence of backdating. (Reproduced with permission, from www.biz.uiowa.edu/faculty/elie/backdating.htm.)



Companies involved in backdating have had to restate past financial statements and have been defendants in class action suits brought by shareholders who claim to have lost money as a result of backdating. For example, McAfee announced in December 2007 that it would restate earnings between 1995 and 2005 by \$137.4 million. In 2006, it set aside \$13.8 million to cover lawsuits.

SUMMARY

Executive compensation has increased very fast in the last 20 years and much of the increase has come from the exercise of stock options granted to the executives. Until 2005 at-the-money stock option grants were a very attractive form of compensation. They had no impact on the income statement and were very valuable to employees. Accounting standards now require options to be expensed.

There are a number of different approaches to valuing employee stock options. A common approach is to use the Black-Scholes-Merton model with the life of the option set equal to the expected time to exercise or expiry of the option. Another approach is to assume that options are exercised as soon as the ratio of the stock price to the strike price reaches a certain barrier. A third approach is to try and estimate the relationship between the probability of exercise, the ratio of the stock price to the strike price, and the time to option maturity. A fourth approach is to create a market for securities that replicate the payoffs on the options.

Academic research has shown beyond doubt that many companies have engaged in the illegal practice of backdating stock option grants in order to reduce the strike price, while still contending that the options were at the money. The first prosecutions for this illegal practice were in 2007.

FURTHER READING

- Carpenter, J., “The Exercise and Valuation of Executive Stock Options,” *Journal of Financial Economics*, 48, 2 (May): 127–58.
- Core, J. E., and W. R. Guay, “Stock Option Plans for Non-Executive Employees,” *Journal of Financial Economics*, 61, 2 (2001): 253–87.
- Heron, R., and E. Lie, “Does Backdating Explain the Stock Price Pattern around Executive Stock Option Grants,” *Journal of Financial Economics*, 83, 2 (February 2007): 271–95.
- Huddart, S., and M. Lang, “Employee Stock Option Exercises: An Empirical Analysis,” *Journal of Accounting and Economics*, 21, 1 (February): 5–43.
- Hull, J., and A. White, “How to Value Employee Stock Options,” *Financial Analysts Journal*, 60, 1 (January/February 2004): 3–9.
- Lie, E., “On the Timing of CEO Stock Option Awards,” *Management Science*, 51, 5 (May 2005): 802–12.
- Rubinstein, M., “On the Accounting Valuation of Employee Stock Options,” *Journal of Derivatives*, 3, 1 (Fall 1996): 8–24.
- Yermack, D., “Good Timing: CEO Stock Option Awards and Company News Announcements,” *Journal of Finance*, 52 (1997): 449–76.

Practice Questions (Answers in Solutions Manual)

- 15.1. Why was it attractive for companies to grant at-the-money stock options prior to 2005? What changed in 2005?
- 15.2. What are the main differences between a typical employee stock option and an American call option traded on an exchange or in the over-the-counter market?
- 15.3. Explain why employee stock options on a non-dividend-paying stock are frequently exercised before the end of their lives, whereas an exchange-traded call option on such a stock is never exercised early.
- 15.4. “Stock option grants are good because they motivate executives to act in the best interests of shareholders.” Discuss this viewpoint.
- 15.5. “Granting stock options to executives is like allowing a professional footballer to bet on the outcome of games.” Discuss this viewpoint.
- 15.6. Why did some companies backdate stock option grants in the US prior to 2002? What changed in 2002?
- 15.7. In what way would the benefits of backdating be reduced if a stock option grant had to be revalued at the end of each quarter?
- 15.8. Explain how you would do the analysis to produce a chart such as the one in Figure 15.2.
- 15.9. On May 31 a company’s stock price is \$70. One million shares are outstanding. An executive exercises 100,000 stock options with a strike price of \$50. What is the impact of this on the stock price?
- 15.10. The notes accompanying a company’s financial statements say: “Our executive stock options last 10 years and vest after 4 years. We valued the options granted this year using the Black–Scholes–Merton model with an expected life of 5 years and a volatility of 20%.” What does this mean? Discuss the modeling approach used by the company.

- 15.11. In a Dutch auction of 10,000 options, bids are as follows:
 A bids \$30 for 3,000
 B bids \$33 for 2,500
 C bids \$29 for 5,000
 D bids \$40 for 1,000
 E bids \$22 for 8,000
 F bids \$35 for 6,000.
 What is the result of the auction? Who buys how many at what price?
- 15.12. A company has granted 500,000 options to its executives. The stock price and strike price are both \$40. The options last for 12 years and vest after 4 years. The company decides to value the options using an expected life of 5 years and a volatility of 30% per annum. The company pays no dividends and the risk-free rate is 4%. What will the company report as an expense for the options on its income statement?
- 15.13. A company's CFO says: "The accounting treatment of stock options is crazy. We granted 10,000,000 at-the-money stock options to our employees last year when the stock price was \$30. We estimated the value of each option on the grant date to be \$5. At our year-end the stock price had fallen to \$4, but we were still stuck with a \$50 million charge to the P&L." Discuss.

Further Questions

- 15.14. What is the (risk-neutral) expected life for the employee stock option in Example 15.2? What is the value of the option obtained by using this expected life in Black–Scholes–Merton?
- 15.15. A company has granted 2,000,000 options to its employees. The stock price and strike price are both \$60. The options last for 8 years and vest after 2 years. The company decides to value the options using an expected life of 6 years and a volatility of 22% per annum. Dividends on the stock are \$1 per year, payable halfway through each year, and the risk-free rate is 5%. What will the company report as an expense for the options on its income statement?
- 15.16. A company has granted 1,000,000 options to its employees. The stock price and strike price are both \$20. The options last 10 years and vest after 3 years. The stock price volatility is 30%, the risk-free rate is 5%, and the company pays no dividends. Use a four-step tree to value the options. Assume that there is a probability of 4% that an employee leaves the company at the end of each of the time steps on your tree. Assume also that the probability of voluntary early exercise at a node, conditional on no prior exercise, when (a) the option has vested and (b) the option is in the money, is

$$1 - \exp[-a(S/K - 1)/T]$$

where S is the stock price, K is the strike price, T is the time to maturity, and $a = 2$.

CHAPTER 16

Options on Stock Indices and Currencies



Options on stock indices and currencies were introduced in Chapter 9. This chapter discusses them in more detail. It explains how they work and reviews some of the ways they can be used. In the second half of the chapter, the valuation results in Chapter 14 are extended to cover European options on a stock paying a known dividend yield. It is then argued that both stock indices and currencies are analogous to stocks paying dividend yields. This enables the results for options on a stock paying a dividend yield to be applied to these types of options as well.

16.1 OPTIONS ON STOCK INDICES

Several exchanges trade options on stock indices. Some of the indices track the movement of the market as a whole. Others are based on the performance of a particular sector (e.g., computer technology, oil and gas, transportation, or telecoms). Among the index options traded on the Chicago Board Options Exchange (CBOE) are American and European options on the S&P 100 (OEX and XEO), European options on the S&P 500 (SPX), European options on the Dow Jones Industrial Average (DJX), and European options on the Nasdaq 100 (NDX). In Chapter 9, we explained that the CBOE trades LEAPS and flex options on individual stocks. It also offers these option products on indices.

One index option contract is on 100 times the index. (Note that the Dow Jones index used for index options is 0.01 times the usually quoted Dow Jones index.) Index options are settled in cash. This means that, on exercise of the option, the holder of a call option contract receives $(S - K) \times 100$ in cash and the writer of the option pays this amount in cash, where S is the value of the index at the close of trading on the day of the exercise and K is the strike price. Similarly, the holder of a put option contract receives $(K - S) \times 100$ in cash and the writer of the option pays this amount in cash.

Portfolio Insurance

Portfolio managers can use index options to limit their downside risk. Suppose that the value of an index today is S_0 . Consider a manager in charge of a well-diversified portfolio whose beta is 1.0. A beta of 1.0 implies that the returns from the portfolio mirror those

from the index. Assuming the dividend yield from the portfolio is the same as the dividend yield from the index, the percentage changes in the value of the portfolio can be expected to be approximately the same as the percentage changes in the value of the index. Since each contract is on 100 times the index, it follows that the value of the portfolio is protected against the possibility of the index falling below K if, for each $100S_0$ dollars in the portfolio, the manager buys one put option contract with strike price K . Suppose that the manager's portfolio is worth \$500,000 and the value of the index is 1,000. The portfolio is worth 500 times the index. The manager can obtain insurance against the value of the portfolio dropping below \$450,000 in the next three months by buying five three-month put option contracts on the index with a strike price of 900.

To illustrate how the insurance works, consider the situation where the index drops to 880 in three months. The portfolio will be worth about \$440,000. The payoff from the options will be $5 \times (900 - 880) \times 100 = \$10,000$, bringing the total value of the portfolio up to the insured value of \$450,000.

When the Portfolio's Beta Is Not 1.0

If the portfolio's beta (β) is not 1.0, β put options must be purchased for each $100S_0$ dollars in the portfolio, where S_0 is the current value of the index. Suppose that the \$500,000 portfolio just considered has a beta of 2.0 instead of 1.0. We continue to assume that the index is 1,000. The number of put options required is

$$2.0 \times \frac{500,000}{1,000 \times 100} = 10$$

rather than 5 as before.

To calculate the appropriate strike price, the capital asset pricing model can be used (see the appendix to Chapter 3). Suppose that the risk free rate is 12%, the dividend yield on both the index and the portfolio is 4%, and protection is required against the value of the portfolio dropping below \$450,000 in the next three months. Under the capital asset pricing model, the expected excess return of a portfolio over the risk-free

Table 16.1 Calculation of expected value of portfolio when the index is 1,040 in three months and $\beta = 2.0$.

| | |
|--|---|
| Value of index in three months: | 1,040 |
| Return from change in index: | 40/1,000, or 4% per three months |
| Dividends from index: | $0.25 \times 4 = 1\%$ per three months |
| Total return from index: | $4 + 1 = 5\%$ per three months |
| Risk-free interest rate: | $0.25 \times 12 = 3\%$ per three months |
| Excess return from index over risk-free interest rate: | $5 - 3 = 2\%$ per three months |
| Expected excess return from portfolio over risk-free interest rate: | $2 \times 2 = 4\%$ per three months |
| Expected return from portfolio: | $3 + 4 = 7\%$ per three months |
| Dividends from portfolio: | $0.25 \times 4 = 1\%$ per three months |
| Expected increase in value of portfolio: | $7 - 1 = 6\%$ per three months |
| Expected value of portfolio: | $\$500,000 \times 1.06 = \$530,000$ |

Table 16.2 Relationship between value of index and value of portfolio for $\beta = 2.0$.

| <i>Value of index in three months</i> | <i>Value of portfolio in three months (\$)</i> |
|---|--|
| 1,080 | 570,000 |
| 1,040 | 530,000 |
| 1,000 | 490,000 |
| 960 | 450,000 |
| 920 | 410,000 |
| 880 | 370,000 |

rate is assumed to equal beta times the excess return of the index portfolio over the risk-free rate. The model enables the expected value of the portfolio to be calculated for different values of the index at the end of three months. Table 16.1 shows the calculations for the case where the index is 1,040. In this case the expected value of the portfolio at the end of the three months is \$530,000. Similar calculations can be carried out for other values of the index at the end of the three months. The results are shown in Table 16.2. The strike price for the options that are purchased should be the index level corresponding to the protection level required on the portfolio. In this case the protection level is \$450,000 and so the correct strike price for the 10 put option contracts that are purchased is 960.¹

To illustrate how the insurance works, consider what happens if the value of the index falls to 880. As shown in Table 16.2, the value of the portfolio is then about \$370,000. The put options pay off $(960 - 880) \times 10 \times 100 = \$80,000$, and this is exactly what is necessary to move the total value of the portfolio manager's position up from \$370,000 to the required level of \$450,000.

The examples in this section show that there are two reasons why the cost of hedging increases as the beta of a portfolio increases. More put options are required and they have a higher strike price.

16.2 CURRENCY OPTIONS

Currency options are primarily traded in the over-the-counter market. The advantage of this market is that large trades are possible, with strike prices, expiration dates, and other features tailored to meet the needs of corporate treasurers. Although currency options do trade on NASDAQ OMX in the United States, the exchange-traded market for these options is much smaller than the over-the-counter market.

An example of a European call option is a contract that gives the holder the right to buy one million euros with US dollars at an exchange rate of 1.2000 US dollars per euro. If the actual exchange rate at the maturity of the option is 1.2500, the payoff is

¹ Approximately 1% of \$500,000, or \$5,000, will be earned in dividends over the next three months. If we want the insured level of \$450,000 to include dividends, we can choose a strike price corresponding to \$445,000 rather than \$450,000. This is 955.

$1,000,000 \times (1.2500 - 1.2000) = \$50,000$. Similarly, an example of a European put option is a contract that gives the holder the right to sell ten million Australian dollars for US dollars at an exchange rate of 0.9000 US dollars per Australian dollar. If the actual exchange rate at the maturity of the option is 0.8700, the payoff is $10,000,000 \times (0.9000 - 0.8700) = \$300,000$.

For a corporation wishing to hedge a foreign exchange exposure, foreign currency options are an alternative to forward contracts. A company due to receive sterling at a known time in the future can hedge its risk by buying put options on sterling that mature at that time. The hedging strategy guarantees that the exchange rate applicable to the sterling will not be less than the strike price, while allowing the company to benefit from any favorable exchange-rate movements. Similarly, a company due to pay sterling at a known time in the future can hedge by buying calls on sterling that mature at that time. This hedging strategy guarantees that the cost of the sterling will not be greater than a certain amount while allowing the company to benefit from favorable exchange-rate movements. Whereas a forward contract locks in the exchange rate for a future transaction, an option provides a type of insurance. This is not free. It costs nothing to enter into a forward transaction, but options require a premium to be paid up front.

Range Forwards

A *range forward contract* is a variation on a standard forward contract for hedging foreign exchange risk. Consider a US company that knows it will receive one million pounds sterling in three months. Suppose that the three-month forward exchange rate is 1.5200 dollars per pound. The company could lock in this exchange rate for the dollars it receives by entering into a short forward contract to sell one million pounds sterling in three months. This would ensure that the amount received for the one million pounds is \$1,520,000.

An alternative is to buy a European put option with a strike price of K_1 and sell a European call option with a strike price K_2 , where $K_1 < 1.5200 < K_2$. This is known as a short range forward contract. The payoff is shown in Figure 16.1a. In both cases, the options are on one million pounds. If the exchange rate in three months proves to be less than K_1 , the put option is exercised and as a result the company is able to sell the one

Figure 16.1 Payoffs from (a) short and (b) long range-forward contract.





Figure 16.2 Exchange rate realized when either (a) a short range-forward contract is used to hedge a future foreign currency inflow or (b) a long range-forward contract is used to hedge a future foreign currency outflow.



million pounds at an exchange rate of K_1 . If the exchange rate is between K_1 and K_2 , neither option is exercised and the company gets the current exchange rate for the one million pounds. If the exchange rate is greater than K_2 , the call option is exercised against the company with the result that the one million pounds is sold at an exchange rate of K_2 . The exchange rate realized for the one million pounds is shown in Figure 16.2.

If the company knew it was due to pay rather than receive one million pounds in three months, it could sell a European put option with strike price K_1 and buy a European call option with strike price K_2 . This is known as a long range forward contract and the payoff is shown in Figure 16.1b. If the exchange rate in three months proves to be less than K_1 , the put option is exercised against the company and as a result the company buys the one million pounds it needs at an exchange rate of K_1 . If the exchange rate is between K_1 and K_2 , neither option is exercised and the company buys the one million pounds at the current exchange rate. If the exchange rate is greater than K_2 , the call option is exercised and the company is able to buy the one million pounds at an exchange rate of K_2 . The exchange rate paid for the one million pounds is the same as that received for the one million pounds in the earlier example and is shown in Figure 16.2.

In practice, a range forward contract is set up so that the price of the put option equals the price of the call option. This means that it costs nothing to set up the range forward contract, just as it costs nothing to set up a regular forward contract. Suppose that the US and British interest rates are both 5%, so that the spot exchange rate is 1.5200 (the same as the forward exchange rate). Suppose further that the exchange rate volatility is 14%. We can use DerivaGem to show that a European put with strike price 1.5000 to sell one pound has the same price as a European call option with a strike price of 1.5413 to buy one pound. (Both are worth 0.03250.) Setting $K_1 = 1.5000$ and $K_2 = 1.5413$ therefore leads to a contract with zero cost in our example.

In the limit, as the strike prices of the call and put options in a range forward contract are moved closer, the range forward contract becomes a regular forward contract. A short range forward contract becomes a short forward contract and a long range forward contract becomes a long forward contract.

16.3 OPTIONS ON STOCKS PAYING KNOWN DIVIDEND YIELDS

In this section we produce a simple rule that enables valuation results for European options on a non-dividend-paying stock to be extended so that they apply to European options on a stock paying a known dividend yield. Later we show how this enables us to value options on stock indices and currencies.

Dividends cause stock prices to reduce on the ex-dividend date by the amount of the dividend payment. The payment of a dividend yield at rate q therefore causes the growth rate in the stock price to be less than it would otherwise be by an amount q . If, with a dividend yield of q , the stock price grows from S_0 today to S_T at time T , then in the absence of dividends it would grow from S_0 today to $S_T e^{qT}$ at time T . Alternatively, in the absence of dividends it would grow from $S_0 e^{-qT}$ today to S_T at time T .

This argument shows that we get the same probability distribution for the stock price at time T in each of the following two cases:

1. The stock starts at price S_0 and provides a dividend yield at rate q .
2. The stock starts at price $S_0 e^{-qT}$ and pays no dividends.

This leads to a simple rule. When valuing a European option lasting for time T on a stock paying a known dividend yield at rate q , we reduce the current stock price from S_0 to $S_0 e^{-qT}$ and then value the option as though the stock pays no dividends.²

Lower Bounds for Option Prices

As a first application of this rule, consider the problem of determining bounds for the price of a European option on a stock paying a dividend yield at rate q . Substituting $S_0 e^{-qT}$ for S_0 in equation (10.4), we see that a lower bound for the European call option price, c , is given by

$$c \geq \max(S_0 e^{-qT} - K e^{-rT}, 0) \quad (16.1)$$

We can also prove this directly by considering the following two portfolios:

Portfolio A: one European call option plus an amount of cash equal to $K e^{-rT}$

Portfolio B: e^{-qT} shares with dividends being reinvested in additional shares.

To obtain a lower bound for a European put option, we can similarly replace S_0 by $S_0 e^{-qT}$ in equation (10.5) to get

$$p \geq \max(K e^{-rT} - S_0 e^{-qT}, 0) \quad (16.2)$$

² This rule is analogous to the one developed in Section 14.12 for valuing a European option on a stock paying known cash dividends. (In that case we concluded that it is correct to reduce the stock price by the present value of the dividends; in this case we discount the stock price at the dividend yield rate.)

This result can also be proved directly by considering the following portfolios:

Portfolio C: one European put option plus e^{-qT} shares with dividends on the shares being reinvested in additional shares

Portfolio D: an amount of cash equal to Ke^{-rT} .

Put–Call Parity

Replacing S_0 by S_0e^{-qT} in equation (10.6) we obtain put–call parity for an option on a stock paying a dividend yield at rate q :

$$c + Ke^{-rT} = p + S_0e^{-qT} \quad (16.3)$$

This result can also be proved directly by considering the following two portfolios:

Portfolio A: one European call option plus an amount of cash equal to Ke^{-rT}

Portfolio C: one European put option plus e^{-qT} shares with dividends on the shares being reinvested in additional shares.

Both portfolios are both worth $\max(S_T, K)$ at time T . They must therefore be worth the same today, and the put–call parity result in equation (16.3) follows. For American options, the put–call parity relationship is (see Problem 16.12)

$$S_0e^{-qT} - K \leq C - P \leq S_0 - Ke^{-rT}$$

Pricing Formulas

By replacing S_0 by S_0e^{-qT} in the Black–Scholes–Merton formulas, equations (14.20) and (14.21), we obtain the price, c , of a European call and the price, p , of a European put on a stock paying a dividend yield at rate q as

$$c = S_0e^{-qT}N(d_1) - Ke^{-rT}N(d_2) \quad (16.4)$$

$$p = Ke^{-rT}N(-d_2) - S_0e^{-qT}N(-d_1) \quad (16.5)$$

Since

$$\ln \frac{S_0e^{-qT}}{K} = \ln \frac{S_0}{K} - qT$$

it follows that d_1 and d_2 are given by

$$d_1 = \frac{\ln(S_0/K) + (r - q + \sigma^2/2)T}{\sigma\sqrt{T}}$$

$$d_2 = \frac{\ln(S_0/K) + (r - q - \sigma^2/2)T}{\sigma\sqrt{T}} = d_1 - \sigma\sqrt{T}$$

These results were first derived by Merton.³ As discussed in Chapter 14, the word *dividend* should, for the purposes of option valuation, be defined as the reduction in the stock price on the ex-dividend date arising from any dividends declared. If the dividend yield rate is known but not constant during the life of the option, equations (16.4)

³ See R. C. Merton, “Theory of Rational Option Pricing,” *Bell Journal of Economics and Management Science*, 4 (Spring 1973): 141–83.

and (16.5) are still true, with q equal to the average annualized dividend yield during the option's life.

Differential Equation and Risk-Neutral Valuation

To prove the results in equations (16.4) and (16.5) more formally, we can either solve the differential equation that the option price must satisfy or use risk-neutral valuation.

When we include a dividend yield of q in the analysis in Section 14.6, the differential equation (14.16) becomes⁴

$$\frac{\partial f}{\partial t} + (r - q)S \frac{\partial f}{\partial S} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 f}{\partial S^2} = rf \quad (16.6)$$

Like equation (14.16), this does not involve any variable affected by risk preferences. Therefore the risk-neutral valuation procedure described in Section 14.7 can be used.

In a risk-neutral world, the total return from the stock must be r . The dividends provide a return of q . The expected growth rate in the stock price must therefore be $r - q$. It follows that the risk-neutral process for the stock price is

$$dS = (r - q)S dt + \sigma S dz \quad (16.7)$$

To value a derivative dependent on a stock that provides a dividend yield equal to q , we set the expected growth rate of the stock equal to $r - q$ and discount the expected payoff at rate r . When the expected growth rate in the stock price is $r - q$, the expected stock price at time T is $S_0 e^{(r-q)T}$. A similar analysis to that in the appendix to Chapter 14 gives the expected payoff for a call option in a risk-neutral world as

$$e^{(r-q)T} S_0 N(d_1) - KN(d_2)$$

where d_1 and d_2 are defined as above. Discounting at rate r for time T leads to equation (16.4).

16.4 VALUATION OF EUROPEAN STOCK INDEX OPTIONS

In valuing index futures in Chapter 5, we assumed that the index could be treated as an asset paying a known yield. In valuing index options, we make similar assumptions. This means that equations (16.1) and (16.2) provide a lower bound for European index options; equation (16.3) is the put-call parity result for European index options; equations (16.4) and (16.5) can be used to value European options on an index; and the binomial tree approach can be used for American options. In all cases, S_0 is equal to the value of the index, σ is equal to the volatility of the index, and q is equal to the average annualized dividend yield on the index during the life of the option.

Example 16.1

Consider a European call option on the S&P 500 that is two months from maturity. The current value of the index is 930, the exercise price is 900, the risk-free interest rate is 8% per annum, and the volatility of the index is 20% per annum. Dividend

⁴ See Technical Note 6 at www.rotman.utoronto.ca/~hull/TechnicalNotes for a proof of this.

yields of 0.2% and 0.3% are expected in the first month and the second month, respectively. In this case $S_0 = 930$, $K = 900$, $r = 0.08$, $\sigma = 0.2$, and $T = 2/12$. The total dividend yield during the option's life is $0.2\% + 0.3\% = 0.5\%$. This corresponds to 3% per annum. Hence, $q = 0.03$ and

$$d_1 = \frac{\ln(930/900) + (0.08 - 0.03 + 0.2^2/2) \times 2/12}{0.2\sqrt{2/12}} = 0.5444$$

$$d_2 = \frac{\ln(930/900) + (0.08 - 0.03 - 0.2^2/2) \times 2/12}{0.2\sqrt{2/12}} = 0.4628$$

$$N(d_1) = 0.7069, \quad N(d_2) = 0.6782$$

so that the call price, c , is given by equation (16.4) as

$$c = 930 \times 0.7069e^{-0.03 \times 2/12} - 900 \times 0.6782e^{-0.08 \times 2/12} = 51.83$$

One contract would cost \$5,183.

The calculation of q should include only dividends for which the ex-dividend dates occur during the life of the option. In the United States ex-dividend dates tend to occur during the first week of February, May, August, and November. At any given time the correct value of q is therefore likely to depend on the life of the option. This is even more true for indices in other countries. In Japan, for example, all companies tend to use the same ex-dividend dates.

If the absolute amount of the dividend that will be paid on the stocks underlying the index (rather than the dividend yield) is assumed to be known, the basic Black–Scholes–Merton formulas can be used with the initial stock price being reduced by the present value of the dividends. This is the approach recommended in Chapter 14 for a stock paying known dividends. However, it may be difficult to implement for a broadly based stock index because it requires a knowledge of the dividends expected on every stock underlying the index.

It is sometimes argued that, in the long run, the return from investing a certain amount of money in a well-diversified stock portfolio is almost certain to beat the return from investing the same amount of money in a bond portfolio. If this were so, a long-dated put option allowing the stock portfolio to be sold for the value of the bond portfolio should not cost very much. In fact, as indicated by Business Snapshot 16.1, it is quite expensive.

Forward Prices

Define F_0 as the forward price of the index for a contract with maturity T . As shown by equation (5.3), $F_0 = S_0 e^{(r-q)T}$. This means that the equations for the European call price c and the European put price p in equations (16.4) and (16.5) can be written

$$c = F_0 e^{-rT} N(d_1) - K e^{-rT} N(d_2) \quad (16.8)$$

$$p = K e^{-rT} N(-d_2) - F_0 e^{-rT} N(-d_1) \quad (16.9)$$

where

$$d_1 = \frac{\ln(F_0/K) + \sigma^2 T/2}{\sigma\sqrt{T}} \quad \text{and} \quad d_2 = \frac{\ln(F_0/K) - \sigma^2 T/2}{\sigma\sqrt{T}}$$

Business Snapshot 16.1 Can We Guarantee that Stocks Will Beat Bonds in the Long Run?

It is often said that if you are a long-term investor you should buy stocks rather than bonds. Consider a US fund manager who is trying to persuade investors to buy, as a long-term investment, an equity fund that is expected to mirror the S&P 500. The manager might be tempted to offer purchasers of the fund a guarantee that their return will be at least as good as the return on risk-free bonds over the next 10 years. Historically stocks have outperformed bonds in the United States over almost any 10-year period. It appears that the fund manager would not be giving much away.

In fact, this type of guarantee is surprisingly expensive. Suppose that an equity index is 1,000 today, the dividend yield on the index is 1% per annum, the volatility of the index is 15% per annum, and the 10-year risk-free rate is 5% per annum. To outperform bonds, the stocks underlying the index must earn more than 5% per annum. The dividend yield will provide 1% per annum. The capital gains on the stocks must therefore provide 4% per annum. This means that we require the index level to be at least $1,000e^{0.04 \times 10} = 1,492$ in 10 years.

A guarantee that the return on \$1,000 invested in the index will be greater than the return on \$1,000 invested in bonds over the next 10 years is therefore equivalent to the right to sell the index for 1,492 in 10 years. This is a European put option on the index and can be valued from equation (16.5) with $S_0 = 1,000$, $K = 1,492$, $r = 5\%$, $\sigma = 15\%$, $T = 10$, and $q = 1\%$. The value of the put option is 169.7. This shows that the guarantee contemplated by the fund manager is worth about 17% of the fund—hardly something that should be given away!

The put–call parity relationship in equation (16.3) can be written

$$c + Ke^{-rT} = p + F_0e^{-rT}$$

or

$$F_0 = K + (c - p)e^{rT} \quad (16.10)$$

If, as is not uncommon in the exchange-traded markets, pairs of puts and calls with the same strike price are traded actively for a particular maturity date, this equation can be used to estimate the forward price of the index for that maturity date. Once the forward prices of the index for a number of different maturity dates have been obtained, the term structure of forward prices can be estimated, and other options can be valued using equations (16.8) and (16.9). The advantage of this approach is that the dividend yield on the index does not have to be estimated explicitly.

Implied Dividend Yields

If estimates of the dividend yield are required (e.g., because an American option is being valued), calls and puts with the same strike price and time to maturity can again be used. From equation (16.3),

$$q = -\frac{1}{T} \ln \frac{c - p + Ke^{-rT}}{S_0}$$

For a particular strike price and time to maturity, the estimates of q calculated from this equation are liable to be unreliable. But when the results from many matched pairs of calls and puts are combined, a clearer picture of the dividend yield being assumed by the market emerges.

16.5 VALUATION OF EUROPEAN CURRENCY OPTIONS

To value currency options, we define S_0 as the spot exchange rate. To be precise, S_0 is the value of one unit of the foreign currency in US dollars. As explained in Section 5.10, a foreign currency is analogous to a stock paying a known dividend yield. The owner of foreign currency receives a yield equal to the risk-free interest rate, r_f , in the foreign currency. Equations (16.1) and (16.2), with q replaced by r_f , provide bounds for the European call price, c , and the European put price, p :

$$c \geq \max(S_0 e^{-r_f T} - K e^{-r T}, 0)$$

$$p \geq \max(K e^{-r T} - S_0 e^{-r_f T}, 0)$$

Equation (16.3), with q replaced by r_f , provides the put–call parity result for European currency options:

$$c + K e^{-r T} = p + S_0 e^{-r_f T}$$

Finally, equations (16.4) and (16.5) provide the pricing formulas for European currency options when q is replaced by r_f :

$$c = S_0 e^{-r_f T} N(d_1) - K e^{-r T} N(d_2) \quad (16.11)$$

$$p = K e^{-r T} N(-d_2) - S_0 e^{-r_f T} N(-d_1) \quad (16.12)$$

where

$$d_1 = \frac{\ln(S_0/K) + (r - r_f + \sigma^2/2)T}{\sigma\sqrt{T}}$$

$$d_2 = \frac{\ln(S_0/K) + (r - r_f - \sigma^2/2)T}{\sigma\sqrt{T}} = d_1 - \sigma\sqrt{T}$$

Both the domestic interest rate, r , and the foreign interest rate, r_f , are the rates for a maturity T . Put and call options on a currency are symmetrical in that a put option to sell currency A for currency B at strike price K is the same as a call option to buy B with currency A at strike price $1/K$ (see Problem 16.8).

Example 16.2

Consider a four-month European call option on the British pound. Suppose that the current exchange rate is 1.6000, the exercise price is 1.6000, the risk-free interest rate in the United States is 8% per annum, the risk-free interest rate in Britain is 11% per annum, and the option price is 4.3 cents. In this case, $S_0 = 1.6$, $K = 1.6$, $r = 0.08$, $r_f = 0.11$, $T = 0.3333$, and $c = 0.043$. The implied volatility can be calculated by trial and error. A volatility of 20% gives an option price of 0.0639; a volatility of 10% gives an option price of 0.0285; and so on. The implied volatility is 14.1%.

Using Forward Exchange Rates

Because banks and other financial institutions trade forward contracts on foreign exchange rates actively, foreign exchange rates are often used for valuing options.

From equation (5.9), the forward rate, F_0 , for a maturity T is given by

$$F_0 = S_0 e^{(r-r_f)T}$$

This relationship allows equations (16.11) and (16.12) to be simplified to

$$c = e^{-rT} [F_0 N(d_1) - KN(d_2)] \quad (16.13)$$

$$p = e^{-rT} [KN(-d_2) - F_0 N(-d_1)] \quad (16.14)$$

where

$$d_1 = \frac{\ln(F_0/K) + \sigma^2 T/2}{\sigma\sqrt{T}}$$

$$d_2 = \frac{\ln(F_0/K) - \sigma^2 T/2}{\sigma\sqrt{T}} = d_1 - \sigma\sqrt{T}$$

Equations (16.13) and (16.14) are the same as equations (16.8) and (16.9). As we shall see in Chapter 17, a European option on the spot price of any asset can be valued in terms of the price of a forward or futures contract on the asset using equations (16.13) and (16.14). The maturity of the forward or futures contract must be the same as the maturity of the European option.

16.6 AMERICAN OPTIONS

As described in Chapter 12, binomial trees can be used to value American options on indices and currencies. As in the case of American options on a non-dividend-paying stock, the parameter determining the size of up movements, u , is set equal to $e^{\sigma\sqrt{\Delta t}}$, where σ is the volatility and Δt is the length of time steps. The parameter determining the size of down movements, d , is set equal to $1/u$, or $e^{-\sigma\sqrt{\Delta t}}$. For a non-dividend-paying stock, the probability of an up movement is

$$p = \frac{a - d}{u - d}$$

where $a = e^{r\Delta t}$. For options on indices and currencies, the formula for p is the same, but a is defined differently. In the case of options on an index,

$$a = e^{(r-q)\Delta t} \quad (16.15)$$

where q is the dividend yield on the index. In the case of options on a currency,

$$a = e^{(r-r_f)\Delta t} \quad (16.16)$$

where r_f is the foreign risk-free rate. Example 12.1 in Section 12.11 shows how a two-step tree can be constructed to value an option on an index. Example 12.2 shows how a three-step tree can be constructed to value an option on a currency. Further examples of the use of binomial trees to value options on indices and currencies are given in Chapter 20.

In some circumstances, it is optimal to exercise American currency and index options prior to maturity. Thus, American currency and index options are worth more than their European counterparts. In general, call options on high-interest currencies and put options on low-interest currencies are the most likely to be exercised prior to maturity. The reason is that a high-interest currency is expected to depreciate and a low-interest currency is expected to appreciate. In addition, call options on indices with high-dividend yields and put options on indices with low-dividend yields are most likely to be exercised early.

SUMMARY

The index options that trade on exchanges are settled in cash. On exercise of an index call option, the holder receives 100 times the amount by which the index exceeds the strike price. Similarly, on exercise of an index put option contract, the holder receives 100 times the amount by which the strike price exceeds the index. Index options can be used for portfolio insurance. If the value of the portfolio mirrors the index, it is appropriate to buy one put option contract for each $100S_0$ dollars in the portfolio, where S_0 is the value of the index. If the portfolio does not mirror the index, β put option contracts should be purchased for each $100S_0$ dollars in the portfolio, where β is the beta of the portfolio calculated using the capital asset pricing model. The strike price of the put options purchased should reflect the level of insurance required.

Most currency options are traded in the over-the-counter market. They can be used by corporate treasurers to hedge a foreign exchange exposure. For example, a US corporate treasurer who knows that the company will be receiving sterling at a certain time in the future can hedge by buying put options that mature at that time. Similarly, a US corporate treasurer who knows that the company will be paying sterling at a certain time in the future can hedge by buying call options that mature at that time. Currency options can also be used to create a range forward contract. This is a zero-cost contract that can be used to provide downside protection while giving up some of the upside for a company with a known foreign exchange exposure.

The Black–Scholes–Merton formula for valuing European options on a non-dividend-paying stock can be extended to cover European options on a stock paying a known dividend yield. The extension can be used to value European options on stock indices and currencies because:

1. A stock index is analogous to a stock paying a dividend yield. The dividend yield is the dividend yield on the stocks that make up the index.
2. A foreign currency is analogous to a stock paying a dividend yield. The foreign risk-free interest rate plays the role of the dividend yield.

Binomial trees can be used to value American options on stock indices and currencies.

FURTHER READING

Amin, K., and R. A. Jarrow. "Pricing Foreign Currency Options under Stochastic Interest Rates," *Journal of International Money and Finance*, 10 (1991): 310–29.

- Biger, N., and J.C. Hull. "The Valuation of Currency Options," *Financial Management*, 12 (Spring 1983): 24–28.
- Bodie, Z. "On the Risk of Stocks in the Long Run," *Financial Analysts Journal*, 51, 3 (1995): 18–22.
- Garman, M.B., and S.W. Kohlhagen. "Foreign Currency Option Values," *Journal of International Money and Finance*, 2 (December 1983): 231–37.
- Giddy, I.H., and G. Dufey. "Uses and Abuses of Currency Options," *Journal of Applied Corporate Finance*, 8, 3 (1995): 49–57.
- Grabbe, J.O. "The Pricing of Call and Put Options on Foreign Exchange," *Journal of International Money and Finance*, 2 (December 1983): 239–53.
- Jorion, P. "Predicting Volatility in the Foreign Exchange Market," *Journal of Finance* 50, 2(1995): 507–28.
- Merton, R.C. "Theory of Rational Option Pricing," *Bell Journal of Economics and Management Science*, 4 (Spring 1973): 141–83.

Practice Questions (Answers in Solutions Manual)

- 16.1. A portfolio is currently worth \$10 million and has a beta of 1.0. An index is currently standing at 800. Explain how a put option on the index with a strike price of 700 can be used to provide portfolio insurance.
- 16.2. "Once we know how to value options on a stock paying a dividend yield, we know how to value options on stock indices and currencies." Explain this statement.
- 16.3. A stock index is currently 300, the dividend yield on the index is 3% per annum, and the risk-free interest rate is 8% per annum. What is a lower bound for the price of a six-month European call option on the index when the strike price is 290?
- 16.4. A currency is currently worth \$0.80 and has a volatility of 12%. The domestic and foreign risk-free interest rates are 6% and 8%, respectively. Use a two-step binomial tree to value (a) a European four-month call option with a strike price of 0.79 and (b) an American four-month call option with the same strike price.
- 16.5. Explain how corporations can use range forward contracts to hedge their foreign exchange risk when they are due to receive a certain amount of a foreign currency in the future.
- 16.6. Calculate the value of a three-month at-the-money European call option on a stock index when the index is at 250, the risk-free interest rate is 10% per annum, the volatility of the index is 18% per annum, and the dividend yield on the index is 3% per annum.
- 16.7. Calculate the value of an eight-month European put option on a currency with a strike price of 0.50. The current exchange rate is 0.52, the volatility of the exchange rate is 12%, the domestic risk-free interest rate is 4% per annum, and the foreign risk-free interest rate is 8% per annum.
- 16.8. Show that the formula in equation (16.12) for a put option to sell one unit of currency A for currency B at strike price K gives the same value as equation (16.11) for a call option to buy K units of currency B for currency A at strike price $1/K$.
- 16.9. A foreign currency is currently worth \$1.50. The domestic and foreign risk-free interest rates are 5% and 9%, respectively. Calculate a lower bound for the value of a six-month call option on the currency with a strike price of \$1.40 if it is (a) European and (b) American.

- 16.10. Consider a stock index currently standing at 250. The dividend yield on the index is 4% per annum, and the risk-free rate is 6% per annum. A three-month European call option on the index with a strike price of 245 is currently worth \$10. What is the value of a three-month put option on the index with a strike price of 245?
- 16.11. An index currently stands at 696 and has a volatility of 30% per annum. The risk-free rate of interest is 7% per annum and the index provides a dividend yield of 4% per annum. Calculate the value of a three-month European put with an exercise price of 700.
- 16.12. Show that, if C is the price of an American call with exercise price K and maturity T on a stock paying a dividend yield of q , and P is the price of an American put on the same stock with the same strike price and exercise date, then

$$S_0 e^{-qT} - K < C - P < S_0 - K e^{-rT},$$

where S_0 is the stock price, r is the risk-free rate, and $r > 0$. (*Hint*: To obtain the first half of the inequality, consider possible values of:

Portfolio A: a European call option plus an amount K invested at the risk-free rate

Portfolio B: an American put option plus e^{-qT} of stock with dividends being reinvested in the stock.

To obtain the second half of the inequality, consider possible values of:

Portfolio C: an American call option plus an amount $K e^{-rT}$ invested at the risk-free rate

Portfolio D: a European put option plus one stock with dividends being reinvested in the stock.)

- 16.13. Show that a European call option on a currency has the same price as the corresponding European put option on the currency when the forward price equals the strike price.
- 16.14. Would you expect the volatility of a stock index to be greater or less than the volatility of a typical stock? Explain your answer.
- 16.15. Does the cost of portfolio insurance increase or decrease as the beta of a portfolio increases? Explain your answer.
- 16.16. Suppose that a portfolio is worth \$60 million and the S&P 500 is at 1,200. If the value of the portfolio mirrors the value of the index, what options should be purchased to provide protection against the value of the portfolio falling below \$54 million in one year's time?
- 16.17. Consider again the situation in Problem 16.16. Suppose that the portfolio has a beta of 2.0, the risk-free interest rate is 5% per annum, and the dividend yield on both the portfolio and the index is 3% per annum. What options should be purchased to provide protection against the value of the portfolio falling below \$54 million in one year's time?
- 16.18. An index currently stands at 1,500. European call and put options with a strike price of 1,400 and time to maturity of six months have market prices of 154.00 and 34.25, respectively. The six-month risk-free rate is 5%. What is the implied dividend yield?
- 16.19. A total return index tracks the return, including dividends, on a certain portfolio. Explain how you would value (a) forward contracts and (b) European options on the index.
- 16.20. What is the put-call parity relationship for European currency options?

- 16.21. Can an option on the yen–euro exchange rate be created from two options, one on the dollar–euro exchange rate, and the other on the dollar–yen exchange rate? Explain your answer.
- 16.22. Prove the results in equations (16.1), (16.2), and (16.3) using the portfolios indicated.

Further Questions

- 16.23. The Dow Jones Industrial Average on January 12, 2007, was 12,556 and the price of the March 126 call was \$2.25. Use the DerivaGem software to calculate the implied volatility of this option. Assume the risk-free rate was 5.3% and the dividend yield was 3%. The option expires on March 20, 2007. Estimate the price of a March 126 put. What is the volatility implied by the price you estimate for this option? (Note that options are on the Dow Jones index divided by 100.)
- 16.24. A stock index currently stands at 300 and has a volatility of 20%. The risk-free interest rate is 8% and the dividend yield on the index is 3%. Use a three-step binomial tree to value a six-month put option on the index with a strike price of 300 if it is (a) European and (b) American?
- 16.25. Suppose that the spot price of the Canadian dollar is US \$0.95 and that the Canadian dollar/US dollar exchange rate has a volatility of 8% per annum. The risk-free rates of interest in Canada and the United States are 4% and 5% per annum, respectively. Calculate the value of a European call option to buy one Canadian dollar for US \$0.95 in nine months. Use put–call parity to calculate the price of a European put option to sell one Canadian dollar for US \$0.95 in nine months. What is the price of a call option to buy US \$0.95 with one Canadian dollar in nine months?
- 16.26. Hedge funds earn a fixed fee plus a percentage of the profits, if any, that they generate (see Business Snapshot 1.2). How is a fund manager motivated to behave with this type of arrangement?
- 16.27. Assume that the price of currency A expressed in terms of the price of currency B follows the process $dS = (r_B - r_A)S dt + \sigma S dz$, where r_A is the risk-free interest rate in currency A and r_B is the risk-free interest rate in currency B. What is the process followed by the price of currency B expressed in terms of currency A?
- 16.28. The USD/euro exchange rate is 1.3000. The exchange rate volatility is 15%. A US company will receive 1 million euros in three months. The euro and USD risk-free rates are 5% and 4%, respectively. The company decides to use a range forward contract with the lower strike price equal to 1.2500.
- What should the higher strike price be to create a zero-cost contract?
 - What position in calls and puts should the company take?
 - Show that your answer to (a) does not depend on interest rates provided that the interest rate differential between the two currencies, $r - r_f$, remains the same.
- 16.29. In Business Snapshot 16.1, what is the cost of a guarantee that the return on the fund will not be negative over the next 10 years?



17

CHAPTER

Futures Options

The options we have considered so far provide the holder with the right to buy or sell a certain asset by a certain date for a certain price. They are sometimes termed *options on spot* or *spot options* because, when the options are exercised, the sale or purchase of the asset at the agreed-on price takes place immediately. In this chapter we move on to consider *options on futures*, also known as *futures options*. In these contracts, exercise of the option gives the holder a position in a futures contract.

The Commodity Futures Trading Commission in the US authorized the trading of options on futures on an experimental basis in 1982. Permanent trading was approved in 1987, and since then the popularity of the contract with investors has grown very fast.

In this chapter we consider how futures options work and the differences between these options and spot options. We examine how futures options can be priced using either binomial trees or formulas similar to those produced by Black, Scholes, and Merton for stock options. We also explore the relative pricing of futures options and spot options.

17.1 NATURE OF FUTURES OPTIONS

A futures option is the right, but not the obligation, to enter into a futures contract at a certain futures price by a certain date. Specifically, a call futures option is the right to enter into a long futures contract at a certain price; a put futures option is the right to enter into a short futures contract at a certain price. Futures options are generally American; that is, they can be exercised any time during the life of the contract.

If a call futures option is exercised, the holder acquires a long position in the underlying futures contract plus a cash amount equal to the most recent settlement futures price minus the strike price. If a put futures option is exercised, the holder acquires a short position in the underlying futures contract plus a cash amount equal to the strike price minus the most recent settlement futures price. As the following examples show, the effective payoff from a call futures option is $\max(F_T - K, 0)$ and the effective payoff from a put futures option is $\max(K - F_T, 0)$, where F_T is the futures price at the time of exercise and K is the strike price.

Example 17.1

Suppose it is August 15 and an investor has one September futures call option contract on copper with a strike price of 240 cents per pound. One futures contract

is on 25,000 pounds of copper. Suppose that the futures price of copper for delivery in September is currently 251 cents, and at the close of trading on August 14 (the last settlement) it was 250 cents. If the option is exercised, the investor receives a cash amount of

$$25,000 \times (250 - 240) \text{ cents} = \$2,500$$

plus a long position in a futures contract to buy 25,000 pounds of copper in September. If desired, the position in the futures contract can be closed out immediately. This would leave the investor with the \$2,500 cash payoff plus an amount

$$25,000 \times (251 - 250) \text{ cents} = \$250$$

reflecting the change in the futures price since the last settlement. The total payoff from exercising the option on August 15 is \$2,750, which equals $25,000(F - K)$, where F is the futures price at the time of exercise and K is the strike price.

Example 17.2

An investor has one December futures put option on corn with a strike price of 400 cents per bushel. One futures contract is on 5,000 bushels of corn. Suppose that the current futures price of corn for delivery in December is 380, and the most recent settlement price is 379 cents. If the option is exercised, the investor receives a cash amount of

$$5,000 \times (400 - 379) \text{ cents} = \$1,050$$

plus a short position in a futures contract to sell 5,000 bushels of corn in December. If desired, the position in the futures contract can be closed out. This would leave the investor with the \$1,050 cash payoff minus an amount

$$5,000 \times (380 - 379) \text{ cents} = \$50$$

reflecting the change in the futures price since the last settlement. The net payoff from exercise is \$1,000, which equals $5,000(K - F)$, where F is the futures price at the time of exercise and K is the strike price.

Expiration Months

Futures options are referred to by the delivery month of the underlying futures contract—not by the expiration month of the option. As mentioned earlier, most futures options are American. The expiration date of a futures option contract is usually on, or a few days before, the earliest delivery date of the underlying futures contract. (For example, the CME Group Treasury bond futures option expires on the latest Friday that precedes by at least five business days the end of the month before the futures delivery month.) An exception is the CME Group mid-curve Eurodollar contract where the futures contract expires either one or two years after the options contract.

Popular contracts trading in the United States are those on corn, soybeans, cotton, sugar-world, crude oil, natural gas, gold, Treasury bonds, Treasury notes, five-year Treasury notes, 30-day federal funds, Eurodollars, one-year and two-year mid-curve Eurodollars, Euribor, Eurobunds, and the S&P 500.

Options on Interest Rate Futures

The most actively traded interest rate options offered by exchanges in the United States are those on Treasury bond futures, Treasury note futures, and Eurodollar futures.

A Treasury bond futures option, which is traded by the CME Group, is an option to enter a Treasury bond futures contract. As mentioned in Chapter 6, one Treasury bond futures contract is for the delivery of \$100,000 of Treasury bonds. The price of a Treasury bond futures option is quoted as a percentage of the face value of the underlying Treasury bonds to the nearest sixty-fourth of 1%.

An option on Eurodollar futures, which is traded by the CME Group, is an option to enter into a Eurodollar futures contract. As explained in Chapter 6, when the Eurodollar futures quote changes by 1 basis point, or 0.01%, there is a gain or loss on a Eurodollar futures contract of \$25. Similarly, in the pricing of options on Eurodollar futures, 1 basis point represents \$25.

Interest rate futures option contracts work in the same way as the other futures options contracts discussed in this chapter. For example, in addition to the cash payoff, the holder of a call option obtains a long position in the futures contract when the option is exercised and the option writer obtains a corresponding short position. The total payoff from the call, including the value of the futures position, is $\max(F - K, 0)$, where F is the futures price at the time of exercise and K is the strike price.

Interest rate futures prices increase when bond prices increase (i.e., when interest rates fall). They decrease when bond prices decrease (i.e., when interest rates rise). An investor who thinks that short-term interest rates will rise can speculate by buying put options on Eurodollar futures, whereas an investor who thinks the rates will fall can speculate by buying call options on Eurodollar futures. An investor who thinks that long-term interest rates will rise can speculate by buying put options on Treasury note futures or Treasury bond futures, whereas an investor who thinks the rates will fall can speculate by buying call options on these instruments.

Example 17.3

It is February and the futures price for the June Eurodollar contract is 93.82 (corresponding to a 3-month Eurodollar interest rate of 6.18% per annum). The price of a call option on the contract with a strike price of 94.00 is quoted as 0.1, or 10 basis points. This option could be attractive to an investor who feels that interest rates are likely to come down. Suppose that short-term interest rates do drop by about 100 basis points and the investor exercises the call when the Eurodollar futures price is 94.78 (corresponding to a 3-month Eurodollar interest rate of 5.22% per annum). The payoff is $25 \times (94.78 - 94.00) \times 100 = \$1,950$. The cost of the contract is $10 \times 25 = \$250$. The investor's profit is therefore \$1,700.

Example 17.4

It is August and the futures price for the December Treasury bond contract is 96-09 (or $96\frac{9}{32} = 96.28125$). The yield on long-term government bonds is about 6.4% per annum. An investor who feels that this yield will fall by December might choose to buy December calls with a strike price of 98. Assume that the price of these calls is 1-04 (or $1\frac{4}{64} = 1.0625\%$ of the principal). If long-term rates

fall to 6% per annum and the Treasury bond futures price rises to 100-00, the investor will make a net profit per \$100 of bond futures of

$$100.00 - 98.00 - 1.0625 = 0.9375$$

Since one option contract is for the purchase or sale of instruments with a face value of \$100,000, the investor's profit is \$937.50 per option contract bought.

17.2 REASONS FOR THE POPULARITY OF FUTURES OPTIONS

It is natural to ask why people choose to trade options on futures rather than options on the underlying asset. The main reason appears to be that a futures contract is, in many circumstances, more liquid and easier to trade than the underlying asset. Furthermore, a futures price is known immediately from trading on the futures exchange, whereas the spot price of the underlying asset may not be so readily available.

Consider Treasury bonds. The market for Treasury bond futures is much more active than the market for any particular Treasury bond. Also, a Treasury bond futures price is known immediately from exchange trading. By contrast, the current market price of a bond can be obtained only by contacting one or more dealers. It is not surprising that investors would rather take delivery of a Treasury bond futures contract than Treasury bonds.

Futures on commodities are also often easier to trade than the commodities themselves. For example, it is much easier and more convenient to make or take delivery of a live-cattle futures contract than it is to make or take delivery of the cattle themselves.

An important point about a futures option is that exercising it does not usually lead to delivery of the underlying asset, as in most circumstances the underlying futures contract is closed out prior to delivery. Futures options are therefore normally eventually settled in cash. This is appealing to many investors, particularly those with limited capital who may find it difficult to come up with the funds to buy the underlying asset when an option on spot is exercised. Another advantage sometimes cited for futures options is that futures and futures options are traded side by side in the same exchange. This facilitates hedging, arbitrage, and speculation. It also tends to make the markets more efficient. A final point is that futures options entail lower transactions costs than spot options in many situations.

17.3 EUROPEAN SPOT AND FUTURES OPTIONS

The payoff from a European call option with strike price K on the spot price of an asset is

$$\max(S_T - K, 0)$$

where S_T is the spot price at the option's maturity. The payoff from a European call option with the same strike price on the futures price of the asset is

$$\max(F_T - K, 0)$$

where F_T is the futures price at the option's maturity. If the futures contract matures at

the same time as the option, then $F_T = S_T$ and the two options are equivalent. Similarly, a European futures put option is worth the same as its spot put option counterpart when the futures contract matures at the same time as the option.

Most of the futures options that trade are American-style. However, as we shall see, it is useful to study European futures options because the results that are obtained can be used to value the corresponding European spot options.

17.4 PUT–CALL PARITY

In Chapter 10, we derived a put–call parity relationship for European stock options. We now consider a similar argument to derive a put–call parity relationship for European futures options. Consider European call and put futures options, both with strike price K and time to expiration T . We can form two portfolios:

Portfolio A: a European call futures option plus an amount of cash equal to Ke^{-rT}

Portfolio B: a European put futures option plus a long futures contract plus an amount of cash equal to F_0e^{-rT} , where F_0 is the futures price

In portfolio A, the cash can be invested at the risk-free rate, r , and grows to K at time T . Let F_T be the futures price at maturity of the option. If $F_T > K$, the call option in portfolio A is exercised and portfolio A is worth F_T . If $F_T \leq K$, the call is not exercised and portfolio A is worth K . The value of portfolio A at time T is therefore

$$\max(F_T, K)$$

In portfolio B, the cash can be invested at the risk-free rate to grow to F_0 at time T . The put option provides a payoff of $\max(K - F_T, 0)$. The futures contract provides a payoff of $F_T - F_0$.¹ The value of portfolio B at time T is therefore

$$F_0 + (F_T - F_0) + \max(K - F_T, 0) = \max(F_T, K)$$

Because the two portfolios have the same value at time T and European options cannot be exercised early, it follows that they are worth the same today. The value of portfolio A today is

$$c + Ke^{-rT}$$

where c is the price of the call futures option. The daily settlement process ensures that the futures contract in portfolio B is worth zero today. Portfolio B is therefore worth

$$p + F_0e^{-rT}$$

where p is the price of the put futures option. Hence

$$c + Ke^{-rT} = p + F_0e^{-rT} \tag{17.1}$$

The difference between this put–call parity relationship and the one for a non-dividend-paying stock in equation (10.6) is that the stock price, S_0 , is replaced by the discounted futures price, F_0e^{-rT} .

¹ This analysis assumes that a futures contract is like a forward contract and settled at the end of its life rather than on a day-to-day basis.

As shown in Section 17.3, when the underlying futures contract matures at the same time as the option, European futures and spot options are the same. Equation (17.1) therefore gives a relationship between the price of a call option on the spot price, the price of a put option on the spot price, and the futures price when both options mature at the same time as the futures contract.

Example 17.5

Suppose that the price of a European call option on spot silver for delivery in six months is \$0.56 per ounce when the exercise price is \$8.50. Assume that the silver futures price for delivery in six months is currently \$8.00, and the risk-free interest rate for an investment that matures in six months is 10% per annum. From a rearrangement of equation (17.1), the price of a European put option on spot silver with the same maturity and exercise date as the call option is

$$0.56 + 8.50e^{-0.1 \times 6/12} - 8.00e^{-0.1 \times 6/12} = 1.04$$

For American futures options, the put–call relationship is (see Problem 17.19)

$$F_0e^{-rT} - K < C - P < F_0 - Ke^{-rT} \quad (17.2)$$

17.5 BOUNDS FOR FUTURES OPTIONS

The put–call parity relationship in equation (17.1) provides bounds for European call and put options. Because the price of a put, p , cannot be negative, it follows from equation (17.1) that

$$c + Ke^{-rT} \geq F_0e^{-rT}$$

or

$$c \geq (F_0 - K)e^{-rT} \quad (17.3)$$

Similarly, because the price of a call option cannot be negative, it follows from equation (17.1) that

$$Ke^{-rT} \leq F_0e^{-rT} + p$$

or

$$p \geq (K - F_0)e^{-rT} \quad (17.4)$$

These bounds are similar to the ones derived for European stock options in Chapter 10. The prices of European call and put options are very close to their lower bounds when the options are deep in the money. To see why this is so, we return to the put–call parity relationship in equation (17.1). When a call option is deep in the money, the corresponding put option is deep out of the money. This means that p is very close to zero. The difference between c and its lower bound equals p , so that the price of the call option must be very close to its lower bound. A similar argument applies to put options.

Because American futures options can be exercised at any time, we must have

$$C \geq F_0 - K$$

and

$$P \geq K - F_0$$

Thus, assuming interest rates are positive, the lower bound for an American option

price is always higher than the lower bound for the corresponding European option price. There is always some chance that an American futures option will be exercised early.

17.6 VALUATION OF FUTURES OPTIONS USING BINOMIAL TREES

This section examines, more formally than in Chapter 12, how binomial trees can be used to price futures options. A key difference between futures options and stock options is that there are no up-front costs when a futures contract is entered into.

Suppose that the current futures price is 30 and that it will move either up to 33 or down to 28 over the next month. We consider a one-month call option on the futures with a strike price of 29 and ignore daily settlement. The situation is as indicated in Figure 17.1. If the futures price proves to be 33, the payoff from the option is 4 and the value of the futures contract is 3. If the futures price proves to be 28, the payoff from the option is zero and the value of the futures contract is -2 .²

To set up a riskless hedge, we consider a portfolio consisting of a short position in one options contract and a long position in Δ futures contracts. If the futures price moves up to 33, the value of the portfolio is $3\Delta - 4$; if it moves down to 28, the value of the portfolio is -2Δ . The portfolio is riskless when these are the same, that is, when

$$3\Delta - 4 = -2\Delta$$


or $\Delta = 0.8$.

For this value of Δ , we know the portfolio will be worth $3 \times 0.8 - 4 = -1.6$ in one month. Assume a risk-free interest rate of 6%. The value of the portfolio today must be

$$-1.6e^{-0.06 \times 1/12} = -1.592$$

The portfolio consists of one short option and Δ futures contracts. Because the value of the futures contract today is zero, the value of the option today must be 1.592.

Figure 17.1 Futures price movements in numerical example.



² There is an approximation here in that the gain or loss on the futures contract is not realized at time T . It is realized day by day between time 0 and time T . However, as the length of the time step in a binomial tree becomes shorter, the approximation becomes better.

A Generalization

We can generalize this analysis by considering a futures price that starts at F_0 and is anticipated to rise to F_0u or move down to F_0d over the time period T . We consider an option maturing at time T and suppose that its payoff is f_u if the futures price moves up and f_d if it moves down. The situation is summarized in Figure 17.2.

The riskless portfolio in this case consists of a short position in one option combined with a long position in Δ futures contracts, where

$$\Delta = \frac{f_u - f_d}{F_0u - F_0d}$$

The value of the portfolio at time T is then always

$$(F_0u - F_0)\Delta - f_u$$

Denoting the risk-free interest rate by r , we obtain the value of the portfolio today as

$$[(F_0u - F_0)\Delta - f_u]e^{-rT}$$

Another expression for the present value of the portfolio is $-f$, where f is the value of the option today. It follows that

$$-f = [(F_0u - F_0)\Delta - f_u]e^{-rT}$$

Substituting for Δ and simplifying reduces this equation to

$$f = e^{-rT}[pf_u + (1 - p)f_d] \quad (17.5)$$

where


$$p = \frac{1 - d}{u - d} \quad (17.6)$$

This agrees with the result in Section 12.9. Equation (17.6) gives the risk-neutral probability of an up movement.

In the numerical example considered previously (see Figure 17.1), $u = 1.1$, $d = 0.9333$, $r = 0.06$, $T = 1/12$, $f_u = 4$, and $f_d = 0$. From equation (17.6),

$$p = \frac{1 - 0.9333}{1.1 - 0.9333} = 0.4$$

Figure 17.2 Futures price and option price in a general situation.



and, from equation (17.5),

$$f = e^{-0.06 \times 1/12} [0.4 \times 4 + 0.6 \times 0] = 1.592$$

This result agrees with the answer obtained for this example earlier.

Multistep Trees

Multistep binomial trees are used to value American-style futures options in much the same way that they are used to value options on stocks. This is explained in Section 12.9. The parameter u defining up movements in the futures price is $e^{\sigma\sqrt{\Delta t}}$, where σ is the volatility of the futures price and Δt is the length of one time step. The probability of an up movement in the future price is that in equation (17.6):

$$p = \frac{1 - d}{u - d}$$

Example 12.3 illustrates the use of multistep binomial trees for valuing a futures option. Example 20.3 in Chapter 20 provides a further illustration.

17.7 DRIFT OF A FUTURES PRICE IN A RISK-NEUTRAL WORLD

There is a general result that allows us to use the analysis in Section 16.3 for futures options. This result is that in a risk-neutral world a futures price behaves in the same way as a stock paying a dividend yield at the domestic risk-free interest rate r .

One clue that this might be so is given by noting that the equation for the probability p in a binomial tree for a futures price is the same as that for a stock paying a dividend yield equal to q when $q = r$ (compare equation (17.6) with equations (16.15) and (16.16)). Another clue is that the put–call parity relationship for futures options prices is the same as that for options on a stock paying a dividend yield at rate q when the stock price is replaced by the futures price and $q = r$ (compare equations (17.1) and (16.3)).

To prove the result formally, we calculate the drift of a futures price in a risk-neutral world. We define F_t as the futures price at time t and suppose the settlement dates to be at times $0, \Delta t, 2\Delta t, \dots$. If we enter into a long futures contract at time 0, its value is zero. At time Δt , it provides a payoff of $F_{\Delta t} - F_0$. If r is the very-short-term (Δt -period) interest rate at time 0, risk-neutral valuation gives the value of the contract at time 0 as

$$e^{-r\Delta t} \hat{E}[F_{\Delta t} - F_0]$$

where \hat{E} denotes expectations in a risk-neutral world. We must therefore have

$$e^{-r\Delta t} \hat{E}(F_{\Delta t} - F_0) = 0$$

showing that

$$\hat{E}(F_{\Delta t}) = F_0$$

Similarly, $\hat{E}(F_{2\Delta t}) = F_{\Delta t}$, $\hat{E}(F_{3\Delta t}) = F_{2\Delta t}$, and so on. Putting many results like this together, we see that

$$\hat{E}(F_T) = F_0$$

for any time T .

The drift of the futures price in a risk-neutral world is therefore zero. From equation (16.7), the futures price behaves like a stock providing a dividend yield q equal to r . This result is a very general one. It is true for all futures prices and does not depend on any assumptions about interest rates, volatilities, etc.³

The usual assumption made for the process followed by a futures price F in the risk-neutral world is

$$dF = \sigma F dz \quad (17.7)$$

where σ is a constant.

Differential Equation

For another way of seeing that a futures price behaves like a stock paying a dividend yield at rate q , we can derive the differential equation satisfied by a derivative dependent on a futures price in the same way as we derived the differential equation for a derivative dependent on a non-dividend-paying stock in Section 14.6. This is⁴

$$\frac{\partial f}{\partial t} + \frac{1}{2} \frac{\partial^2 f}{\partial F^2} \sigma^2 F^2 = rf \quad (17.8)$$

It has the same form as equation (16.6) with q set equal to r . This confirms that, for the purpose of valuing derivatives, a futures price can be treated in the same way as a stock providing a dividend yield at rate r .

17.8 BLACK'S MODEL FOR VALUING FUTURES OPTIONS

European futures options can be valued by extending the results we have produced. Fischer Black was the first to show this in a paper published in 1976.⁵ Assuming that the futures price follows the (lognormal) process in equation (17.7), the European call price c and the European put price p for a futures option are given by equations (16.4) and (16.5) with S_0 replaced by F_0 and $q = r$:

$$c = e^{-rT} [F_0 N(d_1) - KN(d_2)] \quad (17.9)$$

$$p = e^{-rT} [KN(-d_2) - F_0 N(-d_1)] \quad (17.10)$$

where

$$d_1 = \frac{\ln(F_0/K) + \sigma^2 T/2}{\sigma \sqrt{T}}$$

$$d_2 = \frac{\ln(F_0/K) - \sigma^2 T/2}{\sigma \sqrt{T}} = d_1 - \sigma \sqrt{T}$$

and σ is the volatility of the futures price. When the cost of carry and the convenience

³ As we will discover in Chapter 27, a more precise statement of the result is: "A futures price has zero drift in the traditional risk-neutral world where the numeraire is the money market account." A zero-drift stochastic process is known as a martingale. A forward price is a martingale in a different risk-neutral world. This is one where the numeraire is a zero-coupon bond maturing at time T .

⁴ See Technical Note 7 at www.rotman.utoronto.ca/~hull/TechnicalNotes for a proof of this.

⁵ See F. Black, "The Pricing of Commodity Contracts," *Journal of Financial Economics*, 3 (March 1976): 167–79.

yield are functions only of time, it can be shown that the volatility of the futures price is the same as the volatility of the underlying asset.

Example 17.6

Consider a European put futures option on crude oil. The time to the option's maturity is 4 months, the current futures price is \$20, the exercise price is \$20, the risk-free interest rate is 9% per annum, and the volatility of the futures price is 25% per annum. In this case, $F_0 = 20$, $K = 20$, $r = 0.09$, $T = 4/12$, $\sigma = 0.25$, and $\ln(F_0/K) = 0$, so that

$$d_1 = \frac{\sigma\sqrt{T}}{2} = 0.07216$$

$$d_2 = -\frac{\sigma\sqrt{T}}{2} = -0.07216$$

$$N(-d_1) = 0.4712, \quad N(-d_2) = 0.5288$$

and the put price p is given by

$$p = e^{-0.09 \times 4/12} (20 \times 0.5288 - 20 \times 0.4712) = 1.12$$

or \$1.12.

Using Black's Model Instead of Black–Scholes–Merton

The results in Section 17.3 show that European futures options and European spot options are equivalent when the option contract matures at the same time as the futures contract. Equations (17.9) and (17.10) therefore provide a way of calculating the value of European options on the spot price of a asset.

Example 17.7

Consider a six-month European call option on the spot price of gold, that is, an option to buy one ounce of gold in the spot market in six months. The strike price is \$1,200, the six-month futures price of gold is \$1,240, the risk-free rate of interest is 5% per annum, and the volatility of the futures price is 20%. The option is the same as a six-month European option on the six-month futures price. The value of the option is therefore given by equation (17.9) as

$$e^{-0.05 \times 0.5} [1,240N(d_1) - 1,200N(d_2)]$$

where

$$d_1 = \frac{\ln(1,240/1,200) + 0.2^2 \times 0.5/2}{0.2 \times \sqrt{0.5}} = 0.3026$$

$$d_2 = \frac{\ln(1,240/1,200) - 0.2^2 \times 0.5/2}{0.2 \times \sqrt{0.5}} = 0.1611$$

It is \$88.37.

Traders like to use Black's model rather than Black–Scholes–Merton to value European spot options. It has a fairly general applicability. The underlying asset can be a consumption or investment asset and it can provide income to the holder. The variable F_0 in equations (17.9) and (17.10) is set equal to either the futures or the forward price of the underlying asset for a contract maturing at the same time as the option.

Equations (16.13) and (16.14) show Black's model being used to value European options on the spot value of a currency. Equations (16.8) and (16.9) show Black's model being used to value European options on the spot value of an index. The big advantage of Black's model is that it avoids the need to estimate the income (or convenience yield) on the underlying asset. The futures or forward price that is used in the model incorporate the market's estimate of this income.

When considering stock indices in Section 16.4, we explained that put–call parity is used to imply the forward prices for maturities for which there are actively traded options. Interpolation is then used to estimate forward prices for other maturities. The same approach can be used for a wide range of other underlying assets.

17.9 AMERICAN FUTURES OPTIONS vs. AMERICAN SPOT OPTIONS

Traded futures options are in practice usually American. Assuming that the risk-free rate of interest, r , is positive, there is always some chance that it will be optimal to exercise an American futures option early. American futures options are therefore worth more than their European counterparts.

It is not generally true that an American futures option is worth the same as the corresponding American spot option when the futures and options contracts have the same maturity.⁶ Suppose, for example, that there is a normal market with futures prices consistently higher than spot prices prior to maturity. An American call futures option must be worth more than the corresponding American spot call option. The reason is that in some situations the futures option will be exercised early, in which case it will provide a greater profit to the holder. Similarly, an American put futures option must be worth less than the corresponding American spot put option. If there is an inverted market with futures prices consistently lower than spot prices, the reverse must be true. American call futures options are worth less than the corresponding American spot call option, whereas American put futures options are worth more than the corresponding American spot put option.

The differences just described between American futures options and American spot options hold true when the futures contract expires later than the options contract as well as when the two expire at the same time. In fact, the later the futures contract expires the greater the differences tend to be.

17.10 FUTURES-STYLE OPTIONS

Some exchanges trade what are termed *futures-style options*. These are futures contracts on the payoff from an option. Normally a trader who buys (sells) an option, whether on the spot price of an asset or on the futures price of an asset, pays (receives) cash up front. By contrast, traders who buy or sell a futures-style option post margin in the same way that they do on a regular futures contract (see Chapter 2). The contract is settled daily as with any other futures contract and the final settlement price is the payoff from the option. Just as a futures contract is a bet on what the future price of an

⁶ The spot option “corresponding” to a futures option is defined here as one with the same strike price and the same expiration date.

asset will be, a futures-style option is a bet on what the payoff from an option will be.⁷ If interest rates are constant, the futures price in a futures-style option is the same as the forward price in a forward contract on the option payoff. This shows that the futures price for a futures-style option is the price that would be paid for the option if payment were made in arrears. It is therefore the value of a regular option compounded forward at the risk-free rate.

Black's model in equations (17.9) and (17.10) gives the price of a regular European option on an asset in terms of the futures (or forward) price F_0 for a contract maturing at the same time as the option. The futures price in a call futures-style option is therefore

$$F_0N(d_1) - KN(d_2)$$

and the futures price in a put futures-style option is

$$KN(-d_2) - F_0N(-d_1)$$

where d_1 and d_2 are as defined in equations (17.9) and (17.10). These formulas do not depend on the level of interest rates. They are correct for a futures-style option on a futures contract and a futures-style option on the spot value of an asset. In the first case, F_0 is the current futures price for the contract underlying the option; in the second case, it is the current futures price for a futures contract on the underlying asset maturing at the same time as the option.

The put–call parity relationship for a futures-style options is

$$p + F_0 = c + K$$

An American futures-style option can be exercised early, in which case there is an immediate final settlement at the option's intrinsic value. As it turns out, it is never optimal to exercise an American futures-style options on a futures contract early because the futures price of the option is always greater than the intrinsic value. This type of American futures-style option can therefore be treated as though it were the corresponding European futures-style option.

SUMMARY

Futures options require delivery of the underlying futures contract on exercise. When a call is exercised, the holder acquires a long futures position plus a cash amount equal to the excess of the futures price over the strike price. Similarly, when a put is exercised the holder acquires a short position plus a cash amount equal to the excess of the strike price over the futures price. The futures contract that is delivered usually expires slightly later than the option.

A futures price behaves in the same way as a stock that provides a dividend yield equal to the risk-free rate, r . This means that the results produced in Chapter 16 for options on a stock paying a dividend yield apply to futures options if we replace the stock price by the futures price and set the dividend yield equal to the risk-free interest

⁷ For a more detailed discussion of futures-style options, see D. Lieu, "Option Pricing with Futures-Style Margining," *Journal of Futures Markets*, 10, 4 (1990), 327–38. For pricing when interest rates are stochastic, see R.-R. Chen and L. Scott, "Pricing Interest Rate Futures Options with Futures-Style Margining." *Journal of Futures Markets*, 13, 1 (1993) 15–22.

rate. Pricing formulas for European futures options were first produced by Fischer Black in 1976. They assume that the futures price is lognormally distributed at the option's expiration.

If the expiration dates for the option and futures contracts are the same, a European futures option is worth exactly the same as the corresponding European spot option. This result is often used to value European spot options. The result is not true for American options. If the futures market is normal, an American call futures is worth more than the corresponding American spot call option, while an American put futures is worth less than the corresponding American spot put option. If the futures market is inverted, the reverse is true.

FURTHER READING

- Black, F. "The Pricing of Commodity Contracts," *Journal of Financial Economics*, 3 (1976): 167–79.
- Hilliard, J. E., and J. Reis. "Valuation of Commodity Futures and Options under Stochastic Convenience Yields, Interest Rates, and Jump Diffusions in the Spot," *Journal of Financial and Quantitative Analysis*, 33, 1 (March 1998): 61–86.
- Miltersen, K. R., and E. S. Schwartz. "Pricing of Options on Commodity Futures with Stochastic Term Structures of Convenience Yields and Interest Rates," *Journal of Financial and Quantitative Analysis*, 33, 1 (March 1998): 33–59.

Practice Questions (Answers in Solutions Manual)

- 17.1. Explain the difference between a call option on yen and a call option on yen futures.
- 17.2. Why are options on bond futures more actively traded than options on bonds?
- 17.3. "A futures price is like a stock paying a dividend yield." What is the dividend yield?
- 17.4. A futures price is currently 50. At the end of six months it will be either 56 or 46. The risk-free interest rate is 6% per annum. What is the value of a six-month European call option on the futures with a strike price of 50?
- 17.5. How does the put–call parity formula for a futures option differ from put–call parity for an option on a non-dividend-paying stock?
- 17.6. Consider an American futures call option where the futures contract and the option contract expire at the same time. Under what circumstances is the futures option worth more than the corresponding American option on the underlying asset?
- 17.7. Calculate the value of a five-month European put futures option when the futures price is \$19, the strike price is \$20, the risk-free interest rate is 12% per annum, and the volatility of the futures price is 20% per annum.
- 17.8. Suppose you buy a put option contract on October gold futures with a strike price of \$1,200 per ounce. Each contract is for the delivery of 100 ounces. What happens if you exercise when the October futures price is \$1,180?
- 17.9. Suppose you sell a call option contract on April live cattle futures with a strike price of 90 cents per pound. Each contract is for the delivery of 40,000 pounds. What happens if the contract is exercised when the futures price is 95 cents?

- 17.10. Consider a two-month call futures option with a strike price of 40 when the risk-free interest rate is 10% per annum. The current futures price is 47. What is a lower bound for the value of the futures option if it is (a) European and (b) American?
- 17.11. Consider a four-month put futures option with a strike price of 50 when the risk-free interest rate is 10% per annum. The current futures price is 47. What is a lower bound for the value of the futures option if it is (a) European and (b) American?
- 17.12. A futures price is currently 60 and its volatility is 30%. The risk-free interest rate is 8% per annum. Use a two-step binomial tree to calculate the value of a six-month European call option on the futures with a strike price of 60. If the call were American, would it ever be worth exercising it early?
- 17.13. In Problem 17.12, what does the binomial tree give for the value of a six-month European put option on futures with a strike price of 60? If the put were American, would it ever be worth exercising it early? Verify that the call prices calculated in Problem 17.12 and the put prices calculated here satisfy put–call parity relationships.
- 17.14. A futures price is currently 25, its volatility is 30% per annum, and the risk-free interest rate is 10% per annum. What is the value of a nine-month European call on the futures with a strike price of 26?
- 17.15. A futures price is currently 70, its volatility is 20% per annum, and the risk-free interest rate is 6% per annum. What is the value of a five-month European put on the futures with a strike price of 65?
- 17.16. Suppose that a one-year futures price is currently 35. A one-year European call option and a one-year European put option on the futures with a strike price of 34 are both priced at 2 in the market. The risk-free interest rate is 10% per annum. Identify an arbitrage opportunity.
- 17.17. “The price of an at-the-money European call futures option always equals the price of a similar at-the-money European put futures option.” Explain why this statement is true.
- 17.18. Suppose that a futures price is currently 30. The risk-free interest rate is 5% per annum. A three-month American call futures option with a strike price of 28 is worth 4. Calculate bounds for the price of a three-month American put futures option with a strike price of 28.
- 17.19. Show that, if C is the price of an American call option on a futures contract when the strike price is K and the maturity is T , and P is the price of an American put on the same futures contract with the same strike price and exercise date, then

$$F_0 e^{-rT} - K < C - P < F_0 - K e^{-rT}$$

where F_0 is the futures price and r is the risk-free rate. Assume that $r > 0$ and that there is no difference between forward and futures contracts. (*Hint*: Use an analogous approach to that indicated for Problem 16.12.)

- 17.20. Calculate the price of a three-month European call option on the spot value of silver. The three-month futures price is \$12, the strike price is \$13, the risk-free rate is 4% and the volatility of the price of silver is 25%.
- 17.21. A corporation knows that in three months it will have \$5 million to invest for 90 days at LIBOR minus 50 basis points and wishes to ensure that the rate obtained will be at least 6.5%. What position in exchange-traded options should it take to hedge?

Further Questions

- 17.22. A futures price is currently 40. It is known that at the end of three months the price will be either 35 or 45. What is the value of a three-month European call option on the futures with a strike price of 42 if the risk-free interest rate is 7% per annum?
- 17.23. It is February 4. July call options on corn futures with strike prices of 260, 270, 280, 290, and 300 cost 26.75, 21.25, 17.25, 14.00, and 11.375, respectively. July put options with these strike prices cost 8.50, 13.50, 19.00, 25.625, and 32.625, respectively. The options mature on June 19, the current July corn futures price is 278.25, and the risk-free interest rate is 1.1%. Calculate implied volatilities for the options using DerivaGem. Comment on the results you get.
- 17.24. Calculate the implied volatility of soybean futures prices from the following information concerning a European put on soybean futures:

| | |
|-----------------------|--------------|
| Current futures price | 525 |
| Exercise price | 525 |
| Risk-free rate | 6% per annum |
| Time to maturity | 5 months |
| Put price | 20 |

- 17.25. Calculate the price of a six-month European put option on the spot value of the S&P 500. The six-month forward price of the index is 1,400, the strike price is 1,450, the risk-free rate is 5%, and the volatility of the index is 15%.
- 17.26. The strike price of a futures option is 550 cents, the risk-free interest rate is 3%, the volatility of the futures price is 20%, and the time to maturity of the option is 9 months. The futures price is 500 cents.
- What is the price of the option if it is a European call?
 - What is the price of the option if it is a European put?
 - Verify that put–call parity holds.
 - What is the futures price for a futures-style option if it is a call?
 - What is the futures price for a futures-style option if it is a put?



CHAPTER 18

The Greek Letters

A financial institution that sells an option to a client in the over-the-counter markets is faced with the problem of managing its risk. If the option happens to be the same as one that is traded on an exchange, the financial institution can neutralize its exposure by buying on the exchange the same option as it has sold. But when the option has been tailored to the needs of a client and does not correspond to the standardized products traded by exchanges, hedging the exposure is far more difficult.

In this chapter we discuss some of the alternative approaches to this problem. We cover what are commonly referred to as the “Greek letters”, or simply the “Greeks”. Each Greek letter measures a different dimension to the risk in an option position and the aim of a trader is to manage the Greeks so that all risks are acceptable. The analysis presented in this chapter is applicable to market makers in options on an exchange as well as to traders working in the over-the-counter market for financial institutions.

Toward the end of the chapter, we will consider the creation of options synthetically. This turns out to be very closely related to the hedging of options. Creating an option position synthetically is essentially the same task as hedging the opposite option position. For example, creating a long call option synthetically is the same as hedging a short position in the call option.

18.1 ILLUSTRATION

In the next few sections we use as an example the position of a financial institution that has sold for \$300,000 a European call option on 100,000 shares of a non-dividend-paying stock. We assume that the stock price is \$49, the strike price is \$50, the risk-free interest rate is 5% per annum, the stock price volatility is 20% per annum, the time to maturity is 20 weeks (0.3846 years), and the expected return from the stock is 13% per annum.¹ With our usual notation, this means that

$$S_0 = 49, \quad K = 50, \quad r = 0.05, \quad \sigma = 0.20, \quad T = 0.3846, \quad \mu = 0.13$$

The Black–Scholes–Merton price of the option is about \$240,000 (that is, \$2.40 for an

¹ As shown in Chapters 12 and 14, the expected return is irrelevant to the pricing of an option. It is given here because it can have some bearing on the effectiveness of a hedging scheme.

option to buy one share). The financial institution has therefore sold the option for \$60,000 more than its theoretical value. But it is faced with the problem of hedging the risks.²

18.2 NAKED AND COVERED POSITIONS

One strategy open to the financial institution is to do nothing. This is sometimes referred to as a *naked position*. It is a strategy that works well if the stock price is below \$50 at the end of the 20 weeks. The option then costs the financial institution nothing and it makes a profit of \$300,000. A naked position works less well if the call is exercised because the financial institution then has to buy 100,000 shares at the market price prevailing in 20 weeks to cover the call. The cost to the financial institution is 100,000 times the amount by which the stock price exceeds the strike price. For example, if after 20 weeks the stock price is \$60, the option costs the financial institution \$1,000,000. This is considerably greater than the \$300,000 charged for the option.

As an alternative to a naked position, the financial institution can adopt a *covered position*. This involves buying 100,000 shares as soon as the option has been sold. If the option is exercised, this strategy works well, but in other circumstances it could lead to a significant loss. For example, if the stock price drops to \$40, the financial institution loses \$900,000 on its stock position. This is considerably greater than the \$300,000 charged for the option.³

Neither a naked position nor a covered position provides a good hedge. If the assumptions underlying the Black–Scholes–Merton formula hold, the cost to the financial institution should always be \$240,000 on average for both approaches.⁴ But on any one occasion the cost is liable to range from zero to over \$1,000,000. A good hedge would ensure that the cost is always close to \$240,000.


18.3 A STOP-LOSS STRATEGY

One interesting hedging procedure that is sometimes proposed involves a *stop-loss strategy*. To illustrate the basic idea, consider an institution that has written a call option with strike price K to buy one unit of a stock. The hedging procedure involves buying one unit of the stock as soon as its price rises above K and selling it as soon as its price falls below K . The objective is to hold a naked position whenever the stock price is less than K and a covered position whenever the stock price is greater than K . The procedure is designed to ensure that at time T the institution owns the stock if the option closes in the money and does not own it if the option closes out of the money. In the situation illustrated in Figure 18.1, it involves buying the stock at time t_1 , selling it at time t_2 , buying it at time t_3 , selling it at time t_4 , buying it at time t_5 , and delivering it at time T .

² A call option on a non-dividend-paying stock is a convenient example with which to develop our ideas. The points that will be made apply to other types of options and to other derivatives.

³ Put–call parity shows that the exposure from writing a covered call is the same as the exposure from writing a naked put.

⁴ More precisely, the present value of the expected cost is \$240,000 for both approaches assuming that appropriate risk-adjusted discount rates are used.

Figure 18.1 A stop-loss strategy.

As usual, we denote the initial stock price by S_0 . The cost of setting up the hedge initially is S_0 if $S_0 > K$ and zero otherwise. It seems as though the total cost, Q , of writing and hedging the option is the option's initial intrinsic value:

$$Q = \max(S_0 - K, 0) \quad (18.1)$$

This is because all purchases and sales subsequent to time 0 are made at price K . If this were in fact correct, the hedging procedure would work perfectly in the absence of transactions costs. Furthermore, the cost of hedging the option would always be less than its Black–Scholes–Merton price. Thus, an investor could earn riskless profits by writing options and hedging them.

There are two key reasons why equation (18.1) is incorrect. The first is that the cash flows to the hedger occur at different times and must be discounted. The second is that purchases and sales cannot be made at exactly the same price K . This second point is critical. If we assume a risk-neutral world with zero interest rates, we can justify ignoring the time value of money. But we cannot legitimately assume that both purchases and sales are made at the same price. If markets are efficient, the hedger cannot know whether, when the stock price equals K , it will continue above or below K .

As a practical matter, purchases must be made at a price $K + \epsilon$ and sales must be made at a price $K - \epsilon$, for some small positive number ϵ . Thus, every purchase and subsequent sale involves a cost (apart from transaction costs) of 2ϵ . A natural response on the part of the hedger is to monitor price movements more closely, so that ϵ is reduced. Assuming that stock prices change continuously, ϵ can be made arbitrarily small by monitoring the stock prices closely. But as ϵ is made smaller, trades tend to occur more frequently. Thus, the lower cost per trade is offset by the increased frequency of trading. As $\epsilon \rightarrow 0$, the expected number of trades tends to infinity.⁵

⁵ As mentioned in Section 13.2, the expected number of times a Wiener process equals any particular value in a given time interval is infinite.

Table 18.1 Performance of stop-loss strategy. The performance measure is the ratio of the standard deviation of the cost of writing the option and hedging it to the theoretical price of the option.

| Δt (weeks) | 5 | 4 | 2 | 1 | 0.5 | 0.25 |
|--------------------|------|------|------|------|------|------|
| Hedge performance | 1.02 | 0.93 | 0.82 | 0.77 | 0.76 | 0.76 |

A stop-loss strategy, although superficially attractive, does not work particularly well as a hedging procedure. Consider its use for an out-of-the-money option. If the stock price never reaches the strike price K , the hedging procedure costs nothing. If the path of the stock price crosses the strike price level many times, the procedure is quite expensive. Monte Carlo simulation can be used to assess the overall performance of stop-loss hedging. This involves randomly sampling paths for the stock price and observing the results of using the procedure. Table 18.1 shows the results for the option considered in Section 18.1. It assumes that the stock price is observed at the end of time intervals of length Δt .⁶ The hedge performance measure is the ratio of the standard deviation of the cost of hedging the option to the Black–Scholes–Merton option price. Each result is based on 1,000 sample paths for the stock price and has a standard error of about 2%. A perfect hedge would have a hedge performance measure of zero. In this case it appears to be impossible to produce a value for the hedge performance measure below 0.70 regardless of how small Δt is made.

18.4 DELTA HEDGING

Most traders use more sophisticated hedging procedures than those mentioned so far. These involve calculating measures such as delta, gamma, and vega. In this section we consider the role played by delta.


The *delta* (Δ) of an option was introduced in Chapter 12. It is defined as the rate of change of the option price with respect to the price of the underlying asset. It is the slope of the curve that relates the option price to the underlying asset price. Suppose that the delta of a call option on a stock is 0.6. This means that when the stock price changes by a small amount, the option price changes by about 60% of that amount. Figure 18.2 shows the relationship between a call price and the underlying stock price. When the stock price corresponds to point A, the option price corresponds to point B, and Δ is the slope of the line indicated. In general,

$$\Delta = \frac{\partial c}{\partial S}$$

where c is the price of the call option and S is the stock price.

Suppose that, in Figure 18.2, the stock price is \$100 and the option price is \$10. Imagine an investor who has sold 20 call option contracts—that is, options on 2,000

⁶ The precise hedging rule used was as follows. If the stock price moves from below K to above K in a time interval of length Δt , it is bought at the end of the interval. If it moves from above K to below K in the time interval, it is sold at the end of the interval; otherwise, no action is taken.

Figure 18.2 Calculation of delta.

shares. The investor's position could be hedged by buying $0.6 \times 2,000 = 1,200$ shares. The gain (loss) on the stock position would then tend to offset the loss (gain) on the option position. For example, if the stock price goes up by \$1 (producing a gain of \$1,200 on the shares purchased), the option price will tend to go up by $0.6 \times \$1 = \0.60 (producing a loss of \$1,200 on the options written); if the stock price goes down by \$1 (producing a loss of \$1,200 on the shares purchased), the option price will tend to go down by \$0.60 (producing a gain of \$1,200 on the options written).

In this example, the delta of the trader's short position in 2,000 options is

$$0.6 \times (-2,000) = -1,200$$

This means that the trader loses $1,200\Delta S$ on the option position when the stock price increases by ΔS . The delta of one share of the stock is 1.0, so that the long position in 1,200 shares has a delta of +1,200. The delta of the trader's overall position is, therefore, zero. The delta of the stock position offsets the delta of the option position. A position with a delta of zero is referred to as *delta neutral*.

It is important to realize that, since the delta of an option does not remain constant, the trader's position remains delta hedged (or delta neutral) for only a relatively short period of time. The hedge has to be adjusted periodically. This is known as *rebalancing*. In our example, by the end of 1 day the stock price might have increased to \$110. As indicated by Figure 18.2, an increase in the stock price leads to an increase in delta. Suppose that delta rises from 0.60 to 0.65. An extra $0.05 \times 2,000 = 100$ shares would then have to be purchased to maintain the hedge. A procedure such as this, where the hedge is adjusted on a regular basis, is referred to as *dynamic hedging*. It can be contrasted with *static hedging*, where a hedge is set up initially and never adjusted. Static hedging is sometimes also referred to as *hedge-and-forget*.

Delta is closely related to the Black–Scholes–Merton analysis. As explained in Chapter 14, the Black–Scholes–Merton differential equation can be derived by setting up a riskless portfolio consisting of a position in an option on a stock and a position in

the stock. Expressed in terms of Δ , the portfolio is

- 1: option
- + Δ : shares of the stock.

Using our new terminology, we can say that options can be valued by setting up a delta-neutral position and arguing that the return on the position should (instantaneously) be the risk-free interest rate.

Delta of European Stock Options

For a European call option on a non-dividend-paying stock, it can be shown (see Problem 14.17) that

$$\Delta(\text{call}) = N(d_1)$$

where d_1 is defined as in equation (14.20) and $N(x)$ is the cumulative distribution function for a standard normal distribution. The formula gives the delta of a long position in one call option. The delta of a short position in one call option is $-N(d_1)$. Using delta hedging for a short position in a European call option involves maintaining a long position of $N(d_1)$ for each option sold. Similarly, using delta hedging for a long position in a European call option involves maintaining a short position of $N(d_1)$ shares for each option purchased.

For a European put option on a non-dividend-paying stock, delta is given by

$$\Delta(\text{put}) = N(d_1) - 1$$

Delta is negative, which means that a long position in a put option should be hedged with a long position in the underlying stock, and a short position in a put option should be hedged with a short position in the underlying stock. Figure 18.3 shows the variation of the delta of a call option and a put option with the stock price. Figure 18.4 shows the variation of delta with the time to maturity for in-the-money, at-the-money, and out-of-the-money call options.

Figure 18.3 Variation of delta with stock price for (a) a call option and (b) a put option on a non-dividend-paying stock.





Figure 18.4 Typical patterns for variation of delta with time to maturity for a call option.



Example 18.1

Consider again the call option on a non-dividend-paying stock in Section 18.1 where the stock price is \$49, the strike price is \$50, the risk-free rate is 5%, the time to maturity is 20 weeks (= 0.3846 years), and the volatility is 20%. In this case,

$$d_1 = \frac{\ln(49/50) + (0.05 + 0.2^2/2) \times 0.3846}{0.2 \times \sqrt{0.3846}} = 0.0542$$

Delta is $N(d_1)$, or 0.522. When the stock price changes by ΔS , the option price changes by $0.522\Delta S$.

Dynamic Aspects of Delta Hedging

Tables 18.2 and 18.3 provide two examples of the operation of delta hedging for the example in Section 18.1. The hedge is assumed to be adjusted or rebalanced weekly. The initial value of delta for the option being sold is calculated in Example 18.1 as 0.522. This means that the delta of the short option position is initially $-52,200$. As soon as the option is written, \$2,557,800 must be borrowed to buy 52,200 shares at a price of \$49. The rate of interest is 5%. An interest cost of approximately \$2,500 is therefore incurred in the first week.

In Table 18.2 the stock price falls by the end of the first week to \$48.12. The delta of the option declines to 0.458, so that the new delta of the option position is $-45,800$. This means that 6,400 of the shares initially purchased are sold to maintain the hedge. The strategy realizes \$308,000 in cash, and the cumulative borrowings at the end of Week 1 are reduced to \$2,252,300. During the second week, the stock price reduces to \$47.37, delta declines again, and so on. Toward the end of the life of the option, it

Table 18.2 Simulation of delta hedging. Option closes in the money and cost of hedging is \$263,300.

| <i>Week</i> | <i>Stock price</i> | <i>Delta</i> | <i>Shares purchased</i> | <i>Cost of shares purchased (\$000)</i> | <i>Cumulative cost including interest (\$000)</i> | <i>Interest cost (\$000)</i> |
|-------------|--------------------|--------------|-------------------------|---|---|------------------------------|
| 0 | 49.00 | 0.522 | 52,200 | 2,557.8 | 2,557.8 | 2.5 |
| 1 | 48.12 | 0.458 | (6,400) | (308.0) | 2,252.3 | 2.2 |
| 2 | 47.37 | 0.400 | (5,800) | (274.7) | 1,979.8 | 1.9 |
| 3 | 50.25 | 0.596 | 19,600 | 984.9 | 2,966.6 | 2.9 |
| 4 | 51.75 | 0.693 | 9,700 | 502.0 | 3,471.5 | 3.3 |
| 5 | 53.12 | 0.774 | 8,100 | 430.3 | 3,905.1 | 3.8 |
| 6 | 53.00 | 0.771 | (300) | (15.9) | 3,893.0 | 3.7 |
| 7 | 51.87 | 0.706 | (6,500) | (337.2) | 3,559.5 | 3.4 |
| 8 | 51.38 | 0.674 | (3,200) | (164.4) | 3,398.5 | 3.3 |
| 9 | 53.00 | 0.787 | 11,300 | 598.9 | 4,000.7 | 3.8 |
| 10 | 49.88 | 0.550 | (23,700) | (1,182.2) | 2,822.3 | 2.7 |
| 11 | 48.50 | 0.413 | (13,700) | (664.4) | 2,160.6 | 2.1 |
| 12 | 49.88 | 0.542 | 12,900 | 643.5 | 2,806.2 | 2.7 |
| 13 | 50.37 | 0.591 | 4,900 | 246.8 | 3,055.7 | 2.9 |
| 14 | 52.13 | 0.768 | 17,700 | 922.7 | 3,981.3 | 3.8 |
| 15 | 51.88 | 0.759 | (900) | (46.7) | 3,938.4 | 3.8 |
| 16 | 52.87 | 0.865 | 10,600 | 560.4 | 4,502.6 | 4.3 |
| 17 | 54.87 | 0.978 | 11,300 | 620.0 | 5,126.9 | 4.9 |
| 18 | 54.62 | 0.990 | 1,200 | 65.5 | 5,197.3 | 5.0 |
| 19 | 55.87 | 1.000 | 1,000 | 55.9 | 5,258.2 | 5.1 |
| 20 | 57.25 | 1.000 | 0 | 0.0 | 5,263.3 | |

becomes apparent that the option will be exercised and the delta of the option approaches 1.0. By Week 20, therefore, the hedger has a fully covered position. The hedger receives \$5 million for the stock held, so that the total cost of writing the option and hedging it is \$263,300.

Table 18.3 illustrates an alternative sequence of events such that the option closes out of the money. As it becomes clear that the option will not be exercised, delta approaches zero. By Week 20 the hedger has a naked position and has incurred costs totaling \$256,600.

In Tables 18.2 and 18.3, the costs of hedging the option, when discounted to the beginning of the period, are close to but not exactly the same as the Black–Scholes–Merton price of \$240,000. If the hedging worked perfectly, the cost of hedging would, after discounting, be exactly equal to the Black–Scholes–Merton price for every simulated stock price path. The reason for the variation in the cost of hedging is that the hedge is rebalanced only once a week. As rebalancing takes place more frequently, the variation in the cost of hedging is reduced. Of course, the examples in Tables 18.2 and 18.3 are idealized in that they assume that the volatility is constant and there are no transaction costs.

Table 18.3 Simulation of delta hedging. Option closes out of the money and cost of hedging is \$256,600.

| <i>Week</i> | <i>Stock price</i> | <i>Delta</i> | <i>Shares purchased</i> | <i>Cost of shares purchased (\$000)</i> | <i>Cumulative cost including interest (\$000)</i> | <i>Interest cost (\$000)</i> |
|-------------|--------------------|--------------|-------------------------|---|---|------------------------------|
| 0 | 49.00 | 0.522 | 52,200 | 2,557.8 | 2,557.8 | 2.5 |
| 1 | 49.75 | 0.568 | 4,600 | 228.9 | 2,789.2 | 2.7 |
| 2 | 52.00 | 0.705 | 13,700 | 712.4 | 3,504.3 | 3.4 |
| 3 | 50.00 | 0.579 | (12,600) | (630.0) | 2,877.7 | 2.8 |
| 4 | 48.38 | 0.459 | (12,000) | (580.6) | 2,299.9 | 2.2 |
| 5 | 48.25 | 0.443 | (1,600) | (77.2) | 2,224.9 | 2.1 |
| 6 | 48.75 | 0.475 | 3,200 | 156.0 | 2,383.0 | 2.3 |
| 7 | 49.63 | 0.540 | 6,500 | 322.6 | 2,707.9 | 2.6 |
| 8 | 48.25 | 0.420 | (12,000) | (579.0) | 2,131.5 | 2.1 |
| 9 | 48.25 | 0.410 | (1,000) | (48.2) | 2,085.4 | 2.0 |
| 10 | 51.12 | 0.658 | 24,800 | 1,267.8 | 3,355.2 | 3.2 |
| 11 | 51.50 | 0.692 | 3,400 | 175.1 | 3,533.5 | 3.4 |
| 12 | 49.88 | 0.542 | (15,000) | (748.2) | 2,788.7 | 2.7 |
| 13 | 49.88 | 0.538 | (400) | (20.0) | 2,771.4 | 2.7 |
| 14 | 48.75 | 0.400 | (13,800) | (672.7) | 2,101.4 | 2.0 |
| 15 | 47.50 | 0.236 | (16,400) | (779.0) | 1,324.4 | 1.3 |
| 16 | 48.00 | 0.261 | 2,500 | 120.0 | 1,445.7 | 1.4 |
| 17 | 46.25 | 0.062 | (19,900) | (920.4) | 526.7 | 0.5 |
| 18 | 48.13 | 0.183 | 12,100 | 582.4 | 1,109.6 | 1.1 |
| 19 | 46.63 | 0.007 | (17,600) | (820.7) | 290.0 | 0.3 |
| 20 | 48.12 | 0.000 | (700) | (33.7) | 256.6 | |

Table 18.4 shows statistics on the performance of delta hedging obtained from 1,000 random stock price paths in our example. As in Table 18.1, the performance measure is the ratio of the standard deviation of the cost of hedging the option to the Black–Scholes–Merton price of the option. It is clear that delta hedging is a great improvement over a stop-loss strategy. Unlike a stop-loss strategy, the performance of a delta-hedging strategy gets steadily better as the hedge is monitored more frequently.

Table 18.4 Performance of delta hedging. The performance measure is the ratio of the standard deviation of the cost of writing the option and hedging it to the theoretical price of the option.

| Time between hedge rebalancing (weeks): | 5 | 4 | 2 | 1 | 0.5 | 0.25 |
|---|------|------|------|------|------|------|
| Performance measure: | 0.43 | 0.39 | 0.26 | 0.19 | 0.14 | 0.09 |

Delta hedging aims to keep the value of the financial institution's position as close to unchanged as possible. Initially, the value of the written option is \$240,000. In the situation depicted in Table 18.2, the value of the option can be calculated as \$414,500 in Week 9. Thus, the financial institution has lost \$174,500 on its short option position. Its cash position, as measured by the cumulative cost, is \$1,442,900 worse in Week 9 than in Week 0. The value of the shares held has increased from \$2,557,800 to \$4,171,100. The net effect of all this is that the value of the financial institution's position has changed by only \$4,100 between Week 0 and Week 9.

Where the Cost Comes From

The delta-hedging procedure in Tables 18.2 and 18.3 creates the equivalent of a long position in the option. This neutralizes the short position the financial institution created by writing the option. As the tables illustrate, delta hedging a short position generally involves selling stock just after the price has gone down and buying stock just after the price has gone up. It might be termed a buy-high, sell-low trading strategy! The cost of \$240,000 comes from the average difference between the price paid for the stock and the price at which it is sold.

Delta of a Portfolio

The delta of a portfolio of options or other derivatives dependent on a single asset whose price is S is

$$\frac{\partial \Pi}{\partial S}$$

where Π is the value of the portfolio.

The delta of the portfolio can be calculated from the deltas of the individual options in the portfolio. If a portfolio consists of a quantity w_i of option i ($1 \leq i \leq n$), the delta of the portfolio is given by

$$\Delta = \sum_{i=1}^n w_i \Delta_i$$

where Δ_i is the delta of the i th option. The formula can be used to calculate the position in the underlying asset necessary to make the delta of the portfolio zero. When this position has been taken, the portfolio is referred to as being *delta neutral*.

Suppose a financial institution has the following three positions in options on a stock:

1. A long position in 100,000 call options with strike price \$55 and an expiration date in 3 months. The delta of each option is 0.533.
2. A short position in 200,000 call options with strike price \$56 and an expiration date in 5 months. The delta of each option is 0.468.
3. A short position in 50,000 put options with strike price \$56 and an expiration date in 2 months. The delta of each option is -0.508 .

The delta of the whole portfolio is

$$100,000 \times 0.533 - 200,000 \times 0.468 - 50,000 \times (-0.508) = -14,900$$

This means that the portfolio can be made delta neutral by buying 14,900 shares.

Transactions Costs

Derivatives dealers usually rebalance their positions once a day to maintain delta neutrality. When the dealer has a small number of options on a particular asset, this is liable to be prohibitively expensive because of the transactions costs incurred on trades. For a large portfolio of options, it is more feasible. Only one trade in the underlying asset is necessary to zero out delta for the whole portfolio. The hedging transactions costs are absorbed by the profits on many different trades.

18.5 THETA

The *theta* (Θ) of a portfolio of options is the rate of change of the value of the portfolio with respect to the passage of time with all else remaining the same. Theta is sometimes referred to as the *time decay* of the portfolio. For a European call option on a non-dividend-paying stock, it can be shown from the Black–Scholes–Merton formula (see Problem 14.17) that

$$\Theta(\text{call}) = -\frac{S_0 N'(d_1)\sigma}{2\sqrt{T}} - rKe^{-rT}N(d_2)$$

where d_1 and d_2 are defined as in equation (14.20) and

$$N'(x) = \frac{1}{\sqrt{2\pi}}e^{-x^2/2} \quad (18.2)$$

is the probability density function for a standard normal distribution.

For a European put option on the stock,

$$\Theta(\text{put}) = -\frac{S_0 N'(d_1)\sigma}{2\sqrt{T}} + rKe^{-rT}N(-d_2)$$

Because $N(-d_2) = 1 - N(d_2)$, the theta of a put exceeds the theta of the corresponding call by rKe^{-rT} .

In these formulas, time is measured in years. Usually, when theta is quoted, time is measured in days, so that theta is the change in the portfolio value when 1 day passes with all else remaining the same. We can measure theta either “per calendar day” or “per trading day”. To obtain the theta per calendar day, the formula for theta must be divided by 365; to obtain theta per trading day, it must be divided by 252. (DerivaGem measures theta per calendar day.)

Example 18.2


As in Example 18.1, consider a call option on a non-dividend-paying stock where the stock price is \$49, the strike price is \$50, the risk-free rate is 5%, the time to maturity is 20 weeks (= 0.3846 years), and the volatility is 20%. In this case, $S_0 = 49$, $K = 50$, $r = 0.05$, $\sigma = 0.2$, and $T = 0.3846$.

The option’s theta is

$$-\frac{S_0 N'(d_1)\sigma}{2\sqrt{T}} - rKe^{-rT}N(d_2) = -4.31$$

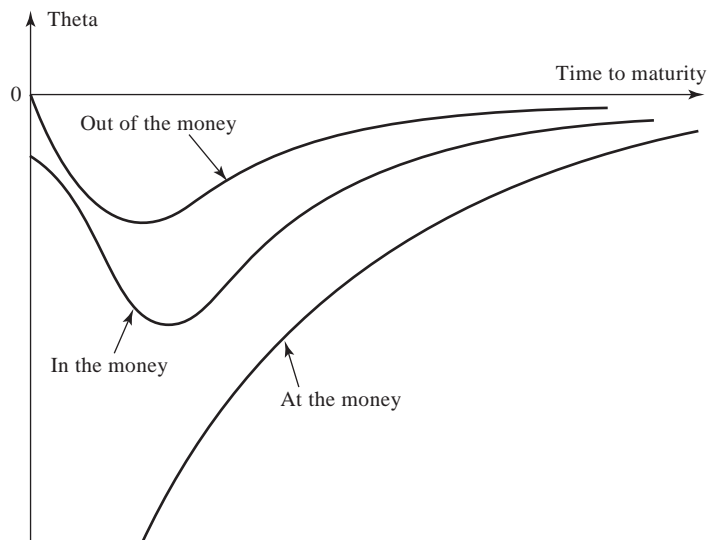
The theta is $-4.31/365 = -0.0118$ per calendar day, or $-4.31/252 = -0.0171$ per trading day.

Figure 18.5 Variation of theta of a European call option with stock price.



Theta is usually negative for an option.⁷ This is because, as time passes with all else remaining the same, the option tends to become less valuable. The variation of Θ with stock price for a call option on a stock is shown in Figure 18.5. When the stock price is very low, theta is close to zero. For an at-the-money call option, theta is large and negative. As the stock price becomes larger, theta tends to $-rKe^{-rT}$. Figure 18.6 shows typical patterns for the variation of Θ with the time to maturity for in-the-money, at-the-money, and out-of-the-money call options.

Figure 18.6 Typical patterns for variation of theta of a European call option with time to maturity.



⁷ An exception to this could be an in-the-money European put option on a non-dividend-paying stock or an in-the-money European call option on a currency with a very high interest rate.

Theta is not the same type of hedge parameter as delta. There is uncertainty about the future stock price, but there is no uncertainty about the passage of time. It makes sense to hedge against changes in the price of the underlying asset, but it does not make any sense to hedge against the passage of time. In spite of this, many traders regard theta as a useful descriptive statistic for a portfolio. This is because, as we shall see later, in a delta-neutral portfolio theta is a proxy for gamma.

18.6 GAMMA

The *gamma* (Γ) of a portfolio of options on an underlying asset is the rate of change of the portfolio's delta with respect to the price of the underlying asset. It is the second partial derivative of the portfolio with respect to asset price:

$$\Gamma = \frac{\partial^2 \Pi}{\partial S^2}$$

If gamma is small, delta changes slowly, and adjustments to keep a portfolio delta neutral need to be made only relatively infrequently. However, if gamma is highly negative or highly positive, delta is very sensitive to the price of the underlying asset. It is then quite risky to leave a delta-neutral portfolio unchanged for any length of time. Figure 18.7 illustrates this point. When the stock price moves from S to S' , delta hedging assumes that the option price moves from C to C' , when in fact it moves from C to C'' . The difference between C' and C'' leads to a hedging error. The size of the error depends on the curvature of the relationship between the option price and the stock price. Gamma measures this curvature.

Suppose that ΔS is the price change of an underlying asset during a small interval of time, Δt , and $\Delta \Pi$ is the corresponding price change in the portfolio. The appendix at the end of this chapter shows that, if terms of order higher than Δt are ignored,

$$\Delta \Pi = \Theta \Delta t + \frac{1}{2} \Gamma \Delta S^2 \quad (18.3)$$

for a delta-neutral portfolio, where Θ is the theta of the portfolio. Figure 18.8 shows the

Figure 18.7 Hedging error introduced by nonlinearity.





Figure 18.8 Relationship between $\Delta\Pi$ and ΔS in time Δt for a delta-neutral portfolio with (a) slightly positive gamma, (b) large positive gamma, (c) slightly negative gamma, and (d) large negative gamma.



nature of this relationship between $\Delta\Pi$ and ΔS . When gamma is positive, theta tends to be negative. The portfolio declines in value if there is no change in S , but increases in value if there is a large positive or negative change in S . When gamma is negative, theta tends to be positive and the reverse is true: the portfolio increases in value if there is no change in S but decreases in value if there is a large positive or negative change in S . As the absolute value of gamma increases, the sensitivity of the value of the portfolio to S increases.

Example 18.3

Suppose that the gamma of a delta-neutral portfolio of options on an asset is $-10,000$. Equation (18.3) shows that, if a change of $+2$ or -2 in the price of the asset occurs over a short period of time, there is an unexpected decrease in the value of the portfolio of approximately $0.5 \times 10,000 \times 2^2 = \$20,000$.

Making a Portfolio Gamma Neutral

A position in the underlying asset has zero gamma and cannot be used to change the gamma of a portfolio. What is required is a position in an instrument such as an option that is not linearly dependent on the underlying asset.

Suppose that a delta-neutral portfolio has a gamma equal to Γ , and a traded option has a gamma equal to Γ_T . If the number of traded options added to the portfolio is w_T , the gamma of the portfolio is

$$w_T \Gamma_T + \Gamma$$

Hence, the position in the traded option necessary to make the portfolio gamma neutral is $-\Gamma/\Gamma_T$. Including the traded option is likely to change the delta of the portfolio, so the position in the underlying asset then has to be changed to maintain delta neutrality. Note that the portfolio is gamma neutral only for a short period of time. As time passes, gamma neutrality can be maintained only if the position in the traded option is adjusted so that it is always equal to $-\Gamma/\Gamma_T$.

Making a portfolio gamma neutral as well as delta-neutral can be regarded as a correction for the hedging error illustrated in Figure 18.7. Delta neutrality provides protection against relatively small stock price moves between rebalancing. Gamma neutrality provides protection against larger movements in this stock price between hedge rebalancing. Suppose that a portfolio is delta neutral and has a gamma of $-3,000$. The delta and gamma of a particular traded call option are 0.62 and 1.50, respectively. The portfolio can be made gamma neutral by including in the portfolio a long position of

$$\frac{3,000}{1.5} = 2,000$$

in the call option. However, the delta of the portfolio will then change from zero to $2,000 \times 0.62 = 1,240$. Therefore 1,240 units of the underlying asset must be sold from the portfolio to keep it delta neutral.

Calculation of Gamma

For a European call or put option on a non-dividend-paying stock, the gamma is given by

$$\Gamma = \frac{N'(d_1)}{S_0 \sigma \sqrt{T}}$$

Figure 18.9 Variation of gamma with stock price for an option.





Figure 18.10 Variation of gamma with time to maturity for a stock option.

where d_1 is defined as in equation (14.20) and $N'(x)$ is as given by equation (18.2). The gamma of a long position is always positive and varies with S_0 in the way indicated in Figure 18.9. The variation of gamma with time to maturity for out-of-the-money, at-the-money, and in-the-money options is shown in Figure 18.10. For an at-the-money option, gamma increases as the time to maturity decreases. Short-life at-the-money options have very high gammas, which means that the value of the option holder's position is highly sensitive to jumps in the stock price.

Example 18.4

As in Example 18.1, consider a call option on a non-dividend-paying stock where the stock price is \$49, the strike price is \$50, the risk-free rate is 5%, the time to maturity is 20 weeks (= 0.3846 years), and the volatility is 20%. In this case, $S_0 = 49$, $K = 50$, $r = 0.05$, $\sigma = 0.2$, and $T = 0.3846$.

The option's gamma is

$$\frac{N'(d_1)}{S_0\sigma\sqrt{T}} = 0.066$$

When the stock price changes by ΔS , the delta of the option changes by $0.066\Delta S$.

18.7 RELATIONSHIP BETWEEN DELTA, THETA, AND GAMMA

The price of a single derivative dependent on a non-dividend-paying stock must satisfy the differential equation (14.16). It follows that the value of Π of a portfolio of such

derivatives also satisfies the differential equation

$$\frac{\partial \Pi}{\partial t} + rS \frac{\partial \Pi}{\partial S} + \frac{1}{2} \sigma^2 S^2 \frac{\partial^2 \Pi}{\partial S^2} = r\Pi$$

Since

$$\Theta = \frac{\partial \Pi}{\partial t}, \quad \Delta = \frac{\partial \Pi}{\partial S}, \quad \Gamma = \frac{\partial^2 \Pi}{\partial S^2}$$

it follows that

$$\Theta + rS\Delta + \frac{1}{2}\sigma^2 S^2 \Gamma = r\Pi \quad (18.4)$$

Similar results can be produced for other underlying assets (see Problem 18.19).

For a delta-neutral portfolio, $\Delta = 0$ and

$$\Theta + \frac{1}{2}\sigma^2 S^2 \Gamma = r\Pi$$

This shows that, when Θ is large and positive, gamma of a portfolio tends to be large and negative, and vice versa. This is consistent with the way in which Figure 18.8 has been drawn and explains why theta can to some extent be regarded as a proxy for gamma in a delta-neutral portfolio.

18.8 VEGA

Up to now we have implicitly assumed that the volatility of the asset underlying a derivative is constant. In practice, volatilities change over time. This means that the value of a derivative is liable to change because of movements in volatility as well as because of changes in the asset price and the passage of time.

The *vega* of a portfolio of derivatives, \mathcal{V} , is the rate of change of the value of the portfolio with respect to the volatility of the underlying asset.⁸

$$\mathcal{V} = \frac{\partial \Pi}{\partial \sigma}$$

If vega is highly positive or highly negative, the portfolio's value is very sensitive to small changes in volatility. If it is close to zero, volatility changes have relatively little impact on the value of the portfolio.

A position in the underlying asset has zero vega. However, the vega of a portfolio can be changed, similarly to the way gamma can be changed, by adding a position in a traded option. If \mathcal{V} is the vega of the portfolio and \mathcal{V}_T is the vega of a traded option, a position of $-\mathcal{V}/\mathcal{V}_T$ in the traded option makes the portfolio instantaneously vega neutral. Unfortunately, a portfolio that is gamma neutral will not in general be vega neutral, and vice versa. If a hedger requires a portfolio to be both gamma and vega neutral, at least two traded derivatives dependent on the underlying asset must usually be used.

Example 18.5

Consider a portfolio that is delta neutral, with a gamma of $-5,000$ and a vega of $-8,000$. The options shown in the table below can be traded. The portfolio can be

⁸ Vega is the name given to one of the "Greek letters" in option pricing, but it is not one of the letters in the Greek alphabet.

made vega neutral by including a long position in 4,000 of Option 1. This would increase delta to 2,400 and require that 2,400 units of the asset be sold to maintain delta neutrality. The gamma of the portfolio would change from $-5,000$ to $-3,000$.

| | <i>Delta</i> | <i>Gamma</i> | <i>Vega</i> |
|-----------|--------------|--------------|-------------|
| Portfolio | 0 | -5000 | -8000 |
| Option 1 | 0.6 | 0.5 | 2.0 |
| Option 2 | 0.5 | 0.8 | 1.2 |

To make the portfolio gamma and vega neutral, both Option 1 and Option 2 can be used. If w_1 and w_2 are the quantities of Option 1 and Option 2 that are added to the portfolio, we require that

$$-5,000 + 0.5w_1 + 0.8w_2 = 0$$

and

$$-8,000 + 2.0w_1 + 1.2w_2 = 0$$

The solution to these equations is $w_1 = 400$, $w_2 = 6,000$. The portfolio can therefore be made gamma and vega neutral by including 400 of Option 1 and 6,000 of Option 2. The delta of the portfolio, after the addition of the positions in the two traded options, is $400 \times 0.6 + 6,000 \times 0.5 = 3,240$. Hence, 3,240 units of the asset would have to be sold to maintain delta neutrality.

For a European call or put option on a non-dividend-paying stock, vega is given by

$$\mathcal{V} = S_0 \sqrt{T} N'(d_1)$$

where d_1 is defined as in equation (14.20). The formula for $N'(x)$ is given in equation (18.2). The vega of a long position in a European or American option is always positive. The general way in which vega varies with S_0 is shown in Figure 18.11.

Example 18.6

As in Example 18.1, consider a call option on a non-dividend-paying stock where the stock price is \$49, the strike price is \$50, the risk-free rate is 5%, the time to maturity is 20 weeks ($= 0.3846$ years), and the volatility is 20%. In this case, $S_0 = 49$, $K = 50$, $r = 0.05$, $\sigma = 0.2$, and $T = 0.3846$.


The option's vega is

$$S_0 \sqrt{T} N'(d_1) = 12.1$$

Thus a 1% (0.01) increase in the volatility from (20% to 21%) increases the value of the option by approximately $0.01 \times 12.1 = 0.121$.

Calculating vega from the Black–Scholes–Merton model and its extensions may seem strange because one of the assumptions underlying the model is that volatility is constant. It would be theoretically more correct to calculate vega from a model in which volatility is assumed to be stochastic. However, it turns out that the vega calculated from a stochastic volatility model is very similar to the Black–Scholes–Merton vega, so the practice of calculating vega from a model in which volatility is constant works reasonably well.⁹

⁹ See J. C. Hull and A. White, "The Pricing of Options on Assets with Stochastic Volatilities," *Journal of Finance* 42 (June 1987): 281–300; J. C. Hull and A. White, "An Analysis of the Bias in Option Pricing Caused by a Stochastic Volatility," *Advances in Futures and Options Research* 3 (1988): 27–61.

Figure 18.11 Variation of vega with stock price for an option.

Gamma neutrality protects against large changes in the price of the underlying asset between hedge rebalancing. Vega neutrality protects against a variable σ . As might be expected, whether it is best to use an available traded option for vega or gamma hedging depends on the time between hedge rebalancing and the volatility of the volatility.¹⁰

When volatilities change, the implied volatilities of short-dated options tend to change by more than the implied volatilities of long-dated options. The vega of a portfolio is therefore often calculated by changing the volatilities of long-dated options by less than that of short-dated options. One way of doing this is discussed in Section 22.6.

18.9 RHO

The *rho* of a portfolio of options is the rate of change of the value of the portfolio with respect to the interest rate:

$$\frac{\partial \Pi}{\partial r}$$

It measures the sensitivity of the value of a portfolio to a change in the interest rate when all else remains the same. For a European call option on a non-dividend-paying stock,

$$\text{rho (call)} = KTe^{-rT}N(d_2)$$

where d_2 is defined as in equation (14.20). For a European put option,

$$\text{rho (put)} = -KTe^{-rT}N(-d_2)$$

Example 18.7

As in Example 18.1, consider a call option on a non-dividend-paying stock where the stock price is \$49, the strike price is \$50, the risk-free rate is 5%, the time to maturity is 20 weeks (= 0.3846 years), and the volatility is 20%. In this case, $S_0 = 49$, $K = 50$, $r = 0.05$, $\sigma = 0.2$, and $T = 0.3846$.

¹⁰ For a discussion of this issue, see J. C. Hull and A. White, "Hedging the Risks from Writing Foreign Currency Options," *Journal of International Money and Finance* 6 (June 1987): 131–52.

Business Snapshot 18.1 Dynamic Hedging in Practice

In a typical arrangement at a financial institution, the responsibility for a portfolio of derivatives dependent on a particular underlying asset is assigned to one trader or to a group of traders working together. For example, one trader at Goldman Sachs might be assigned responsibility for all derivatives dependent on the value of the Australian dollar. A computer system calculates the value of the portfolio and Greek letters for the portfolio. Limits are defined for each Greek letter and special permission is required if a trader wants to exceed a limit at the end of a trading day.

The delta limit is often expressed as the equivalent maximum position in the underlying asset. For example, the delta limit of Goldman Sachs for a stock might be \$1 million. If the stock price is \$50, this means that the absolute value of delta as we have calculated it can be no more than 20,000. The vega limit is usually expressed as a maximum dollar exposure per 1% change in the volatility.

As a matter of course, options traders make themselves delta neutral—or close to delta neutral—at the end of each day. Gamma and vega are monitored, but are not usually managed on a daily basis. Financial institutions often find that their business with clients involves writing options and that as a result they accumulate negative gamma and vega. They are then always looking out for opportunities to manage their gamma and vega risks by buying options at competitive prices.

There is one aspect of an options portfolio that mitigates problems of managing gamma and vega somewhat. Options are often close to the money when they are first sold, so that they have relatively high gammas and vegas. But after some time has elapsed, the underlying asset price has often changed enough for them to become deep out of the money or deep in the money. Their gammas and vegas are then very small and of little consequence. A nightmare scenario for an options trader is where written options remain very close to the money as the maturity date is approached.

The option's rho is

$$KTe^{-rT}N(d_2) = 8.91$$

This means that a 1% (0.01) increase in the risk-free rate (from 5% to 6%) increases the value of the option by approximately $0.01 \times 8.91 = 0.0891$.

18.10 THE REALITIES OF HEDGING

In an ideal world, traders working for financial institutions would be able to rebalance their portfolios very frequently in order to maintain all Greeks equal to zero. In practice, this is not possible. When managing a large portfolio dependent on a single underlying asset, traders usually make delta zero, or close to zero, at least once a day by trading the underlying asset. Unfortunately, a zero gamma and a zero vega are less easy to achieve because it is difficult to find options or other nonlinear derivatives that can be traded in the volume required at competitive prices. Business Snapshot 18.1 provides a discussion of how dynamic hedging is organized at financial institutions.

There are big economies of scale in trading derivatives. Maintaining delta neutrality for a small number of options on an asset by trading daily is usually not economically

Table 18.5 Profit or loss realized in 2 weeks under different scenarios (\$ million).

| <i>Volatility</i> | <i>Exchange rate</i> | | | | | | |
|-------------------|----------------------|-------------|-------------|-------------|-------------|-------------|-------------|
| | <i>0.94</i> | <i>0.96</i> | <i>0.98</i> | <i>1.00</i> | <i>1.02</i> | <i>1.04</i> | <i>1.06</i> |
| 8% | +102 | +55 | +25 | +6 | -10 | -34 | -80 |
| 10% | +80 | +40 | +17 | +2 | -14 | -38 | -85 |
| 12% | +60 | +25 | +9 | -2 | -18 | -42 | -90 |

feasible. The trading costs per option being hedged are high.¹¹ But when a derivatives dealer maintains delta neutrality for a large portfolio of options on an asset, the trading costs per option hedged are likely to be much more reasonable.

18.11 SCENARIO ANALYSIS

In addition to monitoring risks such as delta, gamma, and vega, option traders often also carry out a scenario analysis. The analysis involves calculating the gain or loss on their portfolio over a specified period under a variety of different scenarios. The time period chosen is likely to depend on the liquidity of the instruments. The scenarios can be either chosen by management or generated by a model.

Consider a bank with a portfolio of options on a foreign currency. There are two main variables on which the value of the portfolio depends. These are the exchange rate and the exchange-rate volatility. Suppose that the exchange rate is currently 1.0000 and its volatility is 10% per annum. The bank could calculate a table such as Table 18.5 showing the profit or loss experienced during a 2-week period under different scenarios. This table considers seven different exchange rates and three different volatilities. Because a one-standard-deviation move in the exchange rate during a 2-week period is about 0.02, the exchange rate moves considered are approximately zero, one, two, and three standard deviations.

In Table 18.5, the greatest loss is in the lower right corner of the table. The loss corresponds to the volatility increasing to 12% and the exchange rate moving up to 1.06. Usually the greatest loss in a table such as 18.5 occurs at one of the corners, but this is not always so. Consider, for example, the situation where a bank's portfolio consists of a short position in a butterfly spread (see Section 11.3). The greatest loss will be experienced if the exchange rate stays where it is.

18.12 EXTENSION OF FORMULAS

The formulas produced so far for delta, theta, gamma, vega, and rho have been for a European option on a non-dividend-paying stock. Table 18.6 shows how they change

¹¹ The trading costs arise from the fact that each day the hedger buys some of the underlying asset at the offer price or sells some of the underlying asset at the bid price.

Table 18.6 Greek letters for European options on an asset that provides a yield at rate q .

| <i>Greek letter</i> | <i>Call option</i> | <i>Put option</i> |
|---------------------|---|---|
| Delta | $e^{-qT} N(d_1)$ | $e^{-qT} [N(d_1) - 1]$ |
| Gamma | $\frac{N'(d_1)e^{-qT}}{S_0\sigma\sqrt{T}}$ | $\frac{N'(d_1)e^{-qT}}{S_0\sigma\sqrt{T}}$ |
| Theta | $-S_0N'(d_1)\sigma e^{-qT}/(2\sqrt{T})$ $+qS_0N(d_1)e^{-qT} - rKe^{-rT}N(d_2)$ | $-S_0N'(d_1)\sigma e^{-qT}/(2\sqrt{T})$ $-qS_0N(-d_1)e^{-qT} + rKe^{-rT}N(-d_2)$ |
| Vega | $S_0\sqrt{T}N'(d_1)e^{-qT}$ | $S_0\sqrt{T}N'(d_1)e^{-qT}$ |
| Rho | $KTe^{-rT}N(d_2)$ | $-KTe^{-rT}N(-d_2)$ |

when the stock pays a continuous dividend yield at rate q . The expressions for d_1 and d_2 are as for equations (16.4) and (16.5). By setting q equal to the dividend yield on an index, we obtain the Greek letters for European options on indices. By setting q equal to the foreign risk-free rate, we obtain the Greek letters for European options on a currency. By setting $q = r$, we obtain delta, gamma, theta, and vega for European options on a futures contract. The rho for a call futures option is $-cT$ and the rho for a European put futures option is $-pT$.

In the case of currency options, there are two rhos corresponding to the two interest rates. The rho corresponding to the domestic interest rate is given by the formula in Table 18.6 (with d_2 as in equation (16.11)). The rho corresponding to the foreign interest rate for a European call on a currency is

$$\text{rho}(\text{call, foreign rate}) = -Te^{-r_f T} S_0 N(d_1)$$

For a European put, it is

$$\text{rho}(\text{put, foreign rate}) = Te^{-r_f T} S_0 N(-d_1)$$

with d_1 as in equation (16.11).

The calculation of Greek letters for American options is discussed in Chapter 20.

Delta of Forward Contracts

The concept of delta can be applied to financial instruments other than options. Consider a forward contract on a non-dividend-paying stock. Equation (5.5) shows that the value of a forward contract is $S_0 - Ke^{-rT}$, where K is the delivery price and T is the forward contract's time to maturity. When the price of the stock changes by ΔS , with all else remaining the same, the value of a forward contract on the stock also changes by ΔS . The delta of a long forward contract on one share of the stock is therefore always 1.0. This

means that a long forward contract on one share can be hedged by shorting one share; a short forward contract on one share can be hedged by purchasing one share.¹²

For an asset providing a dividend yield at rate q , equation (5.7) shows that the forward contract's delta is e^{-qT} . For the delta of a forward contract on a stock index, q is set equal to the dividend yield on the index in this expression. For the delta of a forward foreign exchange contract, it is set equal to the foreign risk-free rate, r_f .

Delta of a Futures Contract

From equation (5.1), the futures price for a contract on a non-dividend-paying stock is $S_0 e^{rT}$, where T is the time to maturity of the futures contract. This shows that when the price of the stock changes by ΔS , with all else remaining the same, the futures price changes by $\Delta S e^{rT}$. Since futures contracts are settled daily, the holder of a long futures position makes an almost immediate gain of this amount. The delta of a futures contract is therefore e^{rT} . For a futures position on an asset providing a dividend yield at rate q , equation (5.3) shows similarly that delta is $e^{(r-q)T}$.

It is interesting that daily settlement makes the deltas of futures and forward contracts slightly different. This is true even when interest rates are constant and the forward price equals the futures price. (A related point is made in Business Snapshot 5.2.)

Sometimes a futures contract is used to achieve a delta-neutral position. Define:

T : Maturity of futures contract

H_A : Required position in asset for delta hedging

H_F : Alternative required position in futures contracts for delta hedging

If the underlying asset is a non-dividend-paying stock, the analysis we have just given shows that

$$H_F = e^{-rT} H_A \quad (18.5)$$

When the underlying asset pays a dividend yield q ,

$$H_F = e^{-(r-q)T} H_A \quad (18.6)$$

For a stock index, we set q equal to the dividend yield on the index; for a currency, we set it equal to the foreign risk-free rate, r_f , so that

$$H_F = e^{-(r-r_f)T} H_A \quad (18.7)$$

Example 18.8

Suppose that a portfolio of currency options held by a US bank can be made delta neutral with a short position of 458,000 pounds sterling. Risk-free rates are 4% in the US and 7% in the UK. From equation (18.7), hedging using 9-month currency futures requires a short futures position

$$e^{-(0.04-0.07) \times 9/12} \times 458,000$$

or £468,442. Since each futures contract is for the purchase or sale of £62,500, seven contracts would be shorted. (Seven is the nearest whole number to 468,442/62,500.)

¹² These are hedge-and-forget schemes. Since delta is always 1.0, no changes need to be made to the position in the stock during the life of the contract.

18.13 PORTFOLIO INSURANCE

A portfolio manager is often interested in acquiring a put option on his or her portfolio. This provides protection against market declines while preserving the potential for a gain if the market does well. One approach (discussed in Section 16.1) is to buy put options on a market index such as the S&P 500. An alternative is to create the options synthetically.

Creating an option synthetically involves maintaining a position in the underlying asset (or futures on the underlying asset) so that the delta of the position is equal to the delta of the required option. The position necessary to create an option synthetically is the reverse of that necessary to hedge it. This is because the procedure for hedging an option involves the creation of an equal and opposite option synthetically.

There are two reasons why it may be more attractive for the portfolio manager to create the required put option synthetically than to buy it in the market. First, options markets do not always have the liquidity to absorb the trades required by managers of large funds. Second, fund managers often require strike prices and exercise dates that are different from those available in exchange-traded options markets.

The synthetic option can be created from trading the portfolio or from trading in index futures contracts. We first examine the creation of a put option by trading the portfolio. From Table 18.6, the delta of a European put on the portfolio is

$$\Delta = e^{-qT}[N(d_1) - 1] \quad (18.8)$$

where, with our usual notation,

$$d_1 = \frac{\ln(S_0/K) + (r - q + \sigma^2/2)T}{\sigma\sqrt{T}}$$

The other variables are defined as usual: S_0 is the value of the portfolio, K is the strike price, r is the risk-free rate, q is the dividend yield on the portfolio, σ is the volatility of the portfolio, and T is the life of the option. The volatility of the portfolio can usually be assumed to be its beta times the volatility of a well-diversified market index.

To create the put option synthetically, the fund manager should ensure that at any given time a proportion

$$e^{-qT}[1 - N(d_1)]$$

of the stocks in the original portfolio has been sold and the proceeds invested in riskless assets. As the value of the original portfolio declines, the delta of the put given by equation (18.8) becomes more negative and the proportion of the original portfolio sold must be increased. As the value of the original portfolio increases, the delta of the put becomes less negative and the proportion of the original portfolio sold must be decreased (i.e., some of the original portfolio must be repurchased).

Using this strategy to create portfolio insurance means that at any given time funds are divided between the stock portfolio on which insurance is required and riskless assets. As the value of the stock portfolio increases, riskless assets are sold and the position in the stock portfolio is increased. As the value of the stock portfolio declines, the position in the stock portfolio is decreased and riskless assets are purchased. The cost of the insurance arises from the fact that the portfolio manager is always selling after a decline in the market and buying after a rise in the market.

Example 18.9

A portfolio is worth \$90 million. To protect against market downturns the managers of the portfolio require a 6-month European put option on the portfolio with a strike price of \$87 million. The risk-free rate is 9% per annum, the dividend yield is 3% per annum, and the volatility of the portfolio is 25% per annum. The S&P 500 index stands at 900. As the portfolio is considered to mimic the S&P 500 fairly closely, one alternative, discussed in Section 16.1, is to buy 1,000 put option contracts on the S&P 500 with a strike price of 870. Another alternative is to create the required option synthetically. In this case, $S_0 = 90$ million, $K = 87$ million, $r = 0.09$, $q = 0.03$, $\sigma = 0.25$, and $T = 0.5$, so that

$$d_1 = \frac{\ln(90/87) + (0.09 - 0.03 + 0.25^2/2)0.5}{0.25\sqrt{0.5}} = 0.4499$$

and the delta of the required option is

$$e^{-qT}[N(d_1) - 1] = -0.3215$$

This shows that 32.15% of the portfolio should be sold initially and invested in risk-free assets to match the delta of the required option. The amount of the portfolio sold must be monitored frequently. For example, if the value of the portfolio reduces to \$88 million after 1 day, the delta of the required option changes to 0.3679 and a further 4.64% of the original portfolio should be sold and invested in risk-free assets. If the value of the portfolio increases to \$92 million, the delta of the required option changes to -0.2787 and 4.28% of the original portfolio should be repurchased.

Use of Index Futures

Using index futures to create options synthetically can be preferable to using the underlying stocks because the transaction costs associated with trades in index futures are generally lower than those associated with the corresponding trades in the underlying stocks. The dollar amount of the futures contracts shorted as a proportion of the value of the portfolio should from equations (18.6) and (18.8) be

$$e^{-qT}e^{-(r-q)T^*}[1 - N(d_1)] = e^{q(T^*-T)}e^{-rT^*}[1 - N(d_1)]$$

where T^* is the maturity of the futures contract. If the portfolio is worth A_1 times the index and each index futures contract is on A_2 times the index, the number of futures contracts shorted at any given time should be

$$e^{q(T^*-T)}e^{-rT^*}[1 - N(d_1)]A_1/A_2$$

Example 18.10

Suppose that in the previous example futures contracts on the S&P 500 maturing in 9 months are used to create the option synthetically. In this case initially $T = 0.5$, $T^* = 0.75$, $A_1 = 100,000$, and $d_1 = 0.4499$. Each index futures contract is on 250 times the index, so that $A_2 = 250$. The number of futures contracts shorted should be

$$e^{q(T^*-T)}e^{-rT^*}[1 - N(d_1)]A_1/A_2 = 122.96$$

or 123, rounding to the nearest whole number. As time passes and the index changes, the position in futures contracts must be adjusted.

This analysis assumes that the portfolio mirrors the index. When this is not the case, it is necessary to (a) calculate the portfolio's beta, (b) find the position in options on the index that gives the required protection, and (c) choose a position in index futures to create the options synthetically. As discussed in Section 16.1, the strike price for the options should be the expected level of the market index when the portfolio reaches its insured value. The number of options required is beta times the number that would be required if the portfolio had a beta of 1.0.

18.14 STOCK MARKET VOLATILITY

We discussed in Chapter 14 the issue of whether volatility is caused solely by the arrival of new information or whether trading itself generates volatility. Portfolio insurance strategies such as those just described have the potential to increase volatility. When the market declines, they cause portfolio managers either to sell stock or to sell index futures contracts. Either action may accentuate the decline (see Business Snapshot 18.2). The sale of stock is liable to drive down the market index further in a direct way. The sale of index futures contracts is liable to drive down futures prices. This creates selling pressure on stocks via the mechanism of index arbitrage (see Chapter 5), so that the market index is liable to be driven down in this case as well. Similarly, when the market rises, the portfolio insurance strategies cause portfolio managers either to buy stock or to buy futures contracts. This may accentuate the rise.

In addition to formal portfolio trading strategies, we can speculate that many investors consciously or subconsciously follow portfolio insurance rules of their own. For example, an investor may choose to sell when the market is falling to limit the downside risk.

Whether portfolio insurance trading strategies (formal or informal) affect volatility depends on how easily the market can absorb the trades that are generated by portfolio insurance. If portfolio insurance trades are a very small fraction of all trades, there is likely to be no effect. As portfolio insurance becomes more popular, it is liable to have a destabilizing effect on the market.

SUMMARY

Financial institutions offer a variety of option products to their clients. Often the options do not correspond to the standardized products traded by exchanges. The financial institutions are then faced with the problem of hedging their exposure. Naked and covered positions leave them subject to an unacceptable level of risk. One course of action that is sometimes proposed is a stop-loss strategy. This involves holding a naked position when an option is out of the money and converting it to a covered position as soon as the option moves into the money. Although superficially attractive, the strategy does not provide a good hedge.

The delta (Δ) of an option is the rate of change of its price with respect to the price of the underlying asset. Delta hedging involves creating a position with zero delta (sometimes referred to as a delta-neutral position). Because the delta of the underlying asset

Business Snapshot 18.2 Was Portfolio Insurance to Blame for the Crash of 1987?

On Monday, October 19, 1987, the Dow Jones Industrial Average dropped by more than 20%. Many people feel that portfolio insurance played a major role in this crash. In October 1987 between \$60 billion and \$90 billion of equity assets were subject to portfolio insurance trading rules where put options were created synthetically in the way discussed in Section 18.13. During the period Wednesday, October 14, 1987, to Friday, October 16, 1987, the market declined by about 10%, with much of this decline taking place on Friday afternoon. The portfolio trading rules should have generated at least \$12 billion of equity or index futures sales as a result of this decline. In fact, portfolio insurers had time to sell only \$4 billion and they approached the following week with huge amounts of selling already dictated by their models. It is estimated that on Monday, October 19, sell programs by three portfolio insurers accounted for almost 10% of the sales on the New York Stock Exchange, and that portfolio insurance sales amounted to 21.3% of all sales in index futures markets. It is likely that the decline in equity prices was exacerbated by investors other than portfolio insurers selling heavily because they anticipated the actions of portfolio insurers.

Because the market declined so fast and the stock exchange systems were overloaded, many portfolio insurers were unable to execute the trades generated by their models and failed to obtain the protection they required. Needless to say, the popularity of portfolio insurance schemes has declined significantly since 1987. One of the morals of this story is that it is dangerous to follow a particular trading strategy—even a hedging strategy—when many other market participants are doing the same thing.

is 1.0, one way of hedging is to take a position of $-\Delta$ in the underlying asset for each long option being hedged. The delta of an option changes over time. This means that the position in the underlying asset has to be frequently adjusted.

Once an option position has been made delta neutral, the next stage is often to look at its gamma (Γ). The gamma of an option is the rate of change of its delta with respect to the price of the underlying asset. It is a measure of the curvature of the relationship between the option price and the asset price. The impact of this curvature on the performance of delta hedging can be reduced by making an option position gamma neutral. If Γ is the gamma of the position being hedged, this reduction is usually achieved by taking a position in a traded option that has a gamma of $-\Gamma$.

Delta and gamma hedging are both based on the assumption that the volatility of the underlying asset is constant. In practice, volatilities do change over time. The vega of an option or an option portfolio measures the rate of change of its value with respect to volatility. A trader who wishes to hedge an option position against volatility changes can make the position vega neutral. As with the procedure for creating gamma neutrality, this usually involves taking an offsetting position in a traded option. If the trader wishes to achieve both gamma and vega neutrality, two traded options are usually required.

Two other measures of the risk of an option position are theta and rho. Theta measures the rate of change of the value of the position with respect to the passage of time, with all else remaining constant. Rho measures the rate of change of the value of the position with respect to the interest rate, with all else remaining constant.

In practice, option traders usually rebalance their portfolios at least once a day to maintain delta neutrality. It is usually not feasible to maintain gamma and vega neutrality on a regular basis. Typically a trader monitors these measures. If they get too large, either corrective action is taken or trading is curtailed.

Portfolio managers are sometimes interested in creating put options synthetically for the purposes of insuring an equity portfolio. They can do so either by trading the portfolio or by trading index futures on the portfolio. Trading the portfolio involves splitting the portfolio between equities and risk-free securities. As the market declines, more is invested in risk-free securities. As the market increases, more is invested in equities. Trading index futures involves keeping the equity portfolio intact and selling index futures. As the market declines, more index futures are sold; as it rises, fewer are sold. This type of portfolio insurance works well in normal market conditions. On Monday, October 19, 1987, when the Dow Jones Industrial Average dropped very sharply, it worked badly. Portfolio insurers were unable to sell either stocks or index futures fast enough to protect their positions.

FURTHER READING

Taleb, N. N., *Dynamic Hedging: Managing Vanilla and Exotic Options*. New York: Wiley, 1996.

Practice Questions (Answers in Solutions Manual)

- 18.1. Explain how a stop-loss trading rule can be implemented for the writer of an out-of-the-money call option. Why does it provide a relatively poor hedge?
- 18.2. What does it mean to assert that the delta of a call option is 0.7? How can a short position in 1,000 options be made delta neutral when the delta of each option is 0.7?
- 18.3. Calculate the delta of an at-the-money six-month European call option on a non-dividend-paying stock when the risk-free interest rate is 10% per annum and the stock price volatility is 25% per annum.
- 18.4. What does it mean to assert that the theta of an option position is -0.1 when time is measured in years? If a trader feels that neither a stock price nor its implied volatility will change, what type of option position is appropriate?
- 18.5. What is meant by the gamma of an option position? What are the risks in the situation where the gamma of a position is highly negative and the delta is zero?
- 18.6. “The procedure for creating an option position synthetically is the reverse of the procedure for hedging the option position.” Explain this statement.
- 18.7. Why did portfolio insurance not work well on October 19, 1987?
- 18.8. The Black–Scholes–Merton price of an out-of-the-money call option with an exercise price of \$40 is \$4. A trader who has written the option plans to use a stop-loss strategy. The trader’s plan is to buy at \$40.10 and to sell at \$39.90. Estimate the expected number of times the stock will be bought or sold.
- 18.9. Suppose that a stock price is currently \$20 and that a call option with an exercise price of \$25 is created synthetically using a continually changing position in the stock. Consider the following two scenarios: (a) Stock price increases steadily from \$20 to \$35 during the

- life of the option; (b) Stock price oscillates wildly, ending up at \$35. Which scenario would make the synthetically created option more expensive? Explain your answer.
- 18.10. What is the delta of a short position in 1,000 European call options on silver futures? The options mature in 8 months, and the futures contract underlying the option matures in 9 months. The current 9-month futures price is \$8 per ounce, the exercise price of the options is \$8, the risk-free interest rate is 12% per annum, and the volatility of silver is 18% per annum.
- 18.11. In Problem 18.10, what initial position in 9-month silver futures is necessary for delta hedging? If silver itself is used, what is the initial position? If 1-year silver futures are used, what is the initial position? Assume no storage costs for silver.
- 18.12. A company uses delta hedging to hedge a portfolio of long positions in put and call options on a currency. Which of the following would give the most favorable result?
(a) A virtually constant spot rate
(b) Wild movements in the spot rate
Explain your answer.
- 18.13. Repeat Problem 18.12 for a financial institution with a portfolio of short positions in put and call options on a currency.
- 18.14. A financial institution has just sold 1,000 7-month European call options on the Japanese yen. Suppose that the spot exchange rate is 0.80 cent per yen, the exercise price is 0.81 cent per yen, the risk-free interest rate in the United States is 8% per annum, the risk-free interest rate in Japan is 5% per annum, and the volatility of the yen is 15% per annum. Calculate the delta, gamma, vega, theta, and rho of the financial institution's position. Interpret each number.
- 18.15. Under what circumstances is it possible to make a European option on a stock index both gamma neutral and vega neutral by adding a position in one other European option?
- 18.16. A fund manager has a well-diversified portfolio that mirrors the performance of the S&P 500 and is worth \$360 million. The value of the S&P 500 is 1,200, and the portfolio manager would like to buy insurance against a reduction of more than 5% in the value of the portfolio over the next 6 months. The risk-free interest rate is 6% per annum. The dividend yield on both the portfolio and the S&P 500 is 3%, and the volatility of the index is 30% per annum.
(a) If the fund manager buys traded European put options, how much would the insurance cost?
(b) Explain carefully alternative strategies open to the fund manager involving traded European call options, and show that they lead to the same result.
(c) If the fund manager decides to provide insurance by keeping part of the portfolio in risk-free securities, what should the initial position be?
(d) If the fund manager decides to provide insurance by using 9-month index futures, what should the initial position be?
- 18.17. Repeat Problem 18.16 on the assumption that the portfolio has a beta of 1.5. Assume that the dividend yield on the portfolio is 4% per annum.
- 18.18. Show by substituting for the various terms in equation (18.4) that the equation is true for:
(a) A single European call option on a non-dividend-paying stock
(b) A single European put option on a non-dividend-paying stock
(c) Any portfolio of European put and call options on a non-dividend-paying stock.

- 18.19. What is the equation corresponding to equation (18.4) for (a) a portfolio of derivatives on a currency and (b) a portfolio of derivatives on a futures price?
- 18.20. Suppose that \$70 billion of equity assets are the subject of portfolio insurance schemes. Assume that the schemes are designed to provide insurance against the value of the assets declining by more than 5% within 1 year. Making whatever estimates you find necessary, use the DerivaGem software to calculate the value of the stock or futures contracts that the administrators of the portfolio insurance schemes will attempt to sell if the market falls by 23% in a single day.
- 18.21. Does a forward contract on a stock index have the same delta as the corresponding futures contract? Explain your answer.
- 18.22. A bank's position in options on the dollar/euro exchange rate has a delta of 30,000 and a gamma of $-80,000$. Explain how these numbers can be interpreted. The exchange rate (dollars per euro) is 0.90. What position would you take to make the position delta neutral? After a short period of time, the exchange rate moves to 0.93. Estimate the new delta. What additional trade is necessary to keep the position delta neutral? Assuming the bank did set up a delta-neutral position originally, has it gained or lost money from the exchange-rate movement?
- 18.23. Use the put–call parity relationship to derive, for a non-dividend-paying stock, the relationship between:
- The delta of a European call and the delta of a European put
 - The gamma of a European call and the gamma of a European put
 - The vega of a European call and the vega of a European put
 - The theta of a European call and the theta of a European put.

Further Questions

- 18.24. Consider a 1-year European call option on a stock when the stock price is \$30, the strike price is \$30, the risk-free rate is 5%, and the volatility is 25% per annum. Use the DerivaGem software to calculate the price, delta, gamma, vega, theta, and rho of the option. Verify that delta is correct by changing the stock price to \$30.1 and recomputing the option price. Verify that gamma is correct by recomputing the delta for the situation where the stock price is \$30.1. Carry out similar calculations to verify that vega, theta, and rho are correct. Use the DerivaGem Applications Builder functions to plot the option price, delta, gamma, vega, theta, and rho against the stock price for the stock option.
- 18.25. A financial institution has the following portfolio of over-the-counter options on sterling:

| <i>Type</i> | <i>Position</i> | <i>Delta of option</i> | <i>Gamma of option</i> | <i>Vega of option</i> |
|-------------|-----------------|----------------------------|----------------------------|---------------------------|
| Call | −1,000 | 0.50 | 2.2 | 1.8 |
| Call | −500 | 0.80 | 0.6 | 0.2 |
| Put | −2,000 | −0.40 | 1.3 | 0.7 |
| Call | −500 | 0.70 | 1.8 | 1.4 |

A traded option is available with a delta of 0.6, a gamma of 1.5, and a vega of 0.8.

- (a) What position in the traded option and in sterling would make the portfolio both gamma neutral and delta neutral?

- (b) What position in the traded option and in sterling would make the portfolio both vega neutral and delta neutral?
- 18.26. Consider again the situation in Problem 18.25. Suppose that a second traded option with a delta of 0.1, a gamma of 0.5, and a vega of 0.6 is available. How could the portfolio be made delta, gamma, and vega neutral?
- 18.27. A deposit instrument offered by a bank guarantees that investors will receive a return during a 6-month period that is the greater of (a) zero and (b) 40% of the return provided by a market index. An investor is planning to put \$100,000 in the instrument. Describe the payoff as an option on the index. Assuming that the risk-free rate of interest is 8% per annum, the dividend yield on the index is 3% per annum, and the volatility of the index is 25% per annum, is the product a good deal for the investor?
- 18.28. The formula for the price c of a European call futures option in terms of the futures price F_0 is given in Chapter 17 as

$$c = e^{-rT}[F_0N(d_1) - KN(d_2)]$$

where

$$d_1 = \frac{\ln(F_0/K) + \sigma^2 T/2}{\sigma\sqrt{T}} \quad \text{and} \quad d_2 = d_1 - \sigma\sqrt{T}$$

and K , r , T , and σ are the strike price, interest rate, time to maturity, and volatility, respectively.

(a) Prove that $F_0N'(d_1) = KN'(d_2)$.

(b) Prove that the delta of the call price with respect to the futures price is $e^{-rT}N(d_1)$.

(c) Prove that the vega of the call price is $F_0\sqrt{T}N'(d_1)e^{-rT}$.

(d) Prove the formula for the rho of a call futures option given in Section 18.12.

The delta, gamma, theta, and vega of a call futures option are the same as those for a call option on a stock paying dividends at rate q , with q replaced by r and S_0 replaced by F_0 . Explain why the same is not true of the rho of a call futures option.

- 18.29. Use DerivaGem to check that equation (18.4) is satisfied for the option considered in Section 18.1. (*Note:* DerivaGem produces a value of theta “per calendar day.” The theta in equation (18.4) is “per year.”)
- 18.30. Use the DerivaGem Application Builder functions to reproduce Table 18.2. (In Table 18.2 the stock position is rounded to the nearest 100 shares.) Calculate the gamma and theta of the position each week. Calculate the change in the value of the portfolio each week and check whether equation (18.3) is approximately satisfied. (*Note:* DerivaGem produces a value of theta “per calendar day.” The theta in equation (18.3) is “per year.”)

APPENDIX

TAYLOR SERIES EXPANSIONS AND HEDGE PARAMETERS

A Taylor series expansion of the change in the portfolio value in a short period of time shows the role played by different Greek letters. If the volatility of the underlying asset is assumed to be constant, the value Π of the portfolio is a function of the asset price S , and time t . The Taylor series expansion gives

$$\Delta\Pi = \frac{\partial\Pi}{\partial S}\Delta S + \frac{\partial\Pi}{\partial t}\Delta t + \frac{1}{2}\frac{\partial^2\Pi}{\partial S^2}\Delta S^2 + \frac{1}{2}\frac{\partial^2\Pi}{\partial t^2}\Delta t^2 + \frac{\partial^2\Pi}{\partial S\partial t}\Delta S\Delta t + \dots \quad (18A.1)$$

where $\Delta\Pi$ and ΔS are the change in Π and S in a small time interval Δt . Delta hedging eliminates the first term on the right-hand side. The second term is nonstochastic. The third term (which is of order Δt) can be made zero by ensuring that the portfolio is gamma neutral as well as delta neutral. Other terms are of order higher than Δt .

For a delta-neutral portfolio, the first term on the right-hand side of equation (18A.1) is zero, so that

$$\Delta\Pi = \Theta\Delta t + \frac{1}{2}\Gamma\Delta S^2$$

when terms of order higher than Δt are ignored. This is equation (18.3).

When the volatility of the underlying asset is uncertain, Π is a function of σ , S , and t . Equation (18A.1) then becomes

$$\Delta\Pi = \frac{\partial\Pi}{\partial S}\Delta S + \frac{\partial\Pi}{\partial\sigma}\Delta\sigma + \frac{\partial\Pi}{\partial t}\Delta t + \frac{1}{2}\frac{\partial^2\Pi}{\partial S^2}\Delta S^2 + \frac{1}{2}\frac{\partial^2\Pi}{\partial\sigma^2}\Delta\sigma^2 + \dots$$

where $\Delta\sigma$ is the change in σ in time Δt . In this case, delta hedging eliminates the first term on the right-hand side. The second term is eliminated by making the portfolio vega neutral. The third term is nonstochastic. The fourth term is eliminated by making the portfolio gamma neutral. Traders sometimes define other Greek letters to correspond to later terms in the expansion.

19

CHAPTER

Volatility Smiles



How close are the market prices of options to those predicted by the Black–Scholes–Merton model? Do traders really use the Black–Scholes–Merton model when determining a price for an option? Are the probability distributions of asset prices really log-normal? This chapter answers these questions. It explains that traders do use the Black–Scholes–Merton model—but not in exactly the way that Black, Scholes, and Merton originally intended. This is because they allow the volatility used to price an option to depend on its strike price and time to maturity.

A plot of the implied volatility of an option with a certain life as a function of its strike price is known as a *volatility smile*. This chapter describes the volatility smiles that traders use in equity and foreign currency markets. It explains the relationship between a volatility smile and the risk-neutral probability distribution being assumed for the future asset price. It also discusses how option traders use volatility surfaces as pricing tools.

19.1 WHY THE VOLATILITY SMILE IS THE SAME FOR CALLS AND PUTS

This section shows that the implied volatility of a European call option is the same as that of a European put option when they have the same strike price and time to maturity. This means that the volatility smile for European calls with a certain maturity is the same as that for European puts with the same maturity. This is a particularly convenient result. It shows that when talking about a volatility smile we do not have to worry about whether the options are calls or puts.

As explained in earlier chapters, put–call parity provides a relationship between the prices of European call and put options when they have the same strike price and time to maturity. With a dividend yield on the underlying asset of q , the relationship is

$$p + S_0 e^{-qT} = c + K e^{-rT} \quad (19.1)$$

As usual, c and p are the European call and put price. They have the same strike price, K , and time to maturity, T . The variable S_0 is the price of the underlying asset today, r is the risk-free interest rate for maturity T , and q is the yield on the asset.

A key feature of the put–call parity relationship is that it is based on a relatively simple no-arbitrage argument. It does not require any assumption about the probability

distribution of the asset price in the future. It is true both when the asset price distribution is lognormal and when it is not lognormal.

Suppose that, for a particular value of the volatility, p_{BS} and c_{BS} are the values of European put and call options calculated using the Black–Scholes–Merton model. Suppose further that p_{mkt} and c_{mkt} are the market values of these options. Because put–call parity holds for the Black–Scholes–Merton model, we must have

$$p_{BS} + S_0 e^{-qT} = c_{BS} + K e^{-rT}$$

In the absence of arbitrage opportunities, put–call parity also holds for the market prices, so that

$$p_{mkt} + S_0 e^{-qT} = c_{mkt} + K e^{-rT}$$

Subtracting these two equations, we get

$$p_{BS} - p_{mkt} = c_{BS} - c_{mkt} \tag{19.2}$$

This shows that the dollar pricing error when the Black–Scholes–Merton model is used to price a European put option should be exactly the same as the dollar pricing error when it is used to price a European call option with the same strike price and time to maturity.

Suppose that the implied volatility of the put option is 22%. This means that $p_{BS} = p_{mkt}$ when a volatility of 22% is used in the Black–Scholes–Merton model. From equation (19.2), it follows that $c_{BS} = c_{mkt}$ when this volatility is used. The implied volatility of the call is, therefore, also 22%. This argument shows that the implied volatility of a European call option is always the same as the implied volatility of a European put option when the two have the same strike price and maturity date. To put this another way, for a given strike price and maturity, the correct volatility to use in conjunction with the Black–Scholes–Merton model to price a European call should always be the same as that used to price a European put. This means that the volatility smile (i.e., the relationship between implied volatility and strike price for a particular maturity) is the same for European calls and European puts. It also means that the volatility term structure (i.e., the relationship between implied volatility and maturity for a particular strike) is the same for European calls and European puts.

Example 19.1

The value of the Australian dollar is \$0.60. The risk-free interest rate is 5% per annum in the United States and 10% per annum in Australia. The market price of a European call option on the Australian dollar with a maturity of 1 year and a strike price of \$0.59 is 0.0236. DerivaGem shows that the implied volatility of the call is 14.5%. For there to be no arbitrage, the put–call parity relationship in equation (19.1) must apply with q equal to the foreign risk-free rate. The price p of a European put option with a strike price of \$0.59 and maturity of 1 year therefore satisfies

$$p + 0.60e^{-0.10 \times 1} = 0.0236 + 0.59e^{-0.05 \times 1}$$

so that $p = 0.0419$. DerivaGem shows that, when the put has this price, its implied volatility is also 14.5%. This is what we expect from the analysis just given.