

WILEY FINANCE

University
Edition

Fixed Income Securities

Tools for Today's Markets

THIRD EDITION

BRUCE TUCKMAN
ANGEL SERRAT

Additional Praise for
Fixed Income Securities: Tools for Today's Markets,
3rd Edition

“The coverage of fixed income markets and instruments is even better than in previous editions while the book retains the same clarity of exposition via extensive, carefully worked examples. An outstanding textbook that is extensively used by practitioners is something special. This is indeed the standout text on fixed income.”

—Stephen M. Schaefer, Professor of Finance, London Business School

“This is a terrific reference text that combines a strong conceptual framework with real-world pricing and hedging applications. It is a must-read for any serious investor in fixed income markets.”

—Terry Belton, Global Head of Fixed Income Research, JPMorgan

“This outstanding book achieves the perfect balance between presenting the foundational principles of fixed income markets and providing interesting and insightful practical applications. This classic is required reading for anyone interested in understanding fixed income markets.”

—Francis Longstaff, Allstate Professor of Insurance and Finance,
The Anderson School at UCLA

“This is a great book. It covers the most current issues in fixed income and reflects the authors’ deep understanding of the markets grounded in the theory of finance and many years of practical experience.”

—Ardavan Nozari, Treasurer, Citigroup Global Markets Holding Inc.

Fixed Income Securities

Founded in 1807, John Wiley & Sons is the oldest independent publishing company in the United States. With offices in North America, Europe, Australia and Asia, Wiley is globally committed to developing and marketing print and electronic products and services for our customers' professional and personal knowledge and understanding.

The Wiley Finance series contains books written specifically for finance and investment professionals as well as sophisticated individual investors and their financial advisors. Book topics range from portfolio management to e-commerce, risk management, financial engineering, valuation, and financial instrument analysis, as well as much more.

For a list of available titles, visit our Web site at www.WileyFinance.com.

Fixed Income Securities

Tools for Today's Markets

Third Edition

BRUCE TUCKMAN
ANGEL SERRAT



John Wiley & Sons, Inc.

Copyright © 2012 by Bruce Tuckman and Angel Serrat. All rights reserved.

Published by John Wiley & Sons, Inc., Hoboken, New Jersey.

Published simultaneously in Canada.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning, or otherwise, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 646-8600, or on the Web at www.copyright.com. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at www.wiley.com/go/permissions.

Limit of Liability/Disclaimer of Warranty: While the publisher and author have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. The materials presented herein are the subjective views of the authors and do not represent current or past practices or beliefs of any organization or employer with whom the authors have been affiliated. No warranty may be created or extended by sales representatives or written sales materials. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate. Neither the publisher nor author shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

For general information on our other products and services or for technical support, please contact our Customer Care Department within the United States at (800) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books. For more information about Wiley products, visit our web site at www.wiley.com.

Library of Congress Cataloging-in-Publication Data:

Tuckman, Bruce.

Fixed income securities : tools for today's markets / Bruce Tuckman, Angel Serrat. – 3rd ed.
p. cm. – (Wiley finance series)

Includes index.

ISBN 978-0-470-89169-8 (hardback)

ISBN 978-0-470-90403-9 (paperback); ISBN 978-1-118-13397-2 (ebk);

ISBN 978-1-118-13398-9 (ebk); ISBN 978-1-118-13399-6 (ebk)

1. Fixed-income securities. I. Serrat, Angel. II. Title.

HG4650.T83 2012

332.63'2044–dc23

2011037178

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

Contents

Preface to the Third Edition	xi
Acknowledgments	xiii
An Overview of Global Fixed Income Markets	1
PART ONE	
The Relative Pricing of Securities with Fixed Cash Flows	47
CHAPTER 1	
Prices, Discount Factors, and Arbitrage	51
CHAPTER 2	
Spot, Forward, and Par Rates	69
CHAPTER 3	
Returns, Spreads, and Yields	95
PART TWO	
Measures of Interest Rate Risk and Hedging	119
CHAPTER 4	
One-Factor Risk Metrics and Hedges	123
CHAPTER 5	
Multi-Factor Risk Metrics and Hedges	153
CHAPTER 6	
Empirical Approaches to Risk Metrics and Hedging	171

PART THREE

Term Structure Models	201

CHAPTER 7

The Science of Term Structure Models	207

CHAPTER 8

The Evolution of Short Rates and the Shape of the Term Structure	229

CHAPTER 9

The Art of Term Structure Models: Drift	251

CHAPTER 10

The Art of Term Structure Models: Volatility and Distribution	275

CHAPTER 11

The Gauss+ and LIBOR Market Models	287

PART FOUR

Selected Securities and Topics	325

CHAPTER 12

Repurchase Agreements and Financing	327

CHAPTER 13

Forwards and Futures: Preliminaries	351

CHAPTER 14

Note and Bond Futures	373

CHAPTER 15

Short-Term Rates and Their Derivatives	401

CHAPTER 16

Swaps	435

CHAPTER 17

Arbitrage with Financing and Two-Curve Discounting	457

CHAPTER 18

Fixed Income Options	483

CHAPTER 19	
Corporate Bonds and Credit Default Swaps	527
CHAPTER 20	
Mortgages and Mortgage-Backed Securities	563
CHAPTER 21	
Curve Construction	591
References	607
Exercises	609
Index	623

Preface to the Third Edition

The goal of this book is to present conceptual frameworks for pricing and hedging a broad range of fixed income securities in an intuitive, mathematically simple, and applied manner. Conceptual frameworks are necessary so as to connect ideas across products and to learn new material more easily. An intuitive and mathematically simple approach is certainly useful to students and practitioners without very advanced mathematical training, but it is also really a good way for everyone to learn new material. Finally, an applied approach is crucial for several reasons. First, examples go a long way in solidifying conceptual understanding. The introduction of practically every concept in this book is followed by an example taken from the markets or, at the very least, by an appropriately calibrated example. Second, important details emerge from applications. Third, only by working through real or realistic examples can orders of magnitude be learned and appreciated. For example, a study of *DV01* is not complete without having absorbed that the sensitivity of a 10-year bond is about 8 cents per 100 face amount per basis point, as opposed to 0.8 cents, 80 cents, or 8 dollars.

The book begins with an Overview of global fixed income markets. This section provides institutional descriptions of securities and market participants along with data designed to illustrate absolute and relative sizes of markets and players. A well-informed fixed income market professional has some idea about how central banks around the world have reacted to the financial crisis of 2007–2009 and can say whether the size of the mortgage market in the United States is one-tenth the size of GDP, about equal to GDP, or 10 times GDP.

For securities with fixed cash flows, Part One of the book presents the relationships across prices, spot rates, forward rates, returns, and yields. The fundamental notion of arbitrage pricing is introduced and is central to the analysis. Part Two describes how to measure and hedge interest rate risk, covering one-factor metrics, namely, *DV01*, duration, and convexity (in both their general and yield-based forms); two-factor metrics like key-rate '01s, partial *PV01*s, and forward bucket '01s; and empirical methods like regression and principal component analysis.

Part Three turns to the arbitrage pricing of contingent claims, i.e., of securities with cash flows that depend on interest rates, like options. The science of arbitrage pricing in this context is followed by a framework in which

to think about the shape of the term structure of interest rates in terms of expectations, risk premium, and convexity. One-factor term structure models are then described, to be used both in their own right, when appropriate, and as building blocks toward more sophisticated models. Chapter 11, the last chapter in Part Three, has two parts. First, it presents a multi-factor model for use in relative value applications, along with suggestions for estimating its parameters empirically. Second, it introduces the *LIBOR* Market Model, an extremely popular model for pricing exotic derivatives, in a particularly accessible manner.

Finally, Part Four applies the knowledge gained in the previous three parts to present and analyze a broad and extensive range of fixed income topics and products including repo, bond and note futures, rate futures, swaps, options, corporate bonds and credit default swaps (CDS), and mortgage-backed securities.

This edition substantially revises and expands the second. The only parts of the book that have remained essentially unchanged are Chapters 7 through 10 on pricing contingent claims with one-factor term structure models. The rest of the material that was in the second edition has been updated and, with the exception of a couple of particularly interesting case studies, the numerical illustrations, examples, and applications are all new. In addition, several chapters in this third edition are completely new and others significantly expanded. New chapters include the Overview, Chapter 17 on how the realities of financing have changed the practice of discounting cash flows, and Chapter 19 on corporate bond and CDS markets. Significantly expanded chapters include Chapter 6 on empirical hedging, which now includes principal component analysis; Chapter 11, which was discussed above; Chapter 18 on volatility and fixed income options, which now covers a very broad range of products, Black-Scholes pricing, and a mathematically simple introduction to martingale pricing; and Chapter 20, on mortgages and mortgage-backed securities, which takes a much more market-oriented approach and adds material on pool characteristics, TBAs, and dollar rolls.

Acknowledgments

The authors would like to thank the following people for helpful comments and suggestions: Amitabh Arora, Larry Bernstein, John Feraca, Lawrence Goodman, Jean-Baptiste Homé, Dick Kazarian, Peyron Law, Marco Naldi, Chris Striesow, and Doug Whang. We would like to thank Helen Edersheim for carefully reviewing the manuscript and sparing readers from phrases like “options wroth about \$2” and the like. Finally, we would like to thank Bill Falloon, Meg Freeborn, and Natasha Andrews-Noel at Wiley for their patience and support throughout the process of putting this book together.

An Overview of Global Fixed Income Markets

This overview begins with a snapshot of fixed income markets across the globe and continues with concise reviews of fixed income markets in the United States, Europe, and Japan. These reviews have three goals: one, to describe how households and institutions achieve their borrowing and investing objectives through fixed income markets; two, to highlight the magnitude of the amounts of securities outstanding and of the balance sheets of market participants; and three, to emphasize the themes that are particularly relevant and significant for understanding today's markets.

A SNAPSHOT OF GLOBAL FIXED INCOME MARKETS

While fixed income markets are truly global, the vast majority of securities originate in the United States, Europe, and Japan. Figure O.1 shows the notional amounts outstanding of debt securities by residence of issuer, arranged in order of decreasing size. The largest markets by far are in the United States, the *Eurozone*, Japan, and the United Kingdom. (The Eurozone includes countries that are part of the European Union and also use the Euro as their currency.) The amounts outstanding for many Eurozone countries are shown individually in the graph, and indicated with asterisks, because several of these markets rank among the largest in the world on their own.

Derivative securities do not have an issuer in the same sense as do debt securities, but the distribution of the notional amounts of over-the-counter (OTC) interest rate derivatives across currencies tells a story similar to that of Figure O.1. According to Figure O.2, which shows amounts outstanding of single-currency, OTC interest rate derivatives, markets are dominated by contracts denominated in EUR (Euro), USD (United States dollar), JPY (Japanese Yen), and GBP (British Pound).¹ And with respect

¹The other currencies appearing in the graph are CHF (Swiss Franc), SEK (Swedish krone), CAD (Canadian dollar), AUD (Australian dollar), NOK (Norwegian krone), HKD (Hong Kong dollar), DKK (Danish krone), and NZD (New Zealand dollar).




FIGURE O.1 Debt Securities by Residence of Issuer as of March 2010
Source: Bank for International Settlements.

to exchange-traded derivatives, Table O.1 shows that Europe and North America comprise almost all of the outstanding notional amount.

It is worthwhile noting that Figures O.1, O.2, and Table O.1 report the place of origination of fixed income securities rather than the place of residence of the ultimate owners or counterparties. So, to take one of the more significant examples, China's ownership of nearly \$850 billion of U.S.




FIGURE O.2 Amounts Outstanding of OTC Single-Currency Interest Rate Derivatives as of December 2009
Source: Bank for International Settlements.

TABLE 0.1 Exchange-Traded Interest Rate Derivatives as of March 2010, in Billions of U.S. Dollars

Region	Notional
Europe	27,807
North America	22,604
Asia and Pacific	10
Other	934

Source: Bank for International Settlements.

Treasury securities does not appear anywhere in Figure O.1. Nevertheless, even accounting for such instances, the data presented here justify this book's focus on fixed income securities and markets in the United States, Europe, and Japan.

As a final note before turning to the three overviews, Figure O.3 gives a coarse breakdown of the composition of debt securities in the United States, the Eurozone, Japan, and the United Kingdom. (The totals are the same as those reported in Figure O.1.) While the proportions of debt issued by governments, financial institutions, and corporations are similar in the United States and the Eurozone, debt markets in Japan are dominated by governments while those in the United Kingdom are dominated by the issues of financial institutions.




FIGURE 0.3 Debt Securities by Residence of Issuer and Sector as of March 2010
Source: Bank for International Settlements.




FIGURE O.4 Credit Market Debt in the United States as of March 2010

Source: Flow of Funds Accounts of the United States.

FIXED INCOME MARKETS IN THE UNITED STATES

Securities and Other Assets

Figure O.4 shows the major categories of credit market debt in the United States, along with the size of the market for each, as of March 31, 2010.² Due to the definition of credit market debt in this cut of the data, several assets are not explicitly mentioned here (e.g., deposits, money-market fund shares, security repurchase agreements, insurance and pension reserves, equities), but will be included in the discussions of households and institutions later in this section.

Mortgages The largest single category of debt in the United States is *mortgages*, at a size of \$14.2 trillion. A mortgage is a loan *secured* by property, so that if a borrower fails to make the payments required by a mortgage, the lender has a claim on the property itself. Exercising this claim, the lender could keep proceeds from the sale of the property up to the amount still owed; or the lender could *foreclose* on the property, sell it, and recover the outstanding loan amount that way. In practice there might be

²The data for this figure and for much of this section come from the Board of Governors of the Federal Reserve System, “Flow of Funds Accounts of the United States,” June 10, 2010. See also the accompanying “Guide to the Flow of Funds Accounts.”

restrictions on the immediate or full exercise of this claim, like bankruptcy and other borrower protections or any tax liens on the same property. Finally, depending on the laws of the relevant state, the lender might or might not have *recourse* to the borrower's other assets to collect any remaining amount owed after the sale of the property.

Of the \$14.2 trillion outstanding, \$11.6 trillion is home or other residential mortgages, \$2.5 trillion is commercial mortgages, and \$138 billion is farm mortgages. To put the size of this market into context, two comparative statistics are useful. First, the annual gross domestic product (GDP) of the United States as of the first quarter of 2010 was \$14.6 trillion.³ Hence, it would take almost the entire output of the economy for one year to pay off all mortgage debt. Second, as of March 31, 2010, the public debt of the United States, at a historical high of \$12.8 trillion, was \$1.4 trillion less than the amount of mortgage debt outstanding.

Mortgages and mortgage-backed securities are the subject of Chapter 20.

Corporate and Foreign Bonds The second largest category of debt in Figure O.4 consists of corporate and foreign bonds. Corporate bonds are sold by businesses to finance investment, like the building of a new plant, the purchase of other businesses, or the purchase of investment securities. Bonds are also sold to *refinance* outstanding debt issues, that is, to retire existing debt not with corporate cash, which might have better uses, but with the proceeds raised by selling new debt. Motivations for retiring existing debt include redeeming maturing debt, repurchasing an issue to be rid of bond covenants that have become overly onerous, or exercising an embedded option to repurchase bonds at some prespecified and currently attractive call price.

As of the end of March 2010, \$11.4 trillion of corporate and foreign bonds were outstanding, \$5.6 trillion of which were sold by corporations in the financial sector. Proceeds raised by the financial sector, as will be discussed shortly, are used for the most part to purchase other securities.

Corporate bonds and derivatives on corporate bonds, namely, *credit default swaps*, or CDS, are the subject of Chapter 19.

Treasury Securities The next category is *Treasury securities*, obligations of the U.S. government incurred to finance its spending. U.S. Treasuries are among the most liquid securities in the world, meaning that investors can almost always buy and sell large amounts of Treasuries at prices close to relatively transparent market levels. In addition, while the finances of the U.S. government have deteriorated by historical standards, its debt is still

³Source: Bureau of Economic Analysis, U.S. Department of Commerce.

perceived as one of the safest investments in the world; when world events scare investors and trigger a “flight to quality,” demand for U.S. Treasuries increases and their prices rise. As shown in Figure O.4, at the end of March 2010, there were \$8.3 trillion of such securities outstanding, \$7.7 trillion of which were *marketable*, i.e., traded in markets.

With respect to the credit quality of Treasury securities, it is important to note that Treasuries do not constitute the entire public debt of the United States, which, as mentioned in the discussion of mortgages, is \$12.8 trillion or about 88% of GDP. The public debt includes *intragovernmental holdings*, i.e., debt that one part of the government owes to another in support of third-party claimants (e.g., holdings of government debt in the Medicare and Social Security trust funds). There is a statutory ceiling on the amount of public debt, which, of course, limits the issuance of Treasury securities as well, although this limit has been increased many times. An increase on February 12, 2010, the latest as of the time of this writing, raised the limit to \$14.294 trillion. In any case, the ratio of public debt to GDP is used as an indicator of the credit quality of government debt, although it is widely recognized that certain economies can sustain higher ratios than others. With this caveat, the ratio of 88% in the United States is low compared with Japan but high compared with several, although certainly not all, European countries. Furthermore, 88% is very high relative to recent U.S. history: at the end of 2006, the public debt was \$8.7 trillion and GDP \$13.4 trillion, for a ratio of only 65%.

With the increasing global scrutiny of government financing, the maturity structure of government debt has taken on new importance. Since shorter-term rates are usually lower than longer-term rates, there is always an incentive to reduce borrowing costs by concentrating borrowing at shorter maturities. But this strategy can be dangerous; the more debt with shorter maturities, the greater a government’s difficulties should investors suddenly become reluctant to purchase its new bond issues. While the United States has not had any trouble selling its debt, as Greece and Ireland recently

TABLE 0.2 Maturity Structure of U.S. Treasury Marketable Securities as of March 31, 2010

Maturity Years	Outstanding \$Billions	Percent %
< 2	3,482	45
2–5	1,953	25
5–10	1,528	20
> 10	782	10

Source: Monthly Statement of the Public Debt of the United States, March 31, 2010, and authors’ calculations.

have, in the spirit of this new scrutiny Table O.2 presents the maturity structure of marketable U.S. Treasury securities. In comparison with the maturity structures in Europe and Japan (see Tables O.12 and O.18, respectively), government borrowing in the United States is relatively heavy at the shorter maturities.

Turning now to taxonomy, the U.S. Treasury issues securities in several different forms. Treasury *bills*, or *T-bills*, mature in one year or less and are *discount* securities, meaning that they make no payments until the promised payment at maturity and, consequently, sell for less than, i.e., at a discount from, that promised payment. Treasury *notes* are *coupon-bearing* securities, issued with 10 or fewer years to maturity, that make semiannual interest payments at a fixed rate and then return principal at maturity. Treasury *bonds* are also coupon-bearing securities, but with original maturities greater than 10 years. This separate classification of notes and bonds continues today, but is a historical artifact: bonds used to be subject to a maximum, statutory rate of interest, but this limit was eliminated in 1988.⁴ In any case, this book uses the term “bond” loosely to refer to notes or bonds.

The U.S. Treasury also issues *TIPS*, or *Treasury Inflation Protected Securities*. The principal of TIPS adjusts to reflect changes in the consumer price index so that the coupon, together with the return of indexed principal, preserves a real return, i.e., a return above inflation. The maturing principal of a TIPS, however, will never be set below the original principal, no matter how much deflation might take place. As of March 31, 2010, the amount of TIPS outstanding was a relatively small \$573 billion, less than 7% of the \$8.3 trillion of Treasury issues. Nevertheless, TIPS have a significance beyond their size as their prices reveal market expectations about future inflation.

The last category of U.S. Treasury securities to be mentioned here, simply because they are well known, are U.S. savings bonds, which are nonmarketable, discount securities sold mostly to retail investors. As of March 31, 2010, the amount of savings bonds outstanding was a relatively tiny \$190 billion.

In a largely successful effort to foster the liquidity of Treasury securities, the U.S. Treasury auctions them to the public at preannounced auction dates and quantities. The schedule of which maturities are offered and at what frequencies changes slowly over time with the financing needs of the Treasury. Currently, bills with maturities of 4, 13, and 26 weeks are sold weekly, while bills maturing in 52 weeks are sold every four weeks. Notes with 2-, 3-, 5-, and 7-year maturities are sold monthly. There are quarterly issues of 10-year notes and 30-year bonds, although individual issues are

⁴ Marcia Stigum, *The Money Market*, 3rd Edition, (Dow Jones-Irwin, 1990) p. 37.

reopened, i.e., sold through subsequent auctions.⁵ Finally, 5- and 30-year TIPS are issued annually and reopened twice per year while 10-year TIPS are issued semiannually and reopened four times per year.

Agency- and GSE-Backed Securities *Agency- and GSE-backed securities* are obligations of agencies of the U.S. government and of *GSEs* or *government-sponsored entities*. This category consists of three subcategories:

- Debt issues of U.S. agencies, which comprise only \$24 billion of the \$8.1 trillion total.⁶
- Debt issues of such GSEs as FHLMC (Federal Home Loan Mortgage Corporation or “Freddie Mac”), FNMA (Federal National Mortgage Association or “Fannie Mae”), and FHLB (Federal Home Loan Banks), which comprise \$2.7 trillion of the total. These debt issues are used to finance a portfolio of mortgage-related investments, mostly portfolios of mortgages in the case of FHLMC and FNMA and mostly secured loans to banks making mortgage loans in the case of FHLB.
- Issues of mortgage-backed securities by FHLMC, FNMA, and of the wholly-owned government corporation GNMA (Government National Mortgage Association or “Ginnie Mae”), which comprise \$5.4 trillion of the total. Aside from the *portfolio business* described in the previous bullet point, FHLMC and FNMA have a *guarantee business*, as does GNMA. This business consists of guaranteeing the performance of *conforming* mortgages (i.e., mortgages that meet specified criteria) in exchange for a fee. These mortgages are then bundled into mortgage-backed securities, which, in turn, are sold to investors.

The historical justification for GSEs has been that they serve a public purpose in addition to making profits for their shareholders. In the case of the mortgage-related GSEs, this public purpose has been to facilitate home ownership. As a result of this mix of public and private objectives, there

⁵As an example, consider the issuance and two scheduled reopenings of the 2.625% notes maturing on August 15, 2020. A face amount of \$24 billion of these notes was initially sold to the public on August 11, 2010. Subsequently, in the first reopening auction, on September 8, 2010, another \$21 billion of this issue was sold. Then, in the second and final scheduled reopening, on October 13, 2010, yet another \$21 billion was sold.

⁶This discussion of agency- and GSE-backed securities uses Flow of Funds data as of December 2009 instead of March 2010, as in Figure O.4. As of 2010, mortgage pools were consolidated into the balance sheets of FNMA and FHLMC, blurring the distinction between GSE debt securities and mortgage-backed securities.

has always been furor about the extent to which the U.S. government is responsible for agency or GSE debt that it has not explicitly guaranteed, particularly in the cases of FNMA and FHLMC. These GSEs have been able to borrow at advantageous terms⁷ because the global investment community has believed there is an implicit U.S. government guarantee, despite occasional statements by officials denying that to be the case. And in fact, after September 2008, when FNMA and FHLMC were failing and placed into government conservatorship, the U.S. government did exert considerable effort to protect and calm bondholders.⁸

Municipal Securities Municipal securities or *munis* are for the most part issued by state and local governments. The variation across issues is particularly large in this market, with over 55,000 different issuers⁹ and a staggering number of distinct issues. Shorter-term issues are typically used for cash management purposes, e.g., to manage time gaps between tax collections and expenditures, while longer-term debt issues are often used to finance infrastructure projects. *General obligation (GO)* bonds are backed by the full faith and credit of the issuing municipality while *revenue* bonds are backed by the cash flows from a particular project. Municipal bonds as an investment class have historically had very low rates of default, but perceived creditworthiness does vary dramatically across issues. At the safest extreme are GO bonds of the most creditworthy states while at the other extreme are *revenue* bonds dependent on particularly risky projects. At the time of this writing the credit quality of municipals is under increased scrutiny because spending commitments made in better economic environments are now straining municipal budgets.

An extremely important feature of the municipal bond market is that the interest on the vast majority of issues is exempt from U.S. federal income tax. As a result, municipalities are able to pay much lower rates of interest than would otherwise be the case. Nevertheless, investors subject to the highest marginal federal tax rate earn a higher rate on municipal bonds, particularly those of longer term, than they earn, on an after-tax basis, on otherwise comparable taxable bonds.

Muni investors often enjoy advantages with respect to state income taxes as well, although the exact treatment varies by state. Most commonly,

⁷The Congressional Budget Office has estimated that the implicit government guarantee enables the GSE to raise funds at a rate savings of 0.41% through their debt issues and 0.30% through their mortgage-backed security issues. See "Updated Estimates of the Subsidies to the Housing GSEs," Congressional Budget Office, April 8, 2004.

⁸See, for example, Federal Housing Finance Agency, "U.S. Treasury Support for Fannie Mae and Freddie Mac," Mortgage Market Note 10-1, January 20, 2010.

⁹See the Municipal Securities Rulemaking Board website, www.msrb.org.

a state exempts interest on bonds it has issued while taxing interest on bonds sold by other states.

The much-heralded Build America Bond (BAB) program, created in February 2009, expired at year-end 2010 and, as of the time of this writing, has not been renewed by Congress. Bonds under this program were typically sold as taxable to the investor, with the U.S. government rebating 35% of the interest to the issuing municipality. As of the end of October 2010, only about \$150 billion of BABs had been sold, compared with approximately \$3 trillion of municipal bonds outstanding,¹⁰ but the program was very popular with municipalities. BABs opened the municipal market to investors in low or zero tax brackets who typically buy taxable bonds. On the other hand, the program is costly for the U.S. government to maintain.¹¹

Other Categories Loans and advances of \$4.1 trillion in Figure O.4 include loans made by banks and others (e.g., government, GSEs, finance companies) that are not included in any other category. Almost all of consumer credit of \$2.4 trillion consists of credit card balances and automobile loans. Finally, open market paper of \$1.1 trillion consists almost exclusively of *commercial paper*. Commercial paper issuers borrow money from investors on an unsecured and short-term basis, with maturities extending up to 270 days¹² but averaging about 30 days.

Households and Institutions

Figures O.5 and O.6 show the largest sectors that borrow and that lend funds through credit markets, respectively. These sectors are now discussed in turn, leaving out those that were already described in the securities subsection. (Note that, as in Figure O.4, only securities defined as credit market debt are included in Figures O.5 and O.6. Other assets, however, are included in the balance sheets to follow.)

Households Table O.3 shows the balance sheet for households and non-profit organizations as of March 2010. Note that the percentage of liabilities is exactly that, and not the percentage of liabilities plus net worth. Hence, there is no percentage associated with net worth.

¹⁰ Source: U.S. Build America Bond Issuance, Securities Industry and Financial Markets Association (SIFMA), and Flow of Funds Accounts of the United States.

¹¹ The argument that the program is costly is that since BABs, like all heavily taxed assets, are bought primarily by investors in low or zero tax brackets, the tax paid on BABs is very much below the 35% subsidy. It has been argued by some, however, that the relevant tax rate is higher so that a larger portion of the subsidy is recouped.

¹² Longer maturities would trigger Securities and Exchange Commission (SEC) registration requirements.




FIGURE 0.5 Credit Market Debt Owed as of March 2010
Source: Flow of Funds Accounts of the United States.

The largest asset of households is real estate followed by pension fund savings. Holdings of other assets are spread relatively evenly, with a significant percentage in equity of noncorporate business, e.g., relatively small, family-run businesses. The liabilities of households are predominantly mortgages and consumer credit, the latter consisting mostly of credit card debt and automobile loans. In short then, households own their homes and durable goods and invest in a wide range of financial assets, a significant portion of which are held through pension funds. Households borrow mostly




FIGURE 0.6 Credit Market Assets Held as of March 2010
Source: Flow of Funds Accounts of the United States.

TABLE 0.3 Balance Sheet of Households and Nonprofit Organizations as of March 2010, in Trillions of Dollars

Assets	68.5	100%
Real Estate	18.1	26.5%
Consumer Durables	4.6	6.8%
Deposits and Money Market Funds	7.7	11.2%
Credit Market Instruments	4.2	6.1%
Corporate Equities	7.8	11.4%
Mutual Funds	4.3	6.3%
Pension Fund Reserves	12.3	18.0%
Equity in Noncorporate Business	6.5	9.5%
Life Insurance Reserves	1.3	1.8%
Miscellaneous	1.7	2.5%
Liabilities	14.0	100%
Home Mortgages	10.2	73.3%
Consumer Credit	2.4	17.3%
Miscellaneous	1.3	9.4%
Net Worth	54.6	

to finance their housing and durable purchases, but also to manage their short-term cash requirements. The European market overview, by the way, will discuss pension funds in more detail.

Since the financial crisis of 2007–2009 has had a significant impact on the balance sheets of households and institutions, it is noted here and in subsequent discussions how balance sheets have changed since the end of 2006. With respect to households, net worth has fallen from \$64.4 to \$54.6 trillion, or by more than 15%. And of this \$9.8 trillion drop, \$7.1 trillion or 11% was from a fall in the value of real estate assets and most of the rest from falling values of stocks and noncorporate equity.

Nonfinancial, Nonfarm Businesses Table O.4 gives the balance sheet of corporate and noncorporate businesses, excluding the financial and farm sectors. Businesses in the financial sector will be covered in later subsections and the farm sector is relatively small.

Nonfinancial business assets consist of real estate and equipment, along with a large portion classified as miscellaneous. There is a reasonable amount of trade financing, amounting to 7.2% of assets and 10.8% of liabilities. As for longer-term liabilities, businesses finance property with mortgages while financing other assets with loans and corporate bonds.

Table O.4 is not a snapshot of an individual business but an average across the sector, which obscures the life-cycle of financing a business. Initial capital comes from “friends and family” and bank loans. Then, as a

TABLE 0.4 Balance Sheet of Nonfinancial Nonfarm Businesses as of March 2010, in Trillions of Dollars

Assets	36.4	100%
Real Estate	12.1	33.1%
Equipment and Software	5.0	13.6%
Inventories	1.8	4.9%
Deposits and Credit Market Instruments	2.7	7.3%
Trade Receivables	2.6	7.2%
Miscellaneous	12.4	33.9%
Liabilities	18.9	100%
Corporate Bonds	4.2	22.5%
Mortgages	3.4	18.1%
Trade Payables	2.0	10.8%
Loans and Miscellaneous	9.2	48.7%
Net Worth	17.5	

business grows, it may obtain loans from investor groups and from *private placements* of debt (e.g., negotiating the terms of a loan with one or several insurance companies). Finally, a larger business, with a track record and name recognition, can tap public bond markets.

From year-end 2006 to the end of the first quarter of 2010, the balance sheet of nonfinancial businesses deteriorated along with those of households: liabilities rose and assets fell, the latter predominantly because of real estates' values falling. As a result, net worth fell by about \$5 trillion, or 23%, from \$22.6 to \$17.5 trillion. Or, taking a different perspective, the ratio of liabilities to assets increased from 42% to 52%.

Commercial Banking Table O.5 gives the financial assets and liabilities of the commercial banking sector as of March 31, 2010. Note that unlike the previous balance sheets, this one lists only financial assets and liabilities. This is a reasonable view for financial intermediaries whose nonfinancial assets are relatively insignificant.

The sources of funds for the commercial banking sector as a whole are deposits, *federal funds* (overnight loans between banks in the federal reserve system; see Chapter 15), and *repurchase agreements* or *repo* (usually very short-term loans secured by relatively high-quality collateral; see Chapter 12), bonds, and other sources. These funds are invested in a broad range of assets, although a significant percentage of these are mortgages (26%) or mortgage-related (i.e., agency- and GSE-backed securities at 8.8%). Banks make money by earning spreads between the rates they pay on their sources of funds and the rates of return they earn on their assets. But to earn spreads, banks have to take certain risks. In particular, banks typically take on three types of risk. First, banks take *credit risk* by lending to homeowners and to

TABLE 0.5 Financial Assets and Liabilities of Commercial Banks as of March 2010, in Trillions of Dollars

Financial Assets	14.4	100%
Reserves at Federal Reserve	1.0	6.7%
Agency- and GSE-backed Securities	1.3	8.8%
Corporate and Foreign Bonds	.8	5.4%
Loans	1.8	12.2%
Mortgages	3.8	26.0%
Consumer Credit	1.2	8.1%
Other and Miscellaneous	4.7	32.7%
Financial Liabilities	12.8	100%
Deposits	7.6	59.6%
Federal Funds and Repo (net)	.9	6.7%
Open Market Paper	.2	1.6%
Corporate Bonds	1.4	10.7%
Other Loans and Advances	.4	2.8%
Other and Miscellaneous	2.4	18.6%

businesses that may not repay their borrowings as promised. This source of risk is not a main focus of this book, but will be discussed in Chapter 19. Second, banks may take *interest rate risk* by borrowing with shorter-term securities but investing in longer-term assets. Shorter-term funds can usually be borrowed at relatively low rates of interest but, as these borrowings come due, banks run the risk of having to pay higher rates of interest on new borrowing. At the same time, longer-term lending is usually initiated at relatively high rates of interest, but these rates are fixed for years. Hence, should a bank's shorter-term borrowing costs rise relative to its fixed lending rates, its profit margin or spread will narrow or even turn negative. Interest rate risk is the subject of Part Two of this book.

The third source of risk for banks is *financing risk*. Deposits are regarded as relatively stable sources of funds because of deposit insurance: because the FDIC (Federal Deposit Insurance Corporation) insures deposits, at least up to a limit, depositors do not need to pull deposits at the first breath of rumor about a bank's financial health. Corporate bonds, with their relatively long maturities, also constitute a stable source of funds in the sense that a bank has time between the surfacing of any financial problems and the maturity of its bonds to sort out its difficulties.¹³ Federal funds and repo, however, along with open market paper, are shorter-term sources of funds and are less stable: at the first sign that a bank is in financial difficulties, its ability to finance itself with federal (fed) funds and repo can erode in days, to be

¹³ Of course, corporations typically stagger the maturities of their longer-term debt to ensure that only manageable amounts come due at any one time.

followed over subsequent months by the erosion of its ability to sell open market paper. Depositors with amounts above the insured limit may also withdraw the excess amounts.¹⁴ And should all this happen, a bank cannot simply let its assets mature commensurately because most of the assets are of much longer term. Hence, the only way to meet short-term maturing obligations might very well be a fire-sale of assets at a substantial loss. Such a dramatic and sudden loss of short-term financing is often called a *run on the bank* and can certainly lead to bankruptcy. Chapter 12 revisits financing risk in the context of broker-dealer balance sheets during the financial crisis of 2007–2009.

As discussed with respect to the business sector, the balance sheet of a sector does not show variation across individual entities within that sector. As larger banks have better opportunities to borrow than do smaller banks, they tend to rely less on funding from deposits: the ratio of deposits to liabilities for the largest 25 U.S.-chartered banks was 69% in March 2010, compared with 84% for the smaller banks.¹⁵ Another significant source of variation across banks is the historical reliance of smaller banks on local real estate lending, which reliance resulted in significant losses and failures during the 2007–2009 crisis. For the smaller U.S.-chartered banks, 54% of their assets as of March 2010 were related to real estate, 45% in the form of loans and 9% in the form of mortgage-backed securities. For the 25 largest U.S.-chartered banks, by contrast, 43% of assets were real estate-related, with 30% in loans and 13% in securitized form.¹⁶

A significant difference between household and nonfinancial business balance sheets compared with that of commercial banking is *leverage*, or the amount of assets supported by a given amount of liabilities. From Tables O.3 and O.4, the assets of the household sector are 4.9 times the liabilities, and the assets of the nonfinancial business sector are 1.9 times the liabilities. And at the end of 2006, before the financial crisis, the ratios were somewhat higher, at 5.8 and 2.4 respectively. Nevertheless, the value of the assets of these sectors can fall significantly before assets are insufficient to pay off liabilities. By contrast, the ratio of financial assets to financial liabilities of the commercial banking sector, from Table O.5, is only 1.1. Put another way, according to the table, the *equity* or the cushion of assets over liabilities equals \$1.6 trillion. Thus, an 11.1% drop in the value of the \$14.4 trillion

¹⁴ Lines of credit, which allow customers to draw loans from banks, up to pre-specified amounts, also contribute to financing risk. In times of stress, customers will draw their lines while banks are losing their sources of funds.

¹⁵ Source: Board of Governors of the Federal Reserve System, “Assets and Liabilities of Commercial Banks in the U.S.” Note that the category U.S.-chartered commercial banks is a subset of the larger commercial banking sector described by Table O.5. This explains why these ratios do not bracket the comparable ratio in the table.

¹⁶ Ibid.

in assets would wipe out the equity of the sector. And of course, any agent that had more leverage than its sector average would suffer commensurately larger losses of equity for any given loss of asset value.

This simplified discussion of leverage can explain how the banking sector got into trouble during the 2007–2009 financial crisis. At year-end 2006, 41.9% of assets were in mortgages and in agency- and GSE-backed securities. Taking bank capital, which, roughly speaking, is the cushion between the value of bank assets and liabilities, to be 10% of assets,¹⁷ it takes only a 12% drop in the 41.9% of mortgage-related assets, assuming no other assets fall in value, to reduce asset value by $41.9\% \times 12\%$ or 5% and cut bank capital in half. And to the extent that a particular bank had an even larger fraction of mortgage-related assets, or to the extent that other assets, like loans to troubled businesses also fell in value, the effect would be that much greater.

Monetary Authority, or the Board of Governors of the Federal Reserve System (Fed) The liabilities of the Fed are predominantly the reserves and deposits of banking institutions in the federal reserve system. The conduct of monetary policy is far beyond the scope of this book, but, to review, in the simplest of terms: the Fed's goals are given by the Federal Reserve Act, namely, "to promote effectively the goals of maximum employment, stable prices, and moderate long-term interest rates." Therefore, when the Fed believes that economic growth could be greater without causing inflation, it lowers short-term interest rates in an attempt to encourage borrowing and investment. And the way in which the Fed lowers interest rates is to increase the supply of funds relative to the demand by lending money to banks and taking securities as collateral. This collateralized lending is done through repo, which technically means that the Fed buys securities while simultaneously agreeing to resell them at a fixed price at some short time in the future. (See Chapter 12.) The Fed's balance sheet increases with these operations: its assets increase by the amount of securities taken as collateral (i.e., temporarily bought) from banks and its liabilities increase by the deposits made by banks with the amounts borrowed (i.e., temporarily sold). The process works in reverse when the Fed believes that inflation risks dominate and decides to raise short-term interest rates to discourage borrowing and investment. In this case the Fed borrows money and gives

¹⁷ A fuller discussion of bank capital ratios is beyond the scope of this overview. Measures of capital ratios and their corresponding regulatory thresholds vary depending on which forms of financing count as capital and on how assets are measured. With respect to capital, common equity always counts, but, for example, subordinated debt is included only by the broader definitions in the spectrum. And with respect to assets, the computation is typically either a simple sum or a risk-weighted sum of individual asset values, where the risk-weights are determined by regulators.

securities as collateral, which reduces its balance sheet: assets fall by the amount of securities temporarily sold and liabilities fall by the decrease in bank deposits to pay for securities temporarily purchased. Traditionally, the Fed has invested the proceeds from incurring bank reserve and deposit liabilities in U.S. Treasuries, the safest and most liquid domestic securities available.

The 2007–2009 financial crisis dramatically changed the Fed’s balance sheet. Responding to the crisis and the ensuing economic slowdown, the Fed lowered rates from 5.25% at the end of 2006 to between 0% and 0.25% in December 2008. Then, believing that the traditional lowering of interest rates by supplying the banking system with reserves was not spurring growth as desired, and worried that the real estate and mortgage markets remained dangerously fragile, the Fed also bought mortgage-related securities directly. The idea was to inject cash into the system in a different way while stabilizing the real estate and mortgage markets. The scale of these operations raised the assets on the Fed’s balance sheet from \$908 billion at the end of 2006 to \$2.3 trillion at the end of March 2010. Furthermore, over the same period, the composition of the Fed’s assets changed from over 90% in either Treasuries or loans against Treasury collateral to about 33% in Treasuries and 53% in agency- and GSE-backed securities. Or, from another perspective, the Fed held no agency and GSE-backed securities at the end of 2006 but held about 16% of the amount outstanding of these securities by the end of March 2010. A concern about this situation is that with \$1.2 trillion of mortgage-related securities, the Fed’s balance sheet is subject to an unprecedented amount of risk. From this point of view, one facet of the U.S. government’s intervention through the crisis was to move mortgage-related assets from the private sector’s balance sheet to that of the Fed.

Issuers of Asset-Backed Securities In the boom before the 2007–2009 crisis, there was great demand for securitized assets, i.e., it was profitable to acquire assets, most often mortgage-related but also including student loans, business loans, automobile loans, and credit card receivables, and then sell securities with payouts that depended, sometimes in complex ways, on the performance of those assets. Chapter 20 discusses the securitization process for mortgages.

There are “on-balance sheet” and “off-balance sheet” approaches to securitization. In the on-balance sheet approach, a financial institution acquires the underlying assets outright and then recovers these funds, hopefully at a profit, when selling the securities. In the off-balance sheet approach, a financial institution sets up a separate financial entity, called an *SPV* for *special purpose vehicle* or an *SIV* for *special investment vehicle*, which purchases the securities by issuing short- and long-term debt whose performance ultimately depends on that of the underlying assets. Before the 2007–2009

crisis, financial institutions preferred the off-balance sheet approach for two reasons. First, they wanted to be in the “moving” business rather than the “storage” business, i.e., they wanted to be paid for the acquisition and eventual sale of assets but did not want any of the resulting risks on their books. It turned out, however, that many SPVs could not sell debt without various guarantees from their sponsors. While providing these guarantees made the risk of the off-balance sheet approach similar to that of the on-balance sheet approach, there was the second reason for preferring the former. Regulatory capital requirements and pressure from the investment community discouraged a direct increase in balance-sheet assets and liabilities without commensurate increases in capital while indirect claims on the balance sheet through the guarantees were not penalized as readily.

An inherent problem of asset-backed vehicles that rely on the sale of short-term debt or commercial paper is financing risk. Should the market begin to doubt the quality of the underlying assets, short-term debtholders will refuse to roll their loans at anywhere near the originally contemplated rate levels and new lenders will be equally difficult to find. In that case the SPV might very well not be able to redeem the claims of these short-term debt holders and would either default or fall back on any guarantees provided by the sponsoring financial institution. During the 2007–2009 financial crisis, there were many instances in which SPVs were unwound and put back onto the balance sheets of their sponsors.

At the end of 2006, there were \$4.2 trillion of assets in these special purpose entities, 74.3% of which were mortgage-related. Furthermore, 19.9% of the liabilities of these entities were in the form of commercial paper. The assets in these entities continued to grow for a while, reaching \$4.5 trillion at the end of 2007, but the decline in real estate prices and the resulting effect on mortgage-related securities soon took their toll. By March 31, 2010, the assets in these entities had fallen to only \$2.8 trillion and the fraction of commercial paper in their liabilities was reduced to 4.4%.

Life Insurance Companies Life insurance companies sell insurance and annuity products that investors find particularly attractive for tax reasons, in particular, for tax-free death benefits and tax-deferral of savings. From the sale of these products, life insurance companies collect premiums that they invest so as to meet the obligations incurred and to earn an excess return. They choose to invest a large portion of their portfolios in longer-term assets, both to match the term of their liabilities better and to meet their return hurdles. They also take on default and equity risk to meet these hurdles. As of March 31, 2010, they invested a significant fraction of their assets in corporate and foreign bonds (39.9%) and in equities (27.1%). The former constituted 17% of the amount outstanding of corporate and foreign bonds, making insurance companies significant players in that market.

Broker-Dealers The magnitude of the balance sheet of broker-dealers (B/Ds), at a bit over \$2 trillion at the end of March 2010, did not warrant inclusion in Figures O.5 and O.6. The sector, however, is clearly important to fixed income markets and to the functioning of the financial system.

Broker-dealers have three lines of business, although the three cannot always be cleanly separated from one another. First, investment banking helps customers raise money from capital markets. Second, sales and trading facilitate customer trading in a broad range of securities with the B/D acting as a broker, i.e., buying or selling on behalf of a customer, or as a dealer, i.e., trading on the B/D's own account for later trading with a customer. Third, proprietary trading or positioning, broadly defined trades securities for profit on the B/D's own accounts. Investment banking, strictly defined, does not require much in the way of funding, although there is an associated proprietary side of the business in which, as part of a larger client transaction, a B/D commits its own capital to the client on a short-term or even longer-term basis. Sales and trading often require funding, as broker-dealers find it necessary to hold an inventory of securities to facilitate customer trading. Finally, a proprietary business, which by its nature holds positions, requires longer-term funding.

The asset side of B/D balance sheets consists of a range of securities, consistent with making markets and proprietary positions across different markets. The liability side also has some corporate bonds that serve as a long-term source of funds. Security credit, which is made up of loans to the B/D from banks to finance securities and customer deposits with B/Ds, are a larger part of the liability side. Finally, liabilities include secured borrowing through repos. Repo borrowing can usually be achieved at relatively low credit spreads since the loans are short-term and secured. Precisely because repo borrowing is usually short-term, however, with most being overnight, B/Ds are subject to the same financing risks as discussed in the context of commercial banks. In fact, during the 2007–2009 crisis, when lenders became nervous about the credit quality of B/Ds, repo funding became hard to maintain on any terms and contributed to the stress on B/Ds. (See Chapter 12 for examples of this.) Since then, B/Ds have tried to rely less on short-term repo financing than previously. Even more dramatically with respect to managing financing risk, all of the major investment banks that survived through the fall of 2008 converted to bank holding companies, giving them access to the safety net of the Fed's discount lending window.

FIXED INCOME MARKETS IN EUROPE

An overview of European fixed income markets is particularly challenging. European markets are comprised of many individual country markets which, as mentioned at the start of this chapter, can be divided in many different

TABLE 0.6 Financial Assets of Households, by Asset Class, 2007

	GDP	Financial Assets		Currency and Deposits	Pensions and Insurance	Equities	Other and Misc.
	\$Trillions	% of GDP	\$Trillions	% of Assets	% of Assets	% of Assets	% of Assets
Germany	3.328	188	6.270	36	26	25	13
UK	2.800	296	8.285	27	54	16	3
France	2.598	189	4.905	29	38	27	6
Italy	2.118	241	5.100	27	17	34	22
Spain	1.443	182	2.631	38	14	42	6
Netherlands	.779	280	2.184	22	59	15	4
Belgium	.459	271	1.244	28	23	40	9
Switzerland	.434	375	1.628	24	42	25	9
Denmark	.227	235	0.533	21	43	30	6

Sources: IMF and Eurostat.

ways: politically (i.e., countries in the European Union), by currency (i.e., countries using the Euro), by the intersection of the two, (i.e., countries in the Eurozone), or by other subdivisions (e.g., the Benelux countries, including Belgium, the Netherlands, and Luxembourg). Not surprisingly then, there is no single source of data that looks across all of the relevant countries with consistent classifications of securities or of financial market participants. Consequently, the data presented in this section come from many different sources, neither with perfectly consistent categories nor with a single ‘as of’ date.

Households and Institutions

Households¹⁸ Table O.6 describes the financial assets of the household sector in several European countries. The countries are listed in order of decreasing GDP and financial assets are presented both as a percentage of GDP and in absolute terms. As often the case when several currencies are involved, all absolute quantities have been expressed in U.S. dollars.

The financial assets listed in Table O.6 range from \$1.2 to \$8.3 trillion and sum to about \$33 trillion. In magnitude, then, the financial assets of the household sectors of individual European countries are small relative to those in Japan or the United States. Taken together, however, as a bloc, the financial assets of the household sector in Europe exceed those in Japan,

¹⁸ The data for this subsection come mostly from “Financial Assets and Liabilities of Households in the European Union,” Eurostat, 2009.




FIGURE 0.7 Private Pensions and Insurance Savings of Households in Europe, as a Percentage of Financial Assets, vs. the Replacement Rate of Public Pension Systems, 2006–2007

but still fall short of those in the United States. (See Tables O.13 and O.3, respectively.¹⁹)

The rightmost four columns of Table O.6 give the percentages of financial assets invested in particular asset categories. The first of these categories is currency and deposits. The percentage of financial assets held in this extremely safe and liquid form may be taken as a rough measure of the risk aversion of households with respect to personal investments. These percentages vary from 21% to 38%, which are all large relative to the 17% for U.S. households²⁰ but small relative to the 55% for Japanese households.²¹

The next column of Table O.6 shows that the percentages of European household financial assets held through private pensions and insurance products, two forms of long-term savings, vary widely across countries, from a low of 14% in Spain to a high of 59% in the Netherlands. Much of this cross-sectional variation can be explained by the variation of state-provided retirement benefits across countries. Figure O.7 graphs the pensions and insurance allocations from Table O.6 against the *replacement rate* of state plans, where the replacement rate is a ratio of pre-tax benefits to retirees'

¹⁹ These tables, with 2010 numbers, are not strictly comparable to the 2007 numbers of Table O.6. Comparing the relevant magnitudes with data from a single year, however, would yield the same qualitative results.

²⁰ See Table O.3, noting that real estate, consumer durables, and some of the miscellaneous category are not considered financial assets. Also see footnote 19.

²¹ See Table O.13 and footnote 19.

most recent pre-tax income.²² Clearly, the more generous the state-provided pension benefit, the less the household sector devotes on its own to retirement and long-term savings vehicles. The trend line of this relationship in Europe applies to the United States²³ and Japan²⁴ as well, which are included in Figure O.7 for comparison.

Returning to Table O.6, summing the percentage allocations in currency and deposits with those in pensions and insurance products gives an indicator of the extent of financial intermediation of household savings in Europe. The average of these sums across countries is about 63%, which is high relative to the United States, at about 48%,²⁵ but low relative to Japan, at about 82%.²⁶ The exception to this ordering is Italy, at a sum of 43.5%, partially because of the high replacement rate of its public pension system, through the effect described in the previous paragraph, and partially because households can easily purchase domestic government bonds given the relatively large supply available (see Table O.12).

Pension Funds Pensions provide people with an income when they are older and, most likely, no longer employed. While the structure of pension provisions in Europe varies dramatically across countries, all models are based on a three-pillar system. The first pillar is made up of public pensions, paid by the state; the second pillar is comprised of occupational pension schemes, paid by employers to their retired employees; and the third pillar consists of private retirement plans, through which individuals accumulate savings to provide a pension upon retirement.

In the private plans of the second and third pillars, employers and employees, in some combination, contribute to a fund, often managed by a trustee, and often with certain tax advantages (i.e., tax-deductible contributions and tax-free accumulation of investment income and capital gains). Then, upon retirement, the beneficiary is given a lump-sum payment, an annuity, or a combination of the two. Private pension plans can be divided into three major categories: *defined benefit* plans, *defined contribution* plans, and *hybrid* plans. In a defined benefit plan, the sponsor of the plan promises

²² Source: Allianz, as of 2006 and 2007.

²³ For replacement rates, see, for example, Chart 1 and Table O.1 of Patricia P. Martin, "Comparing Replacement Rates Under Private and Federal Retirement Systems," *Social Security Bulletin*, Vol. 65, No. 1, 2003/2004. For household asset allocations see Table O.3 and footnote 19 in this chapter.

²⁴ For replacement rates, see, for example, Eiji Tajika, "The Public Pension System in Japan: The Consequences of Rapid Expansion," The International Bank for Reconstruction and Development/The World Bank, 2002. For household asset allocation, see Table O.13 and footnote 19.

²⁵ See footnote 20.

²⁶ See footnote 21.

payments to retirees according to some formula, which depends on several factors, e.g., the number of years worked, the level of contributions to the plan, and salary history. Consequently, the sponsor bears the investment risk of the fund. In a defined contribution plan, payments to retirees depend on the accumulated principal plus investment performance of contributions. In these plans, therefore, the beneficiaries of the fund bear its investment risk. (Note that third pillar pensions are, by nature, defined contribution plans.) Lastly, in a hybrid plan, payments typically depend on investment returns, as in defined contribution plans, but the sponsor bears some of the investment risk. Examples include guarantees of paid-in principal (Germany) and minimum guaranteed returns (Switzerland). Historically, almost all pensions were defined benefit plans. For some time now, however, the global trend has been a marked shift to defined contribution and hybrid plans. The effect of this shift, of course, has been to shift the investment risk of pension benefits from employers to employees.

Defined benefit plans can be funded or unfunded. If funded, the plan sponsor uses contributions to buy assets, the income and sales proceeds of which are used to meet pension obligations as they become due. Of course, depending on investment returns, this portfolio of assets may or may not be sufficient to meet promised obligations.

In the case of unfunded defined benefit plans, sometimes called *pay-as-you-go* or *PAYG* plans, the sponsor uses current contributions to meet current obligations, with surpluses or deficits accumulating based on the difference between the two. Assets are bought or sold to manage these accumulated surpluses or deficits.

Most first pillar or public pension plans are *PAYG*. These can be further divided into *Bismarckian* systems, with contributions and benefits linked to pre-retirement earnings (Austria, Belgium, France, Germany, Italy, and Spain) and systems characterized by relatively low contributions and benefits that are designed to prevent poverty in old age (Ireland, the Netherlands, and UK). In any case, following the years of post-war, baby-boom-generation contributions to public pension systems, many plans have accumulated assets. However, any such accumulation of assets by no means implies a fiscally healthy position; most national plans are *underfunded* in the sense that current surpluses plus projected contributions are not nearly sufficient to meet future, promised obligations.

Returning for a moment to the case of a funded, defined benefit plan, the asset management challenge is to predict future obligations and to invest current assets so as to meet those obligations with high probability. Predicting future obligations requires tools outside the area of finance, such as mortality analysis, while the investment of assets is a risk-return problem in the general field of asset-liability management. Not surprisingly, long-dated fixed income assets are particularly suitable to meet projected, long-term obligations. Furthermore, the substantial demand from pension funds in

TABLE 0.7 Asset Allocations of Pension Fund
Assets in Selected Countries as of December 2010

	Equities	Bonds	Other
Germany	40%	45%	15%
UK	60%	31%	9%
Italy	20%	75%	5%
Spain	20%	65%	15%
Netherlands	28%	48%	24%
Switzerland	27%	36%	37%
Denmark	42%	51%	7%

Source: Watson Wyatt, 2008.

Europe for such long-dated assets is critical in the determination of the prices of these assets. (See, for example, the trading case study in Chapter 2.) Having established this, it is also the case that pension funds invest in *real assets* (i.e., assets expected to generate a return that is relatively independent of inflation, like equities and real estate) since pension benefits are often explicitly or implicitly linked to inflation. From this perspective, inflation-linked bonds would be a natural choice for pension investments, but the supply of such bonds is very limited relative to the size of pension portfolios. Table O.7 shows the allocations across equities and bonds for the largest pension systems in Europe and will be referenced further in the discussion of individual country pension systems.

Pension Funds in the UK The largest pension system in Europe is in the UK, with about \$1.75 trillion of assets as of year-end 2009,²⁷ which is about 80% of GDP. About 60% of these assets are part of corporate defined benefit plans, with benefits, subject to some caps, linked to inflation by statute. In fact, the legal requirement that pension benefits increase with inflation in the UK goes far in explaining why inflation-linked fixed income security markets are most developed in the UK.

Historically, UK pension funds have been invested primarily in equities, the legacy of which can be seen in Table O.7. But a combination of poor investment results and relatively high benefit levels have left many corporate defined benefit plans in the UK underfunded.²⁸ Consequently, many companies have closed plans to new employees, curtailed benefits (sometimes even within closed plans), and started to offer new, mostly younger employees, far less generous defined contribution plans.

²⁷ Source: Watson Wyatt, "Global Pension Assets Study," 2008.

²⁸ According to Aon Consulting, the 200 largest companies in the UK faced a combined deficit of close to \$150 billion as of March 2010.

The accumulating deficits of private defined benefit plans over the past decades have also resulted in increased regulation of these plans, like the Minimum Funding Requirement of the late 1990s and the Pensions Act of 2004, which aim to ensure that pension liabilities be valued appropriately, that is, by discounting actuarial projections of benefits at some prevailing market interest rates. In addition, pension plans are required to pay annual, risk-based levies to the Pension Protection Fund, established in 2004, which takes control of plans should they become insolvent or should the sponsoring company declare bankruptcy. The coverage of this fund is still relatively small because, when established, it did not cover existing, underfunded obligations.

The relatively new regulatory and accounting standards for pension funds have caused plan sponsors to hedge more of the interest rate risk as calculated by these standards. This entails increasing the average maturity of bonds in the asset portfolio or *liability hedging*, which, in this case, means receiving fixed in long-dated interest rate swaps. (Part Two describes interest rate risk and hedging.) As a result of this change in portfolio strategy, the asset composition of pension funds in the UK has gradually shifted from equities to fixed income products, including inflation-linked bonds. Furthermore, given the magnitude of pension assets in the UK, this shift has been and continues to be an important explanatory factor for the pricing of fixed income securities in this market.

Pension Funds in the Netherlands The Netherlands is the second largest pension market in Europe, with about \$1 trillion in pension assets as of year-end 2009, about 125% of GDP. Corporate-run pensions are normally mandatory, cover about 90% of the workforce, and are mostly defined-benefit plans. In general, Dutch pension-fund assets are invested more conservatively than in the UK, consistent with the data in Table O.7. As a result, the Dutch pension system stayed relatively well-funded through the 2007–2009 financial crisis, with a *solvency ratio* (i.e., assets divided by the present value of liabilities) falling from 150% before fall 2008 to 107% in 2010.

Dutch pension assets are marked-to-market. Liabilities were historically discounted at a static rate of 4%, but are now discounted at relatively conservative swap rates.²⁹ As a result of this switch to discounting liabilities at a market rate, Dutch pension sponsors, like those in the UK, are now quite active in managing their interest rate risk. In particular, Dutch pension funds tend to be significant receivers in the market for long-dated swaps so as to narrow the gap between the relatively long average duration of their liabilities and the relatively short average duration of their assets (about seven years).

²⁹ International accounting standards allow for the use of a AA corporate curve. See Chapter 19 for a description of bond ratings.

Dutch pension benefits are linked to inflation, but indexation can be suspended when plans are not fully funded. This exemption, combined with a limited supply of products that hedge Dutch-specific inflation, has resulted in inflation hedging being less common for Dutch pension funds than for those in the UK.

Switzerland Switzerland is the third largest pension market in Europe, with about \$550 billion in pension assets as of year-end 2009, about 112% of GDP. Its mandatory, private pension plans have had a minimum guaranteed return of 2.75% since 2008 and are held externally to corporations in a trust-like structure called a *Stiftung*. Of total pension assets, about 58% are held in defined contribution plans, which, in a hybrid component, are usually linked to inflation. A special characteristic of Swiss pension funds, consistent with the data in Table O.7, is a relatively large asset allocation to real estate investment.

Germany³⁰ Germany has the fourth largest pension market in Europe, with more than €400 billion in pension assets as of year-end 2009, about 17% of GDP. Traditionally, companies made direct pension promises to employees, with the resulting liabilities on balance sheet as *book reserves*. With the global trend to segregate and protect pension assets, pensions in Germany have taken on several additional forms along two distinct paths. Along the first path, companies have segregated assets held against pension liabilities through *Contractual Trust Agreements (CTAs)*. In this way, while the pension plan is still fully run by the company, the liabilities move off balance sheet by international accounting standards. These direct, company-run plans in book reserve and CTA forms total €234 billion of pension assets.³¹ Along the second path of segregating pension assets, companies outsource the management of plans to legally independent entities. The most widely-used of these structures, the *Pensionskasse*, accounting for about €96 billion of assets, is technically an insurance company and is subject to funding requirements and guaranteed minimum rates of return. Next, accounting for about €47 billion of assets, is the direct purchase of insurance on behalf of employees, which is followed by *Unterstützungskasse* or support funds, which account for about €37 billion and have wide discretion with respect to investment decisions. Lastly, relatively new structures, called *Pensionsfonds*, offer substantial funding and investment flexibility but account presently for only about €2 billion. The *Pensionskasse* and direct insurance purchases, covered already by existing insurance regulation, do

³⁰ Sources: Allianz Global Investors; “The German Pension System,” Mayer-Brown, 2009; Aegion Global Pensions, “Pension Provision in Germany,” 2010.

³¹ The book reserve entries are technically liabilities, but, as part of corporate balance sheets, they can be thought of as having corresponding assets.

TABLE 0.8 Asset Mix of European Insurers, 2010

Asset class	Portfolio weight
Equities	7%
Government bonds	28%
Corporate bonds	26%
Structured credit	5%
Real estate	4%
Loans	11%
Covered bonds	10%
Other	9%

Source: Deutsche Bank.

not require contributions to the national Pension Insurance Association as do the other structures.

Life Insurance Companies Life insurance products are among the most important savings vehicles in Europe, with assets of over €4 trillion in the four largest markets, which, in descending size, are the UK, France, Germany, and Italy.³² While originally designed to provide a payment upon death, life insurance products normally feature tax-advantaged accumulation of savings, a minimum guaranteed return, and the redemption of proceeds when desired as a lump-sum payment or an annuity. Insurance companies have to invest premiums collected in a manner conducive to meeting both the death benefits and the minimum rates of return promised by the policies.

As with pension funds, insurance companies are being directed by regulators to fund and manage the gap between the market values of their assets and liabilities. A new regulatory regime, Solvency II, effective in 2012, uses a risk-based approach to determine required reserves. And the industry response has been similar to that of the pension funds: more focus on interest rate risk and asset-liability management and a shift of investment allocations from equities to fixed income securities. The portfolio composition of industry assets in Europe as of 2010 is presented in Table O.8.³³ Note that the allocation to equities is very low relative to that of pension funds shown in Table O.7.

As in the case of pension funds, insurance companies can either buy long-term assets to match the long-term nature of their liabilities, or they

³² Source: Allianz, 2008.

³³ Covered bonds are typically AAA, mortgage-related, asset-backed securities with recourse to the issuer.

can choose assets based on various criteria and then receive fixed in swaps to achieve the appropriate risk profile. Not surprisingly then, life-insurance companies are important buyers in the long-dated bond market as well as important receivers in long-term swaps. Furthermore, in balancing risk against return hurdles, insurance companies, consistent with the data in Table O.8, tend to buy significant amounts of corporate bonds. These intermediate- to long-term fixed income securities are a good match against typical liability profiles and, by offering extra return as compensation for credit risk, are more useful than government bonds for meeting return objectives. The amount of credit risk insurance companies can prudently take is limited, however, whether constrained by themselves or by their regulators, so they tend to purchase investment-grade corporate bonds, i.e., those rated BBB or better. In fact, recent data indicate the following allocations by rating class: AAA 3%; AA 20%; A 47%; and BBB 30%.³⁴

Banks In Europe there are a large number of independent banks and a small number of very large banks. Measured by assets, the banking sector is of significant size relative to the economies of the various countries. Columns (2) to (4) of Table O.9 report the GDP of selected countries, the assets of each country's banking sector, and those assets expressed as a percentage of GDP. Assets as a percentage of GDP mostly range from about 225% to 425%, with Ireland an outlier at over 1,000%, as compared with a ratio of about 100% in the United States.³⁵ The size of banking in Europe relative to GDP certainly reflects the structure of its financial system, but is at least in part the result of expansion and leverage leading up to the financial crisis of 2007–2009.

Stable sources of funding enable a banking system to supply credit reliably to the household and corporate sectors and to ensure its own financial soundness. Since deposits are one of the more stable sources of funds, largely because of governmental deposit insurance, one indicator of banking system robustness is the percentage of liabilities in the form of deposits. These ratios, reported in column (5) of Table O.9, which range from 12%, in the case of Ireland, to around 40% for several other countries, are low compared with a ratio of about 60% in the United States.³⁶ While a relatively low reliance on deposits for funding has historically meant a relatively high reliance on other financial institutions for funding, the European Central Bank (ECB)

³⁴ Source: BNP Paribas, August 2010.

³⁵ In this discussion of banks, caution is called for when comparing the balance sheet numbers in Table O.9, which cover Monetary Financial Institutions in Europe as defined by the ECB, with numbers used here from the Flow of Funds Accounts, which cover the commercial banking sector in the United States.

³⁶ Source: Flow of Funds Accounts of the United States for 2009 or Table O.5 for 2010.

TABLE 0.9 Selected Statistics for the Banking Sector in Europe, as of Year-End 2009. The final column gives doubtful and nonperforming loans, net of provisions, as a percent of Tier 1 capital

(1)	(2) GDP	(3) Bank Assets	(4)	(5) Deposits	(6) Housing Loans	(7) Tier 1 Capital Ratio	(8) Net At-Risk Loans / Tier 1 Capital
	€Trillions	€Trillions	% GDP	% Liabilities	% Assets	%Risk-Wtd Assets	% Capital
Germany	2.397	7.438	310	38	13	10.7	31.8
France	1.907	8.148	427	14	9	10.1	30.4
Italy	1.521	3.965	261	25	9	8.3	40.9
Spain	1.054	3.470	329	48	19	9.3	9.2
Netherlands	.572	2.269	397	23	15	12.5	11.1
Belgium	.339	1.167	344	20	7	13.2	36.8
Austria	.274	1.014	370	29	8	9.3	3.2
Greece	.233	.536	230	42	15	10.8	23.7
Finland	.171	.460	268	21	16		
Portugal	.168	.557	332	38	20	7.8	6.0
Ireland	.160	1.648	1,032	12	7		

Sources: GDP from Eurostat; bank balance sheet and capital ratio data from the ECB.

has recently taken on a much expanded role as the supplier of this funding. This will be discussed further in the context of the ECB.

Since the connection between the housing market and bank balance sheets is particularly important in light of the 2007–2009 financial crisis, column (6) of the table gives the percentage of bank assets in the form of housing loans. These data make sense in the context of European banking's participation in the global real estate boom, e.g., German Landesbank's financing of U.S. real estate or Spanish banks' financing of local real estate. While not insignificant, these percentages are low compared with an equivalent percentage of over 36% for U.S. banks, which includes both mortgage loans and various mortgage-related securities.

Columns (7) and (8) of Table O.9 report on bank capital adequacy. Column (7) reports the Tier 1 Capital Ratio, which is essentially common and preferred stock divided by risk-weighted assets. The regulatory, "Basel II" minimum for this ratio is currently 4% but the Bank for International Settlements calls for an increase to 6% by 2015. The banking systems in all of the countries listed in Table O.9 are above these minimum thresholds. Many market participants, however, believe that 10% is a more appropriate threshold, in which case the results in the table are less impressive. The last column reports doubtful and nonperforming loans,³⁷ net of provisions already taken for potential losses, as a function of Tier 1 capital; the proportions are not inconsequential.

To conclude the discussion of Table O.9, the soundness of European banking systems is an important issue for today's markets. Some concerns can be gleaned with respect to banks in particular countries from balance sheet data, e.g., large assets relative to GDP, a relatively small deposit base, just adequate capital ratios, and somewhat high levels of doubtful and nonperforming loans. Meaningful conclusions, however, cannot be drawn from balance sheet data alone. Some additional insights with respect to the robustness of the banking system are discussed next, in the context of the ECB, but a more complete analysis is beyond the scope of this overview.

The final point to be made here about the European banking system is that it plays a crucial role in financing the governments of Europe. Table O.12, to be presented shortly, gives the amounts outstanding of government bonds and other borrowings across Europe. Aggregating these across countries, the totals are €4.563 trillion of government bonds and €2.444 trillion of other borrowings. At the same time, bank balance sheet data³⁸ show that European banks, in aggregate, hold €1.478 trillion of government securities and €1.068 trillion of other loans. Expressed as proportions, then, European banks hold 32% of European government

³⁷ Doubtful loans are those that a bank suspects may not perform. Nonperforming loans are those which are not currently making contractually required payments.

³⁸ Source: ECB June 2010.

TABLE O.10 Balance Sheet of the ECB as of March 1, 2002, in Billions of Euros

Assets	570	Liabilities	570
Autonomous factors, gold, and foreign exchange reserves	387	Autonomous factors, bank-notes in circulation, and other	378
Outstanding operations		Government deposits	57
Main refinancing operations	123	Bank current account holdings	135
Long-term refinancing operations	60	Deposit facility	0

TABLE O.11 Balance Sheet of the ECB as of September 3, 2009, in Billions of Euros

Assets	1,135	Liabilities	1,135
Autonomous factors, gold, and foreign exchange reserves	428	Autonomous factors, bank-notes in circulation, and other	650
Outstanding operations		Government deposits	145
Main refinancing operations	72	Bank current account holdings	200
Long-term refinancing operations	635	Deposit facility	140

securities and 44% of other loans to these governments. These are remarkable proportions, perhaps put in context by noting that the commercial banking sector in the United States, as of the end of 2009, held about 2.4% of U.S. Treasury securities.³⁹

The European Central Bank The ECB conducts monetary policy by adjusting the supply of bank reserves. The ECB makes these adjustments in its *main refinancing operations* with repo agreements that mature from overnight to one week and in its *long-term financing operations* with repo agreements that mature in three or six months. The ECB also runs a deposit facility in which it accepts deposits from banks (in excess of their reserve requirements) at a relatively low rate of interest. But despite the continuity of these generic operating procedures, the policy and role of the ECB has changed dramatically since the financial crisis.

Tables O.10 and O.11 present the balance sheet of the ECB as of March 2002 and as of September 2009, respectively. As of the earlier date, most of the refinancing operations were in the shorter-term category and the order of magnitude of refinancing operations matched the order of magnitude of bank reserves. In other words, the refinancing operations were being used mostly to manage the size of these reserves. Furthermore, the deposit facility was

³⁹ Source: Flow of Funds Accounts of the United States.

not being used by banks. As of September 2009, however, the situation had changed dramatically, with the balance sheet of the ECB having increased from €570 billion to €1.135 trillion.

After the onset of financial crisis, the ECB shifted the purpose of refinancing operations from managing reserves to financing securities for banks. Evidence of this shift appears in the refinancing operations of the ECB's September 2009 balance sheet. First, the magnitude of the operations is significantly larger than the reserves of the banks. Second, the vast majority of the operations are long term, reflecting the fact that the banks are using the ECB to finance their assets on a semi-permanent basis.

Another significant change in the behavior and balance sheet of the ECB after the crisis is the use of the deposit facility. Banks are lending to the ECB through this facility at relatively low rates of interest either because they do not see attractive opportunities in the economy or because their own financial conditions prudently require larger cash balances. In any case, these balances are being accepted by the ECB and lent to other banks through the refinancing operations. In other words, the ECB has taken over some of the financial intermediation previously done within the banking sector.

Since the balance sheet snapshot of Table O.11, the European sovereign debt crisis came to a boil in 2010 with the relatively abrupt deterioration of the perceived creditworthiness of the public debt of Greece, Ireland, Portugal, and Spain. While each of these cases is different from the others, common themes include sharp increases in fiscal deficits during the 2007–2009 crisis, elevated levels of public debt, and skepticism about the post-crisis valuation of assets on bank balance sheets, particularly those related to real estate. In response, the ECB has supported sovereign bond markets with outright purchases through its Security Market Program. By the end of 2010, these purchases totalled about €47 billion for Greece, €15 billion for Ireland, and €12 billion for Portugal. These purchases are, of course, in addition to support through the repo operations described previously. According to Deutsche Bank, the ECB's combined outright and repo holdings of bonds issued by Greece, Ireland, and Portugal was nearly €400 billion at the end of 2010, close to 65% of the total public debt outstanding across these three countries.

Selected Securities

Government Debt Table O.12 reports the amount of government bonds outstanding by maturity, the amount of other government borrowings, and the grand total as a percentage of GDP. As the table shows, government bonds in Europe are relatively well spread out across maturities, which is comforting from the perspective of refinancing risk. On the other hand, overall debt levels for some countries, like Greece, Ireland, Portugal, and Spain, have approached problematic levels relative to their economies.

TABLE 0.12 Outstanding Marketable Government Debt and Other Borrowings as of December 2010

	Maturity in Years				Total Bonds	Other Borrowings	Total/ GDP
	<2	2–5	5–10	>10			
	%	%	%	%	€Billions	€Billions	%
Germany	28	27	27	17	997	1,082	83
France	20	27	28	25	1,125	454	82
Italy	23	25	24	27	1,379	465	119
Spain	28	26	25	20	514	125	60
Netherlands	24	29	28	19	253	118	63
Belgium	25	29	28	18	283	58	97
Austria	7	30	39	24	157	48	72
Greece	18	30	28	24	266	63	143
Finland	11	33	45	11	54	33	48
Portugal	25	28	29	18	120	40	93
Ireland	11	20	60	9	90	58	96

Sources: Bloomberg, Eurostat, Government Financial Statistics and National Accounts.

Money Markets Banks borrow and lend funds in the uncollateralized Euro Overnight Index Average (EONIA) market, which is the European equivalent of the U.S. Fed Funds market. The volume-weighted average rate of transactions in the EONIA market over a day is called the EONIA rate. The majority of transactions are for overnight funds (about 70% in 2007) and the vast majority are for less than one month (96% in 2007). Lending on a secured basis occurs in the repo markets, where the vast majority of terms are less than one month.

The European issuance of short-term, unsecured commercial paper (CP) has grown steadily over time, to about €389 billion as of September 2010. (This is about half the size of the equivalent market in the United States.) Financial institutions and governments are the biggest borrowers in the market, but, at about half the size, non-financial corporations issue here as well. And interestingly, less than half of the total amount of European issued CP is denominated in EUR. Other money markets in Europe include: certificates of deposit (CDs) and other short-term notes, which total \$1.85 trillion, of which 27% is denominated in EUR and 60% in USD; and medium-term notes (MTNs), which total \$8.6 trillion, of which 56% are denominated in EUR and 18.5% in USD.⁴⁰ Taken together, these data show that there is strong demand in Europe for short-term USD financing.

⁴⁰ Source: Euroclear, September 2010.

Money market funds, investment vehicles that invest primarily in the short-term debt of investment grade borrowers, are important for distributing the issuance of short-term, unsecured paper.

FIXED INCOME MARKETS IN JAPAN

Savings-Investment Dynamics and the Macroeconomic Environment

The savings and investment dynamics in Japan have shaped a fixed income market in which individuals save through a single instrument, namely, government debt. This subsection describes the macroeconomic backdrop of these dynamics: corporate deleveraging, an aging population, and expanding government expenditure. Subsequent subsections describe the market participants and the most actively traded markets.

Figure O.8 shows that Japanese households have historically been characterized by high savings rates, hovering around 10% of GDP since the early 1970s. Over the last several years, however, this rate has declined to below 5%. While comparatively high savings rates were associated with generational behavior since World War II, the latest declines are likely attributable to the aging of the population and of the most saving-intensive cohorts, namely, those 45 to 65 years old who are now approaching retirement. Over the past 30 years, these persistent domestic savings surpluses have been sufficient to finance the investment needs of the corporate sector and the expenditures of the government.

Figure O.8 also describes the savings and investment behavior of the corporate sector. For most of the post-war period, this sector was the major




FIGURE O.8 Savings and Investment in Japan, by Sector

user of household savings. In fact, from the mid- to late-1980s, the strong expansion resulting from these high levels of corporate investment increased government revenues to the extent of temporarily turning the public sector into a net supplier of savings, which episode can be clearly seen in Figure O.8. But the loose monetary conditions that accompanied this booming economy—the official discount rate of the Bank of Japan (BoJ) reached 0.25% in 1987—encouraged extremely high levels of corporate indebtedness and leverage. As a consequence, the burst of the economic bubble at the end of the 1980s inaugurated an era of dismal stock returns, low real growth, and falling nominal prices. In reaction, as can be seen clearly from the graph, the corporate sector focused on debt reduction or *balance sheet deleveraging*⁴¹ at the expense of investment to such an extent as to become a net saver.

With the corporate sector's deleveraging, the government took over as the main user of household savings through transfer payments and increased public works (and deteriorating public finances). As shown in Figure O.8, this situation has persisted with the added complication that the gross supply of savings by households is decreasing.

Lastly, the foreign sector's behavior in Figure O.8 implies consistent current account surpluses that have resulted in Japan's accumulation of a sizable, international, net creditor position.

To round out the description of the economic background, GDP growth in Japan dipped into negative territory during the recessions of 1998 and 2001–2002, but then measured between 0 and 3% until the financial crisis that resulted in a –8.4% year-on-year GDP growth in the first quarter of 2009. The main source of information to track the performance of the Japanese economy, by the way, is the *Tankan* survey, published quarterly by the BoJ.

Households and Institutions

Households Household savings and investment behavior in Japan is dominated by demographics; an aging population, facing retirement and medical costs, results in preferences for precautionary savings in cash, for liquid assets, and for capital preservation. To put some perspective on these demographics, the share of the population above 65 years of age reached 20% in 2005, is estimated to be 22.5% in 2010, and is expected to be 35.7% by 2050.⁴² This compares, for example, with 2010 levels of 12.8% in the United States and 20.4% in Germany. From another perspective, more than half of Japanese households are headed by someone

⁴¹ The corporate debt to equity ratio decreased consistently from slightly above 4 in 1990 to about 1.8 at the end 2009. Source: Ministry of Finance, corporate statistics.

⁴² Source: National Institute of Population and Social Security Research.

TABLE 0.13 Financial Assets of Households in Japan
as of March 2010

Assets (¥1,453 Trillion)	% of Financial Assets
Currency and deposits	54.9%
Insurance and pension reserves	27.0%
Stocks	7.1%
Investment trusts	3.8%
Bonds	2.9%
Other	4.3%

Source: Bank of Japan.

over the age of 55. The preferences of these households are far from overwhelmed by younger households, whose investable amounts are limited by the weight of such major expenditures as home purchases and children's education. It is an open question, however, whether younger households, in response to the uncertain prospects of public pensions, will begin to accept higher levels of risk and seek high-yielding alternatives to traditional savings deposits.

As of March 2010, households held financial assets of ¥1,453 trillion and liabilities of ¥369 trillion.⁴³ Table O.13 breaks down these financial assets by instrument. Most of the corresponding liabilities, about 85% of the total, are in the form of loans.

Household risk aversion, in addition to institutional and regulatory features of Japanese financial markets, has led to a high proportion of household wealth in cash and deposits, reported as almost 55% in Table O.13. High risk aversion has been also been evident in the choice between demand deposits, which can be withdrawn without penalty at any time, and time deposits, which have maturities ranging up to three years. In 2000, time deposits constituted about 67% of total bank deposits, but fell to about 20% from 2006. But due to changes in government deposit insurance, this trend reversed itself and time deposits grew back to about 50% of total bank deposits. It should be noted, however, that the most popular savings vehicle has been offered not by any bank, but rather by the Japanese Postal System, which, using its huge network of offices, has acted for years as the largest retail financial institution in the world. Its most popular product is the postal savings certificate, a time deposit that can be withdrawn without penalty after six months.

Moving beyond the 55% held in cash and deposits, another 27% of household financial wealth is in the form of insurance policies and

⁴³ Source: Bank of Japan.

TABLE 0.14 Aggregate Balance Sheet of Depository Institutions as of March 2010

Banks (¥1,522 trillion)			
Assets		Liabilities	
Loans	43.1%	Deposits	75.9%
Bonds	31.0%	Bonds	3.4%
Currency and Deposits	10.3%	Borrowings	10.3%
Stock	3.1%	Equity	5.2%
Other	12.1%	Other	5.2%

Source: Bank of Japan.

pension reserves. Hence, household investments in financial markets are highly intermediated.⁴⁴

Life insurance products in Japan are primarily savings vehicles with death-contingent payoffs. The most popular products are single-premium *endowment* policies in which a lump-sum contribution buys a tax-advantaged investment to be realized as a death benefit or a withdrawal at a guaranteed rate of return.

Banks Traditionally, commercial banks in Japan have been divided into three categories: city, regional, and trust. City banks have historically been the most active in financial markets. City and regional banks are banks in the sense of making commercial loans, while trust banks act mostly as investor agents for pension funds. The largest banks are known as the six “Megs,” with combined holdings approximately equal in size to Japan’s GDP.⁴⁵

Table O.14 reports the composition of the ¥1,522 trillion balance sheet of all depository institutions. The dominance of deposits in the liability structure, relative to the United States and Europe, has been noted previously. With respect to loans, as of May 2010, the composition of bank loans was as follows: 28% in housing and consumer loans; 24% in loans to large companies; and 42% in loans to small and medium-sized companies.

Japanese commercial banks essentially raise funds by taking deposits, making loans to the corporate sector, and, with most of the funds that

⁴⁴ This has not always been the case. Direct stock ownership was high when the *zaibatsu*, the large industrial conglomerates from the prewar era, were broken up after the war and stocks sold to employees. Since then, however, direct ownership has declined as the investment trusts, i.e., stock funds, developed as an investment vehicle in the early 1950s.

⁴⁵ These are Mitsubishi UFJ, Mizuhuo FG, Sumitomo Mistsui FG, Resona holdings, Sumitomo Trust, and Mitsui Trust holdings.

TABLE O.15 Assets of Investment Trusts, Pension Funds, and Life Insurance Companies as of March 2010

	Investment Trusts	Pension Funds	Life Insurance Cos.
Assets	¥90 trillion	¥123 trillion	¥372 trillion
Assets as a % of GDP	19%	26%	78%
	%	%	%
Foreign securities	48.9	19.7	12.2
Domestic bonds	19.4	33.3	56.4
Domestic stock	17.1	11.4	7.8
Loans	4.6	2.2	15.2
Currency and deposits	0.4	6.6	1.8
Other	9.6	26.8	6.6

Source: Bank of Japan.

remain, investing in government bonds. It follows that bank holdings of government bonds vary inversely with loan activity: these holdings fell in 2004–2007 as loan volume picked up, but increased during the financial crisis and the large economic contraction of 2008–2009, reaching 16% of financial assets. (The Bonds asset percentage of 31% in Table O.14 includes assets other than Japanese Government Bonds (JGBs))

Following the burst of the financial bubble in the early 1990s, banks went through a long period of balance sheet restructuring while amortizing bad loans. The proportion of nonperforming loans in the major banks' portfolios, which peaked at 8% in late 2001, fell steadily to slightly less than 2% in 2009, although for regional banks the improvement was not as pronounced. While banks initially received capital injections from the government to deal with the nonperforming-loan problem, they did not receive further assistance after 2003.⁴⁶ And, in another sign of recovery, although the net interest margin on loans has been low, it is positive and somewhat higher than in the 1980s, when funding costs were much higher.

Private, Asset-Management Institutions: Investment Trusts, Pension Funds, and Life Insurance Companies Table O.15 presents the composition of assets of private asset-management institutions. The domestic bond holdings are mostly government bonds with maturities greater than 10 years and, in particular, with maturities around 20 years. The liabilities of these institutions appears for the most part on the asset side of household balance sheets as insurance and pension reserves and as investment trusts.

⁴⁶ Source: Financial Services Authority (FSA); Deposit Insurance Corp.

Investment trusts are affiliates of Japanese securities firms (the largest being Nomura Asset Management) or mutual fund managers that moved into the Japanese market (e.g., Fidelity Investments). The public can invest in these trusts through banks, insurance companies, and securities firms.

As percentages of GDP, Japanese investment trust and pension fund assets are quite low relative to those in the United States, while Japanese insurance assets are quite high. The relevant percentages for Japan are in Table O.15 while the corresponding percentages for the United States are 84%,⁴⁷ 40%,⁴⁸ and 44%,⁴⁹ respectively. But the trend in Japan is for households to allocate less of their savings to insurance companies. This is partially due to an aging population, but also in good part due to insurance companies' failing and not paying eligible claims: between 1997 and 2000, six insurers collapsed. In the past decade the sector underwent a major reorganization, including the entrance of foreign insurers. Nevertheless, the industry still faces the structural problem that investment returns have not kept up with rates that have been guaranteed to policy holders: since the 1990s the average yield on interest-bearing assets has been low, e.g., around 2% in 2009, while the average rate on outstanding policies has been around 3%.

Pension plans in Japan have faced the same issues as those described in Europe: poor performance of financial markets relative to defined benefits and pressure to account properly for the value of liabilities *versus* assets. In any case, the net effect has been that the size of the corporate pension market has not grown significantly since the beginning of the decade.

Semi-Public Asset-Management Institutions: The Postal Savings System and the Public Pensions System

The Postal Savings System The public or semi-public postal savings system has historically attracted between 20% and 25% of household financial wealth, two-thirds in savings accounts (*yū-cho*) and one-third in insurance

⁴⁷ The Flow of Funds accounts as of March 2010 report the following amounts: money market mutual funds, \$1.8 billion; mutual funds (credit instruments), \$2.8 billion; mutual fund shares (equities), \$7.3 billion; closed-end funds, \$140 billion. Dividing the total of about \$12 billion by a 2009 U.S. GDP of \$14.3 billion gives a proportion of 84%.

⁴⁸ The Flow of Funds accounts report private pension fund assets of \$5.7 trillion as of March 2010. Dividing by a 2009 GDP of \$14.3 trillion gives a proportion of about 40%.

⁴⁹ The Flow of Fund accounts as of March 2010 report life insurance company assets of \$4.9 trillion and property and casualty insurance company assets of \$1.4 trillion. Dividing the \$6.3 trillion sum by a 2009 U.S. GDP of \$14.3 trillion gives a proportion of about 44%.

products (*kampo*). Japanese households are attracted by the convenience of obtaining the products at post offices and by the perceived safety of the offerings, originated, as they are, by a quasi-public entity. In addition, postal service life insurance products are particularly accessible (e.g., they require no medical exam) and, like other such policies in Japan, also function as tax-advantaged vehicles for savings.

Until 2003, the system was managed by the Postal Services Agency, which deposited all its assets with the Trust Fund Bureau, an arm of the Ministry of Finance, which, in turn, would buy government debt and extend loans to local government entities. In 2003 the agency was reorganized into Japan Post as preparation for a long-term plan to privatize the system. This had implications for fixed income markets as Japan Post changed the management of postal savings assets by investing them directly in financial markets. But to complete the institutional story, in October 2007, Japan Post was split into four entities, Japan Post Bank, Japan Post Insurance, and two nonfinancial companies, with a view to spin them off as independent entities by 2017. Subsequent policy reversals have delayed these privatization plans, however, and a final privatization is far from certain.

As of September 2009, the combined assets of the postal system were ¥303.6 trillion. And with its ¥177 trillion in deposits, it is the largest savings institution in the world. (The next largest in Japan would be the country's largest publicly-traded bank, Mitsubishi UFJ, with ¥119 trillion in deposits.) Finally, the postal system invests a substantial portion of its funds in Japanese government bonds, holding ¥158 trillion in its banking unit and ¥68 trillion in its insurance unit, which together constitute nearly one third of the amount outstanding.

The Public Pensions System Japan's public pension program has two pillars. The first is the National Pension System (NPS), a program for all Japanese nationals in which they contribute premiums while they work and ultimately receive a fixed pension benefit that is independent of individual contribution levels. The NPS is a PAYG system and, since contributions are no longer sufficient to cover benefits, the Treasury contributes a significant amount to fund the program.

The second pillar is a public pension system for employees of private companies. Complementing the NPS, this system pays benefits based on each participant's income and contributions. The sum of contributions by active workers and the government has, in fact, exceeded cumulative benefits paid, thus building a reserve that amounted to ¥121 trillion in June 2009.⁵⁰ The funds in this reserve are managed by the Government Pension Investment Fund, both directly and, more commonly, through private financial

⁵⁰ Source: Government Pension Investment Fund.

TABLE O.16 Government Revenues and Expenditures as a % of GDP. Nominal GDP in Trillions of Yen.

	2005	2006	2007	2008	2009
Revenue	31.1	33.8	32.2	33.6	31.0
Tax (direct)	15.0	15.6	16.1	14.3	13.0
Tax (indirect)	2.6	2.6	2.5	2.5	2.5
NPS contributions	10.6	10.8	11.1	11.2	11.3
Other	2.9	4.9	2.8	5.6	4.3
Expenditure	37.3	34.8	35.2	38.6	40.6
Debt interest	2.4	2.5	2.5	2.6	2.7
NPS benefits	17.2	17.1	17.5	18.2	19.1
Public works	4.0	3.6	3.4	3.4	4.4
Other	13.7	11.6	11.9	14.4	14.4
Surplus/Deficit	-6.1	-1.0	-3.0	-5.0	-9.6
Nominal GDP	503.2	510.9	515.8	497.7	483.0

Source: Ministry of Finance and Morgan Stanley.

institutions. Its guideline is to invest about 70% of its funds in domestic government bonds. For public employees, the role of this second pillar is assumed by *mutual-aid* pension programs, which are funded individually by the government.

The Government The overall indebtedness of the Japanese government, relative to GDP, is higher than that of any other G20 country. Total gross government debt, securitized and not securitized, including the debt of regional governments, the National Pension System (NPS), and semi-public administrative corporations, reached ¥968 trillion at the end of 2008, 188% of GDP. Public finances have deteriorated consistently since the beginning of the 1990s, explained by a shrinking tax base and increasing expenditure levels. Table O.16 details government revenues and expenditures over the past several years and the resulting annual deficits. Note that over half of government expenditures are used to service the debt and pay NPS benefits.

As of March 2009, the outstanding amount of JGBs was ¥798 trillion. The next subsection will detail the ownership distribution of JGBs, but note here that almost half of all JGBs are owned by public or semi-public institutions. This public debt in public hands should be included in the total debt outstanding as it supports third-party claims (e.g., private deposits in the postal savings system or premia paid into Kampo life insurance policies).⁵¹

⁵¹ By contrast, the BoJ's holdings of JGBs, resulting from its open market purchases, should not be included in total debt outstanding.

The future of Japan's public finances and the sustainability of its promised benefits seem bleak given the problems described in this overview: a large amount of government debt relative to GDP; endemic budget deficits; consistently lackluster economic performance; demographic trends toward an aging society; and a fall in the household savings rate. Providing some room for maneuver, however, is the fact that the tax and social security burden in Japan is lower than that of developed European economies like the UK, Germany, and France. Also, Japan runs a positive current account balance and has accumulated a net position of ¥258 trillion of foreign assets, which is more than half of GDP. In addition to providing resources to improve its fiscal position, this holding of foreign assets hedges Japan against an underperforming domestic economy: about half of the income from foreign assets is interest from bonds purchased abroad, while somewhat more than half of payments to foreigners is a return on direct business investment by foreigners.

The Bank of Japan and Monetary Policy Before 2001 the BoJ conducted monetary policy by choosing a target for the Tokyo Overnight Average Rate (TONAR) consistent with its macroeconomic objectives, where TONAR is the daily average rate on uncollateralized borrowing and lending of reserves held at the BoJ, analogous to the Fed Funds rate in the United States and the EONIA rate in Europe. After having chosen a target level for TONAR, the BoJ pushed the call rate to that target mostly by trading short-term funds to supply or drain cash from bank reserve accounts. The bulk of these operations were conducted with repo (*Gensaki*) with an average maturity of about seven days, as well as with indirect loans against collateral.

From 2001 to 2006, in an attempt to stimulate the economy by injecting large amounts of liquidity into financial markets, the BoJ embarked on a program of *quantitative easing*. The strategy was to expand significantly the reserves that banks have at the BoJ in the hope that these reserves would support loan volumes to businesses that would, in turn, stimulate economic activity. The BoJ used two tactics to implement this strategy. First, the BoJ began using 25-day term repos to inject liquidity, thus expanding the set of participants willing to borrow money from the BoJ. Second, the BoJ began to conduct monthly auctions (*Rinban*) through which it bought JGBs of maturities greater than two years. This allowed the bank to put large amounts of cash into the hands of the public in exchange for bonds over a much longer term than achievable through repo markets. The quantitative easing experiment, however, was not particularly successful. First, banks chose to keep the extra liquidity rather than aggressively expand their extensions of credit. Second, the corporate sector persisted in deleveraging from the very high indebtedness levels of the early 1990s.

After 2006 the BoJ suspended its use of 25-day repos and returned to calibrating TONAR through the more traditional use of short-term repos

and indirect, collateralized loans. However, since the 2007–2009 financial crisis, the BoJ, like the ECB, has been using longer-term repo, with the BoJ conducting a significant portion of its operations with 3- and 6-month repos.

The *Rinban* has continued beyond the quantitative easing period, with auctions now held four times per month across all JGB maturities.

The Markets

The products used for most investments in Japan are relatively simple: deposits, short-term money market products, and JGBs. Swap markets are developed as well, but activity there derives from the risk management needs of investors in JGBs. With respect to money markets, overnight, uncollateralized funds among banks trade in a *call market*, and overnight, collateralized funds trade in the larger repo (*Gensaki*) market in which trust banks, and to a lesser extent regional banks, have been the traditional suppliers of funds. Money markets in Japan also include, in order of importance, government bills, bank certificates of deposit, and commercial paper.

Japanese Government Bonds JGBs are center stage of fixed income markets in Japan. Table O.17 shows the distribution of ownership of JGBs. The amount of JGBs owned by public or semi-public institutions has already been noted. But it is remarkable that about 94% of JGBs are held by domestic investors, a proportion much higher than that of any other major government debt market and, for that matter, of Japanese equity markets

TABLE O.17 Distribution of Ownership of JGBs in 2009

	Share
Public sector	55.6%
Japan Post Bank	23.7%
Post Insurance (Kampo)	10.0%
Public Pensions	11.6%
Bank of Japan	7.7%
Mutual Aid Pensions	2.6%
Other	0.3%
Private sector	44.4%
Banks	12.4%
Life Insurance	6.5%
Foreign	6.0%
Households	5.2%
Other	14.3%

Source: Barclays Bank.

TABLE 0.18 JGBs Maturity Structure
as of July 2010

Maturity Years	Outstanding ¥ Trillion	Percent %
< 2	190	27
2–5	183	26
5–10	191	28
> 10	130	19

Source: Bloomberg.

in which foreign investors have very relevant stakes.⁵² Strong domestic demand for JGBs is the result of many factors discussed earlier, including the life-cycle of Japanese households, the institutional structure of savings and investment, and financial regulation.

Low foreign ownership of JGBs does not correspond to low trading participation: foreigner trading has accounted for between 15% and 25% of JGB trading volume over the last decade and for between 25% and 40% of JGB futures trading volume. Domestic traders of JGBs include many institutions described in the previous subsection, with the most important being public sector entities, regional and city banks, life insurance companies, and pension funds. Putting the participation of these institutions in perspective, city and trust banks account for about four times the trading volume as that of regional banks and life insurance companies.⁵³

JGBs are currently issued in six categories: short-term (6-month and 1-year bills); medium-term (2-year and 5-year bonds, with tickers JN and JS, respectively); long-term (10-year bonds, ticker JB); super-long term (20-year bonds, JL, 30-year bonds, JX, 40-year bonds, JU, and the recently discontinued 15-year floating rate note, JF); JGBs for individual investors (5-year and 10-year); and inflation-indexed bonds (10-year, JBI).⁵⁴ The average maturity of all outstanding JGB issues is about 5.8 years and Table O.18 shows that outstanding volumes are relatively evenly spread across maturities. The resulting percentage of volume in the short end is high relative to Europe (see Table O.12) but low relative to the United States (see Table O.2).

JGB 15-year floaters pay a coupon equal to the greater of zero and the difference between 1) the average yield of the 10-year bond at the most

⁵² In equity markets, overseas investors absorbed approximately all net selling of stocks by domestic banks and insurers from 2000 to 2010. Source: Tokyo Stock Exchange.

⁵³ Source: Japan Securities Dealers Association (JSDA).

⁵⁴ Source: Ministry of Finance.

recent, prior auction; and 2) a fixed margin set at the auction of the floater. Floating rate notes were issued with the intention of enabling large, institutional investors in JGBs to hedge against rising yields. Currently, floaters account for about 6% of total JGBs outstanding.

The government started issuing inflation-linked bonds (JGBi) with a maturity of 10 years in March 2004 and stopped issuing them, buying them back as well, in August 2008. Notional outstanding reached a maximum of about ¥10 trillion, but, as of 2010, is at a relatively small ¥4.8 trillion. The structure of the JGBi is like that of TIPS, discussed in the section on U.S. markets, with the important difference that JGBi principal is not floored at the initial amount. As a result, the nominal return to a JGBi investor can be negative in a deflationary environment.

Asset-Backed Securities While there are markets for asset-backed securities in Japan, they are of much less importance than comparable markets in western economies. The market for residential, mortgage-backed securities has grown the most, but is still relatively small and most mortgages remain on bank balance sheets. A good part of the explanation for this lies in the residential mortgage market itself. First, residential mortgages constitute only about 30% of GDP, which is substantially less than a comparable proportion of about 80% in the United States. Second, the residential mortgage market in Japan has not grown for the last two decades due to the lackluster performance of the real estate market. But despite this backdrop, securitization of residential real estate has grown very rapidly since 2005 and is approaching a ¥200 trillion cumulative total. This growth is in good part due to the insistence of the government that the Government Home Loan Corporation arrange securitizations rather than make loans directly.

The Relative Pricing of Securities with Fixed Cash Flows

Consumers and businesses are willing to pay more than \$1 in the future in exchange for \$1 today. A newly independent adult borrows money to buy a car today, agreeing to repay the loan price plus interest over time; a family takes a mortgage to purchase a new home today, assuming the obligation to make principal and interest payments for years; and a business, which believes it can transform \$1 of investment into \$1.10 or \$1.20, chooses to take on debt and pay the prevailing market rate of interest. At the same time, this willingness of potential borrowers to pay interest attracts lenders and investors to make consumer loans, mortgage loans, and business loans. This fundamental fact of financial markets, that receiving \$1 today is better than receiving \$1 in the future, or, equivalently, that borrowers pay lenders for the use of their funds, is known as the *time value of money*.

Borrowers and lenders meet in fixed income markets to trade funds across time. They do so in very many forms: from one-month U.S. Treasury bills that are almost certain to return principal and interest to the long-term debt of companies that have already filed for bankruptcy; from assets with a simple dependence on rates, like Eurodollar futures, to callable bonds and swaps; from assets whose value depends only on rates, like interest rate swaps, to mortgage-backed securities or inflation-protected securities; and from fully taxable private-sector debt to partially or fully tax-exempt issues of governments and municipalities.

To cope with the challenge of pricing the vast number of existing and potential fixed income securities, market professionals often focus on a limited set of benchmark securities, for which prices are most consistently and reliably available, and then price all similar assets relative to those benchmarks. Sometimes, as when pricing a UK government bond in terms of other UK government bonds, or when pricing an EUR swap in terms of other EUR swaps, relative prices can be determined rigorously and for the most part accurately by *arbitrage pricing*. This methodology is developed in Chapter 1, where it is also shown that discounting, i.e., calculating present values with discount factors, is really just shorthand for arbitrage pricing.

While discount factors in many ways solve the relative pricing problem, they are not very intuitive for understanding the time value of money that is embedded in market prices. For this purpose, markets rely on spot, forward, and par rates. Chapter 2 introduces these rates and derives the relationships linking them to each other and to discount factors. The trading case study in Chapter 2, inspired by an abnormally shaped EUR forward swap curve, illustrates how fixed income analytics, market technicals (due to institutional factors described in the Overview), and a macroeconomic setting all contribute to a trade idea.

While the interest rates of Chapter 2 provide excellent intuition with respect to the time value of money embedded in market prices, other quantities provide intuition with respect to the returns offered by individual securities. The first half of Chapter 3 defines returns, spreads, and yields. Spreads describe the pricing of particular securities relative to benchmark government bond or swap curves and yields are the widely used, although sometimes misunderstood, internal rates of return on individual securities. The second half of Chapter 3 breaks down a security's return into several component parts. First, how does the security perform if rates and spreads stay the same? Second, how does the security perform if rates change? Third, how does the security perform if spreads change?

Given the central role of benchmarks in Part One, it is worth describing which securities are used as benchmarks and why. Until relatively recently, benchmark curves in U.S. and Japanese markets were derived from the historically most liquid markets, that is, from government bond markets. Recently, however, the benchmark has shifted significantly to swap curves. European markets, on the other hand, have for some time relied predominantly on interest rate swap markets for benchmarks because their swap markets have been, on average across the maturity spectrum, more liquid than government bond markets.

It is not hard to understand why government bond and interest rate swap markets are the preferred choices for use as benchmarks. First, they are the most liquid markets, consistently providing prices at which market participants can execute trades in reasonable size. Second, they incorporate information about interest rates that is common to all fixed income markets.

The value of a corporate bond, for example, depends on the interest rate information embedded in the government bond or swap curve in addition to depending on the credit characteristics of the individual corporate issuer.

But what about the choice between government bond and swap curves as benchmarks? Historically, government bonds were the only choice because swaps did not exist until the early 1980s and it took some time for their liquidity to become adequate. But bond markets have a significant disadvantage when used as a benchmark, namely that an individual bond issue is not a commodity in the sense of being a fungible collection of cash flows: bond issues are in fixed supply and have idiosyncratic characteristics. The best-known examples of nonfungibility are *on-the-run* U.S. Treasury bonds that trade at a premium relative to other government bonds because of their superior liquidity and financing characteristics. Put another way, pricing with a curve that is constructed from “similar” bonds, which are not *on-the-run* bonds, will underestimate the prevailing prices of *on-the-runs*. By contrast, an interest rate swap is really a commodity, that is, a fungible collection of cash flows. A 10-year, 4% interest rate swap cannot possibly be in short supply because any willing buyer and seller can create a new contract with exactly those terms. In fact, market practice bears out this distinction between bonds and swaps. While bond traders set prices for each and every bond they trade (although they certainly may use heuristics relating various prices to each other or to related futures markets), swap traders strike a curve that is then used to price their entire book of swaps automatically.

In short, global fixed income markets currently use interest rate swaps as benchmarks or base curves and build other curves from spreads or spread curves on top of swap curves. Even in the liquid U.S. Treasury market, strategists assess relative value using spreads of individual Treasury issues against the USD swap curve.

Prices, Discount Factors, and Arbitrage

This chapter begins by introducing the cash flows of fixed-rate, government coupon bonds. It shows that prices of these bonds can be used to extract *discount factors*, which are the market prices of one unit of currency to be received on various dates in the future.

Relying on a principle known as *the law of one price*, discount factors extracted from a particular set of bonds can be used to price other bonds, outside the original set. A more complex but more convincing relative pricing methodology, known as *arbitrage pricing*, turns out to be mathematically identical to pricing with discount factors. Hence, discounting can rightly be used and regarded as shorthand for arbitrage pricing.

The application of this chapter uses the U.S. Treasury coupon bond and Separate Trading of Registered Interest and Principal of Securities (STRIPS) markets to illustrate that bonds are not commodities, meaning that their prices reflect individual characteristics other than their scheduled cash flows. This idiosyncratic component of bond valuation implies that the predictions of the simplest relative pricing methodologies only approximate the complex reality of bond markets.

The chapter concludes with a discussion of day-counts and accrued interest, pricing conventions used throughout fixed income markets and, consequently, throughout this book.

THE CASH FLOWS FROM FIXED-RATE GOVERNMENT COUPON BONDS

The cash flows from fixed-rate, government coupon bonds are defined by *face amount*, *principal amount*, or *par value*; *coupon rate*; and *maturity date*. For example, in May 2010 the U.S. Treasury sold a bond with a coupon rate of $2\frac{1}{8}\%$ and a maturity date of May 31, 2015. Purchasing \$1 million face amount of these “ $2\frac{1}{8}\%$ s of May 31, 2015,” entitles the buyer to the schedule of

TABLE 1.1 Cash Flows of the U.S. $2\frac{1}{8}$ s of May 31, 2015

Date	Coupon Payment	Principal Payment
11/30/2010	\$10,625	
5/31/2011	\$10,625	
11/30/2011	\$10,625	
5/31/2012	\$10,625	
11/30/2012	\$10,625	
5/31/2013	\$10,625	
11/30/2013	\$10,625	
5/31/2014	\$10,625	
11/30/2014	\$10,625	
5/31/2015	\$10,625	\$1,000,000

payments in Table 1.1. The Treasury promises to make a coupon payment every six months equal to half the note's annual coupon rate of $2\frac{1}{8}\%$ times the face amount, i.e., $\frac{1}{2} \times 2\frac{1}{8}\% \times \$1,000,000$, or \$10,625. Then, on the maturity date of May 31, 2015, in addition to the coupon payment on that date, the Treasury promises to pay the bond's face amount of \$1,000,000. One fact worth mentioning, although too small a detail to receive much attention in this book, is that scheduled payments that do not fall on a business day are made on the following business day. For example, the payments of the $2\frac{1}{8}$ s scheduled for Sunday, May 31, 2015, would be made on Monday, June 1, 2015.

For concreteness and continuity of exposition this chapter restricts attention to U.S. Treasury bonds. But the analytics of the chapter apply easily to bonds issued by other countries because the cash flows of all fixed rate government coupon bonds are qualitatively similar. The most significant difference across issues is the frequency of coupon payments, which can be semiannual or annual; government bond issues in France and Germany make annual coupon payments, while those in Italy, Japan, and the UK make semiannual payments.

Returning to the U.S. Treasury market, then, Table 1.2 reports the coupons and maturity dates of selected U.S. Treasury bonds, along with their prices as of the close of business on Friday, May 28, 2010. Almost all U.S. Treasury trades settle $T + 1$, which means that the exchange of bonds for cash happens one business day after the trade date. In this case, the next business day was Tuesday, June 1, 2010.

The prices given in Table 1.2 are *mid-market*, *full* (or *invoice*) prices per 100 face amount. A mid-market price is an average of a lower *bid* price, at which traders stand ready to buy a bond, and a higher *ask* price, at which

TABLE 1.2 Selected U.S. Treasury Bond
Prices as of May 28, 2010

Coupon	Maturity	Price
$1\frac{1}{4}\%$	11/30/2010	100.550
$4\frac{7}{8}\%$	5/31/2011	104.513
$4\frac{1}{2}\%$	11/30/2011	105.856
$4\frac{3}{4}\%$	5/31/2012	107.966
$3\frac{3}{8}\%$	11/30/2012	105.869
$3\frac{1}{2}\%$	5/31/2013	106.760
2%	11/30/2013	101.552
$2\frac{1}{4}\%$	5/31/2014	101.936
$2\frac{1}{8}\%$	11/30/2014	100.834

traders stand ready to sell a bond. A *full* price is the total amount a buyer pays for a bond, which is the sum of the *flat* or *quoted* price of the bond and *accrued interest*. This division of full price will be explained later in this chapter. In any case, to take an example from Table 1.2, purchasing \$100,000 face amount of the $3\frac{1}{2}$ s of May 31, 2013, costs a total of \$106,760.

The bonds in Table 1.2 were selected from the broader list of U.S. Treasuries because they all mature and make payments on the same *cycle*, in this case at the end of May and November each year. This means, for example, that all of the bonds make a payment on November 30, 2010, and, therefore, that all their prices incorporate information about the value of a dollar to be received on that date. Similarly, all of the bonds apart from the $1\frac{1}{4}$ s of November 30, 2010, which will have already matured, make a payment on May 31, 2011, and their prices incorporate information about the value of a dollar to be received on that date, etc. The next section describes how to extract information about the value of a dollar to be received on each of the payment dates in the May–November cycle from the prices in Table 1.2.

DISCOUNT FACTORS

The *discount factor* for a particular term gives the value today, or the *present value* of one unit of currency to be received at the end of that term. Denote the discount factor for t years by $d(t)$. Then, for example, if $d(.5)$ equals .99925, the present value of \$1 to be received in six months is 99.925 cents. Another security, which pays \$1,050,000 in six months, would have a present value of $.99925 \times \$1,050,000$ or \$1,049,213.

Since Treasury bonds promise future cash flows, discount factors can be extracted from Treasury bond prices. In fact, each of the rows of Table 1.2 can be used to write one equation that relates prices to discount factors. The equation from the 1¼s of November 30, 2010, is

$$100.550 = \left(100 + \frac{1\frac{1}{4}}{2}\right) d(.5) \tag{1.1}$$

In words, equation (1.1) says that the price of the bond equals the present value of its future cash flows, namely its principal plus coupon payment, all times the discount factor for funds to be received in six months. Solving reveals that $d(.5)$ equals .99925.

By the same reasoning, the equations relating prices to discount factors can be written for the other bonds listed in Table 1.2. The next two of these equations are

$$104.513 = \frac{4\frac{7}{8}}{2} \times d(.5) + \left(100 + \frac{4\frac{7}{8}}{2}\right) d(1) \tag{1.2}$$

$$105.856 = \frac{4\frac{1}{2}}{2} \times d(.5) + \frac{4\frac{1}{2}}{2} \times d(1) + \left(100 + \frac{4\frac{1}{2}}{2}\right) d(1.5) \tag{1.3}$$

Given the solution for $d(.5)$ from equation (1.1), equation (1.2) can be solved for $d(1)$. Then, given the solutions for $d(.5)$ and $d(1)$, equation (1.3) can be solved for $d(1.5)$. Continuing in this fashion through the rows of Table 1.2 generates the discount factors, in six-month intervals, out to four and one-half years, which are reported in Table 1.3. Note how these

TABLE 1.3 Discount Factors from
U.S. Treasury Note and Bond Prices
as of May 28, 2010

Term	Discount Factor
11/30/2010	.99925
5/31/2011	.99648
11/30/2011	.99135
5/31/2012	.98532
11/30/2012	.97520
5/31/2013	.96414
11/30/2013	.94693
5/31/2014	.93172
11/30/2014	.91584

discount factors, falling with term, reflect the time value of money: the longer a payment of \$1 is delayed, the less it is worth today.

THE LAW OF ONE PRICE

Another U.S. Treasury bond issue, one not included in the set of base bonds in Table 1.2, is the $\frac{3}{4}$ s of November 30, 2011. How should this bond be priced? A natural answer is to apply the discount factors of Table 1.3 to this bond’s cash flows. After all, the base bonds are all U.S. Treasury bonds and the value to investors of receiving \$1 from a Treasury on some future date should not depend very much on which particular bond paid that \$1. This reasoning is an application of the *law of one price*: absent confounding factors (e.g., liquidity, financing, taxes, credit risk), identical sets of cash flows should sell for the same price.

According to the law of one price, the price of the $\frac{3}{4}$ s of November 30, 2011 should be

$$.375 \times .99925 + .375 \times .99648 + 100.375 \times .99135 = 100.255 \qquad (1.4)$$

where each cash flow is multiplied by the discount factor from Table 1.3 that corresponds to that cash flow’s payment date. As it turns out, the market price of this bond is 100.190, close to, but not equal to, the prediction of 100.255 in equation (1.4).

Table 1.4 compares the market prices of three bonds as of May 28, 2010, to their present values (PVs), i.e., to their prices as predicted by the law of one price. The differences range from −2.8 cents to +6.5 cents per 100 face value, indicating that the law of one price describes the pricing of these bonds relatively well but not perfectly.

According to the last row of Table 1.4, the $\frac{7}{8}$ s of May 31, 2011, trade 2.8 cents *rich* to the base bonds, i.e., its market price is high relative to the discount factors in Table 1.3. In the same sense, the $\frac{3}{4}$ s of November 30, 2011, and the $\frac{3}{4}$ s of May 31, 2012, trade *cheap*. In fact, were these price discrepancies sufficiently large relative to transaction costs, an arbitrageur might consider trying to profit by selling the rich $\frac{7}{8}$ s and

TABLE 1.4 Testing the Law of One Price for Three U.S. Treasury Notes as of May 28, 2010

Bond	$\frac{7}{8}$ s 5/31/11	$\frac{3}{4}$ s 11/30/11	$\frac{3}{4}$ s 5/31/12
PV	100.521	100.255	100.022
Price	100.549	100.190	99.963
PV−Price	−.028	.065	.059

simultaneously buying some combination of the base bonds; by buying either of the cheap bonds and simultaneously selling base bonds; or by selling the rich $\frac{7}{8}$ s and buying both of the cheap bonds in the table. Trades of this type, arising from deviations from the law of one price, are the subject of the next section.

ARBITRAGE AND THE LAW OF ONE PRICE

While the law of one price is intuitively reasonable, its justification rests on a stronger foundation. It turns out that a deviation from the law of one price implies the existence of an *arbitrage opportunity*, that is, a trade that generates profits without any chance of losing money.¹ But since arbitrageurs would rush *en masse* to do any such trade, market prices would quickly adjust to rule out any such opportunity. Hence, arbitrage activity can be expected to do away with significant deviations from the law of one price. And it is for this reason that the law of one price usually describes security prices quite well.

To make this argument more concrete, the discussion turns to an arbitrage trade based on the results of Table 1.4, which showed that the $\frac{3}{4}$ s of November 30, 2011, are cheap relative to the discount factors in Table 1.3 or, equivalently, to the bonds listed in Table 1.2. The trade is to purchase the $\frac{3}{4}$ s of November 30, 2011, and simultaneously sell or *short*² a portfolio of bonds from Table 1.2 that replicates the cash flows of the $\frac{3}{4}$ s. Table 1.5 describes this *replicating portfolio* and the arbitrage trade.

Columns (2) to (4) of Table 1.5 correspond to the three bonds chosen from Table 1.2 to construct the replicating portfolio: the $1\frac{1}{4}$ s of November 30, 2010; the $4\frac{7}{8}$ s of May 31, 2011; and the $4\frac{1}{2}$ s of November 30, 2011. Row (iii) gives the face amount of each bond in the replicating portfolio, so that this portfolio is long 98.166 face amount of the $4\frac{1}{2}$ s, short 1.790 of the $4\frac{7}{8}$ s, and short 1.779 of the $1\frac{1}{4}$ s. Rows (iv) through (vi) show the cash flows from those face amounts of each bond. For example, 98.166 face amount of the $4\frac{1}{2}$ s, which pay a coupon of 2.25% on May 31, 2011, generates a cash flow of $98.166 \times 2.25\%$ or 2.209 on that date. Similarly, -1.779 of the $1\frac{1}{4}$ s, which pay coupon and principal totalling $100 + \frac{1.25}{2}$ or

¹Market participants often use the term arbitrage more broadly to encompass trades that could conceivably lose money, but promise large profits relative to the risks borne.

²To short a security means to sell a security one does not own. The mechanics of short selling bonds will be discussed in Chapter 12. For now, assume that a trader shorting a bond receives the price of the bond and is obliged to pay all its coupon and principal cash flows.

TABLE 1.5 The Replicating Portfolio of the $\frac{3}{4}$ s of November 30, 2011, with Prices as of May 28, 2010

	(1)	(2)	(3)	(4)	(5)	(6)
(i) Coupon		$1\frac{1}{4}$ s	$4\frac{7}{8}$ s	$4\frac{1}{2}$ s		$\frac{3}{4}$ s
(ii) Maturity		11/30/10	5/31/11	11/30/11	Portfolio	11/30/11
(iii) Face Amount		-1.779	-1.790	98.166		100
	Date	Cash Flows				
(iv)	11/30/10	-1.790	-.044	2.209	.375	.375
(v)	5/31/11		-1.834	2.209	.375	.375
(vi)	11/30/11			100.375	100.375	100.375
(vii) Price		100.550	104.513	105.856		100.190
(viii) Cost		-1.789	-1.871	103.915	100.255	100.190
(ix) Net Proceeds		.065				

100.625 per 100 face value on November 30, 2010, produces a cash flow of $-1.779 \times 100.625\%$ or -1.790 on that date. Row (vii) gives the price of each bond per 100 face amount, simply copied from Table 1.2. Row (viii) gives the initial cost of purchasing the indicated face amount of each bond. So, for example, the “cost” of “purchasing” -1.790 face amount of the $4\frac{7}{8}$ s is $-1.790 \times 104.513\%$ or -1.871 . Said more naturally, the proceeds from selling 1.790 face amount of the $4\frac{7}{8}$ s are 1.871.

Column (5) of Table 1.5 sums columns (2) through (4) to obtain the cash flows and cost of the replicating portfolio. Rows (iv) through (vi) of column (5) confirm that the cash flows of the replicating portfolio do indeed match the cash flows of 100 face amount of the $\frac{3}{4}$ s of November 30, 2011, given in the same rows of column (6). Note that most of the work of replicating the $\frac{3}{4}$ s of November 30, 2011, is accomplished by the $4\frac{1}{2}$ s maturing on the same date. The other two bonds in the replicating portfolio are used for minor adjustments to the cash flows in six months and one year. Appendix A in this chapter shows how to derive the face amounts of the bonds in this or any such replicating portfolio.

With the construction of the replicating portfolio completed, the discussion returns to the arbitrage trade. According to row (viii) of Table 1.5, an arbitrageur can buy 100 face amount of the $\frac{3}{4}$ s of November 30, 2011, for 100.190, sell the replicating portfolio for 100.255, pocket the difference or “net proceeds” of 6.5 cents, shown in row (ix), and not owe anything on any future date. And while a 6.5-cent profit may seem small, the trade can be scaled up: for \$500 million face of the $\frac{3}{4}$ s, which would not be an abnormally large position, the riskless profit increases to $\$500,000,000 \times .065\%$ or \$325,000.

As stated at the start of this section, if a riskless and profitable trade like the one just described were really available, arbitrageurs would rush to do the trade and, in so doing, force prices to relative levels that admit no arbitrage opportunities. More specifically, arbitrageurs would drive the prices of the $\frac{3}{4}$ s and of the replicating portfolio together until the two were equal.

The crucial link between arbitrage and the law of one price can now be explained. The total cost of the replicating portfolio, 100.255, given in column (5), row (viii) of Table 1.5, exactly equals the present value of the $\frac{3}{4}$ s of November 30, 2011, computed in Table 1.4. In other words, the law of one price methodology of pricing the $\frac{3}{4}$ s (i.e., discounting with factors derived from the $1\frac{1}{4}$ s, $4\frac{7}{8}$ s, and $4\frac{1}{2}$ s) comes up with exactly the same value as does the arbitrage pricing methodology (i.e., calculating the value of portfolio of the $1\frac{1}{4}$ s, $4\frac{7}{8}$ s, and $4\frac{1}{2}$ s that replicates the cash flows of the $\frac{3}{4}$ s). This is not a coincidence. In fact, Appendix B in this chapter proves that these two pricing methodologies are mathematically identical. Hence, applying the law of one price, i.e., pricing with discount factors, is identical to relying on the activity of arbitrageurs to eliminate relative mispricings, i.e., pricing by arbitrage. Expressed another way, discounting can be justifiably regarded as shorthand for the more complex and persuasive arbitrage pricing methodology.

Despite this discussion, of course, the market price of the $\frac{3}{4}$ s was quoted at a level somewhat below the level predicted by the law of one price. This can be attributed to one or a combination of the following reasons. First, there are transaction costs in doing arbitrage trades which could significantly lower or wipe out any arbitrage profit. In particular, the prices in Table 1.2 are mid-market whereas, in reality, an arbitrageur would have to buy securities at higher ask prices and sell at lower bid prices. Second, bid-ask spreads in the financing markets (see Chapter 12), incurred when shorting securities, might also overwhelm any arbitrage profit. Third, it is only in theory that U.S. Treasury bonds are commodities, i.e., fungible collections of cash flows. In reality, bonds have idiosyncratic differences that are recognized by the market and priced accordingly. And it is this last point that is the subject of the next section.

APPLICATION: STRIPS AND THE IDIOSYNCRATIC PRICING OF U.S. TREASURY NOTES AND BONDS

STRIPS

In contrast to coupon bonds that make payments every six months, *zero coupon* bonds make no payments until maturity. Zero coupon bonds issued by the U.S. Treasury are called STRIPS. For example, \$1,000,000

TABLE 1.6 STRIPS Face Amounts from
1,000,000 Face Amount of the $3\frac{1}{2}$ s of
May 15, 2020

Date	C-STRIP Face Amount	P-STRIP Face Amount
11/15/10	\$17,500	0
5/15/11	\$17,500	0
11/15/11	\$17,500	0
$\vdots$	$\vdots$	$\vdots$
5/15/19	\$17,500	0
11/15/19	\$17,500	0
5/15/20	\$17,500	\$1,000,000

face amount of STRIPS maturing on May 15, 2020, promises only one payment: \$1,000,000 on that date. STRIPS are created when a particular coupon bond is delivered to the Treasury in exchange for its coupon and principal components. Table 1.6 illustrates the stripping of \$1,000,000 face amount of the $3\frac{1}{2}$ s of May 15, 2020, which was issued in May 2010, to create coupon STRIPS maturing on the 20 coupon payment dates and principal STRIPS maturing on the maturity date. Coupon or interest STRIPS are called TINTs, INTs, or C-STRIPS while principal STRIPS are called TP, Ps, or P-STRIPS. Note that the face amount of C-STRIPS on each date is $1/2 \times 3.5\% \times \$1,000,000$ or \$17,500.

The Treasury not only creates STRIPS but retires them as well. For example, upon delivery of the set of STRIPS in Table 1.6 the Treasury would *reconstitute* the \$1,000,000 face amount of the $3\frac{1}{2}$ s of May 15, 2020. But in this context it is crucial to note that C-STRIPS are fungible while P-STRIPS are not. When reconstituting a bond, any C-STRIPS maturing on a particular date may be applied toward the coupon payment of that bond on that date. By contrast, only P-STRIPS that were stripped from a particular bond may be used to reconstitute the principal payment of that bond.³ This feature of the STRIPS program implies that P-STRIPS, and not C-STRIPS, inherit the cheapness or richness of the bonds from which they came, an implication that will be demonstrated in the following subsection.

STRIPS prices are essentially discount factors. If the price of the C-STRIPS maturing on May 31, 2015, is 89.494 per 100 face amount, then the implied discount factor to that date is .89494. With this in

³Making P-STRIPS fungible would not affect either the total or the timing of cash flows owed by the Treasury, but could change the amounts outstanding of particular securities.




FIGURE 1.1 Discount Factors from C-STRIPS Prices as of May 28, 2010

mind, Figure 1.1 graphs the C-STRIPS prices per unit face amount as of May 28, 2010.

The Idiosyncratic Pricing of U.S. Treasury Notes and Bonds

If U.S. Treasury bonds were commodities, with each regarded solely as a particular collection of cash flows, then the price of each would be well approximated by discounting its cash flows with the C-STRIPS discount factors of Figure 1.1. If however individual bonds have unique characteristics that are reflected in pricing, the law of one price would not be as accurate an approximation. Furthermore, since C-STRIPS are fungible while P-STRIPS are not, any such pricing idiosyncrasies would manifest themselves as differences between the prices of P-STRIPS and C-STRIPS of the same maturity. To this end, Figure 1.2 graphs the differences between the prices of P-STRIPS and C-STRIPS that mature on the same date as of May 28, 2010. So, for example, with the price of P-STRIPS and C-STRIPS, both maturing on May 31, 2015, at 89.865 and 89.494, respectively, Figure 1.2 records the difference for May 31, 2015, as $89.865 - 89.494$ or .371. Note that Figure 1.2 shows two sets of P-STRIPS prices, those P-STRIPS originating from Treasury bonds and those originating from Treasury notes.⁴

Inspection of Figure 1.2 shows that there are indeed significant pricing differences between P-STRIPS and C-STRIPS that mature on the same date. This does not necessarily imply the existence of arbitrage opportunities, as discussed at the end of the previous section. However, the results do suggest

⁴The difference between notes and bonds is of historical interest only; see “Fixed Income Markets in the United States, Securities and Other Assets” in the Overview.




FIGURE 1.2 Differences between the Prices of P-STRIPS and C-STRIPS Maturing on the Same Date per 100 Face Amount as of May 28, 2010

that bonds have idiosyncratic pricing differences and that these differences are inherited by their respective P-STRIPS. Of particular interest, for example, is the largest price difference in the figure, the 2.16 price difference between the P-STRIPS and C-STRIPS maturing on May 15, 2020. These P-STRIPS come from the most recently sold or *on-the-run* 10-year note, an issue which, as will be discussed in Chapter 12, traditionally trades rich to other bonds because of its superior liquidity and financing characteristics. In any case, to determine whether idiosyncratic bond characteristics are indeed inherited by P-STRIPS, Table 1.7 analyzes the pricing of selected U.S. Treasury coupon securities in terms of STRIPS prices. The particular securities selected are those on the mid-month, May-November cycle with 10 or more years to maturity as of May 2010.

Columns (1) to (3) of Table 1.7 give the coupon, maturity, and market price of each bond. Column (4) computes a price for each bond by discounting all of its cash flows using the C-STRIPS prices in Figure 1.1, and column (5) gives the difference between the market price and that computed price. By the simplest application of the law of one price, these computed prices should be a good approximation of market prices. There are, however, some very significant discrepancies. The approximation misses the price of the 3½s of May 15, 2020, the 10-year on-the-run security, by a very large 2.076; the 5s of May 15, 2037, by .924; and the 6¼s of 5/15/30 by .708.

Column (6) of Table 1.7 computes the price of each bond by discounting its coupon payments with C-STRIPS prices and its principal payment with the P-STRIPS of that bond. Column (7) gives the difference between the market price and that computed price. To the extent that P-STRIPS prices inherit pricing idiosyncrasies of their respective bonds, these computed prices should be better approximations to market prices than the prices computed

TABLE 1.7 Market Prices Compared with Pricing Using C-STRIPS and with Pricing Using C-STRIPS for Coupon Payments and the Respective P-STRIPS for Principal Payments

(1) Coupon	(2) Maturity	(3) Market Price	(4) C- Pricing	(5) Error	(6) C- and P- Pricing	(7) Error
$3\frac{1}{2}$	5/15/20	101.896	99.820	2.076	101.982	-.086
$8\frac{3}{4}$	5/15/20	146.076	145.738	.338	146.070	.006
$8\frac{1}{8}$	5/15/21	142.438	142.357	.080	142.407	.031
8	11/15/21	141.916	141.750	.167	141.980	-.063
$7\frac{5}{8}$	11/15/22	139.696	139.545	.151	139.805	-.109
$7\frac{1}{2}$	11/15/24	140.971	140.694	.277	141.059	-.087
$6\frac{1}{2}$	11/15/26	131.582	130.894	.687	131.716	-.134
$6\frac{1}{8}$	11/15/27	127.220	126.643	.578	127.291	-.070
$5\frac{1}{4}$	11/15/28	116.118	115.456	.661	116.175	-.058
$6\frac{1}{4}$	5/15/30	130.523	129.815	.708	130.639	-.116
5	5/15/37	113.840	112.916	.924	113.943	-.102
$4\frac{1}{2}$	5/15/38	105.114	104.625	.490	105.214	-.100
$4\frac{1}{4}$	5/15/39	100.681	100.425	.256	100.764	-.083
$4\frac{3}{8}$	11/15/39	102.780	102.638	.143	102.905	-.124
$4\frac{3}{8}$	5/15/40	102.999	102.308	.691	102.969	.030

using C-STRIPS prices alone. And, in fact, this is the case. Comparing the absolute values of the two error columns reveals that the approximation in column (6) is better than the approximation in column (4) for every bond in the table.

In conclusion, then, individual Treasury bonds have idiosyncratic characteristics that are reflected in market prices. Furthermore, since P-STRIPS are not fungible across bonds, their prices inherit the idiosyncratic pricing of their respective bond issues.

ACCRUED INTEREST

This section describes the useful market practice of separating the full price of a bond, which is the price paid by a buyer to a seller, into two parts: a quoted or flat price, which is the price that appears on trading screens and is used when negotiating transactions; and *accrued interest*. The full and quoted prices are also known as the *dirty* and *clean* prices, respectively.

Definition

To make the concepts concrete, consider an investor who purchases \$10,000 face amount of the U.S. Treasury $3\frac{5}{8}$ s of August 15, 2019, for settlement on June 1, 2010. The bond made a coupon payment of $\frac{1}{2} \times 3\frac{5}{8}\% \times \$10,000$ or \$181.25 on February 15, 2010, and will make its next coupon payment of \$181.25 on August 15, 2010. See the time line in Figure 1.3.

Assuming the purchaser holds the bond through the next coupon date, the purchaser will collect the coupon on that date. But it can be argued that the purchaser is not entitled to the full semiannual coupon payment on August 15 because, as of that time, the purchaser will have held the bond for only two and a half months of a six-month coupon period. More precisely, using what is known as the *actual/actual day-count convention*, which will be explained later in this section, and referring again to Figure 1.3, the purchaser should receive only 75 of 181 days of the coupon payment, that is, $\frac{75}{181} \times \$181.25$, or \$75.10. The seller of the bond, whose cash was invested in the bond from February 15 to June 1, should collect the rest of the coupon, i.e., $\frac{106}{181} \times \$181.25$, or \$106.15. A conceivable institutional arrangement is for the seller and purchaser to divide the coupon on the payment date, but this would undesirably require additional arrangements to ensure that this split of the coupon actually takes place. Consequently, market convention dictates instead that the purchaser pay the \$106.15 of accrued interest to the seller on the settlement date and that the purchaser keep the entire coupon of \$181.25 on the coupon payment date.

On May 28, 2010, for delivery on June 1, 2010, the flat or quoted price of the $3\frac{5}{8}$ s was 102-26, meaning $102 + \frac{26}{32}$ or 102.8125. The full or invoice price of the bond per 100 face amount is defined as the quoted price plus accrued interest, which, in this case, is $102.8125 + 1.0615$ or 103.8740. For this particular trade, of \$10,000 face amount, the invoice price is \$10,387.40.

At this point, by the way, it becomes clear why discussion earlier in the chapter had to make reference to the fact that prices were full prices. When trading bonds that make coupon payments on May 31, 2010, for settlement on June 1, 2010, purchasers have to pay one day of accrued interest to sellers.




FIGURE 1.3 Example of Accrued Interest Time Line

Pricing Implications

The present value of a bond's cash flows should be equated or compared with its full price, that is, with the amount a purchaser actually pays to purchase those cash flows. Conceptually, denoting the flat price by p , accrued interest by AI , the present value of the cash flows by PV , and the full price, as before, by P ,

$$P = p + AI = PV \quad (1.5)$$

Equation (1.5) reveals an important point about accrued interest: the particular market convention used in calculating accrued interest does not really matter. Say, for example, that everyone recognizes that the convention in place is too generous to the seller because, instead of being made to wait for a share of the interest until the next coupon date, the seller receives that share at settlement. In that case, by equation (1.5), the flat price would adjust downward to mitigate this advantage. Put another way, the only quantity that matters is the invoice price, which determines the amount of money that changes hands.

Having made this argument, why is the accrued interest convention useful in practice? The answer is told in Figure 1.4, which draws the full and flat prices of the $3\frac{5}{8}\%$ s of August 15, 2019, from February 15, 2010, to September 15, 2010, under several assumptions, with the most important being that 1) the discount function does not change, i.e., $d(t)$ does not change, where t is the number of days from settlement; and 2) the flat price of the bond for settlement on June 1 is 102.8125. In words, then, Figure 1.4 says that the full price changes dramatically over time even when the market is



FIGURE 1.4 Full and Flat Prices for the $3\frac{5}{8}\%$ s of August 15, 2019, Over Time with a Constant Discount Function

unchanged, including a discontinuous jump on coupon payment dates, while the flat price changes only gradually over time. Therefore, when trading bonds day to day, it is more intuitive to track flat prices and negotiate transactions in those terms.

The shapes of the price functions in Figure 1.4 can be understood as follows. Within a coupon period, the full price of the bond, which is just the present value of its cash flows, increases over time as the bond's payments draw near. But from an instant before the coupon payment date to an instant after, the full price falls by the coupon payment: the coupon is included in the present value of the remaining cash flows at the instant before the payment, but not at the instant after. The time pattern of the flat price, supposing that prevailing interest rates do not change, will be discussed in Chapter 3. Basically, however, the flat price of a bond like the $3\frac{5}{8}$ s, which sells for more than its face value, will trend down to its value at maturity, i.e., par.

Day-Count Conventions

Accrued interest equals the coupon times the fraction of the coupon period from the previous coupon payment date to the settlement date. For the $3\frac{5}{8}$ s, as for most government bonds, this fraction is calculated by dividing the actual number of days since the previous coupon date by the actual number of days in the coupon period. Hence the term “actual/actual” for this day-count convention.

Other day-count conventions, however, are applied in other markets. Two of the most common are *actual/360* and *30/360*. The *actual/360* convention divides the actual number of days between two dates by 360, and is commonly used in money markets, i.e., for short-term, *discount* (i.e., zero coupon) securities, and for the floating legs of interest rate swaps. The *30/360* convention assumes that there are 30 days in a month when calculating the difference between two dates and then divides by 360. Applying this convention, the number of days between June 1 and August 15 is 74 (29 days left in June, 30 days in July, and 15 days in August), as opposed to the 75 days using an actual day count. The *30/360* convention is used most commonly for corporate bonds and for the fixed leg of interest rate swaps.

APPENDIX A: DERIVING REPLICATING PORTFOLIOS

To replicate the $\frac{3}{4}$ s of November 30, 2011, Table 1.5 uses the $1\frac{1}{4}$ s due November 30, 2010, the $4\frac{7}{8}$ s due May 31, 2011, and the $4\frac{1}{2}$ s due November 30, 2011. Number these bonds from 1 to 3 and let F^i be the face amount of bond i in the replicating portfolio. Then, the following equations express

the requirement that the cash flows of the replicating portfolio equal those of the $\frac{3}{4}$ s on each of the three cash flow dates.

For the cash flow on November 30, 2010:

$$\left(100\% + \frac{1\frac{1}{4}\%}{2}\right) F^1 + \left(\frac{4\frac{7}{8}\%}{2}\right) F^2 + \left(\frac{4\frac{1}{2}\%}{2}\right) F^3 = \frac{\frac{3}{4}\%}{2} \quad (1.6)$$

For the cash flow on May 31, 2011:

$$0 \times F^1 + \left(100\% + \frac{4\frac{7}{8}\%}{2}\right) F^2 + \left(\frac{4\frac{1}{2}\%}{2}\right) F^3 = \frac{\frac{3}{4}\%}{2} \quad (1.7)$$

And, for the cash flow on November 30, 2011:

$$0 \times F^1 + 0 \times F^2 + \left(100\% + \frac{4\frac{1}{2}\%}{2}\right) F^3 = 100\% + \frac{\frac{3}{4}\%}{2} \quad (1.8)$$

Solving equations (1.6), (1.7), and (1.8) for F^1 , F^2 , and F^3 gives the replicating portfolio's face amounts in Table 1.5. Note that since one bond matures on each date, these equations can be solved one-at-a-time instead of simultaneously, i.e., solve (1.8) for F^3 , then, using that result, solve (1.7) for F^2 , and then, using both results, solve (1.6) for F^1 . In any case, the results are as follows:

$$F^1 = -1.779\% \quad (1.9)$$

$$F^2 = -1.790\% \quad (1.10)$$

$$F^3 = 98.166\% \quad (1.11)$$

Replicating portfolios are easier to describe and manipulate using matrix algebra. To illustrate, equations (1.6) through (1.8) are expressed in matrix form as follows:

$$\begin{pmatrix} 1 + \frac{1.25\%}{2} & \frac{4.875\%}{2} & \frac{4.5\%}{2} \\ 0 & 1 + \frac{4.875\%}{2} & \frac{4.5\%}{2} \\ 0 & 0 & 1 + \frac{4.5\%}{2} \end{pmatrix} \begin{pmatrix} F^1 \\ F^2 \\ F^3 \end{pmatrix} = \begin{pmatrix} \frac{.75\%}{2} \\ \frac{.75\%}{2} \\ 1 + \frac{.75\%}{2} \end{pmatrix} \quad (1.12)$$

Note that each column of the leftmost matrix describes the cash flows of one of the bonds in the replicating portfolio; the elements of the vector to the right of this matrix are the face amounts of each bond for which equation (1.12) has to be solved; and the rightmost vector contains the cash flows of the bond to be replicated. This equation can easily be solved by pre-multiplying each side by the inverse of the leftmost matrix.

In general then, suppose that the bond to be replicated makes payments on T dates. Let C be the $T \times T$ matrix of cash flows, principal plus interest, with the T columns representing the T bonds in the replicating portfolio and the T rows the dates on which those bonds make payments. Let $\vec{F}$ be the $T \times 1$ vector of face amounts in the replicating portfolio and let $\vec{c}$ be the vector of cash flows, principal plus interest, of the bond to be replicated. Then, the replication equation is

$$C \vec{F} = \vec{c} \quad (1.13)$$

with solution

$$\vec{F} = C^{-1} \vec{c} \quad (1.14)$$

The only complication is in ensuring that the matrix C does have an inverse. Essentially, any set of T bonds will do so long as there is at least one bond in the replicating portfolio making a payment on each of the T dates. In this case, the T bonds would be said to *span* the payment dates. So, for example, T bonds all maturing on the last date would work, but T bonds all maturing on the second-to-last date would not work: in the latter case there would be no bond in the replicating portfolio making a payment on date T .

APPENDIX B: THE EQUIVALENCE OF THE DISCOUNTING AND ARBITRAGE PRICING APPROACHES

Proposition: Pricing a bond according to either of the following methods gives the same price:

- Derive a set of discount factors from some set of spanning bonds and price the bond in question using those discount factors.
- Find the replicating portfolio of the bond in question using that same set of spanning bonds and calculate the price of the bond as the price of this portfolio.

Proof: Continue using the notation introduced at the end of Appendix A. Also, let $\vec{d}$ be the $T \times 1$ vector of discount factors for each date and let $\vec{P}$ be the vector of prices of each bond in the replicating portfolio, which is the same as the vector of prices of each bond used to compute the discount factors. Generalizing the “Discount Factors” section of this chapter, one can solve for discount factors using the following equation:

$$\vec{d} = (C')^{-1} \vec{P} \quad (1.15)$$

where the $'$ denotes the transpose. Then, the price of the bond according to the first method is $c' \vec{d}$. The price according to the second method is $\vec{P}' \vec{F}$ where $\vec{F}$ is as derived in equation (1.14).

Hence, the two methods give the same price if

$$\vec{c}' \vec{d} = \vec{P}' \vec{F} \quad (1.16)$$

Expanding the left-hand side of equation (1.16) with (1.15) and the right-hand side with (1.14),

$$\vec{c}' (C')^{-1} \vec{P} = \vec{P}' C^{-1} \vec{c} \quad (1.17)$$

And since both sides of this equation are just numbers, take the transpose of the left-hand side to show that equation (1.17) is true.

Spot, Forward, and Par Rates

It is clear from Chapter 1 that price and cash flows completely describe any fixed-rate investment. Nevertheless, investors and traders almost always find it more intuitive to express the time value of money in terms of interest rates. This chapter, therefore, introduces the most commonly-used interest rates, which are spot rates, forward rates, and par rates. The relationships linking these rates to discount factors and to each other reveal why interest rates are so intuitively appealing.

Given the importance of interest rate swaps as a benchmark of market interest rates, the illustrative examples and the trading case study of this chapter are taken from global swap markets. The valuation of interest rates swaps, however, is not covered by this book until Chapter 16. Until then, the reader is asked to accept the assertion, made here and justified there, that interest rates embedded in the swap market can be properly extracted by treating the fixed side of a swap as if it were a coupon bond and the floating side as if it were a floating rate bond worth par.¹

The trading case study of this chapter begins by highlighting the abnormally downward-sloping forward rates of the EUR swap curve in the second quarter of 2010. Then, in the context of macroeconomic factors and market technicals, a trade is constructed to take advantage of this abnormally-shaped curve.

SIMPLE INTEREST AND COMPOUNDING

Price and cash flows completely describe an investment: a bond might cost 101.980 today and pay 103 in six months; a 100,000,000 1.5-year loan, six months *forward* (i.e., a loan made in six months for 1.5 years) might pay

¹Chapter 16 explains and justifies this widely-used methodology, but money market conditions during the financial crisis of 2007–2009 motivate the relatively advanced material in Chapter 17, which shows when an alternate swap-pricing methodology is more appropriate.

103,797,070 in two years. But investors and traders often prefer to quote and think in terms of interest rates, saying that the bond just described earns 2% and the forward loan 2.5%. Interest rates are more intuitive than prices because they automatically normalize for the amount invested and, expressed as annual rates, normalize for the investment horizon as well. So even though the bond costs 101.98 and matures in six months while the forward loan invests 100,000,000 for 1.5 years, the interest rates on the two investments can be sensibly and intuitively compared.

The purpose of this section is to describe the conventions through which interest rates are quoted given prices and cash flows. The most straightforward convention is *simple interest*, in which interest paid is the quoted, annualized rate times the term of the investment, in years. While the discussion of day-count conventions in Chapter 1 showed that there are many ways to define the term of an investment in years, in the context of this chapter semiannual periods are defined to have a term of half a year. Continuing then with the bond example of the previous paragraph, the six-month bond earns 2% because

$$101.98 + 101.98 \times \frac{2\%}{2} = 101.98 \times \left(1 + \frac{2\%}{2}\right) = 103 \quad (2.1)$$

In words, a simple interest investment is conceptualized as making a single payment at maturity equal to the initial investment amount plus interest on that initial investment. In equation (2.1), the initial investment is 101.980 and the interest earned is that 101.98 times $\frac{2\%}{2}$, where the latter is one-half the quoted, annual rate of 2%. The sum of these two is the bond's total payment of 103.

The forward loan example introduced at the start of this section has a term of 1.5 years or of three semiannual periods, requiring an outlay of 100 million in six months for a terminal payment of 103,797,070 in two years. Under the convention of *semiannual compounding*, an investment is conceptualized as follows. First, simple interest is earned within each six-month period. Second, each six-month period's total proceeds, that is, both principal and interest, are reinvested for the subsequent six-month period. So, in the case of the forward loan earning a rate of 2.5%, the proceeds from earning simple interest over the first six months are

$$100,000,000 \times \left(1 + \frac{2.5\%}{2}\right) = 101,250,000 \quad (2.2)$$

Then, reinvesting this total amount over the subsequent six months at the same rate produces a total of

$$\begin{aligned} 101,250,000 \times \left(1 + \frac{2.5\%}{2}\right) &= 100,000,000 \times \left(1 + \frac{2.5\%}{2}\right)^2 \\ &= 102,515,625 \end{aligned} \quad (2.3)$$

To appreciate the impact of compounding, note that an investment of 100 million earning simple interest of 2.5% over a year would be worth $100,000,000 \times (1 + 1 \times 2.5\%)$ or 102,500,000. The 15,625 difference between the semiannually compounded proceeds in (2.3) and this simple interest amount is exactly equal to the *interest on interest*, that is, the interest earned in the second six-month period on the interest earned over the first six-month period. More specifically, the interest over the first period, from (2.2), is 1,250,000 and the interest on that amount for six months is $1,250,000 \times \frac{2.5\%}{2}$ or 15,625.

Returning now to the forward loan, over the last of its three semiannual periods, the proceeds grow to

$$\begin{aligned} 102,515,625 \times \left(1 + \frac{2.5\%}{2}\right) &= 100,000,000 \times \left(1 + \frac{2.5\%}{2}\right)^3 \\ &= 103,797,070 \end{aligned} \quad (2.4)$$

which is the terminal payoff set out in the example.

Generalizing this discussion, investing F at a rate of $\hat{r}$ compounded semiannually for T years generates

$$F \times \left(1 + \frac{\hat{r}}{2}\right)^{2T} \quad (2.5)$$

at the end of those T years. (Note that the power in this expression is $2T$ since an investment for T years compounded semiannually is, in fact, an investment for $2T$ half-year periods.)

This discussion has been framed in terms of semiannual compounding because coupon bonds and the fixed side of interest rate swaps most commonly pay interest semiannually. Other compounding conventions, including *continuous compounding* (for which interest is assumed to be paid every instant), are useful in other contexts and are presented in Appendix A in this chapter.

EXTRACTING DISCOUNT FACTORS FROM INTEREST RATE SWAPS

As the examples of this chapter are drawn from global swap markets, this section digresses with a very brief introduction to interest rate swaps. Chapters 16 and 17 present a much fuller discussion of these important derivatives.

Two parties might agree, on May 28, 2010, to enter into an *interest rate swap* with the following terms. Starting in two business days, on




FIGURE 2.1 An Example of an Interest Rate Swap

June 2, 2010, party A agrees to pay a fixed rate of 1.235% on a notional amount of \$100 million to party B for two years, who, in return, agrees to pay *three-month LIBOR* (*London Interbank Offered Rate*) on this same notional to Party A. See Figure 2.1. Chapter 15 explains the mechanics of payments based on floating rate indexes and describes LIBOR rates in detail. For the present, suffice it to say that three-month LIBOR is the rate at which the most creditworthy banks can borrow money from each other for three months and that a *fixing* of this rate is published once each trading day.

The \$100 million in the example is called the notional amount of a swap, rather than the face, par, or principal amount, because it is used only to compute the fixed- and floating-rate payments: the \$100 million itself is never paid or received by either party. In any case, party A, who *pays fixed* and *receives floating*, makes fixed payments of $\frac{1.235\%}{2} \times \$100,000,000$, or \$617,500 every six months. Party B, who *receives fixed* and *pays floating*, makes floating rate payments quarterly.

While swap contracts do not include any payment of the notional amount, it is convenient to assume that, at maturity, party A pays the notional amount to party B and that party B pays that same notional amount to party A. Once again, see Figure 2.1. There are three points to be made about these fictional payments. First, since they cancel each other, their inclusion has no effect on the value of the swap.² Second, adding the fictional notional amount to the fixed side makes that leg of the swap look like a coupon bond, i.e., a security with semiannual, fixed coupon payments and a terminal principal payment. Third, adding the fictional notional amount to the floating side makes that leg look like a floating rate bond, i.e., a security with semiannual, floating coupon payments and a terminal principal payment.

Chapter 16 presents the widely-used valuation methodology in which the floating leg of the swap, with its fictional notional amount, is worth par, or, in the example, \$100 million, on payment dates. Taking this as given for the purposes of this chapter, an interest rate swap can be viewed in a very simple way: party B, the fixed receiver, “buys” a 1.235% semiannually-paying

²See Chapter 16 for a more detailed discussion of this point.

TABLE 2.1 Discount Factors, Spot Rates, and Forward Rates Implied by Par USD Swap Rates as of May 28, 2010

Term in Years	Swap Rate	Discount Factor	Spot Rate	Forward Rate
0.5	.705%	.996489	.705%	.705%
1.0	.875%	.991306	.875%	1.046%
1.5	1.043%	.984494	1.045%	1.384%
2.0	1.235%	.975616	1.238%	1.820%
2.5	1.445%	.964519	1.450%	2.301%

coupon bond (i.e., the fixed leg) for \$100 million (i.e., the value of the floating leg). Party A, the fixed payer, “sells” a 1.235% bond for \$100 million. This interpretation of swaps is so useful and commonplace that the phrase, the “fixed leg of a swap,” is almost always meant to include the fictional notional payment at maturity.

Invoking the interpretation of swaps in the previous paragraph, discount factors can be derived from swaps using the methodology of Chapter 1, developed in the context of coupon bonds. To illustrate this, along with the rate calculations of later sections, Table 2.1 presents some data on shorter-maturity, USD interest rate swaps as of May 28, 2010. The second column gives the rates that are quoted and observed in swap market trading. These indicate that counterparties are willing to exchange fixed payments of .875% against three-month LIBOR for one year, 1.043% against three-month LIBOR for 1.5 years, etc. The 2-year swap rate, depicted in Figure 2.1, is 1.235%. In any case, to derive the third column of Table 2.1, the discount factors implied by swap rates, proceed as in Chapter 1. Write an equation for each “bond” that equates the present value of its cash flows to its price of par, i.e.,

$$\left(100 + \frac{.705}{2}\right) d(.5) = 100 \quad (2.6)$$

$$\frac{.875}{2} d(.5) + \left(100 + \frac{.875}{2}\right) d(1) = 100 \quad (2.7)$$

etc. The set of five such equations, corresponding to the maturities .5 through 2.5, allows for the solution of the discount factors given in the third column of the table.

The derivation of spot and forward rates, the fourth and fifth columns of Table 2.1, along with the relationships across all of these rates, is the subject of the rest of the chapter.

DEFINITIONS OF SPOT, FORWARD, AND PAR RATES

Chapter 1 defined a curve of discount factors, $d(t)$, which gives the present values of one unit of currency to be received at various times t . This section expresses the information in the discount curve as a *term structure of interest rates* and, in particular, in terms of semiannually-compounded *spot*, *forward*, and *par* rates. Definitions of continuously compounded spot and forward rates can be found in Appendix B in this chapter.

Spot Rates

A spot rate is the rate on a *spot loan*, an agreement in which a lender gives money to the borrower at the time of the agreement to be repaid at some single, specified time in the future. Denote the semiannually compounded t -year spot rate by $\hat{r}(t)$. Then, following (2.5), investing 1 unit of currency from now to year t will generate proceeds at that time of

$$\left(1 + \frac{\hat{r}(t)}{2}\right)^{2t} \quad (2.8)$$

To link spot rates and discount factors, note that if \$1 grows to the quantity (2.8) in t years, then the present value of that quantity is \$1. Using discount factors to compute that present value,

$$\left(1 + \frac{\hat{r}(t)}{2}\right)^{2t} d(t) = 1 \quad (2.9)$$

Then, solving for $d(t)$ gives

$$d(t) = \frac{1}{\left(1 + \frac{\hat{r}(t)}{2}\right)^{2t}} \quad (2.10)$$

Table 2.1 gives the discount factors from the USD swap curve as of May 28, 2010. Taking the 2-year discount factor of .975616 from that table, equation (2.10) can be used to derive the 2-year, semiannually-compounded spot rate of 1.238%:

$$d(2) = \frac{1}{\left(1 + \frac{1.238\%}{2}\right)^{2 \times 2}} = .975616 \quad (2.11)$$

From (2.8), this rate implies that, in two years, an investment of \$100 grows to

$$\$100 \times \left(1 + \frac{1.238\%}{2}\right)^{2 \times 2} = \$102.499 \quad (2.12)$$

Forward Rates

A *forward* rate is the rate on a *forward loan*, which is an agreement to lend money at some time in the future to be repaid some time after that. There are many possible forward rates: the rate on a loan given in six months for a subsequent term of 1.5 years; the rate in five years for six months; etc. This subsection, however, focuses exclusively on forward rates over sequential, six-month periods. Let $f(t)$ denote the forward rate on a loan from year $t - .5$ to year t . Then, investing 1 unit of currency from year $t - .5$ for six months generates proceeds, at year t , of

$$\left(1 + \frac{f(t)}{2}\right) \quad (2.13)$$

To link forward rates to spot rates, note that a spot loan for $t - .5$ years combined with a forward loan from year $t - .5$ to year t covers the same investment period as a spot loan to year t . To ensure that rates are quoted consistently, that is, to ensure that the proceeds from these identical investments are the same,

$$\begin{aligned} \left(1 + \frac{\hat{r}(t)}{2}\right)^{2t} &= \left(1 + \frac{\hat{r}(t - .5)}{2}\right)^{2(t - .5)} \left(1 + \frac{f(t)}{2}\right) \\ &= \left(1 + \frac{\hat{r}(t - .5)}{2}\right)^{2t-1} \left(1 + \frac{f(t)}{2}\right) \end{aligned} \quad (2.14)$$

This logic can be extended further, to write the spot rate of term t as a function of all forward rates up to $f(t)$:

$$\left(1 + \frac{\hat{r}(t)}{2}\right)^{2t} = \left(1 + \frac{f(.5)}{2}\right) \left(1 + \frac{f(1)}{2}\right) \cdots \left(1 + \frac{f(t)}{2}\right) \quad (2.15)$$

Finally, to express forward rates in terms of discount factors, simply use equation (2.10) to replace the spot rates in (2.14) with discount factors:

$$\left(1 + \frac{f(t)}{2}\right) = \frac{d(t - .5)}{d(t)} \quad (2.16)$$

Continuing with the USD swap data in Table 2.1, use the 2- and 2.5-year spot rates or discount factors from the table, together with (2.14) or (2.16), to derive that $f(2.5) = 2.301\%$. This value implies that an investment of \$100 in 2 years will, in 2.5 years, be worth

$$\$100 \times \left(1 + \frac{2.301\%}{2}\right) = \$101.151 \quad (2.17)$$

In passing, note that if the term structure of spot interest rates is *flat*, so that all spot rates are the same, i.e., $\widehat{r}(t) = \widehat{r}$ for all t , then, from (2.14), each forward rate must equal that same $\widehat{r}$ and the term structure of forward interest rates is flat as well.

Par Rates

Consider 100 face or notional amount of a fixed-rate asset that makes regular semiannual coupon or fixed-rate payments of $100 \times \frac{c}{2}$ and a terminal payment at year T of that 100. The T -year, semiannual *par* rate is the rate $C(T)$ such that the present value of this asset equals par or 100. But that is exactly the definition of swap rates given earlier in this chapter. Hence, swap rates in Table 2.1 are, in fact, par rates. For example, for the 2-year swap rate of 1.235%,

$$\frac{1.235}{2} [d(.5) + d(1) + d(1.5) + d(2)] + 100d(2) = 100 \quad (2.18)$$

This equality can be verified by substituting the discount factors from Table 2.1 into (2.18), but this comes as no surprise: the discount factors from that table are derived from a set of pricing equations that included (2.18).

In general, for an asset with a par amount of one unit that makes semiannual payments and matures in T years,

$$\frac{C(T)}{2} \sum_{t=1}^{2T} d\left(\frac{t}{2}\right) + d(T) = 1 \quad (2.19)$$

The sum in equation (2.19), i.e., the value of one unit of currency to be received on every payment date until maturity in T years, is often called an *annuity factor* and denoted by $A(T)$. For semiannual payments,

$$A(T) = \sum_{t=1}^{2T} d\left(\frac{t}{2}\right) \quad (2.20)$$

Using the discount factors from Table 2.1, for example, $A(2)$ is about 3.948. In any case, substituting the annuity notation of (2.20) into (2.19), the par rate equation can also be written as

$$\frac{C(T)}{2} A(T) + d(T) = 1 \quad (2.21)$$

If the term structure of spot or forward rates is flat at some rate, then the term structure of par rates is flat at that same rate. This is proven in Appendix C in this chapter.

Before closing this subsection it is important to point out that a bond with a price of par, or the fixed leg of a swap worth par, may be valued at par only for a moment. As interest rates and discount factors change, the present values of these bonds or swaps change as well and the assets cease to be “par” bonds or swaps.

Synopsis: Quoting Prices with Semiannual Spot, Forward, and Par Rates

Chapter 1 showed that prices of fixed-rate assets can be expressed in terms of discount factors and this section showed that spot, forward, and par rates can be expressed in terms of discount factors. Hence, prices of fixed-rate assets can be expressed in terms of either discount factors or rates. For review and easy reference, this subsection collects these relationships for a unit par amount of a fixed-rate asset with price P that makes semiannual payments at a rate c for T years and then returns par. Using discount and annuity factors,

$$P = \frac{c}{2} A(T) + d(T) \quad (2.22)$$

Using spot rates,

$$P = \frac{c}{2} \left[\frac{1}{\left(1 + \frac{\hat{r}(.5)}{2}\right)} + \frac{1}{\left(1 + \frac{\hat{r}(1)}{2}\right)^2} + \cdots + \frac{1}{\left(1 + \frac{\hat{r}(T)}{2}\right)^{2T}} \right] + \frac{1}{\left(1 + \frac{\hat{r}(T)}{2}\right)^{2T}} \quad (2.23)$$

Using forward rates,

$$P = \frac{c}{2} \left[\frac{1}{\left(1 + \frac{f(.5)}{2}\right)} + \frac{1}{\left(1 + \frac{f(.5)}{2}\right)\left(1 + \frac{f(1)}{2}\right)} + \cdots + \frac{1}{\left(1 + \frac{f(.5)}{2}\right)\left(1 + \frac{f(1)}{2}\right) \cdots \left(1 + \frac{f(T)}{2}\right)} \right] + \frac{1}{\left(1 + \frac{f(.5)}{2}\right)\left(1 + \frac{f(1)}{2}\right) \cdots \left(1 + \frac{f(T)}{2}\right)} \quad (2.24)$$

And finally, using the par rate, $C(T)$, subtract (2.21) from (2.22) to obtain

$$P = 1 + \frac{c - C(T)}{2} A(T) \quad (2.25)$$

CHARACTERISTICS OF SPOT, FORWARD, AND PAR RATES

The six-month spot rate is identically equal to the corresponding forward rate: both are rates on a six-month loan starting on the settlement date. But an interesting first observation from Table 2.1 is that each of the other spot rates is nearly equal to the average of all the forward rates of equal and lower term. Taking the 2.5-year spot rate, for example,

$$1.450\% \approx \frac{.705\% + 1.046\% + 1.384\% + 1.820\% + 2.301\%}{5} \quad (2.26)$$

Intuitively this is not at all surprising: the interest rate on a 2.5-year loan is approximately equal to the average of the rates on a six-month loan and on six-month loans six months, one year, one and a half years, and two years forward. Mathematically, the proceeds from the 2.5-year spot loan must be the same as those from the five forward loans:

$$\begin{aligned} \left(1 + \frac{\hat{r}(2.5)}{2}\right)^5 &= \left(1 + \frac{f(.5)}{2}\right) \left(1 + \frac{f(1)}{2}\right) \left(1 + \frac{f(1.5)}{2}\right) \\ &\quad \times \left(1 + \frac{f(2)}{2}\right) \left(1 + \frac{f(2.5)}{2}\right) \end{aligned} \quad (2.27)$$

So while the 2.5-year spot rate is, strictly speaking, a complex average of the first five six-month forward rates, the simple average is usually a very good approximation.³

A second observation from Table 2.1 is that spot rates are increasing with term while forward rates are greater than spot rates. This is not a coincidence. It has just been established that spot rates are an average of forward rates. Furthermore, adding a number to an average increases that average if and only if the added number is larger than the pre-existing

³Very precisely, one plus half the spot rate is a geometric average of one plus half of each of the forward rates. But a first-order Taylor series approximation to the geometric average is, in fact, the arithmetic average, and is relatively accurate since interest rates are usually small numbers.

average. Using the data in the table, adding the 2-year forward of 2.301% to the 2-year “average” or spot rate of 1.238%, gives a higher new “average” or 2.5-year spot rate of 1.450%. Appendix E in this chapter proves in general that, for any t , $\hat{r}(t) > \hat{r}(t - .5)$ if and only if $f(t) > \hat{r}(t - .5)$ and that $\hat{r}(t) < \hat{r}(t - .5)$ if and only if $f(t) < \hat{r}(t - .5)$. These are period-by-period statements and, as such, do not necessarily extend to entire spot and forward rate curves. In practice, however, spot rates increase or decrease over relatively wide maturity ranges and therefore forward rates are above or below spot rates over relatively wide maturity ranges. Figures 2.2 and 2.3, of the EUR and GBP swap curves as of May 28, 2010, illustrate typical relationships between spot and forward rate curves. In each currency, the spot rate curve increases with term while forward rates are above spot rates, but, as forward rates cross from above to below the spot rates, the spot rate curve begins to decrease with term.

A third and final observation from Table 2.1 is that while spot rates are increasing with term, par rates are near, but below, spot rates. To understand the intuition here, consider the 2.5-year par and spot rates of 1.445% and 1.450%, respectively. From the discussion earlier in this chapter, were the spot rate curve flat at 1.450%, the par rate would be 1.450% as well. In other words, discounting fixed payments of 1.450% at a flat spot rate curve of 1.450% would give a price of par. But this means that discounting 1.450% payments at the spot rates in Table 2.1, which are all less than or equal to 1.450%, would give a price greater than par. Hence, discounting with the spot rates in the table, the par rate must be below 1.450%. More generally,




FIGURE 2.2 EUR Swap Curves as of May 28, 2010




FIGURE 2.3 GBP Swap Curves as of May 28, 2010

Appendix F in this chapter proves that when the spot rate curve is strictly upward-sloping, par rates are below equal-maturity spot rates and that when spot rates are strictly downward-sloping, par rates are above equal-maturity spot rates. USD swap rate curves as of May 28, 2010, shown in Figure 2.4, illustrate how par rates are below spot rates as spot rates increase over most of the maturity range. By the end of the year 2041 the spot rate curve starts




FIGURE 2.4 USD Swap Rates as of May 28, 2010

to decrease very gradually, but not nearly enough for par rates to exceed spot rates. By contrast, the EUR spot rate curve in Figure 2.2 does decrease rapidly enough at the longer maturities for the par rate curve to rise above the spot rate curve.

Maturity and Price or Present Value

If the term structure of rates remains completely unchanged over a six-month period, will the price of a bond or the present value of the fixed side of a swap increase or decrease over the period?

Table 2.2 explores this question by computing the present value of the fixed sides of swaps paying 1.445% to different maturities using the discount factors or rates from Table 2.1. Since 1.445% is the 2.5-year par rate, the present value of 100 face amount of the fixed side of the 2.5-year swap is 100. Six months later, should the term structure be exactly the same, the swap would be a two-year swap and this present value would rise to 100.41. Then, after another six months, the swap would be a 1.5-year swap and, with the term structure still unchanged, would have a present value of 100.60, etc. The third column of the table simply reproduces the forward rates of Table 2.1.

To understand why the present value behaves as it does, rising and then falling, begin by comparing the six-month and 1-year swaps. Both swaps pay 1.445% over the first six months. But then the 1-year swap pays 1.445% for an additional six months while the forward rate over that additional six-month period is only 1.046%. This paying of an above market rate makes the 1-year swap more valuable than the six-month swap and so its price is higher. And so with the 1.5-year swap: it pays 1.445% for the six months from 1-year to 1.5-years from now while the forward rate over that period is only 1.384%. And so, again, the present value increases as maturity increases from one to 1.5 years. But now consider the 2-year swap relative to the 1.5-year swap. The 2-year swap pays 1.445% for an additional six months while the forward rate for that six months is 1.820%. Hence the 2-year swap pays a below-market rate for the additional six months and

TABLE 2.2 Present Values of 100 Face Amount of the Fixed Sides of 1.445% Swaps as of May 28, 2010

Maturity	Present Value	Forward Rate
.5	100.37	.705%
1	100.57	1.046%
1.5	100.60	1.384%
2	100.41	1.820%
2.5	100.00	2.301%

has a present value less than that of the 1.5-year swap. Finally, the present value of the 2.5-year swap is less than that of the 2-year swap because the 2.5-year swap pays 1.445% for an additional six months while the forward rate is 2.301%.

Returning to the original question, then, if the term structure of rates remains unchanged over a six-month period, the present value will rise as the swap matures if its fixed rate is less than the forward rate corresponding to the expiring six-month period. The present value will fall as the swap matures if its fixed rate is greater than that forward rate. Appendix G in this chapter proves this general result.

TRADING CASE STUDY: TRADING AN ABNORMALLY DOWNWARD-SLOPING 10S-30S EUR FORWARD RATE CURVE IN Q2 2010

Figure 2.5 graphs six-month forward rate curves for USD, EUR, GBP, and JPY as of May 28, 2010. In EUR for example, the six-month rate, 10-years forward, or the 10y6m rate, is about 4.25% while the USD six-month rate, 30-years forward, or the 30y6m rate, is about 4%. By historical standards the EUR curve is remarkable in how the “10s-30s” forward curve, i.e., the curve from 10- to 30-year terms, slopes so steeply downward. The more usual historical shape is more like that of the other curves in the figure,




FIGURE 2.5 Forward Swap Rates in USD, EUR, GBP, and JPY as of May 28, 2010

sloping upward from short- to intermediate-maturities and then flattening out and falling gradually at the long end.

The macroeconomic context at the time was concerned about the fiscal difficulties and economic prospects of EUR countries triggered by fears that Greece and a number of other countries might default on their government debts. These fears were somewhat mitigated by a bailout fund proposed by EUR countries and the International Monetary Fund.

The technical context of these curves at this time was a particular theme of the Overview, namely, the need for European pension funds and insurance companies to invest in long-dated assets, or, in swap language, to receive fixed on the long end, so that their asset profiles better matched their long-term liabilities. This need was particularly acute after the approval of the Solvency II directive, which required additional capital to reflect any asset and liability mismatches. In any case, this institutional pressure to receive fixed on the long end, without any commensurately sized payers on the long end, drove down long-term swap rates and was one factor responsible for the abnormally downward sloping EUR forward curve.

Before moving on to trade ideas, it will be useful to explain some market jargon. Consider the two pairs of abstract term structures of rates depicted in Figure 2.6. Market practitioners use the word *flattening* to describe shifts in which either 1) longer-term rates fall by more than shorter-term rates, or 2) shorter-term rates rise by more than longer-term rates. Therefore, by 1), a shift from either of the solid lines in the figure to its corresponding




FIGURE 2.6 Shifting from Either Solid Line to Its Dashed Line Is Called a “Flattening” of the Term Structure

dashed line would be called a flattening even though everyday usage of the word “flatten” would not apply to the shift from the lower solid line to the lower dashed line. Similarly, market practitioners use the word *steepening* to describe shifts in which either 3) longer-term rates increase by more than shorter-term rates, or 4) shorter-term rates fall by more than longer-term rates. Therefore, by 3), a shift from either of the dashed lines to its corresponding solid line would be called a steepening even though everyday normal usage of “steepen” would not apply to the shift from the lower dashed line to the lower solid line.

Returning now to the case, many market participants wanted to bet that the EUR forward curve in Figure 2.5 would revert to a more normal shape, i.e., that the 10s-30s forward curve would steepen. It was argued that the institutional demand to receive fixed would eventually be absorbed by the market so that a more normally sloped curve could be obtained. Furthermore, the technical factors holding down the long end would soon be overpowered by trading to follow in the wake of the resolution of macroeconomic uncertainty in Europe. More precisely, should the fiscal and economic situation in the EUR seriously deteriorate, the EUR forward curve would converge to the JPY forward curve and 10s-30s would steepen. On the other hand, should the fiscal and economic situation in the EUR improve, the EUR forward curve would converge to the USD and GBP curve and, once again, 10s-30s would steepen.

It might be the case, of course, that 10s-30s does not steepen. First, the institutional demand to receive fixed in the long end might so overwhelm the supply of payers that no amount of trading driven by macroeconomic considerations would drive 10s-30s EUR forwards back to historical norms. In fact, should incremental institutional demand to receive fixed continue to exceed incremental supply, 10s-30s might flatten even more. Also, global macroeconomic forces might flatten 10s-30s across the globe, which may very well have nothing to do with EUR technicals but which would still result in the EUR curve's flattening.

A trader who comes to the conclusion that the risk-return characteristics of the steepening bet are appealing can implement the bet through the following trade: receive fixed in the relatively high EUR 10y6m rate and pay fixed in the abnormally low long end.⁴ Put another way, lock in a rate to receive 10y6m and lock in a rate to pay in the long end as a bet that the 10y6m forward is going to fall relative to the longer-dated forwards. In addition, construct the trade so that if the 10y6m and longer-dated forward rates both increase by one basis point (i.e., .01%), the loss from the 10-year

⁴It is possible that the trade would be implemented in exactly this way, but as six-month forwards at long maturities are not liquid, a much more likely implementation would use portfolios of par swaps. For clarity of exposition, however, the text assumes direct trading in short-term forwards.

TABLE 2.3 Selected EUR and JPY Forward Rates as of May 28, 2010

	10y6m	9y6m	25y6m	24y6m	30y6m	29y6m
EUR	4.254%	4.127%	2.550%	2.724%	2.293%	2.237%
JPY	2.712%	2.594%	2.433%	2.452%	2.219%	2.339%

TABLE 2.4 One-Year Roll-Down from Receiving 10y6m EUR and Paying 30y6m EUR as of May 28, 2010, Assuming an Unchanged Term Structure

	Today		One Year Later		Gain/Loss
	Forward	Rate	Forward	Rate	(bps)
Receive	10y6m	4.254	9y6m	4.127	+12.7
Pay	30y6m	2.293	29y6m	2.237	-5.6
Total					+7.1

leg is offset by the gain from the longer-dated leg and the trade neither makes nor loses money. (Part Two shows how this type of hedge is constructed.)

A final aspect of the trade to consider is *roll-down*,⁵ i.e., how the trade fares if rates do not change much at all, which would be the case, for example, if the forward rate curve remains the same. For if the trade does lose money over time as nothing happens, then the trader may not be able to stay in the trade long enough to realize the anticipated profits. This implied impatience can arise from internal risk management controls that force the closure of trades hitting *stop-losses* (i.e., loss thresholds). Impatience can also arise from the inability or reluctance, as trades lose money, to post more and more collateral to counterparties to ensure the performance of increasingly under-water contracts. (See Chapter 12.) In any case, to analyze the roll-down of the trade discussed thus far, Table 2.3 gives six-month forward rates of various terms in EUR and, for later use, in JPY as of May 28, 2010.

Say that a trader decides to implement the suggested trade by receiving in EUR 10y6m and paying in EUR 30y6m. How does this trade roll-down over a year in which the term structure does not change? Table 2.4, using the forward rates in Table 2.3, outlines the answer. After one year the trader will have a position receiving 4.254% in 9y6m and paying 2.293% in 29y6m, but the market rates for those forwards will have fallen to 4.127% and 2.237%, respectively. As the table shows, this means a gain of 12.7 basis points (i.e., $4.254\% - 4.127\%$) on the receiving leg of the trade and a loss

⁵Some practitioners would call this *carry* or *carry-roll-down*. See the discussion in Chapter 3.

of 5.6 basis points (i.e., $-2.293\% + 2.237\%$) on the paying leg of the trade. Furthermore, since the trade is constructed so that each leg has the same exposure to a change in interest rates, the net result would simply be the sum of the individual results or +7.1 basis points. So, for example, a trade scaled to have an interest rate exposure of €10,000 per basis point would gain €71,000.

But what if, instead of selling the 30y6m forward, the trader pays fixed in the 25y6m forward? This may be harder to transact, as the 30-year maturity is more liquidly traded, but it is a choice to be considered. Table 2.5 computes the roll-down in this case, again using the forward rates in Table 2.3. The receiving leg is unchanged and still gains 12.7 basis points. But the paying leg, since the 24y6m rate is greater than the 25y6m rate, gains as well, in the amount of 17.4 basis points. Hence the total roll-down, the sum of the roll-down of the two legs, is 30.1 basis points. This revised trade, then, has much better roll-down properties than the originally conceived trade.

It was noted above that the proposed trade would lose money if 10s-30s around the globe flattened due to shared macroeconomic shocks. A possible hedge to this losing scenario is to put on the opposite trade in another currency, e.g., to pay fixed in 10y6m and to receive fixed in 25y6m in JPY. It makes sense to put on this hedge only if two conditions are met. One, 10s-30s in that currency is not likely to experience any idiosyncratic moves over the time horizon of the trade; if such idiosyncratic moves were likely, the hedge might very well increase rather than decrease the volatility of the trade's results. Two, the roll-down of the hedge is not so negative as to spoil the appealing risk-return profile of the original trade.

As it turns out, the JPY curve seems very suitable for this hedge, i.e., paying in 10y6m and receiving in 25y6m. First, resolution of Japan's fiscal and economic situation and, therefore, a reshaping of its swap curve, is expected to happen much more slowly than a resolution of the situation in the EUR countries. Second, using the data in Table 2.3, the incremental roll-down of this trade is a negative $2.712\% - 2.594\%$ or -11.8 basis points from the 10-year leg and a negative $2.433\% - 2.453\%$ or -2 basis points

TABLE 2.5 One-Year Roll-Down from Receiving 10y6m EUR and Paying 25y6m EUR as of May 28, 2010, Assuming an Unchanged Term Structure

	Today		One Year Later		Gain/Loss (bps)
	Forward	Rate	Forward	Rate	
Receive	10y6m	4.254	9y6m	4.127	+12.7
Pay	25y6m	2.550	24y6m	2.724	+17.4
Total					+30.1

from the 25-year leg for a total of -13.8 basis points. Noting that the overall roll-down of the trade, the original 30.1 basis points minus the 13.8 basis points of the macroeconomic hedge, is a reasonable $+16.3$ basis points, a trader might very well choose to purchase this insurance by adding the hedge to the original trade.

It is possible, of course, that 10s-30s in EUR becomes more steeply downward sloping at the same time that JPY 10s-30s becomes less steeply downward sloping, in which case both the original trade and the hedge lose money. But nothing in the analysis of the macroeconomic and technical foundations of the trade suggests this eventuality. And, after all, a trade is always a bet on something!

APPENDIX A: COMPOUNDING CONVENTIONS

The text discussed semiannual compounding, which assumes that interest is paid twice a year, and showed that one unit of currency invested at the rate $\hat{r}^{sa}$ for T years would grow to

$$\left(1 + \frac{\hat{r}^{sa}}{2}\right)^{2T} \quad (2.28)$$

Similarly, it is easy to see that one unit of currency invested at an annual rate $\hat{r}^a$, a monthly rate $\hat{r}^m$, or a daily rate $\hat{r}^d$, would grow after T years to the following quantities, respectively,

$$(1 + \hat{r}^a)^T \quad (2.29)$$

$$\left(1 + \frac{\hat{r}^m}{12}\right)^{12T} \quad (2.30)$$

$$\left(1 + \frac{\hat{r}^d}{365}\right)^{365T} \quad (2.31)$$

More generally, if interest at a rate $\hat{r}$ is paid n times per year, the proceeds after T years will be

$$\left(1 + \frac{\hat{r}}{n}\right)^{nT} \quad (2.32)$$

One would expect that, holding all other characteristics of investment constant, the market would offer a single terminal amount for having invested one unit of currency for T years. Given the quantities in equations (2.28) through (2.32), this means that the market could offer many different

rates of interest for that investment, each associated with a different compounding convention. So, for example, if the market offers 2% annually compounded for a one-year investment, so that a unit investment grows to 1.02 at the end of a year, rates of other compounding conventions would be determined by the equations

$$\left(1 + \frac{\hat{r}^{sa}}{2}\right)^2 = \left(1 + \frac{\hat{r}^m}{12}\right)^{12} = \left(1 + \frac{\hat{r}^d}{365}\right)^{365} = 1.02 \quad (2.33)$$

Solving equations (2.33) for each rate, $\hat{r}^{sa} = 1.9901\%$; $\hat{r}^m = 1.9819\%$; and $\hat{r}^d = 1.9803\%$. Note that the more often interest is paid, the more interest can earn interest on interest, and the lower the rate required to earn the fixed amount 1.02 over the year.

Under continuous compounding, interest is paid every instant, resulting in a terminal value equal to the limit of the quantity (2.32) as n approaches infinity. Taking the natural logarithm of both sides of that equation and rearranging terms,

$$nT \ln\left(1 + \frac{r}{n}\right) = \frac{T \ln\left(1 + \frac{r}{n}\right)}{\frac{1}{n}} \quad (2.34)$$

Using l'Hôpital's rule, the limit of the right-hand side of (2.34) as n becomes large is rT . Hence, the limit of (2.32) must be e^{rT} , where $e = 2.71828 \dots$ is the base of the natural logarithm. Therefore, if interest is paid at a rate r^c every instant, an investment of one unit of currency will grow after T years to

$$e^{r^c T} \quad (2.35)$$

Equivalently, the value of one unit of currency to be received in T years is

$$e^{-r^c T} \quad (2.36)$$

APPENDIX B: CONTINUOUSLY COMPOUNDED SPOT AND FORWARD RATES

Let $\hat{r}^c(t)$ be the continuously compounded spot rate from time 0 to t , let $f^c(t)$ be the continuously compounded forward rate at time t , and let $f^c(t - \Delta, t)$ be the forward rate from time $t - \Delta t$ to time t , which approaches $f^c(t)$ as Δt approaches 0. From the discussion on spot rates in the text and the discussion on continuous compounding in Appendix A and equation (2.36) in particular, the continuously compounded spot rate is defined such that

$$d(t) = e^{-\hat{r}^c(t)t} \quad (2.37)$$

With respect to forward rates, the continuously compounded analogue of equation (2.14) of the text is

$$e^{\hat{r}^c(t-\Delta) \times (t-\Delta)} e^{f^c(t-\Delta, t)\Delta} = e^{\hat{r}^c(t)t} \quad (2.38)$$

Substituting for each of the two spot rates using equation (2.37) and rearranging terms,

$$e^{f^c(t-\Delta, t)\Delta} = \frac{d(t-\Delta)}{d(t)} \quad (2.39)$$

Next, taking the natural logarithm of both sides and rearranging terms,

$$f^c(t-\Delta, t) = -\frac{\ln[d(t)] - \ln[d(t-\Delta)]}{\Delta} \quad (2.40)$$

Finally, taking the limit of both sides, recognizing the limit of the right-hand side of (2.40) as the derivative of $\ln[d(t)]$,

$$f^c(t) = -\frac{d'(t)}{d(t)} \quad (2.41)$$

where $d'(t)$ is the derivative of the discount function.

APPENDIX C: FLAT SPOT RATES IMPLY FLAT PAR RATES

Proposition: If spot rates are flat at the rate $\hat{r}$, then par rates are flat at that same rate.

Proof: Write equation (2.19) in terms of the single spot rate, $\hat{r}$:

$$\frac{C(T)}{2} \sum_{t=1}^{2T} \frac{1}{(1 + \frac{\hat{r}}{2})^t} + \frac{1}{(1 + \frac{\hat{r}}{2})^{2T}} = 1 \quad (2.42)$$

Using (2.49) in Appendix 2.D, rewrite this equation as

$$\frac{C(T)}{\hat{r}} \left[1 - \frac{1}{(1 + \frac{\hat{r}}{2})^{2T}} \right] + \frac{1}{(1 + \frac{\hat{r}}{2})^{2T}} = 1 \quad (2.43)$$

But solving (2.43) for $C(T)$ shows that $C(T) = \hat{r}$ for all T .

APPENDIX D: A USEFUL SUMMATION FORMULA

Proposition:

$$\sum_{t=a}^b z^t = \frac{z^a - z^{b+1}}{1 - z} \quad (2.44)$$

Proof: Define S such that

$$S = \sum_{t=a}^b z^t \quad (2.45)$$

Then,

$$zS = \sum_{t=a+1}^{b+1} z^t \quad (2.46)$$

And, subtracting (2.46) from (2.45),

$$S(1 - z) = z^a - z^{b+1} \quad (2.47)$$

Finally, dividing both sides of (2.47) by $1 - z$ gives equation (2.44), as desired.

This proposition is quite useful in fixed income where expressions like the one in equation (2.42) of Appendix C are common:

$$\sum_{t=1}^{2T} \frac{1}{(1 + \frac{\hat{r}}{2})^t} \quad (2.48)$$

Setting $z = \frac{1}{(1 + \frac{\hat{r}}{2})}$ and applying the proposition of this appendix gives the result

$$\begin{aligned} \sum_{t=1}^{2T} \frac{1}{(1 + \frac{\hat{r}}{2})^t} &= \frac{\frac{1}{(1 + \frac{\hat{r}}{2})} - \frac{1}{(1 + \frac{\hat{r}}{2})^{2T+1}}}{1 - \frac{1}{(1 + \frac{\hat{r}}{2})}} \\ &= \frac{1 - \frac{1}{(1 + \frac{\hat{r}}{2})^{2T}}}{\frac{\hat{r}}{2}} = \frac{2}{\hat{r}} \left[1 - \frac{1}{(1 + \frac{\hat{r}}{2})^{2T}} \right] \end{aligned} \quad (2.49)$$

APPENDIX E: THE RELATIONSHIP BETWEEN SPOT AND FORWARD RATES AND THE SLOPE OF THE TERM STRUCTURE

Proposition: For semiannually compounded rates, $f(t) > \hat{r}(t - .5)$ if and only if $\hat{r}(t) > \hat{r}(t - .5)$.

Proof: $f(t) > \hat{r}(t - .5)$ is equivalent to

$$\begin{aligned} \left(1 + \frac{\hat{r}(t - .5)}{2}\right)^{2t-1} \left(1 + \frac{f(t)}{2}\right) &> \left(1 + \frac{\hat{r}(t - .5)}{2}\right)^{2t-1} \left(1 + \frac{\hat{r}(t - .5)}{2}\right) \\ &> \left(1 + \frac{\hat{r}(t - .5)}{2}\right)^{2t} \end{aligned} \quad (2.50)$$

But the left-hand side of this equation can be written in terms of $\hat{r}(t)$ using equation (2.14).

$$\left(1 + \frac{\hat{r}(t)}{2}\right)^{2t} > \left(1 + \frac{\hat{r}(t - .5)}{2}\right)^{2t} \quad (2.51)$$

And this is equivalent to $\hat{r}(t) > \hat{r}(t - .5)$.

Proposition: For semiannually compounded rates, $f(t) < \hat{r}(t - .5)$ if and only if $\hat{r}(t) < \hat{r}(t - .5)$.

Proof: Reverse the inequalities in the previous proof.

Proposition: For continuously compounded rates, $f^c(t) \geq \hat{r}^c(t)$ if and only if $\frac{d\hat{r}^c(t)}{dt} \geq 0$ and $f^c(t) \leq \hat{r}^c(t)$ if and only if $\frac{d\hat{r}^c(t)}{dt} \leq 0$.

Proof: Taking the derivative of equation (2.37),

$$d'(t) = - \left[\hat{r}^c(t) + t \frac{d\hat{r}^c(t)}{dt} \right] d(t) \quad (2.52)$$

Dividing both sides by $-d(t)$ and then substituting for the left-hand side using (2.41),

$$f^c(t) = \hat{r}^c(t) + t \frac{d\hat{r}^c(t)}{dt} \quad (2.53)$$

Rearranging terms,

$$\frac{d\hat{r}^c(t)}{dt} = \frac{f^c(t) - \hat{r}^c(t)}{t} \quad (2.54)$$

By inspection, then, $\frac{d\hat{r}^c(t)}{dt}$ has the same sign as $f^c(t) - \hat{r}^c(t)$.

APPENDIX F: THE RELATIONSHIP BETWEEN SPOT AND PAR RATES AND THE SLOPE OF THE TERM STRUCTURE

Proposition: If $\hat{r}(.5) < \hat{r}(1) < \dots < \hat{r}(T)$ then $C(T) < \hat{r}(T)$.

Proof: By the definition of the par rate, $C(T)$,

$$\frac{C(T)}{2} \left[\frac{1}{\left(1 + \frac{\hat{r}(.5)}{2}\right)} + \dots + \frac{1}{\left(1 + \frac{\hat{r}(T)}{2}\right)^{2T}} \right] + \frac{1}{\left(1 + \frac{\hat{r}(T)}{2}\right)^{2T}} = 1 \quad (2.55)$$

Also, setting all spot rates in (2.55) equal to $C(T)$, it follows from (2.49) of Appendix D that

$$\frac{C(T)}{2} \left[\frac{1}{\left(1 + \frac{C(T)}{2}\right)} + \dots + \frac{1}{\left(1 + \frac{C(T)}{2}\right)^{2T}} \right] + \frac{1}{\left(1 + \frac{C(T)}{2}\right)^{2T}} = 1 \quad (2.56)$$

Furthermore, since $\hat{r}(.5) < \hat{r}(1) < \dots < \hat{r}(T)$, the expression

$$\frac{C(T)}{2} \left[\frac{1}{\left(1 + \frac{\hat{r}(T)}{2}\right)} + \dots + \frac{1}{\left(1 + \frac{\hat{r}(T)}{2}\right)^{2T}} \right] + \frac{1}{\left(1 + \frac{\hat{r}(T)}{2}\right)^{2T}} \quad (2.57)$$

which sets all of the discounting rates to $\hat{r}(T)$, is less than the left-hand side of equation (2.55). But since the left-hand sides of both (2.55) and (2.56) equal 1, this implies that (2.57) is also less than the left-hand side of (2.56). And this, in turn, implies that $C(T) < \hat{r}(T)$, as was to be proved.

Proposition: If $\hat{r}(.5) > \hat{r}(1) > \dots > \hat{r}(T)$ then $C(T) > \hat{r}(T)$.

Proof: In this case, (2.57) is greater than the left-hand side of equation (2.55) and, therefore, of (2.56). But this implies that $C(T) > \hat{r}(T)$, as was to be proved.

APPENDIX G: MATURITY, PRESENT VALUE, AND FORWARD RATES

Proposition: The sign of $P(T) - P(T - .5)$ equals the sign of $c - f(T)$.

Proof: Using equation (2.24) for $P(T)$ and for $P(T - .5)$ it can be shown that

$$P(T) - P(T - .5) = \frac{1 + \frac{c}{2}}{\left(1 + \frac{f(.5)}{2}\right) \left(1 + \frac{f(1)}{2}\right) \cdots \left(1 + \frac{f(T)}{2}\right)} - \frac{1}{\left(1 + \frac{f(.5)}{2}\right) \cdots \left(1 + \frac{f(T-.5)}{2}\right)} \quad (2.58)$$

Or,

$$P(T) - P(T - .5) = \frac{1 + \frac{c}{2} - \left(1 + \frac{f(T)}{2}\right)}{\left(1 + \frac{f(.5)}{2}\right) \left(1 + \frac{f(1)}{2}\right) \cdots \left(1 + \frac{f(T)}{2}\right)} \quad (2.59)$$

Or, again,

$$P(T) - P(T - .5) = \frac{\frac{1}{2}(c - f(T))}{\left(1 + \frac{f(.5)}{2}\right) \left(1 + \frac{f(1)}{2}\right) \cdots \left(1 + \frac{f(T)}{2}\right)} \quad (2.60)$$

Therefore the sign of $P(T) - P(T - .5)$ equals the sign of $c - f(T)$.

Returns, Spreads, and Yields

Spot, forward, and par rates, presented in Chapter 2, intuitively describe the time value of money embedded in market prices. To analyze the *ex-post* performance and the *ex-ante* relative attractiveness of individual securities, however, market participants rely on returns, spreads, and yields.

The first section of this chapter defines these terms. Horizon returns in the fixed income context have to account for intermediate cash flows and are often computed both on a gross basis and net of financing, but are otherwise similar to the returns calculated for any asset. Spreads measure the pricing of an individual fixed income security relative to a benchmark curve, usually of swaps or government bonds. Yield is a practical and intuitive way to quote price and is used extensively for quick insight and analysis. It cannot be used, however, as a precise measure of relative value. This first section concludes with a brief news excerpt about the sale of Greek government bonds that illustrates the convenience of speaking in terms of spreads and yields.

The second section of the chapter shows how the profit-and-loss (P&L) or return of a fixed income security can be decomposed into component parts. Such decompositions are defined differently by different market participants, but this book will define terms as follows. Cash carry is a security's coupon income minus its financing cost, a quantity that will be particularly useful in the context of forwards and futures (see Chapters 13 and 14). Carry-roll-down is the change in the (flat) price of a security if rates move "as expected," where one common interpretation of "as expected" is the scenario of realized forwards and another is the scenario of an unchanged term structure, both of which are described in this chapter.

The third and final section of the chapter presents several carry-roll-down scenarios, partly to complete the discussion of return decompositions, but partly for the insights these scenarios provide with respect to bond returns. Two such insights are the following: 1) if realized forward rates exceed the forward rates embedded in bond prices, a strategy of rolling over short-term bonds outperforms an investment in long-term bonds; 2) a bond's return equals its yield only if its yield stays constant and if all coupons are reinvested at that same yield.

DEFINITIONS

Realized Returns

This section begins the chapter by defining *gross* and *net* realized returns over a single period and over several periods. Very simply, net returns are gross returns minus financing costs. For concreteness and ease of exposition this chapter focuses exclusively on bonds, but the principles and definitions presented can easily be extended to other securities. For the same reasons, this chapter calculates returns only over holding periods equal to the length of time between cash flows, so that, for example, the returns of semiannual coupon bonds are calculated only over six-month holding periods.

Since Chapters 1 and 2 have dealt extensively with the details of semiannual cash flows, this chapter simplifies notation by not explicitly recording the length of each period. Denote the price of a particular bond at time t by P_t per unit face value and the price of that same bond, after one period of unspecified length, as P_{t+1} . Also, denote the bond's periodic coupon payment per unit face value by c . Numerical examples, however, will explicitly incorporate semiannual cash flow conventions and will assume face values of 100.

An investor purchasing a bond at time t pays P_t and then, at time $t + 1$, receives a coupon c and has a bond worth P_{t+1} . The gross realized return on that bond from t to $t + 1$, $R_{t,t+1}$, is defined as the total value at the end of the period minus the starting value all divided by the starting value. Mathematically,

$$R_{t,t+1} = \frac{P_{t+1} + c - P_t}{P_t} \quad (3.1)$$

Continuing with the U.S. Treasury example of Chapter 1, say that an investor bought the U.S. Treasury $4\frac{1}{2}$ s of November 30, 2011, for 105.856 for settlement on June 1, 2010. Then suppose that the price of the bond one coupon-period later, on November 30, 2010, turned out to be 105. The six-month return on that investment would have been

$$\frac{105 + 2.25 - 105.856}{105.856} = 1.317\% \quad (3.2)$$

where the 2.25 in the numerator is the bond's semiannual coupon payment.

Computing a realized return over a longer holding period requires keeping track of the rate at which coupons are reinvested over the holding period. Consider an investment in the same bond for one year, that is, to May 31, 2011. The total proceeds at the end of the year consist not only of the value of the bond and the coupon payment on May 31, 2011, but also of the reinvested proceeds of the coupon paid on November 30, 2010. Assuming

that this November coupon is invested at a semiannually compounded rate of .60% and that the price of the bond on May 31, 2011, is 105, the realized gross holding period return over the year would be

$$\frac{105 + 2.25 + 2.25 \times \left(1 + \frac{.60\%}{2}\right) - 105.856}{105.856} = 3.449\% \tag{3.3}$$

Now consider an investor in the 4½s of November 30, 2011, who *financed* the purchase of the bond, that is, who borrowed cash to make the investment. (The important institutional details of such borrowing are presented in Chapter 12.) While not usually the case, assume for the purposes of this chapter that the investor could borrow the entire purchase price of the bond. Assume a rate of .2% for .5 years on the amount borrowed so that paying off the loan costs $105.856 \times \left(1 + \frac{.2\%}{2}\right)$ or 105.962. Also assume, as before, that the price of the bond is 105 on November 30, 2010. Then, this investment over a six-month horizon is described as in Table 3.1.

One obvious problem in calculating a return on this investment is that it requires no initial cash and the final value cannot be divided by zero. But even if the investor did have to put up some amount of initial cash, so that borrowing was 90% rather than 100% of the purchase price, it would still not be sensible to divide the final value by the amount invested when trying to describe the return on the 4½s of November 30, 2011. After all, another investor might have borrowed 95% of the purchase price and a third investor only 85%. Hence it would be sensible to divide by the investor’s outlay only to calculate a return on capital for that investor. But that is not the exercise here. Therefore, when calculating realized returns on securities, even when those securities are financed, it is conventional to divide that final value by the initial price of the security.

With this choice of a denominator, the net realized return on the security looks almost, but not exactly, like the gross return in (3.2):

$$\frac{105 + 2.25 - 105.962}{105.856} = 1.217\% \tag{3.4}$$

TABLE 3.1 A Financed Purchase of the 4½s of November 30, 2011

Settlement Date	Transaction	Proceeds	Total Proceeds
June 1, 2010	buy bond	−105.856	0
	borrow price	105.856	
November 30, 2010	collect coupon	2.250	1.288
	sell bond	105.000	
	pay off loan	−105.962	

In fact, the net return is simply the previously calculated gross return of 1.317% minus the 0.1% cost of six-month financing. To make this a bit more explicit,

$$\begin{aligned} \frac{105 + 2.25 - 105.856 \times (1 + \frac{.2\%}{2})}{105.856} &= \frac{105 + 2.25 - 105.856}{105.856} - \frac{.2\%}{2} \\ &= 1.317\% - .1\% \end{aligned} \quad (3.5)$$

Without going into further detail here, calculating a multi-period net return requires not only the reinvestment rates of the coupons but the future financing costs as well.

Spreads

As mentioned in the introduction to the chapter, spreads are important measures of relative value and their convergence or divergence is an important component of return.

The market price of any security can be thought of as its value computed using some term structure of interest rates, denoted generically by $\mathbb{R}$, plus a premium or discount, ϵ , relative to that term structure:

$$P \equiv P(\mathbb{R}) + \epsilon \quad (3.6)$$

Furthermore, the premium or discount ϵ is often expressed in terms of a spread to interest rates, s , rather than in terms of price. Mathematically, first write $P(\mathbb{R})$ using forward rates (as in equation (2.24) but with periods of unspecified length):

$$\begin{aligned} P &\equiv \frac{c}{(1 + f(1))} + \frac{c}{(1 + f(1))(1 + f(2))} + \cdots \\ &\quad + \frac{1 + c}{(1 + f(1))(1 + f(2)) \cdots (1 + f(T))} + \epsilon \end{aligned} \quad (3.7)$$

Then, instead of defining the deviation of the market from $P(\mathbb{R})$ through ϵ , define it through a spread. In other words, find s such that the following equation is identically true:

$$\begin{aligned} P &\equiv \frac{c}{(1 + f(1) + s)} + \frac{c}{(1 + f(1) + s)(1 + f(2) + s)} + \cdots \\ &\quad + \frac{1 + c}{(1 + f(1) + s)(1 + f(2) + s) \cdots (1 + f(T) + s)} \end{aligned} \quad (3.8)$$

In words, the market price is recovered by discounting a bond's cash flows using an appropriate term structure plus a spread.

Spreads defined as in equation (3.8) are usually intended to be either bond- or sector-specific. As an example of the former, recall the testing of the law of one price in Table 1.4. The $\frac{3}{4}$ s of November 30, 2011, when priced using the discount curve derived in Chapter 1, gave a present value of 100.255 compared with a market price of 100.190. To express this price deviation or ϵ in terms of spread, express the discount factors in Table 1.3 as forward rates and solve the following equation for s :

$$100.190 \equiv \frac{.375}{\left(1 + \frac{.149\%}{2} + \frac{s}{2}\right)} + \frac{.375}{\left(1 + \frac{.149\%}{2} + \frac{s}{2}\right) \left(1 + \frac{.556\%}{2} + \frac{s}{2}\right)} + \frac{100.375}{\left(1 + \frac{.149\%}{2} + \frac{s}{2}\right) \left(1 + \frac{.556\%}{2} + \frac{s}{2}\right) \left(1 + \frac{1.036\%}{2} + \frac{s}{2}\right)} \quad (3.9)$$

The result is $s = .044\%$ or 4.4 basis points. With this spread result, instead of saying that the $\frac{3}{4}$ s of November 30, 2011, trade 6.5 cents cheap relative to the reference bonds, one could say that they trade 4.4 basis points cheap. Sometimes speaking in terms of price is more useful, as when saying that buying the $\frac{3}{4}$ s and selling its replicating portfolio will produce a P&L of 6.5 cents per \$100. But sometimes speaking in terms of spread is more intuitive, as when saying that the $\frac{3}{4}$ s trade at 4.4 basis points above the Treasury curve. There is also an interpretation of that 4.4 basis points in terms of the bond's return, which will be presented in the third section of this chapter.

Equation (3.8) and the U.S. Treasury note example illustrate bond-specific spreads. A common example of sector-specific spreads would be corporate bonds, the subject of Chapter 19. In that context bonds issued by a particular corporation are thought of as trading at a spread curve to government bonds or swaps, where a spread curve means that the forward spread at each term is different. The pricing equation for a bond in that case might take the following form:

$$P \equiv \frac{c}{(1 + f(1) + s(1))} + \frac{c}{(1 + f(1) + s(1))(1 + f(2) + s(2))} + \cdots + \frac{1 + c}{(1 + f(1) + s(1)) \cdots (1 + f(T) + s(T))} \quad (3.10)$$

Yield-to-Maturity

While par, spot, and forward rates are in many contexts more intuitive than prices, their appeal suffers from needing so many rates to describe the

pricing of a single bond. As a result, *yield-to-maturity* is often quoted when describing a security in terms of rates rather than price.

Yield-to-maturity is the single rate such that discounting a security's cash flows at that rate gives that security's market price. For example, Table 1.2 reported that, with 1.5 years to maturity, the price of the 4½s of November 30, 2011, was 105.856. The yield-to-maturity, y , of this bond is therefore defined such that¹

$$105.856 \equiv \frac{2.25}{(1 + \frac{y}{2})} + \frac{2.25}{(1 + \frac{y}{2})^2} + \frac{102.25}{(1 + \frac{y}{2})^3} \quad (3.11)$$

Juxtaposing equation (3.11) with equations (2.23), (2.24), and (3.8) or (3.10) reveals that yield summarizes both the term structure of interest rates as well as any spread or spread curve for this bond relative to that term structure. In any case, solving (3.11) for y by trial-and-error or some numerical method shows that the yield of the 4½s is about .574%. While it is much easier to solve for price given yield than for yield given price, many calculators and computer applications are readily available to move from price to yield or *vice versa*. Yield is often used as an alternate way to quote price: a trader could bid to buy the 4½s of November 30, 2011, at a price of 105.856 or at a yield of .574%. Needless to say, market practice is not such that a trader can bid to buy the bond with three spot or forward rates instead of a price.

The definition of yield for a coupon bond for settlement on a coupon payment date is²

$$P = \frac{\frac{1}{2}c}{(1 + \frac{y}{2})} + \frac{\frac{1}{2}c}{(1 + \frac{y}{2})^2} + \cdots + \frac{1 + \frac{1}{2}c}{(1 + \frac{y}{2})^{2T}} \quad (3.12)$$

Or, more compactly,

$$P = \frac{c}{2} \sum_{t=1}^{2T} \frac{1}{(1 + \frac{y}{2})^t} + \frac{1}{(1 + \frac{y}{2})^{2T}} \quad (3.13)$$

And simplifying using the summation formula given in Appendix D in Chapter 2,

$$P = \frac{c}{y} \left(1 - \frac{1}{(1 + \frac{y}{2})^{2T}} \right) + \frac{1}{(1 + \frac{y}{2})^{2T}} \quad (3.14)$$

¹This is not perfectly correct since the prices in Table 1.2 were for settlement on June 1, 2010, rather than May 31, 2010. See Appendix A in this chapter for a more precise definition.

²The formula for other settlement dates is given in Appendix A in this chapter.




FIGURE 3.1 Prices of Bonds with Varying Coupons Over Time with Yields Fixed at 3%

Equation (3.14) provides several immediate facts about the price-yield relationship. First, when $c = y$, $P(T) = 1$. In words, when the yield is equal to the coupon rate, the bond sells for its face value. Second, when $c > y$, $P(T) > 1$: when the coupon rate exceeds the yield, the bond sells at a premium to its face value. Third, when $c < y$, $P(T) < 1$: when the yield exceeds the coupon rate, the bond sells at a discount to its face value.

Figure 3.1 illustrates these first three implications of equation (3.14). Fixing all yields at 3%, each curve gives the price of a bond with a particular coupon rate as a function of years remaining to maturity. The bond with a coupon of 3% has a price of 100 at all terms. With 30 years to maturity, the 4% and 5.5% bonds sell at substantial premiums to par. As these bonds mature, however, the value of an above-market coupon falls: receiving a coupon 1% or 2.5% above market for 20 years is not so valuable as receiving those above-market coupons for 30 years. Hence, the prices of these premium bonds fall over time until they are worth par at maturity. Conversely, the .5% and 2% bonds sell at substantial discounts to par with 30 years to maturity and rise in price as they mature. The time trend of bond prices depicted in the figure is known as the *pull to par*. Of course, the realized price paths of these bonds will differ dramatically from those in Figure 3.1 (which fixes all yields at 3%) according to the actual realization of yields.

The fourth lesson from the price-yield relationship of equation (3.14) is the *annuity* formula. An annuity makes annual payments of 1 until date

T with no final principal payment. In this case, the second term of (3.14) vanishes and, with $c = 1$, the value of the annuity, $A(T)$, becomes

$$A(T) = \frac{1}{y} \left(1 - \frac{1}{\left(1 + \frac{y}{2}\right)^{2T}} \right) \quad (3.15)$$

The annuity formula appears frequently in fixed income as the present value factor for a bond's coupons, a swap's fixed rate cash flows, or a mortgage's payments, which are most often structured as a series of equal payments. (See Chapter 20.)

A fifth implication of equation (3.14) is that the value of a *perpetuity*, a security that makes the fixed payment c forever, can be found by letting T approach infinity in (3.14) and multiplying by c , which gives $\frac{c}{y}$.

A sixth and final implication of the definition of yield is that if the term structure is flat, so that all spot rates and all forward rates equal some single rate, then the yield-to-maturity of all bonds equals that rate as well. This is easily seen by observing that, in the case of a flat spot rate curve, the pricing equation for each bond would take exactly the same form as equation (3.12) with the yield equal to the single spot rate.

Yield Curves and the Coupon Effect The phrase “yield curve” is used often, but its meaning is not very precise because the concept of yield is intertwined with the cash flows of a particular bond. Spot, forward, and par rate curves can, as shown in Chapter 2, be used to price any similar security. By contrast, the yield of a particular security derived from (3.14) can be used to price only that security. To illustrate this point, Figure 3.2, using C-STRIPS prices as of May 28, 2010, graphs the yields on hypothetical but fairly priced zero coupon bonds, par bonds, and 9% coupon bonds of various maturities on the mid-month, May–November cycle. In other words, using discount factors derived from C-STRIPS prices, the prices of these hypothetical bonds are computed along the lines of Chapter 1. Then the yields of these bonds are calculated. Figure 3.2 also shows the yields of actual U.S. Treasury notes and bonds on the same payment cycle and as of the same pricing date. Figure 3.3 shows the same data as Figure 3.2, but zooms in on a narrower yield range by focusing on the longer maturities.

These figures show that the “zero coupon yield curve,” the “par yield curve,” and the “9% coupon yield curve,” are indeed all different. In other words, a yield curve is not well defined until particular cash flows have been defined. And securities with a structure different from that of a coupon bond, like an amortizing bond or a fixed-rate mortgage, which spread principal payments out over time, would generate more dramatically different “yield curves.”

In Figures 3.2 and 3.3, for any given maturity, zero coupon yields exceed par yields, which, in turn, exceed the 9% coupon yields. This can be explained by the fact that yield is the one rate that describes how a security's




FIGURE 3.2 Yields of Hypothetical Securities Priced with C-STRIPS as of May 28, 2010

cash flows are being discounted. Since a zero coupon bond has only one cash flow at maturity, its yield is simply the spot rate corresponding to that maturity. A 9% coupon bond, on the other hand, makes cash flows every six months. Its yield, therefore, is a complex average of all of the spot rates from terms of six months to the bond's maturity, although the greatest weight is




FIGURE 3.3 Yields of Long-Term Hypothetical Securities Priced with C-STRIPS as of May 28, 2010

on the spot rate corresponding to the bond's largest present value, namely, that of the final payment of coupon plus principal. Furthermore, since the term structure of interest rates in the figures slopes upward, any weight this complex average places on the shorter-term spot rates lowers that average below the spot rate at maturity. Hence the yield on the 9% bond has to be lower than the yield on the 0% bond. The par bonds, with coupons between 0% and 9%, discount a lot of their present value at the shorter-term spot rates relative to zero coupon bonds, but discount little of their present value at those shorter-term rates relative to the 9% bonds. Hence, the yield of a par bond of a given maturity will be between the yield of the 0% and 9% bonds of that maturity. While not illustrated here, if the term structure slopes downward, then the argument just made would be reversed and the zero coupon yield curve would be below the 9%-coupon yield curve.

The fact that fairly priced bonds of the same maturity but different coupons have different yields-to-maturity is called the *coupon effect*. The implication of this effect is that yield is not a reliable measure of relative value. Just because one fixed income security has a higher yield than another does not necessarily mean that it is a better investment. Any such difference may very well be due to the relationship between the time pattern of the security's cash flows and the term structure of spot rates, as discussed in the previous paragraph.

The yields on the actual notes and bonds are seen most easily in Figure 3.3. Many of the bonds, particularly those of longer term, are closest to the 9% coupon yield curve because those bonds, having been issued relatively long ago when rates were much higher, do indeed have very high coupons. The 6½s of November 15, 2026, the 6⅛s of November 15, 2027, the 5¼s of November 15, 2028, and the 6¼s of May 15, 2030, are all easily seen in the figure to fall into this category. Other bonds, however, were issued more recently at lower coupons and trade closer to the par yield curve. The three bonds in the figure with longest maturities, which were issued relatively recently, fall into this category: the 4¼s of May 15, 2039, the 4⅜s of November 15, 2039, and the 4⅜s of May 15, 2040.

Japanese Simple Yield Before concluding the discussion of yield, it is noted here that Japanese government bonds are quoted on a *simple yield* basis. With a fiat price p per unit face amount, a coupon rate c , and a maturity in years, T , this simple yield, y , is given by $y = c/p + (1/T) \times (1 - p)/p$. So, for example, if $p = 101.45\%$, $c = 2\%$, and $T = 20$, then $y = 1.90\%$.

News Excerpt: Sale of Greek Government Bonds in March, 2010

At the end of March, 2010, investors around the world were concerned that Greece might not be able to meet all its debt obligations. At that time,

Bloomberg reported the following³ about the Greek government's sale of new, seven-year bonds:

Greece priced the 5 billion euros (\$6.7 billion) of seven-year bonds to yield 310 basis points more than the benchmark mid-swap rate, according to a banker involved in the transaction . . .

The bonds' 6 percent yield equates to 334 basis points more than seven-year German bunds, Europe's benchmark government securities. That compares with a yield premium, or spread, of 61 basis points for similar maturity Spanish debt and 114 basis points on Portugal's government bonds due 2017, according to composite prices on Bloomberg. Italy's seven-year bonds yield 45 basis points more than bunds, the prices show.

"Greece's borrowing costs exceed those of Spain and Portugal as it still needs to convince the market that it can roll over existing debt . . ."

COMPONENTS OF P&L AND RETURN

As stated in the introduction to this chapter, breaking down P&L or return into component parts is extremely useful for understanding how money is being made or lost in a trading book or investment portfolio. In addition, many sorts of errors can often be caught by a thorough analysis of *ex-post* profitability or loss.

For expositional ease, this section makes the following choices. First, it decomposes P&L; a return decomposition can then be found by dividing each P&L component by the initial price. Second, the P&L considered is that of a single bond trading at a single spread, but the analysis can be extended to more general portfolios and term structures of spreads. Third, the holding period is assumed to be equal to a coupon payment period. Appendix B of this chapter gives the P&L decomposition for holding periods both within and across coupon payment periods.

P&L is generated by price appreciation plus *cash-carry*, which consists of explicit cash flows like coupon payments and financing costs. This section decomposes price appreciation into three components and then presents a sample return decomposition. The next section focuses on one component of return, namely, carry-roll-down, in more detail.

Set the following notation:

- $P_t(\mathbb{R}_t, s_t)$: the price of a bond at time t , under term structure $\mathbb{R}_t$, and bond-specific spread s_t .

³"Greece Pays Bond Investors 5 Times Spain Yield Spread (Update1)," Bloomberg *BusinessWeek*, Thursday May 27, 2010.

- c : periodic coupon payment of the bond.
- $P_{t+1}(\mathbb{R}_{t+1}, s_{t+1})$: the price of the bond at $t + 1$, with the term structure and bond-specific spreads changing as indicated.
- $\mathbb{R}_{t+1}^e$: some term structure of rates that is not necessarily the term structure at time t or $t + 1$. The choice of this term structure will be discussed shortly.

The total price appreciation and a breakdown of that appreciation into its component parts can be defined as follows.⁴ Note that the sum of the component parts is, by design, identically equal to total price appreciation.

- *Total Price Appreciation*: $P_{t+1}(\mathbb{R}_{t+1}, s_{t+1}) - P_t(\mathbb{R}_t, s_t)$
- *Carry-Roll-Down*: $P_{t+1}(\mathbb{R}_{t+1}^e, s_t) - P_t(\mathbb{R}_t, s_t)$
- *Rate Changes*: $P_{t+1}(\mathbb{R}_{t+1}, s_t) - P_{t+1}(\mathbb{R}_{t+1}^e, s_t)$
- *Spread Change*: $P_{t+1}(\mathbb{R}_{t+1}, s_{t+1}) - P_{t+1}(\mathbb{R}_{t+1}, s_t)$

The first component of the decomposition, called carry-roll-down, is the price change due to the passage of time with rates moving “as expected,” from $\mathbb{R}_t$ to $\mathbb{R}_{t+1}^e$, and with no change in spread. Before proceeding further, however, it is worthwhile to explain the name carry-roll-down by discussing the generic concepts of carry and roll-down, which are invoked often in practice, but tend to generate some confusion.

Most generally, P&L due to carry is meant to convey how much a position earns due to the passage of time, holding everything else constant. A clean example is a par bond when the term structure is flat and unchanging: since the bond’s price is always par, its carry is clearly its coupon income minus its cost of financing. Another clean example is a premium bond when the term structure is, again, flat and unchanging. Since this bond’s price is pulled to par over time (see Figure 3.1), its carry is easily defined as its coupon income minus the decline in its price minus its cost of financing. Note, by the way, that the concept of carry just described, by including pull-to-par P&L, is broader than cash carry, defined earlier as coupon income minus financing costs. As to be discussed in Chapters 13 and 14, cash carry plays an important role in describing bond forward and futures prices.

P&L due to roll-down is meant to convey how much a position earns due to the fact that, as a security matures, its cash flows are priced at earlier points on the term structure. A clean example of this is the 10y6m forward highlighted in the case study of Chapter 2. At the time of that

⁴Defining the breakdown in a different order can change the allocation of the total price appreciation, but the magnitude of this change is usually very small except for securities with values that are very nonlinear in rates or spreads.

case, an investor might agree to lend EUR for six months, 10-years forward, at a rate of 4.254%. That trade has no carry in the sense of the previous paragraph: it pays no coupon, it costs nothing to finance, and, if the market rate of the forward trade remains at 4.254%, then its P&L is zero. But if at the time of the trade the 9y6m rate was 4.127%, then the trade would be said to have roll-down P&L in the following sense. If the term structure does not change, then, after a year, the 10y6m forward trade at 4.254% matures into a 9y6m forward with an appropriate market rate of 4.127%. Hence, the investor would gain the difference between 4.254% and 4.127%, or 12.7 basis points, because the forward trade had “rolled-down” the curve.

The examples in the previous two paragraphs clearly illustrate the concepts of carry and roll-down, but the division of P&L between the two often requires further calculation. Consider a premium bond when the term structure is upward-sloping and unchanging. The resulting P&L over time would be a combination of carry, i.e., pull-to-par plus coupon minus financing costs, and roll-down, as the bond’s cash flows are discounted at lower rates. While an investor could define some separation of this P&L into distinct carry and roll-down components, the separation would not be as clean as in the earlier examples and, more importantly, would probably not be worth the effort. From the perspective of understanding P&L over time, the more important objective is to separate out what happens to a position when rates move “as expected” from what happens as rates and spread change.

Taking all of these considerations into account, this book preserves a separate accounting for cash carry, i.e., coupon income minus financing costs, so as to be consistent with concepts in forward and futures markets. The remaining P&L due to the passage of time, i.e., the P&L due to the passage of time excluding cash carry, is called carry-roll-down. This name reflects the fact that carry-roll-down is a mix of P&L that might otherwise be classified as either carry or roll-down.

Returning then to the P&L decomposition given previously, carry-roll-down P&L is the price appreciation due to the bond’s maturing over the period and rates moving from the original term structure $\mathbb{R}_t$ to some hypothetical, “expected,” or intermediate term structure, $\mathbb{R}_{t+1}^e$. There are many possible choices for $\mathbb{R}_{t+1}^e$ and some common ones are discussed in the next section, but no choice clearly dominates another. In any case, note that carry-roll-down price appreciation assumes that the bond’s individual spread has not changed over the period. Also note that practitioners often calculate carry-roll-down in advance, that is, at time t they are interested in knowing the carry-roll-down from time t to time $t + 1$.

The price appreciation due to rate changes is the price effect of rates changing from the intermediate term structure, $\mathbb{R}_{t+1}^e$, to the term structure that actually prevails at time $t + 1$, namely $\mathbb{R}_{t+1}$. This component of price

appreciation is the subject matter of Part Two of this book. Note that spread is assumed unchanged here as well. Note also that price appreciation due to changes in rates might be calculated in advance as part of a scenario analysis, but is usually reserved for calculations done *ex-post* as part of realized return.

Finally, the price appreciation due to a spread change is the price effect due to the bond's individual spread changing from s_t to s_{t+1} . The spread is, in fact, the focus or bet of many trades. Is this U.S. Treasury too cheap relative to others? Is that corporate bond too expensive relative to swaps? Price appreciation due to a spread change, like that due to rate changes, may be calculated in advance as part of a scenario analysis or *ex-post* in the process of computing realized returns.

Note that dividing each of the components of price appreciation and then cash carry by the initial price, $P_t(\mathbb{R}_t, s_t)$, gives the respective components of bond return.

A Sample P&L Decomposition

This subsection works through an example of decomposing the return of the $\frac{3}{4}$ s of November 30, 2011, over the six months from May 28, 2010, to November 30, 2010. The example assumes that:

- The initial term structure and spreads are as in equation (3.9);
- The carry-roll-down scenario is *realized forwards*, which will be explained shortly;
- The term structure falls in parallel by 10 basis points over the six-month holding period;
- The bond's spread converges from its initial 4.4 basis points to 0 over the holding period.

Table 3.2 shows how forward rates and prices change from their initial values to the values in each step of the decomposition. The initial forwards used to price the $\frac{3}{4}$ s on May 28, 2010, given in row (i) of the table, are the sums of the initial base forwards on that date, row (ii), and the computed spread of the $\frac{3}{4}$ s on that date, row (iii). The price of the bond using these forwards and this spread is 100.190, given in the rightmost column of row (i). See equation (3.9). Rows (iv) through (xii) of the table describe the pricing of the $\frac{3}{4}$ s at the end of the holding period, on November 30, 2010.

The first price change, due to carry-roll-down, is presented in rows (iv) through (vi) of Table 3.2. The assumption of realized forwards means the following. As of the initial date, May 28, 2010, the forward rate curve in row (ii) "anticipated" a rate of .556% from November 30, 2010, to May 31, 2011, and a rate of 1.036% from May 31, 2011, to November 31, 2011.

TABLE 3.2 A Decomposition of the Price Appreciation for the $\frac{3}{4}$ s of November 30, 2011, over a Six-Month Holding Period

	Start Period End Period	5/30/10 11/30/10	11/30/10 5/31/11	5/31/11 11/30/11	Price
Pricing Date 5/28/10					
(i)	Initial Forwards	.193%	.600%	1.080%	100.190
(ii)	Term Structure	.149%	.556%	1.036%	
(iii)	Spreads	.044%	.044%	.044%	
Pricing Date 11/30/10					
(iv)	Carry-Roll-Down Forwards		.600%	1.080%	99.911
(v)	Term Structure		.556%	1.036%	
(vi)	Spreads		.044%	.044%	
(vii)	Rate-Change Forwards		.500%	.980%	100.011
(viii)	Term Structure		.456%	.936%	
(ix)	Spreads		.044%	.044%	
(x)	Spread-Change Forwards		.456%	.936%	100.054
(xi)	Term Structure		.456%	.936%	
(xii)	Spreads		.000	.000	

Then, six months later, these anticipated rates were realized: on November 30, 2010, the forward rate curve in row (v) is taken to be .556% in the first period and 1.036% in the second. The justification for the assumption of realized forwards will be described in the next section. Under these forwards in row (v), however, along with an unchanged spread of .044%, row (vi), the price of the now one-year bond is 99.911, given in the rightmost column of row (iv). Hence, the price appreciation due to carry-roll-down in this example is $99.911 - 100.190$ or $-.279$. (Of course, the bond paid a coupon on November 30, 2010, but that will be handled in the cash carry part of the calculations.)

The next price change, due to rate changes, is presented in rows (vii) through (ix). For this example it is assumed that all forward rates fell by 10 basis points. Therefore, the term structure of forwards falls from .556% and 1.036% in row (v) to .456% and .936% in row (viii). The spreads in row (ix) remain again at 4.4 basis points, so the new forwards for pricing the $\frac{3}{4}$ s in row (vii) are .500% and .980%. These new forwards give a bond price of 100.011 in the rightmost column of row (vii) and a price appreciation due to rate changes of $100.011 - 99.911$ or $.1$.

The final price change, due to the change of the spread from .044% to 0%, is presented in rows (x) through (xii). Keeping the new term structure in row (xi) the same as in row (viii) and using a zero spread in row (xii), the new forwards in row (x) are .456% and .936%, which gives a final

TABLE 3.3 Decomposition of P&L of the $\frac{3}{4}$ s of November 30, 2011, over a Six-Month Holding Period

	\$
Initial Price	100.190
Price Appreciation	−.136
Carry-Roll-Down	−.279
Rates	+.100
Spread	+.043
Cash Carry	.375
Coupon	.375
Financing	0.000
P&L	+.239

bond price of 100.054 in the rightmost column of row (x). Hence, the price appreciation due to spread change is $100.054 - 100.011$ or .043.

Table 3.3 summarizes the components of price appreciation and adds the coupon payment to complete the decomposition of gross dollar return. Were the position financed, the financing cost would be included in the carry so as to compute net dollar returns. Finally, these dollar returns can be divided by the initial price to obtain percentage returns, although, since the initial price is very near 100, percentage returns do not add much insight in this particular example.

CARRY-ROLL-DOWN SCENARIOS

When considering potential trades or investments, many practitioners want to calculate the dollar return of the trade or investment under the expectation or scenario of “no change” in rates. So the question with respect to carry-roll-down is, “What are good choices for no change scenarios?”

One common choice is to assume that forward rates equal expectations of future rates and that, as time passes, these forward rates are realized. So, for example, today’s six-month rate two years forward is the realized six-month rate two years from today. This realized forward assumption was used in the sample P&L decomposition of the previous section. A second common choice assumes that the entire term structure of interest rates remains unchanged over time. So, for example, today’s six-month rate two years forward will be the six-month rate two years forward a week from now, a month from now, a year from now, etc.

This section derives some implications of the realized forward and unchanged term structure assumptions, in addition to the related assumption

of unchanged yields. To conclude, the section considers one alternative assumption which, while conceptually attractive, is hardly used in practice.

Realized Forwards

Given the example of realized forwards in the previous section, this subsection proceeds directly to the mathematics. Recall the pricing equation of a bond in terms of forwards, omitting any spreads to the base curve:

$$P_0(\mathbb{R}_0) = \frac{c}{(1+f(1))} + \frac{c}{(1+f(1))(1+f(2))} + \cdots + \frac{1+c}{(1+f(1))(1+f(2))\cdots(1+f(T))} \quad (3.16)$$

Under the assumption of realized forwards, the price of the bond after one period becomes

$$P_1(\mathbb{R}_1) = \frac{c}{(1+f(2))} + \frac{c}{(1+f(2))(1+f(3))} + \cdots + \frac{1+c}{(1+f(2))(1+f(3))\cdots(1+f(T))} \quad (3.17)$$

Combining equations (3.16) and (3.17) it is easy to see that

$$\frac{P_1(\mathbb{R}_1) + c - P_0(\mathbb{R}_0)}{P_0(\mathbb{R}_0)} = f(1) \quad (3.18)$$

In words, equation (3.18) says that the gross, single-period return of any security is the prevailing one-period rate. A two-year bond and a 10-year bond, over the next period, both earn the short-term rate. As will be made clear in Chapter 8, this result and the underlying assumption of realized forwards is not particularly satisfying. It is more common to assume that, since the 10-year bond has more interest rate risk than the two-year bond (see Part Two), investors demand a higher return for the 10-year bond. In any case, under the reasonable assumption that the one-period financing rate is $f(1)$, subtracting this rate from the gross return in (3.18) shows that the single-period, net return of any security is 0.

In a similar manner it is easy to show that in the presence of a term structure of spreads, i.e., with price given by (3.10), the one-period gross return of a bond in the case of realized forward rates and spreads is $f(1) + s(1)$, i.e., the short-term rate plus the short-term spread.

The gross return under the realized forward assumption can be calculated over many periods as well. In general, it can be shown that the return

to maturity under realized forwards is

$$\begin{aligned} & \frac{c(1+f(2))(1+f(3))\cdots(1+f(T))+\cdots}{P_0(\mathbb{R}_0)} + \frac{P_T(\mathbb{R}_T) - P_0(\mathbb{R}_0)}{P_0(\mathbb{R}_0)} \\ &= (1+f(1))(1+f(2))\cdots(1+f(T)) - 1 \end{aligned} \quad (3.19)$$

In words, the return to a bond held to maturity under the assumption of realized forward rates is the same as rolling a \$1 investment one period at a time at those forward rates.

The discussion of this subsection has interesting implications in the answer to the following question. Which of the following two strategies is more profitable, rolling over one-period bonds or investing in a long-term bond and reinvesting coupons at prevailing short-term rates? As just demonstrated, if forward rates are realized, the two strategies are equally profitable. But if realized forwards are greater than the forwards implicit in the initial bond price, rolling over one-period bonds is more profitable. And if realized forwards are less than those implicit in the initial bond price, investing in the long-term bond is more profitable. Hence, the decision to roll short-term investments or to purchase long-term bonds depends on how the decision maker's forecast of rates compares with market forward rates. Note, however, that while this reasoning provides a good deal of intuition about the returns of short- *versus* long-term bonds, it says nothing about the more realistic case of some forwards being realized above the initial forwards and some being realized below.

Unchanged Term Structure

A very common carry-roll-down assumption is that the term structure stays unchanged. If the six-month rate two years forward is 1.25% today, then, six months from now, the six-month rate two-years forward will still be 1.25%. Under this assumption, the prices of a bond today and after one period are

$$\begin{aligned} P_0(\mathbb{R}_0) &= \frac{c}{(1+f(1))} + \frac{c}{(1+f(1))(1+f(2))} + \cdots \\ &\quad + \frac{1+c}{(1+f(1))(1+f(2))\cdots(1+f(T))} \end{aligned} \quad (3.20)$$

$$\begin{aligned} P_1(\mathbb{R}_1) &= \frac{c}{(1+f(1))} + \frac{c}{(1+f(1))(1+f(2))} + \cdots \\ &\quad + \frac{1+c}{(1+f(1))(1+f(2))\cdots(1+f(T-1))} \end{aligned} \quad (3.21)$$

Combining the two equations (3.20) and (3.21) reveals the one-period gross return under the assumption of an unchanged term structure:

$$\frac{P_1(\mathbb{R}_1) + c - P_0(\mathbb{R}_0)}{P_0(\mathbb{R}_0)} = \left[\frac{f(T) - c}{(1 + f(1)) \cdots (1 + f(T))} + c \right] \frac{1}{P_0(\mathbb{R}_0)} \quad (3.22)$$

While equation (3.22) does not have as neat an interpretation as the analogous equation for realized forwards, it does make the following point. The gross return under the assumption of an unchanged term structure depends most crucially on the last relevant forward rate, that is the forward rate from one-period before maturity to maturity, *versus* the bond's coupon rate. The intuition for this result parallels the discussion in the "Maturity and Price or Present Value" subsection of Chapter 2. Finally, it is easy to show that in the presence of a term structure of spreads, the relevant quantity for determining the return becomes $f(T) + s(T) - c$.

The realized forward assumption implicitly assumes that there is no risk premium built into forward rates. The unchanged term structure implicitly assumes the opposite extreme. If the term structure slopes upward on average and yet remains unchanged on average, it must be that the upward-sloping shape is completely explained by investors' requiring a risk premium that increases with term.

Unchanged Yields

Yet another carry-roll-down assumption is that a bond's yield remains unchanged. This assumption is useful not so much for explicit carry-roll-down calculations but for interpreting yield-to-maturity as a measure of return. The bond pricing equation in terms of yield, equation (3.12) without the explicit semiannual payment convention, is

$$P_0(\mathbb{R}_0) = \frac{c}{(1+y)} + \frac{c}{(1+y)^2} + \cdots + \frac{1+c}{(1+y)^T} \quad (3.23)$$

Under the assumption of unchanged yields,

$$P_1(\mathbb{R}_1) = \frac{c}{(1+y)} + \frac{c}{(1+y)^2} + \cdots + \frac{1+c}{(1+y)^{T-1}} \quad (3.24)$$

And, along the lines of the previous subsections, combine equations (3.23) and (3.24) to see that

$$\frac{P_1(\mathbb{R}_1) + c - P_0(\mathbb{R}_0)}{P_0(\mathbb{R}_0)} = y \quad (3.25)$$

In words, the one-period gross return, assuming that yield remains unchanged, is the yield. It is in this sense that an investor in a bond earns its yield-to-maturity.

Extending this analysis to many periods, it can be shown that, under the assumption of constant yields,

$$\begin{aligned} & \frac{c(1+y)^{T-1} + c(1+y)^{T-2} + \cdots}{P_0(\mathbb{R}_0)} + \frac{P_T(\mathbb{R}_T) - P_0(\mathbb{R}_0)}{P_0(\mathbb{R}_0)} \\ &= (1+y)^T - 1 \end{aligned} \quad (3.26)$$

In words, an investor to maturity earns the bond's yield in the sense that, if the yield does not change and if all coupons are reinvested at that yield, then the return of the bond to maturity equals the return of rolling over a \$1 investment period-by-period at that yield. Now while this sounds similar to the statement made in the context of realized forwards, the unchanged yield scenario is even less satisfying. The assumptions that yield stays unchanged over the life of a bond and that all coupons can be reinvested at that same yield are particularly flawed: the fact that there is a term structure of interest rates implies that a bond's yield will change with maturity and that single-period reinvestment rates should not equal bond yield. The unchanged yield assumption is less problematic for these reasons if the term structure is always flat, but that condition is quite unrealistic as well.

Expectations of Short-Term Rates are Realized

A more conceptually appealing scenario for computing carry-roll-down is that expectations of short-term rates are realized. This is much more difficult to implement than the other scenarios presented in this section because an investor has to specify expectations of rates in the future and then describe how forwards rates are formed relative to those expectations. A framework of this sort is presented in Chapter 8. The outcome, arguably more sensible than others in this section, is that the expected return of a bond, which is not the same as the roll-down return,⁵ is equal to the short-term rate plus a risk premium that depends on the riskiness of the bond.

APPENDIX A: YIELD ON SETTLEMENT DATES OTHER THAN COUPON PAYMENT DATES

To keep the presentation of ideas simple, the "Yield-to-Maturity" subsection earlier in this chapter considered only settlement dates that fall on coupon

⁵The expected return of a bond is not the same as the return of the bond should rates evolve according to expectation. Mathematically, a price at expected rates is not equal to the expected price.

payment dates. This appendix gives the formula for yield-to-maturity when the settlement date does not fall on a coupon payment date. The definition of yield is expressed in the text as equation (3.13) or (3.14):

$$P = \frac{c}{2} \sum_{t=1}^{2T} \frac{1}{(1 + \frac{y}{2})^t} + \frac{1}{(1 + \frac{y}{2})^{2T}} \quad (3.27)$$

$$P = \frac{c}{y} \left(1 - \frac{1}{(1 + \frac{y}{2})^{2T}} \right) + \frac{1}{(1 + \frac{y}{2})^{2T}} \quad (3.28)$$

Equation (3.27) has to change in two ways to take account of a settlement date between coupon dates. First, price has to be interpreted to be the full price of the bond. See the “Accrued Interest” section of Chapter 1. Second, the exponents of equation (3.27) have to be adjusted to reflect the timing of the cash flows. When the coupon payments arrive in semiannual intervals, then, following the semiannual compounding convention, the first payment is discounted by dividing by $1 + \frac{y}{2}$, the second by dividing by $(1 + \frac{y}{2})^2$, etc. But what if the first payment is paid in a fraction τ of a semiannual period? (If the next coupon were paid in five months, for example, then $\tau = \frac{5}{6}$.)⁶

Market convention for the purpose of calculating yield (which cannot really be justified in terms of the logic of semiannual compounding) is to discount the next coupon payment by

$$\frac{1}{(1 + \frac{y}{2})^\tau} \quad (3.29)$$

and a subsequent payment i semiannual periods later by

$$\frac{1}{(1 + \frac{y}{2})^{\tau+i}} \quad (3.30)$$

Under this convention, the price-yield equation for a bond making its next payment in a fraction τ of a semiannual period and then making $2T - 1$ subsequent semiannual payments is

$$P = \frac{c}{2} \sum_{t=0}^{2T-1} \frac{1}{(1 + \frac{y}{2})^{\tau+t}} + \frac{1}{(1 + \frac{y}{2})^{\tau+2T-1}} \quad (3.31)$$

⁶More accurately, τ would be calculated with the day-count convention appropriate for the security in question.

Finally, applying the summation formula in Appendix D of Chapter 2 to (3.31) in order to derive the generalization of equation (3.28) gives the relatively simple

$$P = \left(1 + \frac{y}{2}\right)^{1-\tau} \left[\frac{c}{y} \left(1 - \frac{1}{\left(1 + \frac{y}{2}\right)^{2T}}\right) + \frac{1}{\left(1 + \frac{y}{2}\right)^{2T}} \right] \quad (3.32)$$

APPENDIX B: P&L DECOMPOSITION ON DATES OTHER THAN COUPON PAYMENT DATES

For ease of exposition, the text assumed that dates t and $t + 1$ are both coupon payment dates. To generalize the P&L decomposition, this appendix allows these dates to fall between coupon payment dates. The notation of the text continues here, with the following qualifications and additions. Let P_t denote the full price of a bond, p_t denote its quoted price, and $AI(t)$ denote its accrued interest, so that $P_t = p_t + AI(t)$. The coupon rate is c , as in the text, and let the financing rate be r . Finally, let there be d days between dates t and $t + 1$.

Begin with the case in which there is no coupon paid between dates t and $t + 1$. Then the total P&L of a bond, including the cost of financing the full price of the bond for d days, is

$$P_{t+1}(\mathbb{R}_{t+1}, s_{t+1}) - P_t(\mathbb{R}_t, s_t) \left(1 - \frac{rd}{360}\right) \quad (3.33)$$

Using the breakdown of full price into quoted price plus accrued interest and rearranging terms, the P&L becomes

$$p_{t+1}(\mathbb{R}_{t+1}, s_{t+1}) - p_t(\mathbb{R}_t, s_t) + AI(t+1) - AI(t) - P_t(\mathbb{R}_t, s_t) \frac{rd}{360} \quad (3.34)$$

Applying the breakdown in the text to the quoted price appreciation in (3.34) gives

$$\begin{aligned} & [p_{t+1}(\mathbb{R}_{t+1}^e, s_t) - p_t(\mathbb{R}_t, s_t)] + [p_{t+1}(\mathbb{R}_{t+1}, s_t) - p_{t+1}(\mathbb{R}_{t+1}^e, s_t)] \\ & + [p_{t+1}(\mathbb{R}_{t+1}, s_{t+1}) - p_{t+1}(\mathbb{R}_{t+1}, s_t)] \\ & + \left[AI(t+1) - AI(t) - P_t(\mathbb{R}_t, s_t) \frac{rd}{360} \right] \end{aligned} \quad (3.35)$$

The P&L terms of (3.35) are, in order, the contributions due to carry-roll-down, rates, spread, and cash carry.

In the case that there is a coupon payment between dates t and $t + 1$, then, ignoring the second order amount of interest on the coupon payment from its payment date to $t + 1$, the P&L expression (per unit face amount) changes only with the cash carry term in (3.35) changing to

$$\frac{c}{2} + AI(t + 1) - AI(t) - P_t(\mathbb{R}_t, s_t) \frac{rd}{360} \quad (3.36)$$

Note, however, that in (3.35), $AI(t + 1) > AI(t)$ because there is no coupon paid between t and $t + 1$. By contrast, in (3.36), $AI(t + 1)$ may be greater or less than $AI(t)$ depending on where the two dates fall in the coupon cycle.

Measures of Interest Rate Risk and Hedging

The interest rate risk of a security is measured by how much its price changes as interest rates change. Not surprisingly, measures of interest rate risk are used routinely in fixed income markets. To hedge the interest rate risk of one set of securities with other securities, traders have to compute how the price of each security responds to changes in rates. To take a view on the level or the shape of the term structure of interest rates in the future, investors have to determine how securities perform under various interest rate scenarios. To ensure that a portfolio of assets can continue to support a portfolio of liabilities, asset-liability managers have to compare the interest rate risks of the two portfolios. Lastly, to carry an appropriate amount of risk relative to a mandate or charter, risk managers need to be able to compute the volatility of fixed income portfolios.

Computing the price change of a security given a change in interest rates is straightforward. For example, given an initial and a shifted spot rate curve, the tools of Part One can be used to calculate the price change of securities with fixed cash flows, and the models of Part Three can be used to calculate the price change of interest rate derivatives. Therefore, the challenge of measuring price sensitivity comes not so much from the computation of price changes given changes in interest rates but in defining what is meant by changes in interest rates.

One of the simplest formulations of a change in interest rates is that rates of every term move up or down by the same amount. More general *one-factor* formulations assume that changes in the rates of all terms are fully determined by the change in a single interest rate factor, e.g., the 10-year par rate. *Multi-factor* formulations assume that changes in all rates are a

function of two or more factors. A popular two-factor approach posits that a “level” factor and a “slope” factor are sufficient to describe term structure movements while popular approaches using many more factors take each forward rate as a separate interest rate factor.

Simple assumptions about how interest rates change result in risk metrics that are intuitive and easy to use, but may not accurately capture the dynamics of term structure behavior. For example, yield-based *DV01*, which assumes that all yields move up and down in parallel, is easy to compute, easy to understand, and only requires one security to hedge the risk of a large portfolio of securities. But the resulting hedge does not protect against *curve risk*, that is, changes in the slope of the term structure. On the other hand, more complex, multi-factor formulations capture term structure dynamics more accurately but are more challenging to use and require several securities to hedge a given portfolio. In the end, practitioners choose the simplest model that is appropriate for the application at hand.

Chapter 4 introduces one-factor risk metrics. The first part of the chapter presents one-factor metrics without explicitly describing how the term structure changes as a function of that one factor. The point of this expository approach is to distinguish the concepts underlying risk metrics from the models that describe how the term structure changes. Then, with this distinction established, Chapter 4 continues with the popular yield-based *DV01*, duration, and convexity measures of interest rate risk. More sophisticated models of term structure behavior are deferred to Part Three.

Chapter 5 describes risk metrics that divide the term structure into several regions, make very simple assumptions about how rates change within each region, and then measure risk with respect to each region separately. *Key-rate analysis*, for example, measures a portfolio’s exposure to changes in several “key” rates, e.g., the two-year, five-year, 10-year, and 30-year rates. Because each exposure is measured and then hedged separately, there is no need to make assumptions about how the key rates change relative to one another. However, in this example, four securities would be needed to hedge any fixed income portfolio. Key rates exposures are used mostly for bond portfolios. For portfolios of swaps and portfolios with both bonds and swaps, however, *partial PV01s* and *forward bucket ’01s* tend to be used instead. These approaches are similar in spirit to key rates, but divide the term structure into many more regions.

Chapter 6 turns to methods that rely directly on data and empirical analysis to describe changes in rates and to construct appropriate hedges. While presented in a separate chapter, empirical approaches are not completely distinct from the methods of Chapters 4, 5, and Part Three. Sensible practitioners use yield-based *DV01* to hedge only when it is empirically reasonable to do so. Similarly, the assumptions of any good factor model are

consistent with empirical regularities of the data that are expected to persist in the future. The particular methods described in Chapter 6 are single-factor regression hedging, two-factor regression hedging, and principal components analysis. Substantial effort has been made to enable a reader to appreciate the power of these empirical approaches with a minimum amount of mathematics.

One-Factor Risk Metrics and Hedges

This chapter presents some of the most important concepts used to measure and hedge risk in fixed income markets, namely, $DV01$, duration, and convexity. These concepts are first presented in a very general, one-factor framework, meaning that the only significant assumption made about how the term structure changes is that all rate changes are driven by one factor. As an application used to illustrate concepts, the chapter focuses on a market maker who shorts futures options and hedges with futures, although the reader need not know anything about futures at this point.

The chapter then presents the yield-based equivalents of these more general concepts, i.e., yield-based $DV01$, duration, and convexity. Because these can be expressed through relatively simple formulas, they are very useful for building intuition about the interest rate risk of bonds and are widely used in practice. They cannot, however, be applied to securities with interest-rate contingent payoffs, like options.

The chapter concludes with an application in which a portfolio manager is deciding whether to purchase duration in the form of a bullet or barbell portfolio. As it turns out, the choice depends on the manager's view on future interest rate volatility.

DV01

Denote the price-rate function of a fixed income security by $P(y)$, where y is an interest rate factor. Despite the usual use of y to denote a yield, this factor might be a yield, a spot rate, a forward rate, or a factor in one of the models of Part Three. In any case, since this chapter describes one-factor measures of price sensitivity, the single number y completely describes the term structure of interest rates.

This chapter uses three securities, with prices as of May 28, 2010, to illustrate concepts: the U.S. Treasury $4\frac{1}{2}\%$ s of May 15, 2017; the 10-year U.S.




FIGURE 4.1 $4\frac{1}{2}$ s of 5/15/2017 and TYU0 Price-Rate Curves as of May 28, 2010

note futures contract maturing in September 2010, whose ticker is TYU0; and a call option on TYU0 with a strike of 120 and a maturity of August 27, 2010, whose ticker is TYU0C 120. For the purposes of this chapter, the reader need not know anything about futures and futures options, which are covered in Part Four. As mentioned in the introduction to Part Two, understanding the interest rate risk of a security from its price-rate function can be separated from the creation of that price-rate function. For completeness, however, it is noted here that the price-rate curves of the three illustrative securities were created using a particular calibration of the Vasicek model, described in Chapter 9 and applied to futures in Chapter 14.

Figure 4.1 graphs three price-rate curves as a function of a (hypothetical) seven-year U.S. Treasury par rate, which, on the pricing date, was 2.77%. The three curves are for TYU0, for 100 notional amount of the $4\frac{1}{2}$ s, and for an adjusted notional amount of the $4\frac{1}{2}$ s which, because of the technicalities of the futures contract, is more comparable to TYU0.¹ This adjusted notional position is included in Figure 4.1 to highlight the difference between the shape of a bond's price-rate curve and that of a futures contract. The price-rate curve of the $4\frac{1}{2}$ s is typical of all coupon bonds; it decreases with rates and is very slightly convex,² though that is hard to see from this figure. The

¹The notional amount is 100 divided by the conversion factor of the bond for delivery into TYU0. See Chapter 14.

²A line connecting any two points of a convex curve lies above the curve over that region.




FIGURE 4.2 TYU0C 120 Price-Rate Curve of as of May 28, 2010

price rate curve of TYU0 is typical of futures, decreasing with rates but with both convex and concave³ regions. The convex region is to the left of the graph, for low values of rates, while the concave region is to the right of the graph, most easily recognized in contrast with the convexity of the two bond curves over that same region.

Figure 4.2 graphs the price-rate curve of TYU0C 120. Its shape is typical for a call option on a fixed income security, decreasing to zero as rates increase and highly convex between a decreasing linear segment on the left and a flat, zero-valued segment on the right.⁴

The price-rate curves in Figures 4.1 and 4.2 can be used to compute the price sensitivities of the three securities with respect to interest rates. From Figure 4.2, for example, if rates rise 10 basis points from .95% to 1.05%, the price of the option falls from 13.550 to 12.755, for a slope of $\frac{13.550-12.755}{1.05\%- .95\%}$, which is -795 or -7.95 cents per basis point. If rates rise from 2.45% to 2.55% the same option falls in price from 3.096 to 2.622, for a slope of -474 or -4.74 cents per basis point. And finally, if rates rise from 3.45% to 3.55% the option falls from .310 to .225, for a slope of -85 or $-.85$ cents per basis point. The fact that price sensitivity changes as rates change will be explored in later sections.

³The line connecting any two points of a concave curve lies below the curve over that region.

⁴The typical shape of an option price-price curve is a hockey stick increasing to the right. Figure 4.2, however, is a price-rate curve.

To define a measure of interest rates more generally, let ΔP and Δy denote the changes in price and rate, respectively, and note that the change in rate measured in basis points is $10,000 \times \Delta y$. Then, consider the following measure of price sensitivity:

$$DV01 \equiv -\frac{\Delta P}{10,000 \times \Delta y} \quad (4.1)$$

DV01 is an acronym for *dollar value of an 01* (i.e., of .01%) and gives the change in the value of a fixed income security for a one-basis point decline in rates. The negative sign defines DV01 to be positive if price increases when rates decline and negative if price decreases when rates decline. This convention has been adopted so that DV01 is positive most of the time: all fixed coupon bonds and most other fixed income securities do rise in price when rates decline.

In the discussion of Figure 4.2, the slope of the call is estimated using pairs of option prices valued at rates which are 10 basis points apart: the points (.95%, 13.550) and (1.05%, 12.755) are used to provide an estimate of the slope at a rate of 1%, the points (2.45%, 3.096) and (2.55%, 2.622) are used to provide an estimate at a rate of 2.5%, etc. Since the slope of the call does change with rates, using points closer together, e.g., at 2.49% and 2.51% for an estimate of the slope at 2.50%, would—so long as the price of the call can be computed accurately enough—give a more precise estimate of the slope at a single point on the curve. In the limit of moving these points together, the estimation gives the slope of the line *tangent* to the price-rate curve at the chosen rate level. Figure 4.3 graphs two such tangent lines to TYU0C 120, one tangent at 2.50% and one at 3.50%. That the former is steeper than the latter shows that the option is more sensitive to rates at 2.50% than it is at 3.50%.

In the calculus, the slope of the tangent line at a particular rate level is called the *derivative* of the price-rate function at that rate and is denoted $\frac{dP}{dy}$. In some special cases, e.g., the yield-based metrics discussed later in this chapter or certain model-based metrics of Part Three, the derivative of the price-rate function can be written in closed form, i.e., as a relatively simple mathematical formula. In other cases it has to be calculated numerically as in the calculations for TYU0C 120 shown previously. In either case, in terms of the derivative, equation (4.1) for DV01 becomes

$$DV01 \equiv -\frac{1}{10,000} \frac{dP}{dy} \quad (4.2)$$

Before closing this section, a note on terminology is in order. Most market participants use DV01 to mean yield-based DV01, which is discussed later in this chapter. Yield-based DV01 assumes that the yield-to-maturity of a particular security changes by one basis point while, in the general




FIGURE 4.3 Tangent Lines at 2.50% and 3.50% to the TYU0C 120 Price-Rate Curve as of May 28, 2010

definition of $DV01$ in this section, some factor changes by one basis point, which then propagates in some way across the rest of the term structure. To avoid confusion, some market participants have different names for $DV01$ measures according to the assumed change in rates. For example, the change in price after a parallel shift in forward rates might be called $DVDF$ or $DPDF$ while the change in price after a parallel shift in spot or zero coupon rates might be called $DVDZ$ or $DPDZ$.⁵

A HEDGING APPLICATION, PART I: HEDGING A FUTURES OPTION

Say that in the course of business on May 28, 2010, a market maker sells \$100 million face amount of the option, TYU0C 120, when the seven-year par rate used in the figures of the previous section is 2.77%. How might the market maker hedge the resulting interest rate exposure by trading in the underlying futures contract, TYU0?⁶ Since the market maker has sold

⁵The term $PV01$ will be discussed in the next chapter.

⁶For expositional reasons this application is somewhat contrived. Since futures options are traded on exchanges, a broker-dealer would, in reality, act as an agent to purchase TYU0C for a customer's account rather than act as a principal to sell the option to a customer from its own account. Over-the-counter derivatives, on the other hand, would be more strictly consistent with the spirit of the application.

TABLE 4.1 Selected Model Prices and *DV01*s for TYU0 and TYU0C 120 as of May 28, 2010

7-Year Par Rate	TYU0	<i>DV01</i>	TYU0C 120	<i>DV01</i>
2.72%	120.0780		1.9194	
2.77%	119.7061	.07442	1.7383	.03505
2.82%	119.3338		1.5689	

the option and stands to lose money if rates fall, purchasing futures can hedge the resulting exposure. The *DV01* of the two securities can be used to figure out exactly how many futures should be bought against the short option position.

Table 4.1 gives selected price-rate pairs for TYU0 and for TYU0C 120 along with a calculated *DV01*. Note that, along the lines of the previous section, the calculated *DV01* at 2.77% uses the prices at rates of 2.72% and 2.82%, but not the price at 2.77% itself. In any case, let F be the face amount of futures the market maker needs to hedge the \$100 million short option position. Then, set F such that, after a one basis-point decline in rates, the change in the price of the hedge position plus the change in the price of the option position equals zero. Mathematically,

$$F \frac{.07442}{100} - 100,000,000 \times \frac{.03505}{100} = 0 \quad (4.3)$$

There is a negative sign in front of the second term on the left-hand side because the option position is short \$100 million. Also, since *DV01* values quoted in the text and shown in the figures are for 100 face amount, they have to be divided by 100 before being multiplied by face amounts. Rearranging terms of (4.3) shows that

$$F = 100,000,000 \times \frac{.03505}{.07442} \quad (4.4)$$

Solving (4.4) for F , the market maker should purchase \$47.098 million face amount of TYU0.

To summarize this hedging strategy, the change in value of the short option position for each basis point decline in rates is

$$-\$100,000,000 \times \frac{.03505}{100} = -\$35,050 \quad (4.5)$$

The change in the value of the hedge, the \$47 million face amount of TYU0, offsets this loss:

$$\$47,098,000 \times \frac{.07442}{100} = \$35,050 \quad (4.6)$$

Generally, if $DV01$ is expressed in terms of a fixed face amount, hedging a position of F^A face amount of security A requires a position of F^B of security B where

$$F^B = -\frac{F^A \times DV01^A}{DV01^B} \quad (4.7)$$

To avoid careless trading mistakes, it is worth emphasizing the simple implications of equation (4.7), assuming for the moment that, as usually is the case, each $DV01$ is positive. First, hedging a long position in security A requires a short position in security B and *vice versa*. In the example, the market maker sells futures options and buys futures. Second, the security with the higher $DV01$ is traded in smaller quantity than the security with the lower $DV01$. In the example, the market maker buys only \$47.098 million futures against the sale of \$100 million options.

There are securities for which $DV01$ is negative, most notably in mortgage derivatives. See Chapter 20. Hedging such a security with a positive- $DV01$ security would, by (4.7), require both sides of the trade to be long or short.

Return to the market maker who sells \$100 million of TYU0C 120 and buys \$47.098 million TYU0 when rates are 2.77%. Using the prices in Table 4.1, the value of the hedged position immediately after the trades is

$$-\$100,000,000 \times \frac{1.7383}{100} + \$47,098,000 \times \frac{119.7061}{100} = \$54,640,879 \quad (4.8)$$

Now say that rates fall by 5 basis points to 2.72%. Using the prices in Table 4.1 at the new rate level, the value of the position becomes

$$-\$100,000,000 \times \frac{1.9194}{100} + \$47,098,000 \times \frac{120.0780}{100} = \$54,634,936 \quad (4.9)$$

The hedge has succeeded in that the value of the position has hardly changed even though rates have changed.

To avoid misconceptions about market making, note that the market maker in this example makes no money. In reality, the market maker would

sell the options at some premium to their fair value. Taking half a tick, for example, the market maker would take an immediate value gain of half of $\frac{1}{32}$ or .015625 on the \$100 million options for a total of \$15,625. This spread compensates the market maker for executing the original trade and for managing the hedge of the position over the time. Some of the challenges of hedging the option after the initial trade are discussed in the continuation of this application later in this chapter.

DURATION

DV01 measures the dollar change in the value of a security for a basis point change in interest rates. Another measure of interest rate sensitivity, *duration*, measures the percentage change in the value of a security for a unit change in rates. Mathematically, letting D denote duration,

$$D \equiv -\frac{1}{P} \frac{\Delta P}{\Delta y} \quad (4.10)$$

As in the case of *DV01*, when an explicit formula for the price-rate function is available, the derivative of the price-rate function may be used for the change in price divided by the change in rate:

$$D \equiv -\frac{1}{P} \frac{dP}{dy} \quad (4.11)$$

Otherwise, prices at various rates must be substituted into (4.10) to estimate duration.

Table 4.2 gives the same rate levels and prices as Table 4.1 but computes duration instead of *DV01*. Once again, rates above and below the rate level in question are used to compute changes. The duration of TYU0 at 2.77% is given by

$$D = -\frac{1}{119.7061} \frac{(119.3338 - 120.0780)}{2.82\% - 2.72\%} = 6.217 \quad (4.12)$$

TABLE 4.2 Selected Model Prices and Durations for TYU0 and TYU0C 120 as of May 28, 2010

7-Year Par Rate	TYU0	Duration	TYU0C 120	Duration
2.72%	120.0780		1.9194	
2.77%	119.7061	6.217	1.7383	201.6
2.82%	119.3338		1.5689	

One way to interpret the duration number of 6.217 is to multiply both sides of definition (4.10) by Δy :

$$\frac{\Delta P}{P} = -D\Delta y \quad (4.13)$$

In the case of TYU0, equation (4.13) says that the percentage change in price equals -6.217 times the change in rate. Therefore, a one-basis point decrease in rate will result in a percentage price change of $6.217 \times .0001$ or $.06217\%$. Since the price of TYU0 at 2.77% is 119.7061 , this percentage change translates into a dollar change of $.06217\% \times 119.7061$ or $.07442$ per basis point, which is, of course, the $DV01$ of the futures at that rate level.

When speaking about duration, it is conventional to normalize for a 100 basis-point change in rates. In the present case, for example, practitioners would say that TYU0's price changes by 6.217% for a 100 basis-point change in rates. This is a convention of language, not of practice, because duration, like $DV01$, changes with the level of rates so that the actual price change for a move as large as 100 basis points will not be particularly well approximated by 6.217% .

Duration tends to be more convenient than $DV01$ in the investing context, as opposed to the trading context. If an institutional investor has funds to invest when rates are 2.77% , the fact that the duration of TYU0C 120 vastly exceeds that of TYU0 alerts the investor to the far greater risk of investing money in options. With a duration of 6.215 , the funds invested in TYU0 will change in value by about $.62\%$ for a 10-basis point change in rates. However, with a duration of 201.381 , the same funds invested in the option will gain or lose about 20.1% for the same 10-basis point change in rates!

By contrast, in a trading or hedging problem percentage changes are not particularly useful because the dollar amounts of the two sides of the trade are usually not the same. In the example of the previous section, the market maker sells options worth about $\$1.74$ million and buys futures with a bond-equivalent value of $\$56.38$ million. Hence it is much more useful to compute the dollar sensitivity of each position, as in equations (4.5) and (4.6).

Another difference between $DV01$ and duration is their behavior as rates change. Figure 4.3 showed that the $DV01$ of TYU0C 120 decreases as rates increase. As it turns out, however, the duration of the option increases as rates increase because the value of the option, which appears in the denominator of the definition of duration, decreases rapidly with rates. For example, at a rate of 2.77% , the option's $DV01$ is $.0351$ (Table 4.1) and its duration is 201.6 (Table 4.2). At a rate of 3.50% , however, the $DV01$ (calculated earlier in this chapter) is lower, at $.0085$, while its duration is higher, at $.0085 \times 10,000/.265$ or about 321 , where $.265$ is the option price at 3.50% .

Like the section on *DV01*, this section closes with a note on terminology. As defined in this chapter, duration may be computed for any assumed change in the term structure of interest rates. This very general definition is sometimes also called *effective duration*. In any case, note that when using the term duration many market participants mean yield-based duration, which is discussed later in this chapter.

CONVEXITY

As first mentioned in the discussion of Figure 4.3, interest rate sensitivity changes with the level of rates. *Convexity* measures this sensitivity. To start the discussion, Figure 4.4 graphs the *DV01* of the adjusted notional amount of the 4½s of May 15, 2017, TYU0, and TYU0C 120, all as a function of the level of rates. The *DV01* of the bond declines relatively gently as rates rise. The *DV01* of the futures changes gently as well, although it first declines with rates, then increases, and then declines again. (This shape is usual for futures contracts and will be explained in Chapter 14.) Finally, the *DV01* of the futures option declines gradually or steeply, depending on the level of rates.

Mathematically, convexity is defined as

$$C \equiv \frac{1}{P} \frac{d^2 P}{dy^2} \quad (4.14)$$




FIGURE 4.4 DV01-Rate Curves for the Adjusted Notional of the 4½s of 5/15/2017, TYU0, and TYU0C 120 as of May 28, 2010

where the second multiplicand is the *second derivative* of the price-rate function. While the first derivative measures how price changes with rates, the second derivative measures how the first derivative changes with rates. As with DV01 and duration, if there is an explicit formula for the price-rate function then (4.14) may be used to compute convexity. Without such a formula, convexity must be estimated numerically.

Tables 4.3, 4.4, and 4.5 show how to estimate the convexity of the adjusted notional of the 4½s, TYU0, and TYU0C 120, respectively, at three rate levels, namely, 1.77%, 2.77%, and 3.77%. Prices have been recorded to three decimal places, but calculations have been performed using greater accuracy. (This does make a difference in the calculations of second derivatives which divide twice by a small number, namely, by the .05% difference between rates.)

The convexity of the futures contract at 1.77%, as reported in Table 4.4, is estimated as follows. Start by estimating the first derivative between 1.72% and 1.77%, i.e., at 1.745%, by dividing the change in price by the change in rate:

$$\frac{127.172545 - 127.552549}{1.77\% - 1.72\%} = -760.008 \tag{4.15}$$

Then estimate the derivative between 1.77% and 1.82%, i.e., at 1.795%, in the same way to get -757.956. Next, estimate the second derivative

TABLE 4.3 Model Convexity Calculations for the Adjusted Notional Amount of the 4½s of May 15, 2017, as of May 28, 2010

Rate	Price	1st Derivative	Convexity
1.72%	129.043		
1.745%		-755.304	
1.77%	128.665		41.5
1.795%		-752.637	
1.82%	128.289		
2.72%	121.737		
2.745%		-703.902	
2.77%	121.385		40.6
2.795%		-701.436	
2.82%	121.035		
3.72%	114.927		
3.745%		-656.370	
3.77%	114.599		39.8
3.795%		-654.090	
3.82%	114.272		

TABLE 4.4 Model Convexity Calculations for TYU0 as of May 28, 2010

Rate	Price	1st Derivative	Convexity
1.72%	127.553		
1.745%		−760.008	
1.77%	127.173		32.3
1.795%		−757.956	
1.82%	126.794		
2.72%	120.078		
2.745%		−743.792	
2.77%	119.706		−14.3
2.795%		−744.648	
2.82%	119.334		
3.72%	112.505		
3.745%		−773.593	
3.77%	112.119		−19.2
3.795%		−774.669	
3.82%	111.731		

TABLE 4.5 Model Convexity Calculations for TYU0C 120 as of May 28, 2010

Rate	Price	1st Derivative	Convexity
1.72%	7.657		
1.745%		−715.275	
1.77%	7.299		2,575.0
1.795%		−705.878	
1.82%	6.946		
2.72%	1.919		
2.745%		−362.117	
2.77%	1.738		26,860.0
2.795%		−338.771	
2.82%	1.569		
3.72%	.126		
3.745%		−41.434	
3.77%	.105		113,382.0
3.795%		−35.480	
3.82%	.087		

at 1.77% by dividing the change in the first derivative by the change in rates:

$$\frac{-757.956 + 760.008}{1.795\% - 1.745\%} = 4,104 \quad (4.16)$$

Finally, to estimate the convexity, divide the estimate of the second derivative by the price of the futures contract at 1.77%:

$$\frac{1}{127.172545} \times 4,104 = 32.3 \quad (4.17)$$

In Tables 4.3 and 4.5 the second derivatives of the bond and option are always positive so that convexity is always positive. These securities would be said to exhibit *positive convexity*. Graphically this means that their price-rate curves are convex and that, as shown in Figure 4.4, their DV01s fall as rates increase.

The futures contract, by contrast, is convex over part but not all of its range: in Table 4.4 TYU0 exhibits positive convexity at 1.77% but negative convexity at 2.77% and at 3.77%. In terms of Figure 4.4, the DV01 of the futures contract is falling at 1.77% but is rising at 2.77% and also at 3.77%.

The convexity values for the option calculated in Table 4.5 are relatively large. At intermediate rate levels this is certainly due in part to the rapid fall in DV01 as seen in Figure 4.4. At low and high levels of rates, however, the relatively large convexity values are mostly due to the relatively low price of the option. At 3.77%, for example, the change in the first derivative is about 2.3 for the bond and 6.0 for the option. But because the option price at 3.77% is .105, compared with 114.599 for the bond, the convexity of the option is thousands of times bigger. In short, a price factor distinguishes convexity from the second derivative just as a price factor distinguishes duration from DV01.

A HEDGING APPLICATION, PART II: A SHORT CONVEXITY POSITION

In the first part of this hedging application the market maker buys \$47.098 million face amount of TYU0 against a short position of \$100 million TYU0C 120. Figure 4.5 shows the profit and loss, or P&L, of a long position of \$47.098 million futures and of a *long* position of \$100 million options as rates change. Since the market maker is actually short the options, the P&L of the position at any rate level is the P&L of the long futures position minus the P&L of the long option position.




FIGURE 4.5 P&L-Rate Curve for a \$100 Million Long in TYU0C 120 and a DV01-Equivalent Long in TYU0 as of May 28, 2010

By construction, the *DV01* of the long futures and option positions are the same at a rate of 2.77%. In other words, for small rate changes, the change in the value of one position equals the change in the value of the other. Graphically, the P&L curves are tangent at 2.77%.

The first part of this hedging application showed that the hedge performs well in that the market maker neither makes nor loses money after a five-basis point change in rates. At first glance it may appear from Figure 4.5 that the hedge works well after moves of 25 or even 50 basis points. The values on the vertical axis, however, are measured in millions of dollars. After a move of only 25 basis points the hedge is off by about \$150,000, which is a very large number in light of the approximately \$15,625 the market maker collected in spread. Worse yet, since the P&L of the long option is always above that of the long futures position, the market maker loses this \$150,000 whether rates rise or fall by 25 basis points.

The hedged position loses whether rates rise or fall because the option is more convex than the bond. In market jargon, the hedged position is *short convexity*. For small rate changes away from 2.77% the values of the futures and option positions change by the same amount. Due to its greater convexity, however, the sensitivity of the option changes by more than the sensitivity of the bond. When rates increase, the *DV01* of the option falls by more. Hence, after further rate increases, the option falls in value less than the futures, and the P&L of the option position stays above that of the futures position. Similarly, when rates decline below 2.77%, the *DV01* of

both the futures and option rise, but the $DV01$ of the option rises by more. Hence, after further rate declines the option rises in value more than the futures and the P&L of the option position again stays above that of the futures position.

This discussion reveals that $DV01$ hedging is local, that is, valid in a particular neighborhood of rates. As rates move, the quality of the hedge deteriorates. Consequently, the market maker will need to re-hedge the position. If rates rise above 2.77% so that the $DV01$ of the option position falls by more than the $DV01$ of the futures position, the market maker will have to sell futures to re-equate $DV01$ s at the higher level of rates. If, on the other hand, rates fall below 2.77% so that the $DV01$ of the option position rises by more than the $DV01$ of the futures position, the market maker will have to buy futures to re-equate $DV01$ s at the lower level of rates.

An erroneous conclusion might be drawn at this point. Figure 4.5 shows that the value of the option position exceeds the value of the futures position at any rate level. Nevertheless, it is not correct to conclude that the option position is a superior holding to the futures position. Anticipating the discussion in Chapter 8, the market price of an option will be set high enough relative to the price of the futures to reflect its convexity advantages. In particular, if rates do not change by very much, then as time passes the futures will perform better than the option, a disadvantage of the long option position that is not captured in Figure 4.5. In summary, the long option position will outperform the long futures position if rates move a lot while the long futures position will outperform if rates stay about the same. It is in this sense, by the way, that a long convexity position is long volatility while a short convexity position is short volatility.

ESTIMATING PRICE CHANGES AND RETURNS WITH $DV01$, DURATION, AND CONVEXITY

Price changes and returns as a result of changes in rates can be estimated with the measures of price sensitivity used in previous sections. Despite the abundance of calculating machines that, strictly speaking, makes these approximations unnecessary, an understanding of these estimation techniques builds intuition about the behavior of fixed income securities and, with practice, allows for some rapid mental calculations.

A *second-order Taylor approximation* of the price-rate function with respect to rates gives the following approximation for the price of a security after a small change in rate:

$$P(y + \Delta y) \approx P(y) + \frac{dP}{dy} \Delta y + \frac{1}{2} \frac{d^2P}{dy^2} \Delta y^2 \quad (4.18)$$

Equation (4.18) can be rewritten in several useful ways. First, subtracting P from both sides gives an approximation for the change in price:

$$\Delta P \approx \frac{dP}{dy} \Delta y + \frac{1}{2} \frac{d^2P}{dy^2} \Delta y^2 \quad (4.19)$$

Second, dividing (4.19) by P gives an approximation for the percentage change in price:

$$\frac{\Delta P}{P} \approx \frac{1}{P} \frac{dP}{dy} \Delta y + \frac{1}{2} \frac{1}{P} \frac{d^2P}{dy^2} \Delta y^2 \quad (4.20)$$

Third, using the definitions of duration and convexity in equations (4.11) and (4.14), (4.20) can be rewritten as

$$\frac{\Delta P}{P} \approx -D\Delta y + \frac{1}{2} C\Delta y^2 \quad (4.21)$$

As an example, given data on the price and interest rate sensitivity of TYU0C 120 at 2.77% from previous sections, what is an estimate of the price at 2.50%? Any of equations (4.18) through (4.21) could be applied, but choose (4.18) for now. Table 4.1 reports that at 2.77% the price of the option is 1.738 and its $DV01$ is .03505, which, multiplying by $-10,000$, implies a first derivative of -350.5 . Table 4.5 reports that at 2.77% the convexity of the option is 26,860.0, which, multiplied by its price of 1.738, implies a second derivative of 46,682.7. Substituting all these quantities into (4.18) gives the following price estimate at 2.50%:

$$\begin{aligned} P(2.50\%) &\approx P(2.77\%) + \frac{dP}{dy}(2.50\% - 2.77\%) + \frac{1}{2} \frac{d^2P}{dy^2}(2.50\% - 2.77\%)^2 \\ &\approx 1.738 - 350.5 \times (-.27\%) + \frac{1}{2} \times 46,682.7 \times (-.27\%)^2 \\ &\approx 1.738 + .946 + .170 = 2.854 \end{aligned} \quad (4.22)$$

To three decimals the price of TYU0C 120 at 2.50% is 2.854, so the approximation given by (4.22) is quite accurate.

Note that the first derivative or $DV01$ -like term of (4.22), .946, is much larger than the second derivative term, .170. Or, were the approximation (4.21) used instead, the duration term is much larger than the convexity term. This is generally true for individual securities because, while convexity is usually a larger number than duration, the change in rate is so much




FIGURE 4.6 Price-Rate Curve for TYU0C 120 and its First- and Second-Order Approximations as of May 28, 2010

larger than the change in rate squared that the duration effect dominates.⁷ This fact suggests that it may sometimes be safe to drop the convexity term completely and to use the *first-order approximation* for the change in price or the percentage change in price, which follow from (4.19) and (4.21), respectively:

$$\Delta P \approx \frac{dP}{dy} \Delta y \quad (4.23)$$

$$\frac{\Delta P}{P} \approx -D \Delta y \quad (4.24)$$

Figure 4.6 graphs the option price along with the first-order and second-order approximations at a starting rate of 2.77%. Both approximations work very well for very small changes in rate. For larger changes the second-order approximation still works well, but for very large changes it eventually fails. The figure makes clear that approximating price changes with DV01 or duration alone ignores the curvature or convexity of the price-rate function while adding the convexity term captures a good deal of this curvature.

In the case of a bond or futures price, with price-rate curves that exhibit much less convexity than that of the option—compare Figure 4.1 with

⁷This need not be true, of course, for manufactured securities or positions, e.g., hedged positions constructed to have zero duration.

Figure 4.2—both first- and second-order approximations work so well that they would be difficult to distinguish graphically over a relevant range of interest rates.

CONVEXITY IN THE INVESTMENT AND ASSET-LIABILITY MANAGEMENT CONTEXTS

It was mentioned earlier in this chapter, in the discussion of Figure 4.5, that the option, as the more positively convex security, outperforms a *DV01*-matched position in futures if rates move a lot. This effect, that convexity is an exposure to volatility, can be seen directly from the approximation (4.21). Since Δy^2 is always positive, positive convexity increases return so long as interest rates move. The bigger the move in either direction, the greater the gains from positive convexity. Negative convexity works in the reverse. If C is negative, then rate moves in either direction reduce returns. In the investment context, choosing among securities with the same duration expresses a view on interest rate volatility. Choosing a very positively convex security would essentially be choosing to be long volatility, while choosing a negatively convex security would essentially be choosing to be short volatility.

Figure 4.6 suggests that asset-liability managers (or hedgers, more generally) can achieve greater protection against interest rate changes by hedging duration and convexity instead of duration alone. Consider an asset-liability manager who sets both the duration and convexity of assets equal to those of liabilities. Since both the first- and second-derivative terms of the asset and liability price-rate functions match, changes in the value of assets will more closely resemble changes in the value of liabilities than had their durations alone been matched. Furthermore, since matching convexity also sets the initial change in interest rate sensitivity of the assets equal to that of the liabilities, the sensitivity of the assets will be very close to the sensitivity of the liabilities even after a small change in rate. Put another way, the asset-liability manager need not rebalance so often as in the case of matching duration alone.

MEASURING THE PRICE SENSITIVITY OF PORTFOLIOS

This section shows how measures of a portfolio's price sensitivity are related to the measures of its component securities. Computing price sensitivities can be a time-consuming process, especially when using the term structure models of Part Three. Since a typical investor or trader focuses on a particular set of securities at one time and constantly searches for desirable

portfolios from that set, it is often inefficient to compute the sensitivity of every portfolio from scratch. A better solution is to compute sensitivity measures for all the individual securities and then to use the rules of this section to compute portfolio sensitivity measures.

A price or a measure of sensitivity for security i is indicated by the superscript i , while quantities without superscripts denote portfolio quantities. By definition, the value of a portfolio equals the sum of the value of the individual securities in the portfolio:

$$P = \sum P^i \quad (4.25)$$

Recall that in this chapter y has been a single rate or factor sufficient to determine the prices of all securities. Therefore, one can compute the derivative of price with respect to this rate or factor for all securities in the portfolio and, from (4.25),

$$\frac{dP}{dy} = \sum \frac{dP^i}{dy} \quad (4.26)$$

Then, dividing both sides by 10,000 and using the definition of $DV01$ in (4.1) shows that the $DV01$ of a portfolio equals the sum of the individual security $DV01$ s:

$$DV01 = \sum DV01^i \quad (4.27)$$

The rule for duration is only a bit more complex. Starting from equation (4.26), divide both sides by $-P$:

$$-\frac{1}{P} \frac{dP}{dy} = \sum -\frac{1}{P} \frac{dP^i}{dy} \quad (4.28)$$

Now multiply each term in the summation by one in the form of $\frac{P^i}{P}$.

$$-\frac{1}{P} \frac{dP}{dy} = \sum -\frac{P^i}{P} \frac{1}{P^i} \frac{dP^i}{dy} \quad (4.29)$$

Finally, using the definition of duration in (4.11),

$$D = \sum \frac{P^i}{P} D^i \quad (4.30)$$

In words, the duration of a portfolio equals a weighted sum of individual durations, where each security's weight is its value as a percentage of portfolio value.

Since the formula for the convexity of a portfolio can be derived along the same lines as the duration of a portfolio, it is given here without proof:

$$C = \sum \frac{P^i}{P} C^i \quad (4.31)$$

YIELD-BASED RISK METRICS

As a special case of the metrics defined so far in this chapter, this section defines yield-based measures of price sensitivity. These measures have two important weaknesses. First, they are defined only for securities with fixed cash flows. Second, as will be seen shortly, their use implicitly assumes parallel shifts in yield, which is not a particularly good assumption. Despite these weaknesses, however, there are several reasons fixed income professionals must understand these measures. First, these measures of price sensitivity are simple to compute, easy to understand, and, in many situations, perfectly reasonable to use. Second, these measures are widely used in the financial industry. Third, much of the intuition gained from a full understanding of these measures carries over to more general measures of price sensitivity.

Yield-Based *DV01* and Duration

Yield-based *DV01* and duration are special cases of the metrics introduced earlier in this chapter. In particular, these yield-based measures assume that the yield of a security is the interest rate factor and that the price-rate relationship is the price-yield function introduced in equations (3.13) and (3.14). For convenience, these equations are reproduced here for a face value of 100 and with price written explicitly as a function of that security's yield, y :

$$P(y) = \frac{100c}{2} \sum_{t=1}^{2T} \frac{1}{\left(1 + \frac{y}{2}\right)^t} + \frac{100}{\left(1 + \frac{y}{2}\right)^{2T}} \quad (4.32)$$

$$P(y) = \frac{100c}{y} \left(1 - \frac{1}{\left(1 + \frac{y}{2}\right)^{2T}} \right) + \frac{100}{\left(1 + \frac{y}{2}\right)^{2T}} \quad (4.33)$$

Taking the negative of the derivative of the two pricing expressions, (4.32) and (4.33), dividing by 10,000, and applying the definition of *DV01* in (4.2), gives two expressions for yield-based *DV01*. Note that, to avoid

clutter, this section will use the simple notations $DV01$ and D even when referring to the special cases of yield-based $DV01$ and duration.

$$DV01 = \frac{1}{10,000} \frac{1}{1 + \frac{y}{2}} \left[\frac{100c}{2} \sum_{t=1}^{2T} \frac{t}{2} \frac{1}{(1 + \frac{y}{2})^t} + T \frac{100}{(1 + \frac{y}{2})^{2T}} \right] \quad (4.34)$$

$$DV01 = \frac{1}{10,000} \left[\frac{100c}{y^2} \left(1 - \frac{1}{(1 + \frac{y}{2})^{2T}} \right) + T \left(1 - \frac{c}{y} \right) \frac{100}{(1 + \frac{y}{2})^{2T+1}} \right] \quad (4.35)$$

Similarly, applying the definition of duration in (4.11) to the pricing equations (4.32) and (4.33) gives the special cases of yield-based duration:

$$D = \frac{1}{P} \frac{1}{1 + \frac{y}{2}} \left[\frac{100c}{2} \sum_{t=1}^{2T} \frac{t}{2} \frac{1}{(1 + \frac{y}{2})^t} + T \frac{100}{(1 + \frac{y}{2})^{2T}} \right] \quad (4.36)$$

$$D = \frac{1}{P} \left[\frac{100c}{y^2} \left(1 - \frac{1}{(1 + \frac{y}{2})^{2T}} \right) + T \left(1 - \frac{c}{y} \right) \frac{100}{(1 + \frac{y}{2})^{2T+1}} \right] \quad (4.37)$$

These special cases are also known in the industry as *modified* or *adjusted duration*.⁸

There is a certain structure to equations (4.34) and (4.36). Each term in the brackets is the present value of a bond payment multiplied by the time to receipt of that payment, $\frac{t}{2}$. The contribution of a payment to the interest rate risk of a bond varies directly with its present value and with its time to receipt. In addition, duration can be viewed as a weighted-sum of times to receipt, with each weight equal to the corresponding present value divided by the total of the present values, i.e., the price. Viewed this way, duration is a weighted-sum of times to receipt of payments and can be said to be measured in years. Hence, practitioners often refer to a duration of six as *six years*.

Table 4.6 calculates the $DV01$ and duration of the U.S. Treasury $2\frac{1}{8}\%$ due May 31, 2015, as of May 28, 2010, using equations (4.34) and (4.36) and the market yield of the bond on that date, namely 2.092%.⁹ The cash flow dates and cash flows of the bond are as described in Part One. The

⁸This terminology is used because the first metric of this sort was *Macaulay Duration*. But the definition of the text, which divided Macaulay Duration by $1 + \frac{y}{2}$, became the industry standard.

⁹The use of these equations in this case is actually an approximation since the settlement date is June 1, 2010, and not May 31. See Appendix A in Chapter 3.

TABLE 4.6 *DV01 and Duration Calculations for the 2½s of May 31, 2015, as of May 28, 2010, at a Yield of 2.092 Percent*

Date	Term	Cash Flow	Present Value	Time-Wtd. PV	% of Wtd. Sum
11/30/10	0.5	1.0625	1.0515	.5258	.1%
5/31/11	1.0	1.0625	1.0406	1.0406	.2%
11/30/11	1.5	1.0625	1.0298	1.5448	.3%
5/31/12	2.0	1.0625	1.0192	2.0384	.4%
11/30/12	2.5	1.0625	1.0086	2.5216	.5%
5/31/13	3.0	1.0625	.9982	2.9946	.6%
11/30/13	3.5	1.0625	.9879	3.4575	.7%
5/31/14	4.0	1.0625	.9776	3.9105	.8%
11/30/14	4.5	1.0625	.9675	4.3538	.9%
5/31/15	5.0	101.0625	91.0749	455.3746	95.3%
Total			100.1559	477.7621	
DV01		.04728			
Duration		4.7208			

present value of each payment is computed using the market yield. For example, the present value of the coupon payment due on May 31, 2014, is

$$\frac{1.0625}{\left(1 + \frac{2.092\%}{2}\right)^8} = .97763 \quad (4.38)$$

The time-weighted present value of each cash flow is its present value times its term. For the cash flow on May 31, 2014, the time-weighted present value is $.97763 \times 4.0$ or 3.9105.

From equation (4.34), the DV01 of the bond is the sum of the time-weighted present values divided by one plus half the yield and divided by 10,000. Using the total from the table, this bond's DV01 is

$$\frac{1}{10,000} \times \frac{1}{\left(1 + \frac{2.092\%}{2}\right)} \times 477.7621 = .04728 \quad (4.39)$$

From equation (4.36), the duration of the bond is the sum of the time-weighted present values divided by one plus half the yield and divided by the price, the price just being the sum of the present values:

$$\frac{1}{100.1559} \times \frac{1}{\left(1 + \frac{2.092\%}{2}\right)} \times 477.7621 = 4.7208 \quad (4.40)$$

The rightmost column of Table 4.6 gives the time-weighted present value of each cash flow as a percent of the total of these weighted values.

Given the definitions of $DV01$ and duration in equations (4.34) and (4.36), these percentages are also the contribution of each cash flow to the interest rate risk of the bond. Far and away the largest contributor is the large cash flow at maturity. But considering the coupon flows alone, the contribution increases with term. Even though the present values of the longer-term coupon payments decline with term, their contributions to interest rate risk increase with term. Longer-dated cash flows are more sensitive to interest rate changes because they are discounted over longer periods of time.

Having defined and illustrated yield-based measures of interest rate sensitivity, an important limitation of their use becomes clear. Constructing a hedge so that the yield-based $DV01$ of a bond bought equals the yield-based $DV01$ of a bond sold will work as intended only if the two bond yields change by the same amount, i.e., only if their yields move in parallel. Of course, the efficacy of any hedge depends on the validity of its assumptions. In the examples of the previous sections, an underlying pricing model was used to relate the prices of the various securities to the seven-year par rate, and the quality of those hedges depends on that relationship being valid. Nevertheless, a well-thought-out model along the lines of those in Part Three, or well-researched empirical relationships along the lines of Chapter 6, are more likely to produce valid pricing relationships and hedges than the assumption of parallel yield shifts.

Yield-Based $DV01$ and Duration for Zero Coupon Bonds, Par Bonds, and Perpetuities

Yield-based measures are particularly useful because of the intuition furnished by their easy-to-derive formulas. This and the next several subsections exploit this usefulness to compare and contrast the risk profiles of bonds with different cash-flow characteristics.

The yield-based $DV01$ and duration of a zero coupon bond can be found by setting the coupon rate c equal to zero in equations (4.35) and (4.37) and noting for the latter that, for a T -year zero coupon bond with 100 face amount,

$$P = \frac{100}{(1 + \frac{y}{2})^{2T}} \quad (4.41)$$

Hence,

$$DV01_{c=0} = \frac{T}{100(1 + \frac{y}{2})^{2T+1}} = \frac{TP}{10,000(1 + \frac{y}{2})} \quad (4.42)$$

$$D_{c=0} = \frac{T}{(1 + \frac{y}{2})} \quad (4.43)$$

From (4.43), the duration of a zero coupon bond is its years to maturity divided by a factor only slightly greater than one. Also, the duration of a zero, for a fixed yield, always increases with maturity. From (4.42), however, for long maturity zero coupon bonds, the $DV01$ may not increase with maturity because a falling price may outweigh the increase in maturity. This last point will be illustrated in the next subsection.

The yield-based $DV01$ and duration of par bonds are useful formulae as relatively simple approximations for bonds with prices close to par. For a par bond (see Chapter 3), $P = 100$ and $c = y$. Substituting these values into equations (4.35) and (4.37) shows that

$$DV01_{c=y} = \frac{1}{100y} \left(1 - \frac{1}{\left(1 + \frac{y}{2}\right)^{2T}} \right) \quad (4.44)$$

$$D_{c=y} = \frac{1}{y} \left(1 - \frac{1}{\left(1 + \frac{y}{2}\right)^{2T}} \right) \quad (4.45)$$

The last cases to be considered here are the $DV01$ and duration of perpetuities, which are sometimes useful as rough approximations for the risk of extremely long-term fixed income securities. Letting T approach infinity in equations (4.35) and (4.37) and recalling from Chapter 3 that the price of a perpetuity with 100 face amount is $\frac{100c}{y}$,

$$DV01_{T=\infty} = \frac{1}{100} \frac{c}{y^2} \quad (4.46)$$

$$D_{T=\infty} = \frac{1}{y} \quad (4.47)$$

Duration, $DV01$, Maturity, and Coupon: A Graphical Analysis

Figure 4.7 uses the equations in this section to show how duration varies across bonds. For the purposes of this figure, all yields are fixed at 3.50%. At this yield, the duration of a perpetuity is 28.6. Since a perpetuity has no maturity, this duration is shown in Figure 4.7 as a horizontal line. Also, since by equation (4.47) the duration of a perpetuity does not depend on coupon, this line is a benchmark for the duration of any coupon bond with a sufficiently long maturity.

From equation (4.43), and as evident from Figure 4.7, the duration of zero coupon bonds is linear in maturity.

The duration of the par bond in Figure 4.7 increases with maturity. Inspection of equation (4.45) makes it clear that this is always the case and




FIGURE 4.7 Duration Across Bonds Yielding 3.50%

that the duration of a par bond rises from zero at a maturity of zero and steadily approaches the duration of a perpetuity.

Considering all of the curves of Figure 4.7 together reveals that for any given maturity duration falls as coupon increases. (Recognize that the par bond in the figure has a coupon equal to the yield of 3.50%.) The intuition behind this fact is that higher-coupon bonds have a greater fraction of their value paid earlier. The higher the coupon, the larger the weights on the duration terms of early years relative to those of later years. Hence, higher-coupon bonds are effectively shorter-term bonds and therefore have lower durations.

A little-known fact about duration can be extracted from Figure 4.7. The duration of a bond with a very low, near zero, coupon would be just below the zero coupon line of the figure. Furthermore, the coupon could be set low enough such that the bond's duration is still just below the zero coupon line but above the duration of a perpetuity.¹⁰ Eventually, however, as maturity increases, the low coupon bond must approach the duration of a perpetuity, i.e., its duration must fall with maturity. This fact is somewhat of a mathematical curiosity if—as at the time of this writing—yields are low relative to the coupons of outstanding bonds so that few if any bonds exist with the prerequisite long maturities and deep discounts.

¹⁰ In the example of the text, a bond with a coupon of .5% would have a duration that peaked above the duration of a perpetuity.

The next figure will show how $DV01$ varies across bonds. For this discussion it is useful to combine explicitly the definitions of $DV01$ and duration from (4.2) and (4.11) to write that

$$DV01 = \frac{P \times D}{10,000} \quad (4.48)$$

As discussed in the context of Figure 4.7, duration almost always increases with maturity. According to equation (4.48), however, the effect of maturity on $DV01$ is more complex since it depends not only on how duration changes with maturity but also on how price changes with maturity. What will be called the *duration effect* tends to increase $DV01$ with maturity while what will be called the *price effect* can either increase or decrease $DV01$ with maturity.

Figure 4.8 graphs $DV01$ as a function of maturity under the same assumptions used in Figure 4.7. Since the $DV01$ of a perpetuity, unlike its duration, depends on the coupon rate, the perpetuity line is removed.

Inspection of equation (4.44) reveals that the $DV01$ of par bonds always increases with maturity. Since the price of par bonds is always 100, the price effect does not come into play, and, as in the case of duration, longer par bonds have greater price sensitivity. The curve approaches .286, the $DV01$ of a par perpetuity at a yield of 3.50%.

As discussed in Chapter 3, extending the maturity of a premium bond increases its price. As a result, the price and duration effects combine so that the $DV01$ of a premium bond increases with maturity faster than the $DV01$




FIGURE 4.8 $DV01$ Across Bonds Yielding 3.50%

of a par bond. Of course, at some maturity beyond the range of the graph, the price of the bond increases very slowly and the price effect becomes less important. The $DV01$ of the 7% bonds eventually approaches that of a perpetuity with a coupon of 7% (i.e., .571).

The $DV01$ of a zero behaves initially like that of a coupon bond, but it eventually falls to zero. With no coupon payments the present value of a zero with a longer and longer maturity approaches zero, and so does its $DV01$.

Figure 4.8 also shows that, unlike duration, $DV01$ rises with coupon. This fact is immediately evident from equation (4.34).

Duration, $DV01$, and Yield

Inspection of equation (4.34) reveals that increasing yield lowers $DV01$. This fact was already introduced when showing that coupon bonds display positive convexity, that is, that their $DV01$ s fall as interest rates increase. As it turns out, increasing yield also lowers duration. The intuition behind this fact is that increasing yield lowers the present value of all payments but lowers the present value of the longer payments the most. Therefore, the value of the longer payments falls relative to the value of the whole bond. But since the duration of these longer payments is greatest, lowering their corresponding weights in the duration calculation must lower the duration of the whole bond.

To illustrate the effect of yield on duration, return to the example in Table 4.6. At a yield of 2.092%, the duration of the $2\frac{1}{8}$ s of May 31, 2015, is 4.7208. Also, the time-weighted present value of the payment at maturity, as a percentage of the sum of those values, is 95.3%. Reworking the calculations at a yield of 6%, the percentage of the sum attributable to the payment at maturity falls to 95% which, along with the increased relative importance of the shorter coupon payments, drives the duration down to 3.8375.

Yield-Based Convexity

Following the general definition of convexity in (4.14), yield-based convexity can be derived by taking the second derivative of (4.32) and dividing by price. The resulting formula is

$$C = \frac{1}{P} \frac{1}{(1 + \frac{y}{2})^2} \left[\frac{100c}{2} \sum_{t=1}^{2T} \frac{t}{2} \frac{t+1}{2} \frac{1}{(1 + \frac{y}{2})^t} + T(T + .5) \frac{100}{(1 + \frac{y}{2})^{2T}} \right] \quad (4.49)$$

The structure of this equation is similar to those of the expressions for yield-based $DV01$ and duration, but the time weights are $\frac{t}{2} \times \frac{t+1}{2}$ instead

of $\frac{t}{2}$, or, loosely speaking, more like t^2 than like t . With this in mind, the convexity of the $2\frac{1}{8}$ s due May 31, 2015, can be calculated using the first four columns of Table 4.6 but then substituting the weighted present value terms from (4.49) for those appropriate for the duration calculation. Doing this, the sum of the weighted present values, corresponding to the bracketed term in (4.49), is about 2,586 and, therefore, the bond's convexity is

$$\frac{1}{100.1559} \times \frac{1}{\left(1 + \frac{2.092\%}{2}\right)^2} \times 2,586 = 25.29 \quad (4.50)$$

For intuition, a useful special case of (4.49) is that of a zero coupon bond. Setting $c = 0$ and $P = 100 \left(1 + \frac{y}{2}\right)^{-2T}$,

$$C_{c=0} = \frac{T(T + .5)}{\left(1 + \frac{y}{2}\right)^2} \quad (4.51)$$

Applying (4.51), a five-year zero coupon bond yielding 2.092% would have a convexity of $5 \times 5.5 \times \left(1 + \frac{2.092\%}{2}\right)^{-2}$ or 26.93.

This exceeds the convexity of the five-year $2\frac{1}{8}$ s yielding 2.092%: since a coupon bond has some of its present value in earlier payments, and since the convexity contributions of those payments are less than that of the final payment at maturity, a coupon bond will have a lower convexity than a maturity- and yield-equivalent zero.

From (4.51) it is clear that longer-maturity zeros have greater convexity. In fact, the convexity of a zero increases with the square of maturity. Furthermore, thinking of a coupon bond as a portfolio of zeros, longer-maturity coupon bonds usually have greater convexity than shorter-maturity coupon bonds.

For easy reference, another useful special case of convexity is presented here, namely, the convexity of a par bond. This is obtained by differentiating equation (4.33) twice with respect to yield, evaluating the result at $y = c$, and dividing by the price, which, for par bonds is 100:

$$C_{c=y} = \frac{2}{y^2} \left[1 - \frac{1}{\left(1 + \frac{y}{2}\right)^{2T}} \right] - \frac{2T}{y \left(1 + \frac{y}{2}\right)^{2T+1}} \quad (4.52)$$

APPLICATION: THE BARBELL VERSUS THE BULLET

On May 28, 2010, a portfolio manager is considering the purchase of \$100 million face amount of the U.S. Treasury $3\frac{3}{8}$ s due November 15,

TABLE 4.7 Data on Three U.S. Treasury Bonds as of May 28, 2010

Coupon	Maturity	Price	Yield	Duration	Convexity
$2\frac{1}{2}$	3/31/15	102.5954	2.025%	4.520	23.4
$3\frac{3}{8}$	11/15/19	100.8590	3.288%	8.033	74.8
$4\frac{3}{8}$	11/15/39	102.7802	4.221%	16.611	389.7

2019, at a cost of \$100,859,000. After an analysis of the interest rate environment, the manager is comfortable with the pricing of the bond at a yield of 3.288% and with its duration of 8.033. But, after considering the data on two other Treasury bonds in Table 4.7, the manager wishes to consider an alternate investment.

The three bonds in the table have maturities of approximately five years, 10 years, and 30 years, respectively. Thus, an alternative to purchasing a *bullet* investment in the 10-year $3\frac{3}{8}$ s is to purchase a *barbell* portfolio of the shorter maturity, 5-year $2\frac{1}{2}$ s, and the longer maturity, 30-year $4\frac{3}{8}$ s. In particular, the barbell portfolio would be constructed to cost the same and have the same duration as the bullet investment. The advantages and disadvantages of this *barbell* relative to this *bullet* will be discussed after deriving the composition of the barbell portfolio.

Let V^5 and V^{30} be the value in the barbell portfolio of the 5-year and 30-year bonds, respectively. Then, for the barbell to have the same value as the bullet,

$$V^5 + V^{30} = 100,859,000 \quad (4.53)$$

Furthermore, using the data in Table 4.7 and equation (4.30), which describes how to compute the duration of a portfolio, the duration of the barbell equals the duration of the bullet if

$$\frac{V^5}{100,859,000} \times 4.520 + \frac{V^{30}}{100,859,000} \times 16.611 = 8.033 \quad (4.54)$$

Solving equations (4.53) and (4.54) shows that V^5 is \$71.555 million or 70.95% of the portfolio and that V^{30} is \$29.304 million or 29.05% of the portfolio. Finally, the convexity of the portfolio, using the data in Table 4.7 and equation (4.31), which describes how to compute the convexity of a portfolio, is

$$70.95\% \times 23.4 + 29.05\% \times 389.7 = 129.8 \quad (4.55)$$

The barbell has greater convexity than the bullet because duration increases linearly with maturity while convexity increases with the square of maturity. If a combination of short and long durations, essentially maturities, equals the duration of the bullet, that same combination of the two convexities, essentially maturities squared, must be greater than the convexity of the bullet. In the current context, the particularly high convexity of the $4\frac{3}{8}$ more than compensates for the lower convexity of the $2\frac{1}{2}$. As a result, the convexity of the portfolio exceeds the convexity of the $3\frac{3}{8}$. The general lesson is that spreading out the cash flows of a portfolio, without changing duration, raises convexity.

Return now to the decision of the portfolio manager. For the same amount of duration risk, the barbell portfolio has greater convexity, which means that its value will increase more than the value of the bullet when rates rise or fall. This is completely analogous to the price-rate profile of the option TYU0C 120 relative to the DV01-equivalent position in the futures TYU0 depicted in Figure 4.5: the barbell portfolio benefits more from interest rate volatility than does the bullet portfolio. What then is the disadvantage of the barbell portfolio? The weighted yield of the barbell portfolio is

$$70.95\% \times 2.025\% + 29.05\% \times 4.221\% = 2.663\% \quad (4.56)$$

compared with the yield of the bullet of 3.288%. Hence, the barbell will not do as well as the bullet portfolio if yields remain at current levels while, as just argued, the barbell will outperform if rates move sufficiently higher or lower.

In short, the manager's work in choosing to bear a level of interest rate risk consistent with a portfolio duration of about eight is not sufficient to complete the investment decision. A manager believing that rates will be particularly volatile will prefer the barbell portfolio while a manager believing that rates will not be particularly volatile will prefer the bullet portfolio. Of course, further calculations can establish exactly how volatile rates have to be for the barbell portfolio to outperform.

Multi-Factor Risk Metrics and Hedges

A major weakness of the approach in Chapter 4, and of several of the models of Part Three, is the assumption that movements in the entire term structure can be described by one interest rate factor. To make the case in the extreme, because the six-month rate is unrealistically assumed to predict perfectly the change in the 30-year rate, a (naive) *DV01* analysis leads to hedging a 30-year bond with a six-month bill. In reality, of course, it is widely recognized that rates in different regions of the term structure are far from perfectly correlated. Put another way, predicted changes in the 30-year rate relative to changes in the 6-month rate can be wildly off target, whether these predicted changes come from a model, like the one implicitly used in the first part of Chapter 4, or from the implicit assumption when using yield-based *DV01* that the two rates move by the same amount. The risk that rates along the term structure move by different amounts is known as *curve risk*.

This chapter discusses how to measure and hedge the risks of a security or portfolio in terms of several other securities, where each hedging security is most sensitive to a different part of the term structure. The more securities used in the hedge, the less important are any assumptions linking the behavior of one rate with another. At the extreme discussed in the previous paragraph, hedging with one security requires extremely strong assumptions about how rates move together. At the other extreme, a hedge that uses one security for every cash flow being hedged requires no assumptions about how rates move together because risk will have been *immunized* against any and all interest rate scenarios. Such a hedge, however, is almost certainly to be excessively costly. The methods presented in this chapter have been found to strike a sensible balance between hedging effectiveness and cost or practicality.

Key-rate exposures are used for measuring and hedging the risk of bond portfolios in terms of a relatively small number of the most liquid bonds available, usually the most recently issued, near-par, government bonds.

TABLE 5.1 Key Rate Duration
Profile of the U.S. Lehman Aggregate
Bond Index as of December 31, 2004

Key Rate	Duration
6-Month	0.145
2-Year	0.655
5-Year	1.151
10-Year	1.239
20-Year	0.800
30-Year	0.349
Total	4.339

Source: The Lehman Brothers Global Risk Model: A Portfolio Manager's Guide, April 2005.

Partial '01s are used for measuring and hedging the risk of portfolios of swaps or portfolios that contain both bonds and swaps in terms of the most liquid money market and swap instruments. As these instruments are almost always those whose prices are used to build a swap curve, the number of securities used in this methodology is usually greater than the number used in a key-rate framework. Finally, *forward-bucket '01s*, mostly used in the swap or combined bond and swap contexts as well, measure the risk of a portfolio not in terms of other securities but in terms of direct changes in the shape of the term structure. As a result, forward-bucket '01s are often the most intuitive way to understand the curve risks of a portfolio, but not the quickest way to see which hedges are required to neutralize such risks. This chapter concludes with a comment on the use of these methods to measure the volatility of a portfolio.

KEY-RATE '01s AND DURATIONS

Key-rate exposures are designed to describe how the risk of a bond portfolio is distributed along the term structure and how to implement any desired hedge, all in terms of some set of benchmark bonds, usually the more liquid government securities.¹ Table 5.1, as an example, shows a key-rate exposure report for the U.S. Lehman Aggregate Bond Index,² a benchmark portfolio of U.S. governments, agencies, mortgages, and corporates. The duration of

¹The idea was proposed in Thomas Ho, "Key Rate Duration: A Measure of Interest Rate Risk," *Journal of Fixed Income*, September, 1992.

²This set of indexes is now run by Barclays Capital.

the portfolio with respect to U.S. government rates is 4.339, as reported in the last row of the table. While this one number certainly quantifies interest rate risk, along the lines explained in Chapter 4, the rest of the table adds information about the distribution of this risk across the curve. For example, more than half of the portfolio's duration risk is closely related to—and could be hedged with—5- and 10-year bonds.

Continuing with this example for a moment, consider a portfolio manager whose performance is judged against the performance of this index. And say in addition that the manager's portfolio has the same duration as the index but is concentrated in 30-year bonds. If rates move up or down in parallel, the manager's performance will match that of the index. But if the government bond curve steepens the manager's portfolio will underperform, while if it flattens the manager's portfolio will outperform.³

The next three subsections discuss defining key-rate shifts, computing key-rate exposures, and then hedging with these exposures.

Key-Rate Shifts

The crucial assumption of the key-rate approach is that all rates can be determined as a function of a relatively small number of key rates. Therefore, the following decisions have to be made in order to implement the methodology: the number of key rates, the type of the key rates (e.g., spot rates, par yields), the terms of the key rates, and the rule for computing all other rates given the key rates.

In order to cover risk across the term structure, to keep the number of key rates as few as reasonable, and to rely only on the most liquid government securities, one popular choice of key rates for the U.S. Treasury and related markets are the 2-, 5-, 10-, and 30-year par yields. Then, motivated mostly by simplicity, the change in the term structure of par yields given a one-basis point change in each of the key rates is assumed to be as in Figure 5.1. Each of the four shapes is called a *key-rate shift*. Each key rate affects par yields from the term of the previous key rate (or zero) to the term of the next key rate (or the last term). For example, the 10-year key rate affects par yields of terms 5 to 30 years only. Furthermore, the impact of each key rate is normalized to be one basis point at its own maturity and then assumed to decline linearly, reaching zero at the terms of the adjacent key rates. For the two-year shift at terms of less than 2 years and for the 30-year shift at terms greater than 30 years, however, the assumed change is constant at one basis point.

By construction, the four key-rate shifts sum to a constant shift of one basis point. This allows for the interpretation of key-rate exposures as a

³For a definition of steepening and flattening, see Figure 2.6 and the surrounding discussion.




FIGURE 5.1 A Specification of Key-Rate Shifts

decomposition of the total $DV01$ or duration of a security or a portfolio into exposures to four different regions of the term structure.

While the key-rate shifts in Figure 5.1 turn out to be very tractable and useful, they implicitly make quite strong assumptions about the behavior of the term structure. Consider the assumption that the rate of a given term is affected only by its neighboring key rates. The 7-year rate, for example, is assumed to be a function of changes in the 5- and 10-year rates only. Empirically, however, were the 2-year rate to change while the 5- and 10-year rates stayed the same, the 7-year rate would probably change as well so as to maintain reasonable curvature across the term structure. The linearity of the shifts is also not likely to be an empirically valid assumption. All in all, however, the great tractability of working with the shifts in Figure 5.1 has been found to compensate for these theoretical and empirical shortcomings.

Calculating Key-Rate '01s and Durations

As a simple introduction to the calculation of key-rate '01s and duration, this subsection takes the example of a 30-year zero coupon bond. While the exposure of a 30-year zero to spot rates is very simple, its exposure to par yields and, therefore, to key rates (as defined in the previous subsection), is somewhat complicated. Basically, the risk along the curve of a 30-year zero is not equivalent to that of a 30-year par bond because of the latter's coupon payments.

TABLE 5.2 Key Rate DV01s and Durations of the May 15, 2040, C-STRIP as of May 28, 2010

	(1) Value	(2) Key-Rate '01	(3) Key-Rate Duration
Initial Curve	26.22311		
2-year Shift	26.22411	−.0010	−.38
5-year Shift	26.22664	−.0035	−1.35
10-year Shift	26.25763	−.0345	−13.16
30-year Shift	26.10121	.1219	46.49
Total		.0829	31.60

Table 5.2 illustrates the calculations of key-rate DV01s and durations for 100 face amount of the C-STRIP due May 15, 2040, as of May 28, 2010. The C-STRIP curve on that day was taken as the base pricing curve, with the key-rate shifts of Figure 5.1 superimposed as appropriate.

Column (1) of Table 5.2 gives the initial price of the C-STRIP and its present value after applying the key-rate one-basis point shifts of Figure 5.1. Column (2) gives the key-rate '01s. Denoting the key-rate '01 with respect to key rate y^k as $DV01^k$, these are defined analogously to $DV01$ as

$$DV01^k = -\frac{1}{10,000} \frac{\partial P}{\partial y^k} \quad (5.1)$$

where $\frac{\partial P}{\partial y^k}$ denotes the partial derivative of price with respect to that one key rate. Applying this definition to the C-STRIP described in Table 5.2, the key-rate '01 with respect to the 5-year shift is

$$-\frac{1}{10,000} \frac{26.22664 - 26.22311}{.01\%} = -.0035 \quad (5.2)$$

Or, in words, the C-STRIP *increases* in price by .0035 per 100 face amount for a *positive* one-basis-point five-year shift. The intuition behind the sign of this '01 will be explained in a moment.

The key-rate durations, denoted here as D^k , are also defined analogously to duration so that,

$$D^k = -\frac{1}{P} \frac{\partial P}{\partial y^k} \quad (5.3)$$

Continuing with the five-year shift in Table 5.2, the key-rate duration is

$$-\frac{1}{26.22311} \frac{26.22664 - 26.22311}{.01\%} = -1.35 \quad (5.4)$$

Turning now to interpreting the results, the key-rate profile in Table 5.2 shows that the interest rate exposure of the 30-year C-STRIPS is equivalent to that of a long position in a 30-year par bond, a smaller, short position in a 10-year par bond, and even smaller short positions in 5- and 2-year par bonds. This accords with the intuition stated at the beginning of this subsection, that the 30-year par bond's coupons create exposures at shorter terms that have to be offset by shorts of short-term par bonds.

In addition to this replicating portfolio intuition, it is useful to understand the precise mechanics by which the shorter-term key-rate '01s and durations in Table 5.2 turn out to be negative. To this end, Figure 5.2 graphs the 10-year key-rate (i.e., par yield) shift along with the resulting, implied shift of spot rates. (An analogous figure could be constructed for the five- and two-year key-rate shifts.)

From the implied spot rate shift in Figure 5.2 it is immediately apparent why the 10-year, key-rate sensitivities of the 30-year C-STRIPS in Table 5.2 are negative. By definition, the 30-year par yield is unchanged by the 10-year key-rate shift. But, according to the figure, the 30-year spot rate declines by about .45 basis points, meaning the 30-year C-STRIPS increases in value. Hence, the DV01 or duration of the 30-year STRIPS with respect to the 10-year par yield is negative. Since this spot rate declines by only .45 basis points per basis-point increase in the 10-year par rate, however, the absolute magnitude of this sensitivity is relatively small.

As for the intuition behind the shape of the implied spot-rate shift in Figure 5.2, the interested reader can note that with par yields with from




FIGURE 5.2 The Assumed 10-Year Key-Rate Shift of Par Yields and Its Implied Shift of Spot Rates

zero- to five-year terms remaining unchanged, spot rates of those terms have to remain unchanged as well. Therefore, any increases in par rates of terms between 5 and 10 years cannot be spread out across the spot rate curve but have to be concentrated in spot rates with terms greater than 5 years. But this implies that spot rates of terms between 5 and 10 years have to increase by more than par rates. Similarly, as par rates with terms greater than 10 years decrease, all spot rates with terms up to 10 years have already been fixed, implying that all of the decrease in par rates with terms greater than 10 years has to be concentrated in spot rates with terms beyond 10 years. Thus, the decline in spot rates has to be steeper than the decline in par rates. Finally, note that it would be impossible for the change in the 30-year par yield to be zero if all of the spot rates with terms from 5 to 30 years have increased. Hence, the longest-term spot rates have to decline as part of this key-rate shift of par yields.

A final technical point should be made about the last row of Table 5.2, namely, the sum of the key-rate '01s and durations. Since the sum of the key-rate shifts is a parallel shift of par yields, the sums of the key-rate '01s and durations are, as mentioned earlier, conceptually comparable to the one-factor, yield-based *DV01* and duration metrics, respectively. But key-rate exposures, which shift par yields, will not exactly match yield-based metrics, which shift security-specific yields.⁴

Hedging with Key-Rate Exposures

This subsection illustrates how to hedge with key-rate exposures using a stylized example of a trader making markets in U.S. Treasury bonds. On May 28, 2010, the trader executed two large trades:

1. The trader shorted \$100 million face amount of a 30-year STRIPS to a customer, buying about \$47 million face of the 30-year bond to hedge the resulting interest rate risk.
2. The trader facilitated a customer 5s-10s curve trade by shorting \$40 million face of the 10-year note and buying about \$72 million of the 5-year note.

Table 5.3 lists these trades in column (2), with two hedges, to be discussed presently, in the other columns. The coupon bonds featured in the rows of the table are the on-the-run 2-, 5-, 10-, and 30-year U.S. Treasuries, which, consistent with the motivation of key rates, are used by the trader to

⁴For example, it turns out that the sum of the changes in the 30-year spot rate across all the key-rate shifts is 1.08 basis points. Therefore, the sum of the key-rate exposures of a 30-year zero is about 1.08 times its exposure to the 30-year spot rate, which is the same as 1.08 times its yield-based exposure.

TABLE 5.3 Stylized Market Maker Positions and Hedges as of May 28, 2010

(1) Bond	(2) Face Amount (\$ millions)		(3)	(4)
	Position	Hedge	Alternate Hedge	
.75s of 5/31/12		−5.190		
2.125s of 5/31/15	72.446	−80.006	−80.008	
3.5s of 5/15/20	−40	−.487		
0s of 5/15/40	−100			
4.375s of 5/15/40	47.077	22.633	21.806	

hedge risk. The other bond in the table is the STRIPS due May 15, 2040, discussed in the previous subsection. Table 5.4 gives the key-rate '01 profiles for 100 face amount of these bonds in rows (i) through (v) and the '01 profiles for particular portfolios, again, to be discussed presently, in rows (vi) through (ix).

If the maturity of a coupon bond were exactly equal to the term of a key rate and if the price of that bond were exactly par, then that bond's yield would be identical to that key rate. By definition, then, that bond's key-rate '01 with respect to that key rate would equal its yield-based *DV01* while its key-rate '01 with respect to all other key rates would be zero. Since the on-the-run bonds profiled in Table 5.4 are close to 2-, 5-, 10-, and 30-year maturities, and since they do sell for about par, their key-rate exposures in

TABLE 5.4 Key-Rate 01 Profile of a Stylized Market Maker's Position and Hedges as of May 28, 2010

Bond	Key-Rate '01 (per 100 face amount)				
	2-year	5-year	10-year	30-year	Sum
(i) .75s of 5/31/12	.0199	.0000	.0000	.0000	.0199
(ii) 2.125s of 5/31/15	.0000	.0480	.0000	.0000	.0480
(iii) 3.5s of 5/15/20	.0000	−.0001	.0870	.0000	.0869
(iv) 0s of 5/15/40	−.0010	−.0035	−.0345	.1219	.0829
(v) 4.375s of 5/15/40	.0000	.0001	.0010	.1749	.1760
(vi) Total Position (\$)	1,000	38,377	198	−39,578	0
(vii) Hedge (\$)	−1,000	−38,377	−198	39,578	0
(viii) Alternate Hedge (\$)	31	−38,379	217	38,131	0
(ix) Total+Alt. Hedge (\$)	1,031	−2	415	−1,447	0

rows (i), (ii), (iii), and (v) are heavily concentrated in the respective buckets. In row (iii), for example, the 10-year, key-rate '01 of the 3.5s of May 15, 2020, is .0870, while the rest of its key-rate '01s are near zero. Note that the key-rate profile of the 30-year STRIPS in row (iv) is as presented in Table 5.2.

The sums of the key-rate '01s for each of the bonds in rows (i) through (v) are given in the rightmost column of Table 5.4. The trader uses these sums for initial hedging, which, as discussed previously, is very much like single-factor, *DV01* hedging. So, the trader bought \$72.4 million of the five-year against the \$40mm short of the 10-year because

$$\frac{.0869}{.0480} \times \$40\text{mm} = \$72.4\text{mm} \quad (5.5)$$

Similarly, the trader bought \$47.1 million of 30-year bonds against the \$100 million short of 30-year STRIPS because

$$\frac{.0829}{.1760} \times \$100\text{mm} = \$47.1\text{mm} \quad (5.6)$$

Row (vi) of Table 5.4 gives the key-rate '01 profile, in dollars, of the trader's position recorded in column (2) of Table 5.3. The five-year key-rate '01 in millions of dollars, for example, is calculated as

$$\begin{aligned} & 72.446 \times \frac{.048}{100} - 40 \times \left(\frac{-.0001}{100} \right) - 100 \times \left(\frac{-.0035}{100} \right) + 47.077 \times \frac{.0001}{100} \\ & = .038361 \end{aligned} \quad (5.7)$$

which is \$38,361. (The small difference between this number and the \$38,377 in Table 5.4 is due to the rounding of the '01s and the position amounts.)

Because the trader's initial hedges were constructed to be *DV01*-neutral, the trader has no net *DV01*-type exposure, i.e., the sum of the '01s across row (vi) of Table 5.4 is zero. As can be seen from the rest of that row, however, the key-rate profile of the trader's book is not flat. In fact, the trader essentially has on a substantial 5s-30s steepener, meaning a position that will make money if 30-year yields rise relative to 5-year yields but lose money if the opposite occurs. But this accumulated steepener is a by-product of market making and not the result of deliberate risk taking. The trader, therefore, will construct a hedge to flatten out the key-rate profile in row (vi).

The hedging problem is to find the face amount of each of the key-rate securities such that the net key-rate '01s of the overall position are

all zero. Let the face amount of each of the hedging securities be F^2 , F^5 , F^{10} , and F^{30} for the 2-, 5-, 10-, and 30-year bonds, respectively. Then the equations for setting the overall 2-, 5-, 10-, and 30-key-rate '01s to zero are, respectively,

$$\frac{.0199}{100} F^2 + 0 \times F^5 + 0 \times F^{10} + 0 \times F^{30} + \$1,000 = 0 \quad (5.8)$$

$$\frac{.048}{100} \times F^5 - \frac{.0001}{100} \times F^{10} + \frac{.0001}{100} \times F^{30} + \$38,377 = 0 \quad (5.9)$$

$$\frac{.0870}{100} \times F^{10} + \frac{.001}{100} \times F^{30} + \$198 = 0 \quad (5.10)$$

$$\frac{.1749}{100} \times F^{30} - \$39,578 = 0 \quad (5.11)$$

Solving equations (5.8) through (5.11) gives the face amounts in column (3) of Table 5.3. By construction, then, the key-rate profile of the hedging portfolio, shown in row (vii) of Table 5.4, is the negative of the profile of row (vi) so that these two rows sum to zero.

This precisely constructed hedge, with its four equations and four unknowns, may look somewhat daunting. But this should not obscure the essentials of the hedge. The five-year key-rate '01 to be hedged is \$38,377 and the five-year key-rate '01 of the five-year on-the-run bond is .048, so the approximate face amount of the five-year bond that has to be sold is $\frac{\$38,377}{.048\%}$ or about \$79.95 million. Similarly, the 30-year '01 to be hedged is -\$39,578 and the 30-year '01 of the 30-year on-the-run is .1749, so the face amount of the 30-year bond that has to be bought is about $\frac{\$39,578}{.1749\%}$ or about \$22.63 million. These results are very close to the precise results reported in column (3) of Table 5.3.

In practice, a trader might very well recognize that the biggest risk of the position, from row (vi) of Table 5.4, is the 5s-30s steepener. The trader might then sell the \$80 million of the five-year on-the-run, as computed in the previous paragraph. Then, to keep a flat overall DV01, the trader might purchase an amount F^{30} such that

$$F^{30} \frac{.1760}{100} = \$80\text{mm} \times \frac{.0480}{100} \quad (5.12)$$

And solving, F^{30} is \$21.8 million. This quicker, alternate hedge is recorded in column (4) of Table 5.3. Its key-rate profile is given in row (viii) of Table 5.4 and the net key-rate profile of the original position and this alternate hedge is given in row (ix). This net profile is very close to flat, although the residual is a very small 2s-30s steepener!

PARTIAL '01s AND PV01

As mentioned in the introduction to Part One, swaps have become the most popular interest rate benchmark. Interest rate risk is measured in terms of swap curves not only by swaps traders, but also by asset managers that run portfolios that combine bonds and swaps.⁵ Examples of such managers include these: a life insurance company or pension fund that selects attractive corporate credits or mortgage exposures and hedges interest rate risk with swaps; a manager who supervises several traders or portfolio managers, some of whom trade bonds and some of whom trade swaps; or a relative value government bond investor who hedges curve risk with swaps. In any case, when swaps are taken as the benchmark for interest rates, risk along the curve is usually measured with *Partial '01s* or *Partial PV01s* rather than with key-rate '01s. This section discusses these swap-based methodologies without introducing additional numerical examples since the underlying concepts are very similar to those of key-rate exposures.

Swap market participants fit a par swap rate curve every day, if not more frequently, from a set of traded or observable par swap rates and shorter-term money market and futures rates. (See Chapter 21.) Leveraging this curve-fitting machinery, sensitivities of a portfolio or trading book are measured in terms of changes in the rates of the fitting securities. More specifically, the partial '01 with respect to a particular fitted rate is defined as the change in the value of the portfolio after a one-basis-point decline in that fitted rate and a refitting of the curve. All other fitted rates are unchanged. So, for example, if a curve fitting algorithm fits the three-month *London Interbank Offered Rate (LIBOR)* rate and par rates at 2-, 5-, 10-, and 30-year maturities, then the two-year partial '01 would be the change in the value of a portfolio for a one-basis point decline in the two-year par rate and a refitting of the curve, where the three-month LIBOR and the par 5-, 10-, and 30-year rates are kept the same. Note how the details of calculating partial '01s are intertwined with the details of constructing the swap curve.

Given the partial '01 profile of a portfolio, hedges to zero-out this profile are particularly easy to calculate. As pointed out in the previous section, with key-rate shifts defined in terms of par yields, the key-rate profile of the 10-year bond, for example, would be its *DV01* for the 10-year shift and zero for all other shifts only if the 10-year bond matured in exactly 10 years and were priced at exactly par. By contrast, in the case of partial '01s, the shifts are defined precisely in terms of the fitting securities. Therefore, by construction, all of the '01 of a fitting security is concentrated in the partial '01 calculated by shifting its rate, making calculating hedges

⁵In addition to managing interest rate risk, these managers must also manage *spread risk*, i.e., the risk that spreads between bond and swap rates change.

particularly easy. Nevertheless, since there are typically many fitting securities, market practice is to trade enough of the fitting securities so as to achieve an acceptable profile of partial '01s rather than trading every single fitting security so as to zero-out all partial '01s.

The *PV01* of a security is defined as the change in the value of the security if the rates of all fitting securities decline by one basis point. Hence *PV01* is conceptually equivalent to *DV01*, where the underlying curve-fitting methodology defines rates at all terms given the changes in the rates of the fitting securities. Furthermore, since the sum of all the partial '01 shifts is the *PV01* shift—with one caveat to be raised presently—the partial '01s may be thought of as a decomposition of the *PV01* into risks along the curve. The technical caveat is that money market rates and swap rates are quoted under different day-count conventions, namely, actual/360 for LIBOR-related rates and 30/360 for the fixed side of swaps. So, if money market rates and swap rates are mixed when fitting swap curves, as they usually are, changing each market rate by a basis point is not the same as changing all actual/360 rates by a basis point or all 30/360 rates by a basis point. To ensure that the sum of the partial '01s does equal the *PV01*, all rates could be converted into a single day-count convention. This normalization, however, sacrifices the desirable property that the '01 of each fitting security equals its '01 with respect to its own quoted rate.

In passing, it is worth noting that the *CV01* of a swap is the change in value of a swap for a one basis-point decrease in its coupon rate. A moment's reflection reveals that this quantity is proportional to the annuity factor to the swap's maturity. See equation (2.21). The two metrics, *CV01* and *PV01*, are sometimes used interchangeably, and sometimes confused, because the two are essentially equal for par swaps. To see this, note that the expression for the annuity factor in equation (3.15) is 100 times the expression for the *DV01* of a par swap in equation (4.44).

FORWARD-BUCKET '01s

While key rates and partial '01s conveniently express the exposures of a position in terms of hedging securities, *forward-bucket '01s* convey the exposures of a position to different parts of the curve in a much more direct and intuitive way. Basically, forward-bucket '01s are computed by shifting the forward rate over each of several defined regions of the term structure, one region at a time.

The starting point of the methodology is the division of the term structure into buckets. For the illustration of this section, the term structure is divided into five buckets: 0-2 years, 2-5 years, 5-10 years, 10-15 years, and 20-30 years. The best choice of buckets depends, of course, on the application at hand. A financing desk that does most of its trading in very short-term

securities would define many, narrow buckets in the short end and relatively few, wide buckets in the long end. A swaps market-making desk, with business across the curve, might use the buckets defined for this section, although it would likely prefer a greater number of narrower buckets and, particularly in Europe, might need buckets to cover maturities beyond 30 years.

Forward-Bucket Shifts and '01 Calculations

Each forward-bucket '01 is computed by shifting the forward rates in that bucket by one basis point. Depending on how rate curves are stored, this may mean shifting all of a bucket's semiannual forward rates, quarterly forward rates, or rates of even shorter term. This section shifts seminannual rates.

As a first example, consider a 2.12% five-year swap as of May 28, 2010. Table 5.5 lists the cash flows of the fixed side of 100 notional amount of the swap, the "Current" forward rates as of the pricing date, and the three shifted forward curves. For the "0-2 Shift," forward rates of term .5 to 2.0 years are shifted up by one basis point while all other forward rates stay the same. For the "2-5 Shift," forward rates in that bucket, and that bucket only, are shifted. Lastly, for "Shift All," the entire forward curve is shifted.

The row of Table 5.5 labeled "Present Value" gives the present value of the cash flows under the initial forward rate curve and under each of the shifted curves. The forward-bucket '01 for each shift is then the negative of the difference between the shifted and initial present values. For the 2-5-year shift, for example, the '01 is $-(99.9679 - 99.9955)$, or .0276.

TABLE 5.5 Computation of the Forward-Bucket '01s of a Five-Year 2.12 Percent EUR Swap as of May 28, 2010

Term	Cash Flow	Forward Rates (%)			
		Current	0-2 Shift	2-5 Shift	Shift All
.5	1.06	1.012	1.022	1.012	1.022
1.0	1.06	1.248	1.258	1.248	1.258
1.5	1.06	1.412	1.422	1.412	1.422
2.0	1.06	1.652	1.662	1.652	1.662
2.5	1.06	1.945	1.945	1.955	1.955
3.0	1.06	2.288	2.288	2.298	2.298
3.5	1.06	2.614	2.614	2.624	2.624
4.0	1.06	2.846	2.846	2.856	2.856
4.5	1.06	3.121	3.121	3.131	3.131
5.0	101.06	3.321	3.321	3.331	3.331
Present Value		99.9955	99.9760	99.9679	99.9483
'01			.0196	.0276	.0472




FIGURE 5.3 An Example of Spot-Starting and Forward-Starting Swaps

The '01 of the “Shift All” scenario is analogous to a $DV01$. The forward-bucket analysis decomposes this total '01 into .0196 due to the 0-2-year part of the curve and .0276 due to the 2-5-year part of the curve. The factors that determine the exact distribution of a total '01 across buckets are described in the next section.

Understanding Forward-Bucket '01s: a Payer Swaption

This subsection analyzes the forward-bucket '01s of a payer swaption. Swaptions are treated in greater detail in Chapter 18, but, for the present, a payer swaption gives the purchaser the right to pay a fixed rate on a swap at some time in the future. More specifically, consider an EUR 5x10 *payer swaption* struck at 4.044% as of May 28, 2010, which gives the purchaser the right to pay a fixed rate of 4.044% on a 10-year EUR swap in five years, that is, at the end of May 2015. The underlying security of this option is a 10-year swap for settlement in five years, otherwise known as a “5x10” swap. See Figure 5.3.⁶ (Forward starting swaps are discussed further in Chapter 13.) As of May 28, 2010, the rate on the EUR 5x10 swap was 4.044%, so the swaption of this application was at-the-money.

Table 5.6 gives the forward-bucket '01s of the EUR 5x10 payer swaption, along with the forward-bucket '01s of an EUR 5-year swap, 10-year swap, 15-year swap, and 5x10 swap. The column labeled “All” gives the '01 from shifting all forward rates.

Computing the '01s of the swaption requires a pricing model, which is not covered here.⁷ The intuition behind the results, however, is straightforward. The overall '01 of the payer swaption is negative: as rates increase, the value of the option to pay a fixed rate of 4.044% in exchange for a floating side worth par increases. Furthermore, since the swaption gives the right to pay fixed on a 5x10 swap, the '01 of the swaption will be most concentrated in the buckets that determine the value of that 5x10 swap, i.e., the 5-10

⁶This forward swap is a contract to enter into a 10-year swap in five years. Note from the figure that, since swaps settle $T + 2$, the spot-starting swaps begin on June 2, 2010, and the forward starting swap begins on June 2, 2015.

⁷See Part Three and Chapter 18.

TABLE 5.6 Forward-Bucket Exposures of Selected EUR-Denominated Securities as of May 28, 2010

Security	Rate	Forward-Bucket Exposures					
		0-2	2-5	5-10	10-15	20-30	All
5x10 Payer Swaption	4.044%	.0010	.0016	-.0218	-.0188	.0000	-.0380
5-Year Swap	2.120%	.0196	.0276	.0000	.0000	.0000	.0472
10-Year Swap	2.943%	.0194	.0269	.0394	.0000	.0000	.0857
15-Year Swap	3.290%	.0194	.0265	.0383	.0323	.0000	.1164
5x10 Swap	4.044%	.0000	.0000	.0449	.0366	.0000	.0815

and 10-15 buckets. The swaption has some positive '01 in the 0-2 and 2-5 buckets, as well, because the forward rates in that part of the curve affect the present value of the option's payoff at its expiration in five years' time.

The bucket '01 profiles of the 5-, 10-, and 15-year swaps are determined by several effects. First, and most obvious, each swap is exposed to all parts of the curve up to, but not past, its maturity. Second, the wider buckets, which shift the forward curve over a wider range, tend to generate larger '01s. For example, the 10-year swap's 5-10 bucket '01, which shifts forward rates over five years, is greater than its 2-5 bucket '01, which shifts rates over three years. Third, the further a shift is along the curve, the fewer of a swap's coupon payments are affected. This tends to lower the '01s of the longer-term buckets relative to the shorter-term buckets. Fourth, the larger the forward rate in a bucket, the lower the '01, for the same reason that *DV01* falls with rate, as shown in Chapter 4. In Table 5.6 the term structure of forward rates is, in fact, upward-sloping,⁸ so this effect, combined with the third, lowers the 15-year swap's 10-15 bucket '01 relative to its 5-10 bucket '01.

The 5x10 swap has no exposure to forward rates with a term less than 5 years or greater than 15 years, which is easily apparent from Figure 5.3. Its total '01 of .0815 is divided between the 5-10 and 10-15-year buckets, according to the third and fourth effects described in the previous paragraph.

The Appendix in this chapter presents a very simple demonstration of the third and fourth effects just invoked.

Hedging with Forward-Bucket '01s: a Payer Swaption

Table 5.7 shows the forward-bucket exposure of the payer swaption hedged in three different ways: with a 10-year swap, with a 5x10 swap, and with a combination of 5- and 15-year swaps.

⁸This follows from the upward-sloping par rates in the table, or, more directly, from the graph of the EUR forward rates in Figure 2.2.

TABLE 5.7 Forward-Bucket Exposures of Three Hedges of a Payer Swaption as of May 28, 2010

Security or Portfolio	Forward-Bucket Exposures				
	0-2	2-5	5-10	10-15	All
(i) 5x10 Payer Swaption	.0010	.0016	-.0218	-.0188	-.0380
Hedge #1:					
(ii) Long 44.34% 10-Year Swaps	.0086	.0119	.0175	.0000	.0380
(iii) Net Position	.0096	.0135	-.0043	-.0188	.0000
Hedge #2:					
(iv) Long 46.66% 5x10 Swaps			.0209	.0171	.0380
(v) Net Position	.0010	.0016	-.0009	-.0017	.0000
Hedge #3:					
(vi) Long 57.55% 15-Year Swaps	.0112	.0153	.0220	.0186	.0670
(vii) Short 61.55% 5-Year Swaps	-.0120	-.0170			-.0290
(viii) Net Position	.0002	-.0001	.0002	-.0002	.0000

The full '01 of the payer swaption and the 10-year swap are, from Table 5.6, $-.0380$ and $.0857$, respectively. Therefore, hedging the payer swaption requires a long position of $\frac{.0380}{.0857}$ or approximately 44.34% of the 10-year. Multiplying each of the forward-bucket exposures of the 10-year swap in Table 5.6 by this face amount gives row (ii) of Table 5.7. Then, adding the '01s of this hedge to those of the payer swaption gives the net bucket exposures in row (iii). So, while buying 10-year swaps in a DV01-neutral way may be a good first pass at a hedge, that is, a quick way to neutralize the rate risk of the payer swaption with the most liquid security available, the net bucket exposures show that the resulting position is at risk of a flattening.

Hedging the payer swaption by receiving in a DV01-weighted 5x10 swap, depicted in rows (iv) and (v) of Table 5.7, is a better hedge than receiving in the 10-year swap. This is not particularly surprising since the swaption is the right to pay fixed on that very swap. In any case, the resulting hedged position has a very slight exposure to flattening, but, for the most part, is neutral to rates and the term structure.

Since forward swaps are, in practice, not as easy to execute as par swaps, the final hedge of Table 5.7 considers hedging the swaption with 5 and 15-year par swaps. This hedge, depicted in rows (vi) through (viii) of the table, chooses a long face amount of the 15-year swap to neutralize the 5-10 and 10-15 bucket exposures of the payer swaption and a short face amount of the five-year swap to neutralize the 0-2 and 2-5 bucket exposures arising in small part from the original payer position but in large part from the 15-year swap bought as a hedge. The result,

given in row (viii), shows that this hedge neutralizes the risk of each bucket quite closely.

MULTI-FACTOR EXPOSURES AND MEASURING PORTFOLIO VOLATILITY

The facts that a portfolio has a $DV01$ of \$10,000 and that interest rates have a volatility of 100 basis points per year leads to the conclusion that the portfolio has an annual volatility of $\$10,000 \times 100$ or \$1 million. But this measure has the same drawback as one-factor measures of price sensitivity: the volatility of the entire term structure cannot be adequately summarized with just one number. As to be discussed in Part Three, just as there is a term structure of interest rates, there is a term structure of volatility. The 10-year par rate, for example, is usually more volatile than the 30-year par rate.

In general, portfolios are exposed to interest rates all along the curve and changes in these rates are not perfectly correlated. The frameworks of this chapter, therefore, can be used to estimate volatility more precisely. The presentation here will be in terms of key rates; the discussion would be similar in terms of partial '01s or forward bucket '01s.

First, estimate a volatility for each of the key rates and estimate a correlation for each pair of key rates. Second, compute the key-rate 01s of the portfolio. Third, compute the variance and volatility of the portfolio. This computation is quite straightforward given the required inputs. Say that there are only two key rates, C_1 and C_2 , that the key rates of the portfolio are $KR01_1$ and $KR01_2$, that the value of the portfolio is P , and that changes are denoted by Δ . Then, by the definition of key rates,

$$\Delta P = KR01_1 \times \Delta C_1 + KR01_2 \times \Delta C_2 \quad (5.13)$$

Furthermore, letting σ_P^2 , σ_1^2 , and σ_2^2 denote the variances of the portfolio and of the key rates and letting ρ denote the correlation of the key rates, equation (5.13) implies that

$$\sigma_P^2 = KR01_1^2 \sigma_1^2 + KR01_2^2 \sigma_2^2 + 2KR01_1 KR01_2 \rho \sigma_1 \sigma_2 \quad (5.14)$$

The standard deviation of the portfolio, of course, is just σ_P . While, as mentioned, this reasoning can be applied equally well to partial '01s or forward-bucket '01s, those two frameworks tend to have more reference rates than a typical key-rate framework and, therefore, would require the estimation of a greater number of volatilities and a much greater number of correlation pairs.

APPENDIX: SELECTED DETERMINANTS OF FORWARD-BUCKET '01s

Write the price of a two-year bond or fixed leg of a swap, with its fictional notional, in terms of forward rates, as

$$P = \frac{c}{1 + f_1} + \frac{1 + c}{(1 + f_1)(1 + f_2)} \quad (5.15)$$

Differentiating with respect to each of the forward rates and multiplying by -1 ,

$$-\frac{\partial P}{\partial f_1} = \frac{c}{(1 + f_1)^2} + \frac{1 + c}{(1 + f_1)^2(1 + f_2)} \quad (5.16)$$

$$-\frac{\partial P}{\partial f_2} = \frac{1 + c}{(1 + f_1)(1 + f_2)^2} \quad (5.17)$$

To consider the effects of the term of the bucket alone, let $f_1 = f_2$. Then,

$$-\frac{\partial P}{\partial f_1} > -\frac{\partial P}{\partial f_2} \quad (5.18)$$

showing that the '01 of the first bucket, from date 0 to date 1, is greater than the '01 of the second bucket, from date 1 to date 2, precisely because f_1 is used to discount more cash flows than is f_2 .

To consider the effects of the term structure alone, let $c = 0$. Then, the second bucket risk is less than the first if

$$-\frac{\partial P}{\partial f_2} = \frac{1}{(1 + f_1)(1 + f_2)^2} < -\frac{\partial P}{\partial f_1} = \frac{1}{(1 + f_1)^2(1 + f_2)} \quad (5.19)$$

which simplifies to

$$f_1 < f_2 \quad (5.20)$$

Hence, an upward-sloping term structure, because of discounting, lowers the second bucket risk relative to the first.

Empirical Approaches to Risk Metrics and Hedging

Central to the *DV01*-style metrics and hedges of Chapter 4 and the multi-factor metrics and hedges of Chapter 5 are implicit assumptions about how rates of different term structures change relative to one another. In this chapter, the necessary assumptions are derived directly from data on rate changes.

The chapter begins with single-variable hedging based on regression analysis. In the example of the section, a trader tries to hedge the interest rate risk of U.S. nominal *versus* real rates. This example shows that empirical models do not always describe the data very precisely and that this imprecision expresses itself in the volatility of the profit and loss of trades that depend on the empirical analysis.

The chapter continues with two-factor hedging based on multiple regression. The example for this section is that of an EUR swap market maker who hedges a customer trade of 20-year swaps with 10- and 30-year swaps. The quality of this hedge is shown to be quite a bit better than that of nominal *versus* real rates. Before concluding the discussion of regression techniques, the chapter comments on level *versus* change regressions.

The final section of the chapter introduces principal component analysis, which is an empirical description of how rates move together across the curve. In addition to its use as a hedging tool, the analysis provides an intuitive description of the empirical behavior of the term structure. The data illustrations for this section are taken from USD, EUR, GBP, and JPY swap markets. Considerable effort has been made to present this material at as low a level of mathematics as possible.

A theme across the illustrations of the chapter is that empirical relationships are far from static and that hedges estimated over one period of time may not work very well over subsequent periods.

SINGLE-VARIABLE REGRESSION-BASED HEDGING

This section considers the construction of a relative value trade in which a trader sells a U.S. Treasury bond and buys a U.S. Treasury TIPS (Treasury Inflation Protected Securities). As mentioned in the Overview, TIPS make real or inflation-adjusted payments by regularly indexing their principal amount outstanding for inflation. Investors in TIPS, therefore, require a relatively low real rate of return. By contrast, investors in U.S. Treasury bonds—called nominal bonds when distinguishing them from TIPS—require a real rate of return plus compensation for expected inflation plus, perhaps, an inflation risk premium. Thus the spread between rates of nominal bonds and TIPS reflects market views about inflation. In the relative value trade of this section, a trader bets that this inflation-induced spread will increase.

The trader plans to short \$100 million of the (nominal) 3 $\frac{5}{8}$ s of August 15, 2019, and, against that, to buy some amount of the TIPS 1 $\frac{7}{8}$ s of July 15, 2019. Table 6.1 shows representative yields and *DV01*s of the two bonds. The TIPS sells at a relatively low yield, or high price, because its cash flows are protected from inflation while the *DV01* of the TIPS is relatively high because its yield is low (see “Duration, *DV01*, and Yield” in Chapter 4). In any case, what face amount of the TIPS should be bought so that the trade is hedged against the level of interest rates, i.e., to both rates moving up or down together, and exposed only to the spread between nominal and real rates?

One choice is to make the trade *DV01*-neutral, i.e., to buy F^R face amount of TIPS such that

$$F^R \times \frac{.081}{100} = 100\text{mm} \times \frac{.067}{100}$$

$$F^R = 100\text{mm} \times \frac{.067}{.081} = \$82.7\text{mm} \quad (6.1)$$

This hedge ensures that if the yield on the TIPS and the nominal bond both increase or decrease by the same number of basis points, the trade will

TABLE 6.1 Yields and *DV01*s of a TIPS and a Nominal U.S. Treasury as of May 28, 2010

Bond	Yield (%)	<i>DV01</i>
TIPS 1 $\frac{7}{8}$ s of 7/15/19	1.237	.081
3 $\frac{5}{8}$ s of 8/15/19	3.275	.067




FIGURE 6.1 Regression of Changes in the Yield of the Treasury 3 $\frac{5}{8}$ s of August 15, 2019, on Changes in the Yield of the TIPS 1.875s of July 15, 2019, from August 17, 2009, to July 2, 2010

neither make nor lose money. But the trader has doubts about this choice because changes in yields on TIPS and nominal bonds may very well not be one-for-one. To investigate, the trader collects data on daily changes in yield of these two bonds from August 17, 2009, to July 2, 2010, which are then graphed in Figure 6.1, along with a regression line, to be discussed shortly. It is immediately apparent from the graph that, for example, a five basis-point change in the yield of the TIPS does not imply, with very high confidence, a unique change in the nominal yield, nor even an average change of five basis points. In fact, while the daily change in the real yield was about five basis points several times over the study period, the change in the nominal yield over those particular days ranged from 2.2 to 8.4 basis points. This lack of a one-to-one yield relationship calls the *DV01* hedge into question. For context, by the way, it should be noted that graphing the changes in the yield of one nominal Treasury against changes in the yield of another, of similar maturity, would result in data points much more tightly surrounding the regression line.

With respect to improving on the *DV01* hedge, there is not much the trader can do about the dispersion of the change in the nominal yield for a given change in the real yield. That is part of the risk of the trade and will be discussed later. But the trader can estimate the average change in the nominal yield for a given change in the real yield and adjust the *DV01* hedge accordingly. For example, were it to turn out—as it will—that the nominal

yield in the data changes by 1.0189 basis points per basis-point change in the real yield, the trader could adjust the hedge such that

$$F^R \times \frac{.081}{100} = 100\text{mm} \times \frac{.067}{100} \times 1.0189$$

$$F^R = \$100\text{mm} \times \frac{.067}{.081} \times 1.0189 = \$84.3\text{mm} \quad (6.2)$$

Relative to the DV01 hedge of \$82.7 million in (6.1), the hedge in (6.2) increases the amount of TIPS to compensate for the empirical fact that, on average, the nominal yield changes by more than one basis point for every basis-point change in the real yield.

The next subsection introduces *regression analysis*, which is used both to estimate the coefficient 1.0189, used in equation (6.2), and to assess the properties of the resulting hedge.

Least-Squares Regression Analysis

Let Δy_t^N and Δy_t^R be the changes in the yields of the nominal and real bonds, respectively, and assume that

$$\Delta y_t^N = \alpha + \beta \Delta y_t^R + \epsilon_t \quad (6.3)$$

According to equation (6.3), changes in the real yield, the *independent variable*, are used to predict changes in the nominal yield, the *dependent variable*. The intercept, α , and the slope, β , need to be estimated from the data. The error term ϵ_t is the deviation of the nominal yield change on a particular day from the change predicted by the model. *Least-squares estimation* of (6.3), to be discussed presently, requires that the model be a true description of the dynamics in question and that the errors have the same probability distribution, are independent of each other, and are uncorrelated with the independent variable.¹

¹Since the nominal rate is the real rate plus the inflation rate, the error term in equation (6.3) contains the change in the inflation rate. Therefore, the assumption that the independent variable be uncorrelated with the error term requires here that the real rate be uncorrelated with the inflation rate. This is a tolerable, though far from ideal, assumption: the inflation rate can have effects on the real economy and, consequently, on the real rate.

If the regression were specified such that the real rate were the dependent variable and the nominal rate the independent variable, the requirement that the error and the dependent variable be uncorrelated would certainly not be met. In that case, the error term contains the inflation rate and there is no credible argument that the nominal rate is even approximately uncorrelated with the inflation rate. Consequently, a more advanced estimation procedure would be required, like that of *instrumental variables*.

As an example of the relationship between the nominal and real yields in (6.3), say that the parameters estimated with the data, denoted $\hat{\alpha}$ and $\hat{\beta}$, are 0 and 1.02 respectively. Then, if Δy_t^R is 5 basis points on a particular day, the predicted change in the nominal yield, written $\Delta \hat{y}_t^N$, is

$$\begin{aligned}\Delta \hat{y}_t^N &= \hat{\alpha} + \hat{\beta} \Delta y_t^R \\ &= 0 + 1.02 \times 5 = 5.1\end{aligned}\tag{6.4}$$

Furthermore, should it turn out that the nominal yield changes by 5.5 basis points on that day, then the realized error that day, written $\hat{\epsilon}_t$, following equation (6.3), is defined as

$$\begin{aligned}\hat{\epsilon}_t &= \Delta y_t^N - \hat{\alpha} - \hat{\beta} \Delta y_t^R \\ &= \Delta y_t^N - \Delta \hat{y}_t^N\end{aligned}\tag{6.5}$$

In this example,

$$\hat{\epsilon}_t = 5.5 - 5.1 = .4\tag{6.6}$$

Least-squares estimation of α and β finds the estimates $\hat{\alpha}$ and $\hat{\beta}$ that minimize the sum of the squares of the realized error terms over the observation period,

$$\sum_t \hat{\epsilon}_t^2 = \sum_t (\Delta y_t^N - \hat{\alpha} - \hat{\beta} \Delta y_t^R)^2\tag{6.7}$$

where the equality follows from (6.5). The squaring of the errors ensures that offsetting positive and negative errors are not considered as acceptable as zero errors and that large errors in absolute values are penalized substantially more than smaller errors.

Least-squares estimation is available through many statistical packages and spreadsheet add-ins. A typical summary of the regression output from estimating equation (6.3) using the data in Figure 6.1 is given in Table 6.2. The $\hat{\beta}$ reported in the table is 1.0189, which says that, over the sample period, the nominal yield increases by 1.0189 basis points per basis-point increase in real yields. The constant term of the regression, $\hat{\alpha}$, is not very different from zero, which is typically the case in regressions of changes in a yield on changes in a comparable yield. The economic interpretation of this regularity is that a yield does not usually trend up or down while a comparable yield is not changing.

Table 6.2 reports standard errors of $\hat{\alpha}$ and $\hat{\beta}$ of .2529 and .0525, respectively. Under the assumptions of least squares and the availability of sufficient data, the parameters $\hat{\alpha}$ and $\hat{\beta}$ are normally distributed with means

TABLE 6.2 Regression Analysis of Changes in the Yield of the $3\frac{5}{8}\%$ s of August 15, 2019, on the Changes in Yield of the TIPS $1\frac{7}{8}\%$ s of July 15, 2019, From August 17, 2009, to July 2, 2010

No. of Observations	229	
R-Squared	56.3%	
Standard Error	3.82	
Regression Coefficients	Value	Std. Error
Constant ($\hat{\alpha}$)	0.0503	.2529
Change in Real Yield ($\hat{\beta}$)	1.0189	.0595

equal to the true model values, α and β , respectively, and with standard deviations that can be estimated as the standard errors given in the table. Therefore, relying on the properties of the normal distribution, the confidence interval $.0503 \pm 2 \times .2529$ or $(-.4555, .5561)$ has a 95% chance of falling around the true value α . And since this confidence interval does include the value zero, one cannot reject the statistical hypothesis that $\alpha = 0$. Similarly, the 95% confidence interval with respect to β is $1.0189 \pm 2 \times .0595$, or $(.8999, 1.1379)$. So, while regression hedging makes heavy use of the point estimate $\hat{\beta} = 1.0189$, the true value of β may very well be somewhat higher or lower.

Substituting the estimated coefficients from Table 6.2 into the predicted regression equation in the first line of (6.4),

$$\begin{aligned}\Delta \hat{y}_t^N &= \hat{\alpha} + \hat{\beta} \Delta y_t^R \\ \Delta \hat{y}_t^N &= .0503 + 1.0189 \times \Delta y_t^R\end{aligned}\quad (6.8)$$

This relationship is known as the *fitted regression line* and is the straight line through the data that appears in Figure 6.1.

Table 6.2 reports two other useful statistics, the R-squared and the standard error of the regression. The R-squared in this case is 56.3%, which means that 56.3% of the variance of changes in the nominal yield can be explained by the model. In a one-variable regression, the R-squared is just the square of the correlation of the two changes, so the correlation between changes in the nominal and real yields is the square root of 56.3% or about 75%. This is a relatively low number compared with typical correlations between changes in two nominal yields, echoing the comment made in reference to the relatively wide dispersion of the points around the regression line in Figure 6.1.

The second useful statistic reported in Table 6.2 is the standard error of the regression, denoted here by $\hat{\sigma}$ and given as 3.82 basis points. Algebraically, $\hat{\sigma}$ is essentially the standard deviation of the realized error terms

$\hat{\epsilon}_t$,² defined in equation (6.5). Graphically, each $\hat{\epsilon}_t$ is the vertical line from a data point directly down or up to the regression line and $\hat{\sigma}$ is essentially the standard deviation of these distances. Either way, $\hat{\sigma}$ measures how well the model fits the data in the same units as the dependent variable, which, in this case, are basis points.

The Regression Hedge

The use of the regression coefficient in the hedging example of this section was discussed in the development of equation (6.2). More formally, denoting the face amounts of the real and nominal bonds by F^R and F^N and their DV01s by $DV01^R$ and $DV01^N$, the regression-based hedge, characterized earlier as the DV01 hedge adjusted for the average change of nominal yields relative to real yields, can be written as follows:

$$F^R = -F^N \times \frac{DV01^N}{DV01^R} \times \hat{\beta} \quad (6.9)$$

It turns out, however, that this regression hedge has an even stronger justification. The profit and loss (P&L) of the hedged position over a day is

$$-F^R \times \frac{DV01^R}{100} \Delta y_t^R - F^N \times \frac{DV01^N}{100} \Delta y_t^N \quad (6.10)$$

Appendix A in this chapter shows that the hedge of equation (6.9) minimizes the variance of the P&L in (6.10) over the data set shown in Figure 6.1 and used to estimate the regression parameters of Table 6.2.

In the example of this section, $F^N = -\$100\text{mm}$, $\hat{\beta} = 1.0189$, $DV01^N = .067$, and $DV01^R = .081$, so, from (6.9), as derived before, $F^R = \$84.279\text{mm}$. Because the estimated β happens to be close to one, the regression hedge of about \$84.3 million is not very different from the DV01 hedge of \$82.7 million calculated earlier. In fact, some practitioners would describe this hedge in terms of the DV01 hedge. Rearranging the terms of (6.9),

$$\frac{-F^R \times DV01^R}{F^N \times DV01^N} = \hat{\beta} = 101.89\% \quad (6.11)$$

²If the number of observations is n , the standard error of the regression is actually defined as the square root of $\frac{\sum_t \hat{\epsilon}_t^2}{(n-2)}$. The average of the $\hat{\epsilon}_t$ in a regression with a constant is zero by construction, so the standard error of the regression differs from the standard deviation of the errors only because of the division by $n-2$ instead of $n-1$.

In words, the risk of the (TIPS) hedging portfolio, measured by $DV01$, is 101.89% of the risk of the underlying (nominal) position, measured by $DV01$. Alternatively, the *risk weight* of the hedge portfolio is 101.89%. This terminology does connect the hedge to the common $DV01$ benchmark but is somewhat misleading because the whole point of the regression-based hedge is that the risks of the two securities cannot properly be measured by the $DV01$ alone. It should also be noted at this point that the regression-based and $DV01$ hedges are certainly not always this close in magnitude, even in other cases of hedging TIPS *versus* nominals, as will be illustrated in the next subsection.

An advantage of the regression framework for hedging is that it automatically provides an estimate of the volatility of the hedged portfolio. To see this, substitute F^R from (6.9) into the P&L expression (6.10) and rearrange terms to get the following expression for the P&L of the hedged position:

$$-F^N \times \frac{DV01^N}{100} (\Delta y_t^N - \hat{\beta} \Delta y_t^R) \quad (6.12)$$

From the definition of $\hat{\epsilon}_t$ in (6.5), the term in parentheses equals $\hat{\epsilon}_t + \hat{\alpha}$. But since $\hat{\alpha}$ is typically not very important, the standard error of the regression $\hat{\sigma}$ can be used to approximate the standard deviation of $\Delta y_t^N - \hat{\beta} \Delta y_t^R$. Hence, the standard deviation of the P&L in (6.12) is approximately

$$F^N \times \frac{DV01^N}{100} \times \hat{\sigma} \quad (6.13)$$

In the present example, recalling that the standard error of the regression can be found in Table 6.2, the daily volatility of the P&L of the hedged portfolio is approximately

$$\$100\text{mm} \times \frac{.067}{100} \times 3.82 = \$255,940 \quad (6.14)$$

The trader would have to compare this volatility with an expected gain to decide whether or not the risk-return profile of the trade is attractive.

The Stability of Regression Coefficients Over Time

An important difficulty in using regression-based hedging in practice is that the hedger can never be sure that the hedge coefficient, β , is constant over time. Put another way, the errors around the regression line might be random outcomes around a stable relationship, as described by equation (6.3), or they might be manifestations of a changing relationship. In the former

situation a hedger can safely continue to use a previously estimated $\hat{\beta}$ for hedging while, in the latter situation, the hedger should re-estimate the hedge coefficient with more recent data, if available, or with data from a past, more relevant time period. But how can the hedger know which situation prevails?

A useful start for thinking about the stability of an estimated regression coefficient is to estimate that coefficient over different periods of time and then observe if the result is stable or not. To this end, with the same data as before, Figure 6.2 graphs $\hat{\beta}$ for regressions over rolling 30-day windows. This means that the full data set of changes from August 18, 2009, to July 2, 2010, is used in 30-day increments, as follows: the first $\hat{\beta}$ comes from a regression of changes from August 18, 2009, to September 28, 2009; the second $\hat{\beta}$ from that regression from August 19, 2009, to September 29, 2009, etc.; and the last $\hat{\beta}$ from May 24, 2010, to July 2, 2010. The estimates of β in the figure certainly do vary over time, but the range of .75 to 1.29 is not extremely surprising given the previously computed 95% confidence interval with respect to β of (.8999, 1.1379). More troublesome, perhaps, is the fact that the most recent values of $\hat{\beta}$ have been trending up, which may indicate a change in regime in which even higher values of β characterize the relationship between nominal and real rates.

For a bit more perspective before closing this subsection, the period February 15, 2000, to February 15, 2002, when rates were substantially higher, was characterized by significantly higher levels of $\hat{\beta}$ and higher levels of uncertainty with respect to the regression relationship. The two bonds




FIGURE 6.2 Rolling 30-Day Regression Coefficient for the Change in Yield of the Treasury 3 $\frac{5}{8}$ s of August 15, 2019, on the Change in Yield of the TIPS 1 $\frac{7}{8}$ s of July 15, 2019

TABLE 6.3 Regression Analysis of Changes in the Yield of the 6½s of February 15, 2010, on the Changes in Yield of the TIPS 4¼s of January 15, 2010, From February 15, 2000, to February 15, 2002

No. of Observations	519	
R-Squared	43.0%	
Standard Error	4.70	
Regression Coefficients	Value	Std. Error
Constant ($\hat{\alpha}$)	-.0267	.2067
Change in Real Yield ($\hat{\beta}$)	1.5618	.0790

used in this analysis are the TIPS 4¼s of January 15, 2010, and the Treasury 6½s of February 15, 2010. Summary statistics for the regression of changes in yields of the nominal 6½s on the real 4¼s are given in Table 6.3.

Compared with Table 6.2, the estimated β here is 50% larger and the precision of this regression, measured by the R-squared or the standard error of the regression, is substantially worse. The contrast across periods again emphasizes the potential pitfalls of relying on estimated relationships persisting over time. This does not imply, of course, that blindly assuming a β of one, as in DV01 hedging, is a generally superior approach.

TWO-VARIABLE REGRESSION-BASED HEDGING

To illustrate regression hedging with two independent variables, this section considers the case of a market maker in EUR interest rate swaps. An algebraic introduction is followed by an empirical analysis.

The market maker in question has bought or received fixed in relatively illiquid 20-year swaps from a customer and needs to hedge the resulting interest rate exposure. Immediately paying fixed or selling 20-year swaps would sacrifice too much if not all of the spread paid by the customer, so the market maker chooses instead to sell a combination of 10- and 30-year swaps. Furthermore, the market maker is willing to rely on a two-variable regression model to describe the relationship between changes in 20-year swap rates and changes in 10- and 30-year swap rates:

$$\Delta y_t^{20} = \alpha + \beta^{10} \Delta y_t^{10} + \beta^{30} \Delta y_t^{30} + \epsilon_t \quad (6.15)$$

Equation (6.15) can be estimated by least squares, analogously to the single-variable case, by minimizing

$$\sum_t (\Delta y_t^{20} - \hat{\alpha} - \hat{\beta}^{10} \Delta y_t^{10} - \hat{\beta}^{30} \Delta y_t^{30})^2 \quad (6.16)$$

with respect to the parameters $\hat{\alpha}$, $\hat{\beta}^{10}$, and $\hat{\beta}^{30}$. The estimation of these parameters then provides a predicted change for the 20-year swap rate:

$$\Delta \hat{y}_t^{20} = \hat{\alpha} + \hat{\beta}^{10} \Delta y_t^{10} + \hat{\beta}^{30} \Delta y_t^{30} \quad (6.17)$$

To derive the notional face amount of the 10- and 30-year swaps, F^{10} and F^{30} , respectively, required to hedge F^{20} face amount of the 20-year swaps, generalize the reasoning given in the single-variable case as follows. Write the P&L of the hedged position as

$$-F^{20} \frac{DV01^{20}}{100} \Delta y_t^{20} - F^{10} \frac{DV01^{10}}{100} \Delta y_t^{10} - F^{30} \frac{DV01^{30}}{100} \Delta y_t^{30} \quad (6.18)$$

Then substitute the predicted change in the 20-year rate from (6.17) into (6.18), retaining only the terms depending on Δy_t^{10} and Δy_t^{30} , to obtain

$$\begin{aligned} & \left[-F^{20} \frac{DV01^{20}}{100} \hat{\beta}^{10} - F^{10} \frac{DV01^{10}}{100} \right] \Delta y_t^{10} \\ & + \left[-F^{20} \frac{DV01^{20}}{100} \hat{\beta}^{30} - F^{30} \frac{DV01^{30}}{100} \right] \Delta y_t^{30} \end{aligned} \quad (6.19)$$

Finally, choose F^{10} and F^{30} to set the terms in brackets equal to zero, i.e., to eliminate the dependence of the predicted P&L on changes in the 10- and 30-year rates. This leads to two equations with the following solutions:

$$F^{10} = -F^{20} \frac{DV01^{20}}{DV01^{10}} \hat{\beta}^{10} \quad (6.20)$$

$$F^{30} = -F^{20} \frac{DV01^{20}}{DV01^{30}} \hat{\beta}^{30} \quad (6.21)$$

As in the single-variable case, this 10s-30s hedge of the 20-year can be expressed in terms of risk weights. More specifically, the $DV01$ risk in the 10-year part of the hedge and the $DV01$ risk in the 30-year part of the hedge can both be expressed as a fraction of the $DV01$ risk of the 20-year. Mathematically, these risk weights can be found by rearranging (6.20) and (6.21):

$$\frac{-F^{10} \times DV01^{10}}{F^{20} \times DV01^{20}} = \hat{\beta}^{10} \quad (6.22)$$

$$\frac{-F^{30} \times DV01^{30}}{F^{20} \times DV01^{20}} = \hat{\beta}^{30} \quad (6.23)$$

TABLE 6.4 Regression Analysis of Changes in the Yield of the 20-Year EUR Swap Rate on Changes in the 10- and 30-Year EUR Swap Rates From July 2, 2001, to July 3, 2006

No. of Observations	1281	
R-Squared	99.8%	
Standard Error	.14	
Regression Coefficients	Value	Std. Error
Constant ($\hat{\alpha}$)	-.0014	.0040
Change in 10-Year Swap Rate ($\hat{\beta}^{10}$)	.2221	.0034
Change in 30-Year Swap Rate ($\hat{\beta}^{30}$)	.7765	.0037

Proceeding now to the empirical analysis, the market maker, as of July 2006, performs an initial regression analysis using data on changes in the 10-, 20-, and 30-year EUR swap rates from July 2, 2001, to July 3, 2006. Summary statistics for the regression of changes in the 20-year EUR swap rate on changes in the 10- and 30-year EUR swap rates are given in Table 6.4. The statistical quality of these results, characteristic of all regressions of like rates, are far superior to those of the nominal against real yields of the previous section: the R-squared or percent variance explained by the regression is 99.8%; the standard error of the regression is only .14 basis points; and the 95% confidence intervals with respect to the two coefficients are extremely narrow, i.e., (.2153, .2289) for the 10-year and (.7691, .7839) for the 30-year. Lastly, in a result similar to those of the regressions of the previous section, the constant is insignificantly different from zero.

Applying the risk-weight interpretation of the regression coefficients given in equations (6.22) and (6.23), the results in Table 6.4 say that 22.21% of the DV01 of the 20-year swap should be hedged with a 10-year swap and 77.65% with a 30-year swap. The sum of these weights, 99.86%, happens to be very close to one, meaning that the DV01 of the regression hedge very nearly matches the DV01 of the 20-year swap, although this certainly need not be the case: minimizing the variance of the P&L of a hedged position, when rates are not assumed to move in parallel, need not result in a DV01-neutral portfolio.

Tight as the in-sample regression relationship seems to be, the real test of the hedge is whether it works out-of-sample.³ To this end, Figure 6.3 tracks the errors of the hedge over time. All of these errors are computed as

³The phrase *in-sample* refers to behavior within the period of estimation, in this case July 2, 2001, to July 3, 2006. The phrase *out-of-sample* refers to behavior outside the period of estimation, usually after but possibly before that period as well.




FIGURE 6.3 In- and Out-of-Sample Errors for a Regression of Changes of 20-Year and 10- and 30-Year EUR Swap Rates with Estimation Period July 2, 2001, to July 3, 2006

the realized change in the 20-year yield minus the predicted change for that yield based on the estimated regression in Table 6.4:

$$\hat{\epsilon}_t = \Delta y_t^{20} - (-.0014 + .2221\Delta y_t^{10} + .7765\Delta y_t^{30}) \quad (6.24)$$

The errors to the left of the vertical dotted line are in-sample in that the same Δy_t^{20} used to compute $\hat{\epsilon}_t$ in (6.24) were also used to compute the coefficient estimates $-.0014$, $.2221$, and $.7765$. In other words, it is not that surprising that the $\hat{\epsilon}_t$ to the left of the dotted line are small because the regression coefficients were estimated to minimize the sum of squares of these errors. By contrast, the errors to the right of the dotted line are out-of-sample: these $\hat{\epsilon}_t$ are computed from realizations of Δy_t^{20} after July 3, 2006, but using the regression coefficients estimated over the period from July 2, 2001, to July 3, 2006. It is, therefore, the size and behavior of these out-of-sample errors that provide evidence as to the stability of the estimated coefficients over time.

From inspection of Figure 6.3 the out-of-sample errors are indeed small, for the most part, until August and September 2008, a peak in the financial crisis of 2007–2009. After then the daily errors ran as high as about four basis points and as low as about -5.3 basis points. And while the accuracy of the relationship seems to have recovered somewhat to the far right-end of the graph, by the summer of 2009, the errors there are not nearly so well behaved as at the start of the out-of-sample period.

It is obvious and easy to say that the market maker, during the turbulence of a financial crisis, should have replaced the regression of Table 6.4 and the resulting hedging rule. But replace these with what? What does the market maker do at that time, before there exist sufficient post-crisis data points? And what does the market maker do after the worst of the crisis: estimate a regression from data during the crisis or revert to some earlier, more stable period? These are the kinds of issues that make regression hedging an art rather than a science. In any case, it should again be emphasized that avoiding these issues by blindly resorting to a one-security *DV01* hedge, or a two-security *DV01* hedge with arbitrarily assigned risk weights, like 50%-50%, is even less satisfying.

LEVEL VERSUS CHANGE REGRESSIONS

When estimating regression-based hedges, some practitioners regress changes in yields on changes in yields, as in the previous sections, while others prefer to regress yields on yields. Mathematically, in the single-variable case, the level-on-level regression with dependent variable y and independent variable x is

$$y_t = \alpha + \beta x_t + \epsilon_t \quad (6.25)$$

while the change-on-change regression is⁴

$$y_t - y_{t-1} = \Delta y_t = \beta \Delta x_t + \Delta \epsilon_t \quad (6.26)$$

By theory that is beyond the scope of this book, if the error terms ϵ_t are independently and identically distributed random variables with mean zero and are uncorrelated with the independent variable, then so are the $\Delta \epsilon_t$, and least squares on either (6.25) or (6.26) will result in coefficient estimators that are *unbiased*,⁵ *consistent*,⁶ and *efficient*, i.e., of *minimum variance*, in the class of linear estimators. If the error terms of either specification are not independent of each other, however, then the least-squares coefficients of that specification are not necessarily efficient, but retain their unbiasedness and consistency.

⁴It is usual to include a constant term in the change-on-change regression, but for the purposes of this section, to maintain consistency across the two specifications, this constant term is omitted.

⁵An unbiased estimator of a parameter is such that its expectation equals the true value of that parameter.

⁶A consistent estimator of a parameter, with enough data, becomes arbitrarily close to the true value of the parameter.

To illustrate the economics behind the assumption that error terms are independent of each other, say that $\alpha = 0$, that $\beta = 1$, that y is the yield on a coupon bond, and that x is the yield on another, near-maturity coupon bond. Say further that the yield on the x -bond was 5% yesterday and 5% again today while the yield on the y -bond was 1% yesterday. Because the yield on the x -bond is 5% today, the level equation (6.25) predicts that the yield on the y -bond will be 5% today, despite its being 1% yesterday. But if the market yield was so far off yesterday's prediction, with a realized error of -4% , then it is more likely that the error today will be not far from -4% and that the yield of the y -bond yield will be closer to 1% than the 5% predicted by (6.25). Put another way, the errors in (6.25) are not likely to be independent of each other, as assumed, but rather persistent, or correlated over time.

The change regression (6.26) assumes the opposite extreme with respect to the errors, i.e., that they are completely persistent. Continuing with the example of the previous paragraph, with the yield on the y -bond at 1% yesterday and the yield on the x -bond unchanged from yesterday, the change regression predicts that y -bond will remain at 1%. But, as reasoned above, it is more likely that the y -bond yield will move some of the way back from 1% to 5%. Hence, the error terms in (6.26) are also unlikely to be independent of each other.

The first lesson to be drawn from this discussion is that because the error terms in both (6.26) and (6.25) are likely to be correlated over time, i.e., *serially correlated*, their estimated coefficients are not efficient. But, with nothing to gainsay the validity of the other assumptions concerning the error terms, the estimated coefficients of both the level and change specifications are still unbiased and consistent.

The second lesson to be drawn from the discussion of this section is that there is a more sensible way to model the relationship between two bond yields than either (6.26) or (6.25). In particular, model the behavior that the y -bond's yield will, on average, move somewhat closer from 1% to 5%. Mathematically, assume (6.25) with the error dynamics

$$\epsilon_t = \rho\epsilon_{t-1} + \nu_t \quad (6.27)$$

for some constant $\rho < 1$. Assumption (6.27) says that today's error consists of some portion of yesterday's error plus a new random fluctuation. In terms of the numerical example, if $\rho = 75\%$, then yesterday's error of -4% would generate an average error today of $75\% \times -4\%$ or -3% and, therefore, an expected y -bond yield of $5\% - 3\%$ or 2% . In this way the error structure (6.27) has the yield of the y -bond converging to its predicted value of 5% given the yield of the x -bond at 5%. While beyond the scope of this book, the procedure for estimating (6.25) with the error structure (6.27) is presented in many statistical texts.

PRINCIPAL COMPONENTS ANALYSIS

Overview

Regression analysis tries to explain the changes in the yield of one bond relative to changes in the yields of a small number of other bonds. It is often useful, however, to have a single, empirical description of the behavior of the term structure that can be applied across all bonds. *Principal Components* (PCs) provide such an empirical description.

To fix ideas, consider the set of swap rates from 1 to 30 years at annual maturities. One way to describe the time series fluctuations of these rates is through the variances of the rates and their pairwise covariances or correlations. Another way to describe the data, however, is to create 30 interest rate factors or components, where each factor describes a change in each of the 30 rates. So, for example, one factor might be a simultaneous change of 5 basis points in the 1-year rate, 4.9 basis points in the 2-year rate, 4.8 basis points in the 3-year rate, etc. Principal Components Analysis (PCA) sets up these 30 such factors with the following properties:

1. The sum of the variances of the PCs equals the sum of the variances of the individual rates. In this sense the PCs capture the volatility of this set of interest rates.
2. The PCs are uncorrelated with each other. While changes in the individual rates are, of course, highly correlated with each other, the PCs are constructed so that they are uncorrelated.
3. Subject to these two properties or constraints, each PC is chosen to have the maximum possible variance given all earlier PCs. In other words, the first PC explains the largest fraction of the sum of the variances of the rates; the second PC explains the next largest fraction, etc.

PCs of rates are particularly useful because of an empirical regularity: the sum of the variances of the first three PCs is usually quite close to the sum of variances of all the rates. Hence, rather than describing movements in the term structure by describing the variance of each rate and all pairs of correlation, one can simply describe the structure and volatility of each of only three PCs.

The next subsections illustrate PCs and their uses in the context of USD and then global swap markets. For interested readers, Appendix B in this chapter describes the construction of PCs with slightly more mathematical detail, using the simpler context of three interest rates and three PCs. Fully general and more mathematical descriptions are available in numerous other books and articles.

PCAs for USD Swap Rates

Figure 6.4 graphs the first three principal components from daily data on USD swap rates while Table 6.5 provides a selection of the same information in tabular form. Thirty different data series are used, one series for each annual maturity from one to 30 years, and the observation period spans from October 2001 to October 2008. (Data from more recent dates will be presented and discussed later in this section.)

Columns (2) to (4) in Table 6.5 correspond to the three PC curves in Figure 6.4. These components can be interpreted as follows. A one standard-deviation increase in the “Level” PC, given in column (2), is a simultaneous 3.80 basis point increase in the one-year swap rate, a 5.86 basis-point increase in the 2-year, etc., and a 5.38 basis-point increase in the 30-year. This PC is said to represent a “level” change in rates because rates of all maturities move up or down together by, very roughly, the same amount. A one standard-deviation increase in the “Slope” PC, given in column (3), is a simultaneous 2.74 basis-point *drop* in the 1-year rate, a 3.09 basis-point drop in the 2-year rate, etc., and a 6.74 basis-point *increase* in the 30-year rate. This PC is said to represent a “slope” change in rates because short-term rates fall while longer-term rates increase, or *vice versa*. Finally, a one standard-deviation increase in the “Short Rate” PC, given in column (4), is made up of simultaneous increases in short-term rates (e.g., one- and two-year terms), small decreases in intermediate-term rates




FIGURE 6.4 The First Three Principal Components from USD Swap Rates from October 2001 to October 2008

TABLE 6.5 Selected Results of Principal Components for the USD Swap Curve from October 1, 2001, to October 2, 2008. Units are basis points or percentages

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Term	PCs			PC Vol	Total Vol	% of PC Variance			PC Vol / Total Vol (%)
	Level	Slope	Short Rate			Level	Slope	Short Rate	
1	3.80	-2.74	1.48	4.91	4.96	59.8	31.0	9.1	99.05
2	5.86	-3.09	0.59	6.65	6.67	77.7	21.5	0.8	99.74
...	...	...	...	...	...	...	...	...	...
5	6.85	-1.53	-0.57	7.04	7.06	94.7	4.7	0.7	99.85
...	...	...	...	...	...	...	...	...	...
10	6.35	0.06	-0.34	6.36	6.37	99.7	0.0	0.3	99.83
...	...	...	...	...	...	...	...	...	...
20	5.69	0.82	0.14	5.75	5.75	97.9	2.0	0.1	99.95
...	...	...	...	...	...	...	...	...	...
30	5.38	1.09	0.39	5.51	5.52	95.6	3.9	0.5	99.79
Total	32.47	6.74	2.28	33.25	33.29	95.4	4.1	0.5	99.87

(e.g., 5- and 10-year terms), and small increases in long-term rates (e.g., 20- and 30-year terms). While this PC is often called a “curvature” change, because intermediate-term rates move in the opposite direction from short- and long-term rates, the short-term rates moves dominate. Hence, the third PC is interpreted here as an additional factor to describe movements in short-term rates.

One feature of the shape of the level PC warrants additional discussion. Short-term rates might be expected to be more volatile than longer-term rates because changes in short-term rates are determined by current economic conditions, which are relatively volatile, while longer-term rates are determined mostly by expectations of future economic conditions, which are relatively less volatile. But since the Board of Governors of the Federal Reserve System, like many other central banks, anchors the very short-term rate at some desired level, the volatility of very short-term rates is significantly dampened. The level factor, which, as will be discussed shortly, explains the vast majority of term structure movements, and reflects this behavior on the part of central banks: very short-term rates move relatively little. Then, at longer maturities, the original effect prevails and longer-term rates move less than intermediate and shorter-term rates.

Column (5) of Table 6.5 gives the combined standard deviation or volatility of the three principal components for a given rate, and column (6) gives the total or empirical volatility of that rate. For the one-year rate,

for example, recalling that the principal components are uncorrelated, the combined volatility, in basis points, from the three components is

$$\sqrt{3.80^2 + (-2.74)^2 + 1.48^2} = 4.91 \quad (6.28)$$

The total or empirical volatility of the one-year rate, however, computed directly from the time series data, is 4.96 basis points. Column (10) of the table gives the ratio of columns (5) and (6), which, for the 1-year rate is $\frac{4.9099}{4.9572}$ or 99.05%. (For readability, many of the entries of Table 6.5 are rounded although calculations are carried out to higher precision.)

Columns (7) through (9) of Table 6.5 give the ratios of the variance of each PC component to the total PC variance. For the 1-year rate, these ratios are $\frac{3.80^2}{4.91^2} = 59.9\%$; $\frac{(-2.74)^2}{4.91^2} = 31.1\%$; and $\frac{1.48^2}{4.91^2} = 9.1\%$.

Finally, the last row of the table gives statistics on the square root of the sum of the variances across rates of different maturities. The sum of the variances is not a particularly interesting economic quantity—it does not, for example, represent the variance of any interesting portfolio—but, as mentioned in the overview of PCA, this sum is used to ensure that the PCs capture all of the volatility of the underlying interest rate series.

Having explained the calculations of Figure 6.4 and Table 6.5, the text can turn to interpretation. First and foremost, column (10) of Table 6.5 shows that, for rates of all maturities, the three principal components explain over 99% of rate volatility. And, across all rates, the three PCs explain 99.87% of the sum of the variability of these rates. While these findings represent relatively recent data on U.S. swap rates, similarly high explanatory powers characterize the first three components of other kinds of rates, like U.S. government bond yields and rates in fixed income markets in other countries. These results provide a great deal of comfort to hedgers: while in theory many factors (and, therefore, securities) might be required to hedge the interest rate risk of a particular portfolio, in practice, three factors cover the vast majority of the risk.

Columns (7) through (9) of Table 6.5 show that the level component is far and away the most important in explaining the volatility of the term structure. The construction of principal components, described in the overview, does ensure that the first component is the most important component, but the extreme dominance of this component is a feature of the data. This finding is useful for thinking about the costs and benefits of adding a second or third factor to a one-factor hedging framework. Interestingly too, the dominance of the first factor is significantly muted in the very short end of the curve. This implies that hedging one short-term bond with another will not be so effective as hedging one longer-term bond with another. Or, put another way, relatively more factors or hedging

securities are needed to hedge portfolios that are concentrated at the short end of the curve. This makes intuitive sense in the context of the extensive information market participants have about near-term events and their effects on rates relative to the information they have on events further into the future.

Hedging with PCA and an Application to Butterfly Weights

A PCA-based hedge for a portfolio would proceed along the lines of the multi-factor approaches described in Chapter 5. Start with the current price of the portfolio under the current term structure. Then, shift each principal component in turn to obtain new term structures and new portfolio prices. Next, calculate an '01 with respect to each principal component using the difference between the respective shifted price and the original price. Finally, using these portfolio '01s and analogously constructed '01s for a chosen set of hedging securities, find the portfolio of hedging securities that neutralizes the risk of the portfolio to the movement of each PC.

PCA is particularly useful for constructing empirically-based hedges for large portfolios; it is impractical to perform and assess individual regressions for every security in a large portfolio. For illustration purposes, however, this subsection will illustrate how PCA is used, in practice, to hedge a *butterfly* trade. Most typically, butterfly trades use three securities and either buy the security of intermediate maturity and short the *wings* or short the intermediate security and buy the wings.

To take a relatively common butterfly, consider a trader who believes that the 5-year swap rate is too high relative to the 2- and 10-year swap rates and is, therefore, planning to receive in the 5-year and pay in the 2- and 10-year. As of May 28, 2010, the par swap rates and *DV01*s of the swaps of relevant terms are listed in Table 6.6. (The 30-year data will be used shortly.) To calculate the PCA hedge ratios, assume that the trader will receive on 100 notional amount of 5-year swaps and will trade F^2 and F^{10} notional amount of 2- and 10-year swaps. Using the data from Tables 6.5

TABLE 6.6 Par Swap Rates and *DV01*s
as of May 28, 2010

Term	Rate	<i>DV01</i>
2	1.235%	.0197
5	2.427%	.0468
10	3.388%	.0842
30	4.032%	.1731

and 6.6, the equation that neutralizes the overall portfolio's exposure to the level PC is

$$-F^2 \frac{.0197}{100} \times 5.86 - F^{10} \frac{.0842}{100} \times 6.35 - 100 \times \frac{.0468}{100} \times 6.85 = 0 \quad (6.29)$$

Similarly, the equation that neutralizes the overall exposure to the slope PC is

$$-F^2 \frac{.0197}{100} \times (-3.09) - F^{10} \frac{.0842}{100} \times .06 - 100 \times \frac{.0468}{100} \times (-1.53) = 0 \quad (6.30)$$

Solving, $F^2 = -120.26$ and $F^{10} = -34.06$ or, in terms of risk weights relative to the DV01 of the five-year swap,

$$\frac{120.26 \times \frac{.0197}{100}}{.0468} = 50.6\% \quad (6.31)$$

$$\frac{34.06 \times \frac{.0842}{100}}{.0468} = 61.3\% \quad (6.32)$$

In words, the DV01 of the five-year swap is hedged 50.6% by the two-year swap and 61.3% by the 10-year swap. Note that the sum of the risk weights is not 100%: the hedge neutralizes exposures to the level and slope PCs, not exposures to parallel shifts. To the extent that the term structure changes as assumed, i.e., as some combination of the first two PCs, then the hedge will work exactly. On the other hand, to the extent that the actual change deviates from a combination of these two PCs, the hedge will not, *ex post*, have fully hedged interest rate risk.

Hedging the interest rate risk of the five-year swap with two other swaps is not uncommon, a practice supported by the large fraction of rate variance explained by the first two PCs. A trader might also decide, however, to hedge the third PC as well. A hedge against the first three PCs, found by generalizing the two-security hedge just discussed, gives rise to risk weights of 28.1%, 139.1%, and -67.4% in the 2-, 10-, and 30-year swaps, respectively, i.e., pay in the 2- and 10-year, but receive in the 30-year.

Is hedging the third PC worthwhile? The answer depends on the trader's risk preferences, but the following analysis is useful. Say that the trader hedges the first two components alone and then the third component

experiences a one standard-deviation decrease. The P&L of the trade, per 100 face amount of the 5-year swap, would be

$$\left[-120.26 \times \frac{.0197}{100} \times .59 + 100 \times \frac{.0468}{100} \times (-.57) - 34.06 \times \frac{.0842}{100} \times (-.34) \right] \\ = -.031 \quad (6.33)$$

or, for a two standard-deviation move, a loss of a bit more than 6 cents per 100 face amount of the 5-year swap. As these two standard deviations of short rate risk equates to not even 1.5 basis points of convergence of the 5-year swap, a trader might very well not bother with this third leg of the hedge.

Principal Component Analysis of EUR, GBP, and JPY Swap Rates

Figures 6.5 to 6.7 show the first three PCs for the EUR, GBP, and JPY swap rate curves over the same sample period as the USD PCs in Figure 6.4. The striking fact about these graphs is that the shape of the PCs are very much the same across USD, EUR, and GBP. The only significant difference is in magnitudes, with the USD level component entailing larger-sized moves than the level components of EUR and GBP. The PCs of the JPY curve are certainly similar to those of these other countries, but the level component in JPY does not have the same hump: in JPY the first PC does not peak at the five-year maturity point as do the other curves, but increases monotonically with




FIGURE 6.5 The First Three Principal Components from EUR Swap Rates from October 2001 to October 2008




FIGURE 6.6 The First Three Principal Components from GBP Swap Rates from October 2001 to October 2008

maturity before ultimately leveling off. The significance of this difference in shape will be discussed in the next subsection.

The Shape of PCs Over Time

As with any empirically based hedging methodology, a decision has to be made about the relevant time period over which to estimate parameters.




FIGURE 6.7 The First Three Principal Components from JPY Swap Rates from October 2001 to October 2008




FIGURE 6.8 The First Principal Component in USD and EUR Swap Rates Estimated from October 2001 to October 2008 and from October 2008 to October 2010

This is an issue for regression-based methods, as discussed in this chapter, and it is no less an issue for PCA. As will be discussed in this subsection, the qualitative shapes of PCs have, until very recently, remained remarkably stable. This does not imply, however, that differences in PCs estimated over different time periods can be ignored in the sense that they have no important effects on the quality of hedges. But having made this point, the text focuses on the relatively recent changes in the shapes of PCs around the world.

Figure 6.4 showed the first three USD PCs computed over the period 2001 to 2008, but, for quite some time, the qualitative shapes of these PCs was pretty much the same.⁷ The volatility of rates has changed over time, and with it the magnitude or height of the PC curves, but the qualitative shapes have not changed much. Most recently, however, there has been a qualitative change to the shape of the first PC in USD, EUR, and GBP. In fact, these shapes have become more like the past shape of the first PC in JPY!

Figures 6.8 and 6.9 contrast the level PC over the historical period October 2001 to October 2008 with that of the post-crisis period, October 2008 to October 2010. Figure 6.8 makes the comparison for USD and EUR while

⁷See, for example, Figure 2 of Bulent Baygun, Janet Showers, and George Cherpelis, Salomon Smith Barney, “Principles of Principal Components,” January 31, 2000. The shapes of the three PCs in that graph, covering the period from January 1989 to February 1998, are qualitatively extremely similar to those of Figure 6.4 in this chapter.




FIGURE 6.9 The First Principal Component in GBP and JPY Swap Rates Estimated from October 2001 to October 2008 and from October 2008 to October 2010

Figure 6.9 does the same for GBP and JPY. The historical maximum of the level PC at a term of about five years in USD, EUR, and GBP has been pushed out dramatically to 10 years and beyond. In fact, these shapes now more closely resemble the level PC of JPY over the earlier estimation period. One explanation for this is the increasing certainty that central banks will maintain easy monetary conditions and low rates for an extended period of time. This dampens the volatility of short- and intermediate-term rates relative to that of longer-term rates, lowers the absolute volatility of short-term rates, and increases the volatility of long-term rates, reflecting the uncertainty of the ultimate results of central bank policy. Meanwhile, the level PC for JPY in the most recent period has become even more pronouncedly upward-sloping, consistent with an even longer period of central-bank control over the short-term rate.

APPENDIX A: THE LEAST-SQUARES HEDGE MINIMIZES THE VARIANCE OF THE P&L OF THE HEDGED POSITION

The P&L of the hedged position, given in (6.10) and repeated here, is

$$-F^R \times \frac{DV01^R}{100} \Delta y_t^R - F^N \times \frac{DV01^N}{100} \Delta y_t^N \quad (6.34)$$

Let $V(\cdot)$ and $Cov(\cdot, \cdot)$ denote the variance and covariance functions. The variance of the P&L expression in (6.34) is

$$\begin{aligned} & \left(F^R \times \frac{DV01^R}{100} \right)^2 V(\Delta y_t^R) + \left(F^N \times \frac{DV01^N}{100} \right)^2 V(\Delta y_t^N) \\ & + 2 \left(F^R \times \frac{DV01^R}{100} \right) \left(F^N \times \frac{DV01^N}{100} \right) Cov(\Delta y_t^R, \Delta y_t^N) \end{aligned} \quad (6.35)$$

To find the face amount F^R that minimizes this variance, differentiate (6.35) with respect to F^R and set the result to zero:

$$2F^R \left(\frac{DV01^R}{100} \right)^2 V(\Delta y_t^R) + 2F^N \frac{DV01^R}{100} \frac{DV01^N}{100} Cov(\Delta y_t^R, \Delta y_t^N) = 0 \quad (6.36)$$

Then, rearranging terms,

$$F^N \times DV01^N \times \frac{Cov(\Delta y_t^R, \Delta y_t^N)}{V(\Delta y_t^R)} = -F^R \times DV01^R \quad (6.37)$$

But, by the properties of least squares, not derived in this text,

$$\hat{\beta} \equiv \frac{Cov(\Delta y_t^R, \Delta y_t^N)}{V(\Delta y_t^R)} \quad (6.38)$$

Therefore, substituting (6.38) into (6.37) gives the regression hedging rule (6.9) of the text.

APPENDIX B: CONSTRUCTING PRINCIPAL COMPONENTS FROM THREE RATES

The goal of this appendix is to demonstrate the construction and properties of PCs with a minimum of mathematics. To this end, consider three swap rates, the 10-year, 20-year, and 30-year. Over some sample period, the volatilities of these rates, in basis points per day, are 4.25, 4.20, and 4.15. Furthermore, the correlations among these rates are given in the correlation matrix of Table 6.7.

The combination of data on volatilities and correlations are usefully combined into a *variance-covariance* matrix, denoted by V , where the element in the i^{th} row and j^{th} column gives the covariance of the rate of term

TABLE 6.7 Correlation Matrix for Swap Rate Example

Term	10-Year	20-Year	30-Year
10-Year	1.00	0.95	0.90
20-Year	0.95	1.00	0.99
30-Year	0.90	0.99	1.00

i with the rate of term j , or, the correlation of i and j times the standard deviation of i times the standard deviation of j . For example, the covariance of the 20-year swap rate with the 30-year swap rate is $.99 \times 4.20 \times 4.15$, or 17.26. The variance-covariance matrix for the example of this appendix is

$$V = \begin{pmatrix} 18.06 & 16.96 & 15.87 \\ 16.96 & 17.64 & 17.26 \\ 15.87 & 17.26 & 17.22 \end{pmatrix} \quad (6.39)$$

One use of a variance-covariance matrix is to write succinctly the variance of a particular portfolio of the relevant securities. Consider a portfolio with a total DV01 of .50 in the 10-year swap, -1.0 in the 20-year swap, and .60 in the 30-year swap. Without matrix notation, then, the dollar variance of the portfolio, denoted by σ^2 would be given by

$$\begin{aligned} \sigma^2 &= .5^2 4.25^2 + (-1)^2 4.20^2 + .6^2 4.15^2 \\ &\quad + 2 \times .5 \times (-1) \times .95 \times 4.25 \times 4.20 \\ &\quad + 2 \times .5 \times .6 \times .90 \times 4.25 \times 4.15 \\ &\quad + 2 \times (-1) \times .6 \times .99 \times 4.20 \times 4.15 \\ &= .464^2 \end{aligned} \quad (6.40)$$

With matrix notation, letting the transpose of the vector $\mathbf{w}$ be $\mathbf{w}' = (.5, -1, .6)$, the dollar variance of the portfolio is given more compactly by

$$\mathbf{w}'V\mathbf{w} = \begin{pmatrix} .5 & -1 & .6 \end{pmatrix} \begin{pmatrix} 18.06 & 16.96 & 15.87 \\ 16.96 & 17.64 & 17.26 \\ 15.87 & 17.26 & 17.22 \end{pmatrix} \begin{pmatrix} .5 \\ -1 \\ .6 \end{pmatrix} \quad (6.41)$$

Finally, note that the sum of the variances of the rates is $4.25^2 + 4.20^2 + 4.15^2 = 52.925$, or, for a measure of total volatility, take the square root of that sum to get 7.27 basis points.

Returning now to principal components, the idea is to create three factors that capture the same information as the variance-covariance matrix. The procedure is as follows. Denote the first principal component by the vector $\mathbf{a} = (a_1, a_2, a_3)'$. Then find the elements of this vector by maximizing $\mathbf{a}'\mathbf{V}\mathbf{a}$ such that $\mathbf{a}'\mathbf{a} = 1$. As mentioned in the PCA overview, this maximization ensures that, among the three PCs to be found, the first PC explains the largest fraction of the variance. The constraint, $\mathbf{a}'\mathbf{a} = 1$, along with a similar constraint placed on the other PCs, will ensure that the total variance of the PCs equals the total variance of the underlying data. Performing this maximization, which can be done with the solver in Excel, $\mathbf{a} = (.5758, .5866, .5696)$. Note that the variance of this first component is $\mathbf{a}'\mathbf{V}\mathbf{a} = 51.041$ which is $\frac{51.041}{52.925}$ or 96.44% of the total variance of the rates.

The second principal component, denoted by the vector $\mathbf{b} = (b_1, b_2, b_3)$ is found by maximizing $\mathbf{b}'\mathbf{V}\mathbf{b}$ such that $\mathbf{b}'\mathbf{b} = 1$ and $\mathbf{b}'\mathbf{a} = 0$. The maximization and the first constraint are analogous to those for finding the first principal component. The second constraint requires that the PC $\mathbf{b}$ is uncorrelated with the first PC, $\mathbf{a}$. Solving, gives $\mathbf{b} = (-.7815, .1902, .5941)$. Note that $\mathbf{b}'\mathbf{V}\mathbf{b} = 1.867$ which explains $\frac{1.867}{52.925}$ or 3.53% of the total variance of the rates.

Finally, the third PC, denoted by $\mathbf{c} = (c_1, c_2, c_3)$ is found by solving the three equations, $\mathbf{c}'\mathbf{c} = 1$; $\mathbf{c}'\mathbf{a} = 0$; and $\mathbf{c}'\mathbf{b} = 0$. The solution is $\mathbf{c} = (.2402, -.7872, .5680)$.

As will be clear in a moment, it turns out to be more intuitive to work with a different scaling of the PCs, namely, by multiplying each by its volatility. In the example, this means multiplying the first PC by $\sqrt{51.041}$ or 7.14; the second PC by $\sqrt{1.867}$ or 1.37; and the third by $\sqrt{.017}$ or .13. This gives the PCs, to be denoted $\tilde{\mathbf{a}}, \tilde{\mathbf{b}}, \tilde{\mathbf{c}}$, as recorded in Table 6.8.

Under this scaling the PCs have a very intuitive interpretation: a one standard-deviation increase of the first PC or factor is a 4.114 basis-point increase in the 10-year rate, a 4.191 basis-point increase in the 20-year rate, and a 4.069 basis-point increase in the 30-year rate. Similarly, a one standard-deviation increase of the second PC is a 1.068 basis-point drop in the 10-year rate, a .260 basis-point increase in the 20-year rate, and a .812-basis point increase in the 30-year rate. Finally, a one standard-deviation

TABLE 6.8 Transformed PCs for the Swap Rate Example

Term	1st PC	2nd PC	3rd PC
10-Year	4.114	-1.068	.032
20-Year	4.191	.260	-.103
30-Year	4.069	.812	.075

increase of the third PC constitutes changes of .032, $-.103$, and .075 basis points in each of the rates, respectively.

To appreciate the scaling of the PCs in Table 6.8, note the following implications:

- By construction, the PCs are uncorrelated. Hence, the volatility of the 10-year rate can be recovered from Table 6.8 as

$$\sqrt{4.114^2 + (-1.068)^2 + .032^2} = 4.25 \quad (6.42)$$

And the volatilities of the 20- and 30-year rates can be recovered equivalently.

- The variance of each PC is the sum of squares of its elements, or, its volatility is the square root of that sum of squares. For the three PCs,

$$\sqrt{4.114^2 + 4.191^2 + 4.069^2} = 7.14 \quad (6.43)$$

$$\sqrt{(-1.068)^2 + .260^2 + .812^2} = 1.37 \quad (6.44)$$

$$\sqrt{.032^2 + (-.103)^2 + .075^2} = .13 \quad (6.45)$$

- The square root of the sum of the variances of the PCs is the square root of the sum of the variances of the rates, which quantity was given above as 7.27 basis points:

$$\sqrt{7.14^2 + 1.37^2 + .13^2} = \sqrt{52.925} = 7.27 \quad (6.46)$$

- The volatility of any portfolio can be found by computing its volatility with respect to each of the PCs and then taking the square root of the sum of the resulting variances. Returning to the portfolio with DV01 weights of $\mathbf{w}' = (.5, -1, .6)$, its volatility with respect to each of the PCs can be computed as in equations (6.47) through (6.49). Then, adding the sum of these squares and taking the square root, gives a portfolio volatility of .464, as computed earlier from the variances and covariances.

$$\sqrt{(\mathbf{w}'\tilde{\mathbf{a}})^2} = \sqrt{(.5 \times 4.114 - 1 \times 4.191 + .6 \times 4.069)^2} = .3074 \quad (6.47)$$

$$\sqrt{(\mathbf{w}'\tilde{\mathbf{b}})^2} = \sqrt{(.5 \times (-1.068) - 1 \times .260 + .6 \times .812)^2} = .3068 \quad (6.48)$$

$$\sqrt{(\mathbf{w}'\tilde{\mathbf{c}})^2} = \sqrt{(.5 \times .032 - 1 \times (-.103) + .6 \times .075)^2} = .1640 \quad (6.49)$$

In summary, the PCs in Table 6.8 contain the same information as the variances and covariances, but have the interpretation of one standard-deviation changes in the level, slope, and short rate factors. Of course, the power of the methodology is evident not in a simple example like this, but when, as in the text, changes in 30 rates can be adequately expressed with changes in three factors.

Term Structure Models

Part One of this book showed how to price securities with fixed cash flows relative to the prices of other such securities. Part Two showed how to measure and hedge the interest rate risk of securities with fixed cash flows, and also described the hedging of more complex securities, provided that a pricing model for those securities had been made available. Term structure models, the subject of Part Three, are used for more general and complex pricing and hedging problems, including the following:

- Pricing a security with fixed cash flows relative to the prices of other securities with fixed cash flows when the security in question cannot, using the methods of Part One, be priced by arbitrage. An interesting example is pricing a 50-year swap when market swap rates are available out to only 30 years.
- Pricing a generic *interest rate contingent claim* relative to the prices of securities with fixed cash flows, where an interest rate contingent claim is a security whose cash flows depend on the level of interest rates. An option to purchase a bond at a fixed price at some future expiration date would be an example of this application: the value of the option depends on the price of the underlying bond and, therefore, on interest rates, as of the expiration date.
- Pricing an *exotic derivative* relative to the prices of *vanilla* derivatives.¹ An example here would be a derivatives desk pricing a Bermudan-style *swaption* (i.e., an option on a swap that may be exercised on a

¹Vanilla derivatives typically refer to the relatively liquid and simple LIBOR-based derivative products. These include caps and floors, European-style swaptions, and Eurodollar or Euribor futures, all of which are discussed in Part Four.

predetermined set of dates), while market prices are observable only for European-style swaptions.

- Expressing the risks of portfolios, which include contingent claims, in a simple and intuitive manner, or, equivalently, determining a practical way of hedging away the interest rate risk of such portfolios. For example, a risk manager responsible for the combined exposures of several trading desks has to express overall interest rate risk. And as another example, a relative value hedge fund, which buys securities it believes are cheap and sells securities it believes are rich, often has to hedge away any residual exposure to general movements in interest rates.
- Expressing a scenario of changes in fixed income markets in a simple and intuitive manner. A portfolio manager might want to know how a particular portfolio would perform should markets behave as they did in reaction to the Russian debt crisis of 1998. Since it would be difficult and not particularly instructive to replicate the change in every security's value over that historical event, some effective summary of those market changes is useful.

Pricing a security with fixed cash flows by arbitrage along the lines of Part One is relatively simple. Find the portfolio that replicates the cash flows of that security and conclude that the security's price has to equal the price of its replicating portfolio. Importantly, no assumptions about future interest rates have to be made in order to construct this replicating portfolio; since all cash flows are fixed, a portfolio that replicates a set of cash flows when rates are 10% also replicates those cash flows when rates are at 2%. And by the same logic, the composition of these replicating portfolios is *static*, i.e., it does not change over time.

For many securities, however, there does not exist a static, replicating portfolio. Consider an option on a bond. The cash flows of the option depend on the level of interest rates, so any replicating portfolio has to match the option's cash flows for any possible future evolution of interest rates. As will be seen in this part, to no surprise of readers who are familiar with the pricing of equity options, finding such a replicating portfolio typically depends on some assumptions about the future evolution of interest rates. Furthermore, the resulting replicating portfolio is *dynamic*, i.e., its composition changes over time.

The goal of Part Three is to show how to price contingent claims given assumptions about the future evolution of interest rates and then to show how to make reasonable assumptions in the first place. A particular collection of these assumptions is called a *term structure model*. Models that postulate only one source of interest rate risk, i.e., that one random variable determines the entire term structure of interest rates, are called one-factor models. Models that postulate two or more sources of risk are called multi-factor models.

Short-rate models, which can have one or more factors, posit an evolution of the short-term rate which, by arbitrage arguments, allows for the pricing of interest rate contingent claims. While not suitable or sufficiently general for many purposes, one-factor short-rate models are perfectly appropriate for certain applications. Also, one-factor models are traditional and very suitable for pedagogical use. Chapters 7 to 10, therefore, focus on this class of models.

Chapter 7 lays out the science of short-rate models, that is, how to price contingent claims by arbitrage once assumptions about the evolution of rates have been made. This material can be viewed both as a generalization of the pricing methodology of Part One and as an application to fixed income securities of the option-pricing methodology originally developed for equity options.

Chapter 8 uses the short-rate model framework to connect the observed shape of the term structure with expectations of future rates, the volatility of rates, and a risk premium on fixed income investments. This is a useful starting point for thinking about the term structure, to which can be added more sophisticated theories, mentioned briefly ahead, as well as market technicals that create buying and selling pressures—both persistent and temporary—at different parts of the curve.

Chapters 9 and 10 focus on the art of short-rate term structure modeling, that is, how to make reasonable assumptions about the future evolution of the short rate. The approach of these chapters is to start with the simplest of models and to generalize the component parts one at a time. These chapters are useful both for some of the models themselves, which can be used and implemented where appropriate, as well as for introducing the various model components that are the building blocks of more elaborate and sophisticated term structure models.

Chapter 11 leaves the class of one-factor short-rate models to present significantly more advanced material. The first half of the chapter introduces the *Gauss+* model, a multi-factor short-rate model that is particularly popular among relative value traders as a good compromise between the goal of capturing the realities of term structure behavior and the goal of building a model that is intuitive to use and relatively easy to implement. The purpose of including this model in the book is not only to introduce multi-factor short-rate models, but also to provide a determined reader with a useful workhorse for many applications that might arise. To that end, the presentation includes all necessary formulae along with a suggestion for empirically estimating the parameters of the model.

The second half of Chapter 11 introduces the (multi-factor) *LIBOR Market Model (LMM)*. Unlike short-rate models, LMM directly models the evolution of non-overlapping forward rates in order to price contingent claims. Because LMM does have many factors, and because the forward rates themselves are factors, the model automatically matches the term structure

of interest rates and, by design, is particularly easy to calibrate to the prices of vanilla derivatives. As a result, LMM is a particular favorite for pricing and hedging exotic derivatives. Since presentations of LMM tend to be extremely mathematical and technical, LMM was included in this book so as to introduce this popular modeling approach in a manner accessible to a broader audience. While the description here is, in theory, sufficient for a determined reader to implement a simple version of the model, it must be noted that vastly more complex specifications are used in practice and that much research has been done on the calibration and implementation of these models.

The various pricing models presented in the book can now be linked to practical problems and applications, including the list at the start of this introduction:

- **Black-Scholes style models**, presented in Chapter 18. The Black-Scholes model can be applied to interest rate contingent claims that have a relatively simple payoffs and that require calibration only to the price or rate of the underlying security and to its volatility. For example, the model is used by market makers and other traders to price and hedge caplets and floorlets, as well as European-style swaptions, bond options, and futures options.
- **One-factor short-rate models**. As will be seen in this part, these models can be implemented so as to accommodate relatively complex contingent claims. Because of the limitations of a single factor, however, these models are used only when simplicity of implementation is at a premium compared with precision, and when calibration to the term structure might be required but not calibration to more than a handful of volatility products. One example is the interpolation or extrapolation of prices or rates of liquid bonds or swaps so as to price less liquid bonds or swaps. Since this approach brings arbitrage pricing discipline to bear on the problem, some think it superior to purely mathematical approaches, like cubic splines. Another common application would be pricing callable bonds (excluding any handling of credit risk). First, the embedded options in these bonds are relatively complex, with, for example, calls that can be exercised on several dates at a declining schedule of call prices. Second, typical uses, like pricing the embedded call in new bond issues or assessing the risk of callable bonds in the context of a large portfolio, do not require great precision. Third, pricing is usually calibrated to the underlying issuer's curve and to a small number of swaptions that most resemble the features of the embedded call.
- **Multi-factor short-rate models**. These models are used where precision, ease of interpretation, and consistency of approach across a broad range of products are all important considerations, but where calibration to only a relatively small number of securities is required. Examples include

the following: relative value trading, in which security value is assessed and positions hedged relative to a fixed number of points on the term structure and a handful of volatility products; risk management, when the risk of large, cross-product portfolios is most desirably expressed in terms of a limited number of risk factors; and scenario analysis, where historic or hypothesized changes in fixed income markets have to be summarized in terms of a relatively small number of variables.

- **LMM.** Consistent with the description above, LMM is used by derivatives market makers and other derivatives traders to price and hedge exotics. For these applications precision is at a premium, as is calibration to a relatively large set of vanilla derivatives.

Before concluding this introduction, it should be emphasized that the models presented in this book are useful for the relative pricing of fixed income securities, not for determining whether bond or swap rates are, in some sense, too high or too low. In terms of some of the models presented, the assumption that the risk premium earned on fixed income investments is constant means that these models make no interesting predictions about expected returns. A substantial body of empirical evidence, however, contradicts the assumption of a constant risk premium. Investigators have shown that the risk premium changes as a function of a combination of interest rate factors² or, in addition, as a function of macroeconomic factors that are not captured by information in the term structure.³ Consequently, models incorporating a time-varying risk premium enable predictions about expected returns based on both interest rates and on other macroeconomic data. Some practitioners, particularly interest rate strategists, have considered such models in their work. The approach has not seen much use, however, for the pricing of contingent claims. In any case, this book makes no further mention of this promising line of research.

²See, for example, by J. H. Cochrane, and M. Piazzesi, "Bond Risk Premia" *American Economic Review*, 95, 2005, pp. 138–60.

³See, for example, by S. C. Ludvigson and S. Ng, "Macro Factors in Bond Risk Premia," *The Society for Financial Studies*, Oxford University Press, 2009.


The Science of Term Structure Models

This chapter uses a very simple setting to show how to price interest rate contingent claims relative to a set of underlying securities by arbitrage arguments. Unlike the arbitrage pricing of securities with fixed cash flows in Part One, the techniques of this chapter require strong assumptions about how interest rates evolve in the future. This chapter also introduces *option-adjusted spread* (OAS) as the most popular measure of deviations of market prices from those predicted by models.

RATE AND PRICE TREES


Assume that the six-month and one-year spot rates are 5% and 5.15% respectively. Taking these market rates as given is equivalent to taking the prices of a six-month bond and a one-year bond as given. Securities with assumed prices are called underlying securities to distinguish them from the contingent claims priced by arbitrage arguments.

Next, assume that six months from now the six-month rate will be either 4.50% or 5.50% with equal probability. This very strong assumption is depicted by means of a *binomial tree*, where “binomial” means that only two future values are possible:



Note that the columns in the tree represent dates. The six-month rate is 5% today, which will be called date 0. On the next date six months from now, which will be called date 1, there are two possible outcomes or *states of the world*. The 5.50% state will be called the *up state* while the 4.50% state will be called the *down state*.


Given the current term structure of spot rates (i.e., the current six-month and one-year rates), trees for the prices of six month and one-year zero-coupon bonds may be computed. The price tree for \$1,000 face value of the six-month zero is



since $\$1,000 / (1 + \frac{.05}{2}) = \$975.61$. (For easy readability, currency symbols are not included in price trees).

Note that in a tree for the value of a particular security, the maturity of the security falls with the date. On date 0 of the preceding tree the security is a six-month zero, while on date 1 the security is a maturing zero.

The price tree for \$1,000 face value of a one-year zero is the following:



The three date 2 prices of \$1,000 are, of course, the maturity values of the one-year zero. The two date 1 prices come from discounting this certain \$1,000 at the then-prevailing six-month rate. Hence, the date 1 up-state price is $\$1,000 / (1 + \frac{.05}{2})$ or \$973.2360, and the date 1 down-state price is $\$1,000 / (1 + \frac{.045}{2})$ or \$977.9951. Finally, the date 0 price is computed using the given date 0 one-year rate of 5.15%: $\$1,000 / (1 + \frac{.0515}{2})^2$ or 950.423.

The probabilities of moving up or down the tree may be used to compute the average or expected values. As of date 0, the expected value of the one-year zero's price on date 1 is

$$\frac{1}{2} \$973.24 + \frac{1}{2} \$978.00 = \$975.62 \quad (7.1)$$

Discounting this expected value to date 0 at the date 0, six-month rate gives an *expected discounted value*¹ of


$$\frac{\frac{1}{2} \$973.24 + \frac{1}{2} \$978.00}{(1 + \frac{.05}{2})} = \$951.82 \quad (7.2)$$

¹Over one period, discounting the expected value and taking the expectation of discounted values are the same. But, as shown in Chapter 13, over many periods the two are different and, with the approach taken by the short rate models in Part Three, taking the expectation of discounted values is correct—hence the choice of the term “expected discounted value.”

Note that the one-year zero's expected discounted value of \$951.82 does not equal its given market price of \$950.42. These two numbers need not be equal because investors do not price securities by expected discounted value. Over the next six months the one-year zero is a risky security, worth \$973.24 half of the time and \$978 the other half of the time for an average or expected value of \$975.62. If investors do not like this price uncertainty they would prefer a security worth \$975.62 on date 1 with certainty. More specifically, a security worth \$975.62 with certainty after six months would, by the arguments of Part One, sell for $\$975.62 / (1 + \frac{0.5}{2})$ or \$951.82 as of date 0. By contrast, investors penalize the risky one-year zero coupon bond with an average price of \$975.62 after six months by pricing it at \$950.42. The next chapter elaborates further on investor *risk aversion* and how large an impact it might be expected to have on bond prices.

ARBITRAGE PRICING OF DERIVATIVES

The text now turns to the pricing of a derivative security. What is the price of a call option, maturing in six months, to purchase \$1,000 face value of a then six-month zero at \$975? Begin with the price tree for this call option:



If on date 1 the six-month rate is 5.50% and a six-month zero sells for \$973.23, the right to buy that zero at \$975 is worthless. On the other hand, if the six-month rate turns out to be 4.50% and the price of a six-month zero is \$978, then the right to buy the zero at \$975 is worth \$978 – \$975 or \$3. This description of the option's terminal payoffs emphasizes the derivative nature of the option: its value depends on the value of an underlying security.

As shown in Chapter 1, a security is priced by arbitrage by finding and pricing its replicating portfolio. When, as in that context, cash flows do not depend on the levels of rates, the construction of the replicating portfolio is relatively simple. The derivative context is more difficult because cash flows do depend on the levels of rates, and the replicating portfolio must replicate the derivative security for any possible interest rate scenario.

To price the option by arbitrage, construct a portfolio on date 0 of underlying securities, namely six-month and one-year zero coupon bonds, that will be worth \$0 in the up state on date 1 and \$3 in the down state. To solve this problem, let $F^{.5}$ and F^1 be the face values of six-month and one-year zeros in the replicating portfolio, respectively. Then, these values

must satisfy the following two equations:

$$F^{\cdot 5} + .97324F^1 = \$0 \quad (7.3)$$

$$F^{\cdot 5} + .97800F^1 = \$3 \quad (7.4)$$

Equation (7.3) may be interpreted as follows. In the up state, the value of the replicating portfolio's now maturing six-month zero is its face value. The value of the once one-year zeros, now six-month zeros, is .97324 per dollar face value. Hence, the left-hand side of equation (7.3) denotes the value of the replicating portfolio in the up state. This value must equal \$0, the value of the option in the up state. Similarly, equation (7.4) requires that the value of the replicating portfolio in the down state equal the value of the option in the down state.

Solving equations (7.3) and (7.4), $F^{\cdot 5} = -\$613.3866$ and $F^1 = \$630.2521$. In words, on date 0 the option can be replicated by buying about \$630.25 face value of one-year zeros and simultaneously shorting about \$613.39 face amount of six-month zeros. Since this is the case, the law of one price requires that the price of the option equal the price of the replicating portfolio. But this portfolio's price is known and is equal to

$$\begin{aligned} .97561F^{\cdot 5} + .95042F^1 &= -.97561 \times \$613.3866 + .95042 \times \$630.2521 \\ &= \$.58 \end{aligned} \quad (7.5)$$

Therefore, the price of the option must be \$.58.

Recall that pricing based on the law of one price is enforced by arbitrage. If the price of the option were less than \$.58, arbitrageurs could buy the option, short the replicating portfolio, keep the difference, and have no future liabilities. Similarly, if the price of the option were greater than \$.58, arbitrageurs could short the option, buy the replicating portfolio, keep the difference, and, once again, have no future liabilities. Thus, ruling out profits from riskless arbitrage implies an option price of \$.58.

It is important to emphasize that the option cannot be priced by expected discounted value. Under that method, the option price would appear to be

$$\frac{.5 \times \$0 + .5 \times \$3}{1 + \frac{.05}{2}} = \$1.46 \quad (7.6)$$

The true option price is less than this value because investors dislike the risk of the call option and, as a result, will not pay as much as its expected discounted value. Put another way, the risk penalty implicit in the call option price is inherited from the risk penalty of the one-year zero, that is, from the property that the price of the one-year zero is less than its expected

discounted value. Once again, the magnitude of this effect is discussed in the next chapter.

This section illustrates arbitrage pricing with a call option, but it should be clear that arbitrage can be used to price any security with cash flows that depend on the six-month rate. Consider, for example, a security that, in six months, requires a payment of \$200 in the up state but generates a payment of \$1,000 in the down state. Proceeding as in the option example, find the portfolio of six-month and one-year zeros that replicates these two terminal payoffs, price this replicating portfolio as of date 0, and conclude that the price of the hypothetical security equals the price of the replicating portfolio.

A remarkable feature of arbitrage pricing is that the probabilities of up and down moves never enter into the calculation of the arbitrage price. See equations (7.3) to (7.5). The explanation for this somewhat surprising observation follows from the principles of arbitrage. Arbitrage pricing requires that the value of the replicating portfolio matches the value of the option in both the up and the down states. Therefore, the composition of the replicating portfolio is the same whether the probability of the up state is 20%, 50%, or 80%. But if the composition of the portfolio does not depend directly on the probabilities, and if the prices of the securities in the portfolio are given, then the price of the replicating portfolio and hence the price of the option cannot depend directly on the probabilities either.

Despite the fact that the option price does not depend directly on the probabilities, these probabilities must have some impact on the option price. After all, as it becomes more and more likely that rates will rise to 5.50% and that bond prices will be low, the value of options to purchase bonds must fall. The resolution of this apparent paradox is that the option price depends indirectly on the probabilities through the price of the one-year zero. Were the probability of an up move to increase suddenly, the current value of a one-year zero would decline. And since the replicating portfolio is long one-year zeros, the value of the option would decline as well. In summary, a derivative like an option depends on the probabilities only through current bond prices. Given bond prices, however, probabilities are not needed to derive arbitrage-free prices.

RISK-NEUTRAL PRICING

Risk-neutral pricing is a technique that modifies an assumed interest rate process, like the one assumed at the start of this chapter, so that any contingent claim can be priced without having to construct and price its replicating portfolio. Since the original interest rate process has to be modified only once, and since this modification requires no more effort than pricing a single contingent claim by arbitrage, risk-neutral pricing is an

extremely efficient way to price many contingent claims under the same assumed rate process.

In the example of this chapter, the price of a one-year zero does not equal its expected discounted value. The price of the one-year zero is \$950.42, computed from the given one-year spot rate of 5.15%. At the same time, the expected discounted value of the one-year zero is \$951.82, as derived in equation (7.2) and reproduced here:

$$\frac{\frac{1}{2}\$973.24 + \frac{1}{2}\$978.00}{1 + \frac{.05}{2}} = \$951.82 \quad (7.7)$$

The probabilities of $\frac{1}{2}$ for the up and down states are the assumed true or real-world probabilities. But there are other probabilities, called *risk-neutral* probabilities, that do cause the expected discounted value to equal the market price. To find these probabilities, let the risk-neutral probabilities in the up and down states be p and $(1 - p)$, respectively. Then, solve the following equation:

$$\frac{\$973.24p + \$978.00(1 - p)}{1 + \frac{.05}{2}} = \$950.42 \quad (7.8)$$

The solution is $p = .8024$. In words, under the risk-neutral probabilities of .8024 and .1976 the expected discounted value equals the market price.

In later chapters the difference between true and risk-neutral probabilities is described in terms of the *drift* in interest rates. Under the true probabilities there is a 50% chance that the six-month rate rises from 5% to 5.50% and a 50% chance that it falls from 5% to 4.50%. Hence the expected change in the six-month rate, or the drift of the six-month rate, is zero. Under the risk-neutral probabilities there is an 80.24% chance of a 50-basis point increase in the six-month rate and a 19.76% chance of a 50-basis point decline for an expected change of 30.24 basis points. Hence the drift of the six-month rate under these probabilities is 30.24 basis points.

As pointed out in the previous section, the expected discounted value of the option payoff is \$1.46, while the arbitrage price is \$.58. But what if expected discounted value is computed using the risk-neutral probabilities? The resulting option value would be

$$\frac{.8024 \times \$0 + .1976 \times \$3}{1 + \frac{.05}{2}} = \$.58 \quad (7.9)$$

The fact that the arbitrage price of the option equals its expected discounted value under the risk-neutral probabilities is not a coincidence. In

general, to value contingent claims by risk-neutral pricing, proceed as follows. First, find the risk-neutral probabilities that equate the price of the underlying securities with their expected discounted values. (In the simple example of this chapter the only risky, underlying security is the one-year zero.) Second, price the contingent claim by expected discounted value under these risk-neutral probabilities. The remainder of this section will describe intuitively why risk-neutral pricing works. Since the argument is a bit complex, it is broken up into four steps.

Step 1: Given trees for the underlying securities, the price of a security that is priced by arbitrage does not depend on investors' risk preferences. This assertion can be supported as follows.

A security is priced by arbitrage if one can construct a portfolio that replicates its cash flows. Under the assumed process for interest rates in this chapter, for example, the sample bond option is priced by arbitrage. By contrast, it is unlikely that a specific common stock can be priced by arbitrage because no portfolio of underlying securities can mimic the idiosyncratic fluctuations in a single common stock's market value.

If a security is priced by arbitrage and everyone agrees on the price evolution of the underlying securities, then everyone will agree on the replicating portfolio. In the option example, both an extremely risk-averse, retired investor and a professional gambler would agree that a portfolio of \$630.25 face of one-year zeros and $-\$613.39$ face of six-month zeros replicates the option. And since they agree on the composition of the replicating portfolio and on the prices of the underlying securities, they must also agree on the price of the derivative.


Step 2: Imagine an economy identical to the true economy with respect to current bond prices and the possible value of the six-month rate over time but different in that the investors in the imaginary economy are risk neutral. Unlike investors in the true economy, investors in the imaginary economy do not penalize securities for risk and, therefore, price securities by expected discounted value. It follows that, under the probabilities in the imaginary economy, the expected discounted value of the one-year zero equals its market price. But these probabilities satisfy equation (7.8), namely the risk-neutral probabilities of .8024 and .1976.

Step 3: The price of the option in the imaginary economy, like any other security in that economy, is computed by expected discounted value. Since the probability of the up state in that economy is .8024, the price of the option in that economy is given by equation (7.9) and is, therefore, \$.58.

Step 4: Step 1 implies that given the prices of the six-month and one-year zeros, as well as possible values of the six-month rate, the price of an option does not depend on investor risk preferences. It follows that since the real and imaginary economies have the same bond prices and the same possible values for the six-month rate, the option price must be the same in both economies. In particular, the option price in the real economy must equal \$.58, the option price in the imaginary economy. More generally, the price of a derivative in the real economy may be computed by expected discounted value under the risk-neutral probabilities.

ARBITRAGE PRICING IN A MULTI-PERIOD SETTING

Maintaining the binomial assumption, the tree of the previous section might be extended for another six months as follows:




When, as in this tree, an up move followed by a down move does not give the same rate as a down move followed by an up move, the tree is said to be *nonrecombining*. From an economic perspective, there is nothing wrong with this kind of tree. To justify this particular tree, for example, one might argue that when short rates are 5% or higher they tend to change in increments of 50 basis points. But when rates fall below 5%, the size of the change starts to decrease. In particular, at a rate of 4.50%, the short rate may change by only 45 basis points. A volatility process that depends on the level of rates exhibits *state-dependent* volatility.

Despite the economic reasonableness of nonrecombining trees, practitioners tend to avoid them because such trees are difficult or even impossible to implement. After six months there are two possible states, after one year there are four, and after N semiannual periods there are 2^N possibilities. So, for example, a tree with semiannual steps large enough to price 10-year securities will, in its rightmost column alone, have over 500,000 nodes, while a tree used to price 20-year securities will in its rightmost column have over 500 billion nodes. Furthermore, as discussed later in the chapter, it is often desirable to reduce substantially the time interval between dates. In short, even with modern computers, trees that grow this quickly are computationally unwieldy. This doesn't mean, by the way, that the effects that

give rise to nonrecombining trees, like state-dependent volatility, have to be abandoned. It simply means that these effects must be implemented in a more efficient way.


Trees in which the up-down and down-up states have the same value are called *recombining trees*. An example of this type of tree that builds on the two-date tree of the previous sections is



Note that there are two nodes after six months, three after one year, and so on. A tree with weekly rather than semiannual steps capable of pricing a 30-year security would have only $52 \times 30 + 1$ or 1,561 nodes in its rightmost column. Evidently, recombining trees are much more manageable than nonrecombining trees from a computational viewpoint.

As trees grow it becomes convenient to develop a notation with which to refer to particular nodes. One convention is as follows. The dates, represented by columns of the tree, are numbered from left to right starting with 0. The states, represented by rows of the tree, are numbered from bottom to top, also starting from 0. For example, in the preceding tree the six-month rate on date 2, state 0 is 4%. The six-month rate on state 1 of date 1 is 5.50%.

Continuing where the option example left off, having derived the risk-neutral tree for the pricing of a one-year zero, the goal is to extend the tree for the pricing of a 1.5-year zero assuming that the 1.5-year spot rate is 5.25%. Ignoring the probabilities for a moment, several nodes of the 1.5-year zero price tree can be written down immediately:



On date 3 the zero with an original term of 1.5 years matures and is worth its face value of \$1,000. On date 2 the value of the then six-month

zero equals its face value discounted for six months at the then-prevailing spot rates of 6%, 5%, and 4% in states 2, 1, and 0, respectively:

$$\frac{\$1,000}{1 + \frac{.06}{2}} = \$970.87 \quad (7.10)$$

$$\frac{\$1,000}{1 + \frac{.05}{2}} = \$975.61 \quad (7.11)$$


$$\frac{\$1,000}{1 + \frac{.04}{2}} = \$980.39 \quad (7.12)$$

Finally, on date 0 the 1.5-year zero equals its face value discounted at the given 1.5-year spot rate:

$$\frac{\$1,000}{\left(1 + \frac{.0525}{2}\right)^3} = \$925.21 \quad (7.13)$$

The prices of the zero on date 1 in states 1 and 0 are denoted $P_{1,1}$ and $P_{1,0}$ respectively. The then one-year zero prices are not known because, at this point in the development, possible values of the one-year rate in six months are not available.

The previous section showed that the risk-neutral probability of an up move on date 0 is .8024. Letting q be the risk-neutral probability of an up move on date 1,² the tree becomes



²For simplicity alone this example assumes that the probability of moving up from state 0 equals the probability of moving up from state 1. Choosing among the many possible interest rate processes is discussed in Chapters 9 through 11.

By definition, expected discounted value under risk-neutral probabilities must produce market prices. With respect to the 1.5-year zero price on date 0, this requires that


$$\frac{.8024P_{1,1} + .1976P_{1,0}}{1 + \frac{.05}{2}} = \$925.21 \quad (7.14)$$

With respect to the prices of a then one-year zero on date 1,

$$P_{1,1} = \frac{\$970.87q + \$975.61(1 - q)}{1 + \frac{.055}{2}} \quad (7.15)$$

$$P_{1,0} = \frac{\$975.61q + \$980.39(1 - q)}{1 + \frac{.045}{2}} \quad (7.16)$$


While equations (7.14) through (7.16) may appear complicated, substituting (7.15) and (7.16) into (7.14) results in a linear equation in the one unknown, q . Solving this resulting equation reveals that $q = .6489$. Therefore, the risk-neutral interest rate process may be summarized by the following tree:



Furthermore, any derivative security that depends on the six-month rate in six months and in one year may be priced by computing its discounted expected value along this tree. An example appears in the next section.

The difference between the true and risk-neutral probabilities may once again be described in terms of drift. From dates 1 to 2, the drift under the true probabilities is zero. Under the risk-neutral probabilities the drift is computed from a 64.89% chance of a 50-basis point increase in the six-month rate and a 35.11% chance of a 50-basis point decline in the rate. These numbers give a drift or expected change of 14.89 basis points.

Substituting $q = .6489$ back into equations (7.15) and (7.16) completes the tree for the price of the 1.5-year zero:



It follows immediately from this tree that the one-year spot rate six months from now may be either 5.5736% or 4.5743% since

$$\$946.51 = \frac{\$1,000}{\left(1 + \frac{5.5736\%}{2}\right)^2} \quad (7.17)$$

$$\$955.78 = \frac{\$1,000}{\left(1 + \frac{4.5743\%}{2}\right)^2} \quad (7.18)$$

The fact that the possible values of the one-year spot rate can be extracted from the tree is at first surprising. The starting point of the example is the date 0 values of the .5-, 1-, and 1.5-year spot rates as well as an assumption about the evolution of the six-month rate over the next year. But since this information, in combination with arbitrage or risk-neutral arguments, is sufficient to determine the price tree of the 1.5-year zero, it is sufficient to determine the possible values of the one-year spot rate in six months. Considering this fact from another point of view, having specified initial spot rates and the evolution of the six-month rate, a modeler may not make any further assumptions about the behavior of the one-year rate.

The six-month rate process completely determines the one-year rate process because the model presented here has only one factor. Writing down a tree for the evolution of the six-month rate alone implicitly assumes that prices of all fixed income securities can be determined by the evolution of that rate. Multi-factor models for which this is not the case will be introduced in Chapter 11.

Just as some replicating portfolio can reproduce the cash flows of a security from date 0 to date 1, some other replicating portfolios can reproduce the cash flows of a security from date 1 to date 2. The composition of these replicating portfolios depends on the date and state. More specifically, the replicating portfolios held on date 0, on state 0 of date 1, and on state 1 of date 1 are usually different. From the trading perspective, the replicating portfolio must be adjusted as time passes and as interest rates change. This process is known as *dynamic replication*, in contrast to the *static replication strategies* of Part One. As an example of static replication, the portfolio of zero coupon bonds that replicates a coupon bond does not change over time nor with the level of rates.

Having built a tree out to date 2 it should be clear how to extend the tree to any number of dates. Assumptions about the future possible values of the short-term rate have to be extrapolated further into the future and risk-neutral probabilities have to be calculated to recover a given set of bond prices.

EXAMPLE: PRICING A CONSTANT-MATURITY TREASURY SWAP

Equipped with the last tree of interest rates in the previous section, this section prices a particular derivative security, namely \$1,000,000 face value of a stylized *constant-maturity Treasury (CMT) swap* struck at 5%. This swap pays

$$\$1,000,000 \frac{y_{CMT} - 5\%}{2} \quad (7.19)$$

every six months until it matures, where y_{CMT} is a semiannually compounded yield, of a predetermined maturity, on the payment date. The text prices a one-year CMT swap on the six-month yield. In practice, CMT swaps trade most commonly on the yields of the most liquid maturities, i.e., on 2-, 5- and 10-year yields.

Since six-month semiannually compounded yields equal six-month spot rates, rates from the tree of the previous section can be substituted into (7.19) to calculate the payoffs of the CMT swap. On date 1, the state 1 and state 0 payoffs are, respectively,

$$\$1,000,000 \frac{5.50\% - 5\%}{2} = \$2,500 \quad (7.20)$$

$$\$1,000,000 \frac{4.50\% - 5\%}{2} = -\$2,500 \quad (7.21)$$

Similarly on date 2, the state 2, 1, and 0 payoffs are, respectively,

$$\$1,000,000 \frac{6\% - 5\%}{2} = \$5,000 \quad (7.22)$$

$$\$1,000,000 \frac{5\% - 5\%}{2} = \$0 \quad (7.23)$$

$$\$1,000,000 \frac{4\% - 5\%}{2} = -\$5,000 \quad (7.24)$$

The possible values of the CMT swap at maturity, on date 2, are given by equations (7.22) through (7.24). The possible values on date 1 are given by the expected discounted value of the date 2 payoffs under the risk-neutral probabilities plus the date 1 payoffs given by (7.20) and (7.21). The resulting date 1 values in states 1 and 0, respectively, are


$$\frac{.6489 \times \$5,000 + .3511 \times \$0}{1 + \frac{.055}{2}} + \$2,500 = \$5,657.66 \quad (7.25)$$

$$\frac{.6489 \times 0 + .3511 \times (-\$5,000)}{1 + \frac{.045}{2}} - \$2,500 = -\$4,216.87 \quad (7.26)$$

Finally, the value of the swap on date 0 is the expected discounted value of the date-1 payoffs, given by (7.25) and (7.26), under the risk-neutral probabilities:

$$\frac{.8024 \times \$5,657.66 + .1976 \times (-\$4,216.87)}{1 + \frac{.05}{2}} = \$3,616.05 \quad (7.27)$$

The following tree summarizes the value of the stylized CMT swap over dates and states:



A value of \$3,616.05 for the CMT swap might seem surprising at first. After all, the cash flows of the CMT swap are zero at a rate of 5%, and 5% is, under the real probabilities, the average rate on each date. The explanation, of course, is that the risk-neutral probabilities, not the real probabilities, determine the arbitrage price of the swap. The expected discounted value of

the swap under the real probabilities can be computed by following the steps leading to (7.25) through (7.27) but using .5 for all up and down moves. The result of these calculations does give a value close to zero, namely $-\$5.80$.

The expected cash flow of the CMT swap on both dates 1 and 2, under the real probabilities, is zero. It follows immediately that the discounted value of these expected cash flows is zero. At the same time, the expected discounted value of the CMT swap is $-\$5.80$. Why are these values different? The answer to this question is deferred to Chapter 13.

OPTION-ADJUSTED SPREAD

Option-adjusted spread (OAS) is a widely-used measure of the relative value of a security, that is, of its market price relative to its model value. OAS is defined as the spread such that the market price of a security equals its model price when discounted values are computed at risk-neutral rates plus that spread. To illustrate, say that the market price of the CMT swap in the previous section is $\$3,613.25$, $\$2.80$ less than the model value. In that case, the OAS of the CMT swap turns out to be 10 basis points. To see this, add 10 basis points to the discounting rates of 5.5% and 4.5% in equations (7.25) and (7.26), respectively, to get new swap values of

$$\frac{.6489 \times \$5,000 + .3511 \times \$0}{1 + \frac{.056}{2}} + \$2,500 = \$5,656.13 \quad (7.28)$$

$$\frac{.6489 \times 0 + .3511 \times (-\$5,000)}{1 + \frac{.046}{2}} - \$2,500 = -\$4,216.03 \quad (7.29)$$

Note that, when calculating value with an OAS spread, rates are only shifted for the purpose of discounting. Rates are not shifted for the purposes of computing cash flows. In the CMT swap example, cash flows are still computed using equations (7.20) through (7.24).

Completing the valuation with an OAS of 10 basis points, use the results of (7.28) and (7.29) and a discount rate of 5% plus the OAS spread of 10 basis points, or 5.10%, to obtain an initial CMT swap value of

$$\frac{.8024 \times \$5,656.13 + .1976 \times (-\$4,216.03)}{1 + \frac{.051}{2}} = \$3,613.25 \quad (7.30)$$

Hence, as claimed, discounting at the risk-neutral rates plus an OAS of 10 basis points produces a model price equal to the given market price of $\$3,613.25$.

If a security's OAS is positive, its market price is less than its model price, so the security trades cheap. If the OAS is negative, the security trades rich.

Another perspective on the relative value implications of an OAS spread is the fact that the expected return of a security with an OAS, under the risk-neutral process, is the short-term rate plus the OAS per period. Very simply, discounting a security's expected value by a particular rate per period is equivalent to that security's earning that rate per period. In the example of the CMT swap, the expected return of the fairly-priced swap under the risk-neutral process over the six months from date 0 to date 1 is

$$\frac{.8024 \times \$5,657.66 - .1976 \times \$4,216.87 - \$3,616.05}{\$3,616.05} = 2.5\% \quad (7.31)$$

which is six month's worth of the initial rate of 5%. On the other hand, the expected return of the cheap swap, with an OAS of 10 basis points, is

$$\frac{.8024 \times \$5,656.13 - .1976 \times \$4,216.03 - \$3,613.25}{\$3,613.25} = 2.55\% \quad (7.32)$$

which is six month's worth of the initial rate of 5% plus the OAS of 10 basis points, or half of 5.10%.

PROFIT AND LOSS ATTRIBUTION WITH AN OAS

Chapter 3 introduced profit and loss (P&L) attribution. This section gives a mathematical description of attribution in the context of term structure models and of securities that trade with an OAS.

By the definition of a one-factor model, and by the definition of OAS, the market price of a security at time t and a factor value of x can be written as $P_t(x, \text{OAS})$. Using a first-order Taylor approximation, the change in the price of the security is

$$dP = \frac{\partial P}{\partial x} dx + \frac{\partial P}{\partial t} dt + \frac{\partial P}{\partial \text{OAS}} d\text{OAS} \quad (7.33)$$

Dividing by the price and taking expectations,

$$E \left[\frac{dP}{P} \right] = \frac{1}{P} \frac{\partial P}{\partial x} E[dx] + \frac{1}{P} \frac{\partial P}{\partial t} dt \quad (7.34)$$

Since the OAS calculation assumes that OAS is constant over the life of the security, moving from (7.33) to (7.34) assumes that the expected change in the OAS is zero.

As mentioned in the previous section, if expectations are taken with respect to the risk-neutral process,³ then, for any security priced according to the model,

$$E \left[\frac{dP}{P} \right] = r dt \quad (7.35)$$

But equation (7.35) does not apply to securities that are not priced according to the model, that is, to securities with an OAS not equal to zero. For these securities, by definition, the cash flows are discounted not at the short-term rate but at the short-term rate plus the OAS. Equivalently, as argued in the previous section, the expected return under the risk-neutral probabilities is not the short-term rate but the short-term rate plus the OAS. Hence, the more general form of (7.35), is

$$E \left[\frac{dP}{P} \right] = (r + \text{OAS}) dt \quad (7.36)$$

Combining these pieces, substitute (7.34) and (7.36) into (7.33) and rearrange terms to break down the return of a security into its component parts:

$$\frac{dP}{P} = (r + \text{OAS}) dt + \frac{1}{P} \frac{\partial P}{\partial x} (dx - E[dx]) + \frac{1}{P} \frac{\partial P}{\partial \text{OAS}} d\text{OAS} \quad (7.37)$$

Finally, multiplying through by P ,

$$dP = (r + \text{OAS}) P dt + \frac{\partial P}{\partial x} (dx - E[dx]) + \frac{\partial P}{\partial \text{OAS}} d\text{OAS} \quad (7.38)$$

In words, the return of a security or its P&L may be divided into a component due to the passage of time, a component due to changes in the factor, and a component due to the change in the OAS. In the language of Chapter 3, the terms on the right-hand side of (7.38) represent, in order, carry-roll-down,⁴ gains or losses from rate changes, and gains or losses from spread change. For models with predictive power, the OAS converges or

³Taking expected values with respect to the true probabilities would add a risk premium term to the right-hand side of this equation. See Chapter 8.

⁴For expositional simplicity no explicit coupon or other direct cash flows have been included in this discussion.

tends to zero, or, equivalently, the security price converges or tends toward its fair value according to the model.

The decompositions (7.37) and (7.38) highlight the usefulness of OAS as a measure of the value of a security with respect to a particular model. According to the model, a long position in a cheap security earns superior returns in two ways. First, it earns the OAS over time intervals in which the security does not converge to its fair value. Second, it earns its sensitivity to OAS times the extent of any convergence.

The decomposition equations also provide a framework for thinking about relative value trading. When a cheap or rich security is identified, a relative value trader buys or sells the security and hedges out all interest rate or factor risk. In terms of the decompositions, $\frac{\partial P}{\partial x} = 0$. In that case, the expected return or P&L depends only on the short-term rate, the OAS, and any convergence. Furthermore, if the trader finances the trade at the short-term rate, i.e., borrows P at a rate r to purchase the security, the expected return is simply equal to the OAS plus any convergence return.

REDUCING THE TIME STEP

To this point this chapter has assumed that the time elapsed between dates of the tree is six months. The methodology outlined previously, however, can be easily adapted to any time step of Δt years. For monthly time steps, for example, $\Delta t = \frac{1}{12}$ or .0833, and one-month rather than six-month interest rates appear on the tree. Furthermore, discounting must be done over the appropriate time interval. If the rate of term Δt is r , then discounting means dividing by $1 + r\Delta t$. In the case of monthly time steps, discounting with a one-month rate of 5% means dividing by $1 + \frac{.05}{12}$.

In practice there are two reasons to choose time steps smaller than six months. First, a security or portfolio of securities rarely makes all of its payments in even six-month intervals from the starting date. Reducing the time step to a month, a week, or even a day can ensure that all cash flows are sufficiently close in time to some date in the tree. Second, assuming that the six-month rate can take on only two values in six months, three values in one year, and so on, produces a tree that is too coarse for many practical pricing problems. Reducing the step size can fill the tree with enough rates to price contingent claims with sufficient accuracy. Figure 7.1 illustrates this point by showing a relatively realistic-looking probability distribution of the six-month rate in six months from a tree with daily time steps, a drift of zero, and a horizon standard deviation of 65 basis points.

While smaller time steps generate more realistic interest rate distributions, it is not the case that smaller time steps are always desirable. First, the greater the number of computations in pricing a security, the more




FIGURE 7.1 Sample Probability Distribution of the Six-Month Rate in Six Months with Daily Time Steps

attention must be paid to numerical issues like round-off error. Second, since decreasing the time step increases computation time, practitioners requiring quick results cannot make the time step too small. Customers calling market makers in options on swaps, or *swaptions*, for example, expect price quotations within minutes if not sooner. Hence, the time step in a model used to price swaptions must be consistent with the market maker's required response time.

The best choice of step size ultimately depends on the problem at hand. When pricing a 30-year callable bond, for example, a model with monthly time steps may provide a realistic enough interest rate distribution to generate reliable prices. The same monthly steps, however, will certainly be inadequate to price a one-month bond option: that tree would imply only two possible rates on the option expiration date.

While the trees in this chapter assume that the step size is the same throughout the tree, this need not be the case. Sophisticated implementations of trees allow step size to vary across dates in order to achieve a balance between realism and computational concerns.

FIXED INCOME VERSUS EQUITY DERIVATIVES

While the ideas behind pricing fixed income and equity derivatives are similar in many ways, there are important differences as well. In particular, it is

worth describing why models created for the stock market cannot be adopted without modification for use in fixed income markets.

The famous Black-Scholes-Merton pricing analysis of stock options can be summarized as follows. Under the assumption that the stock price evolves according to a particular random process and that the short-term interest rate is constant, it is possible to form a portfolio of stocks and short-term bonds that replicates the payoffs of an option. Therefore, by arbitrage arguments, the price of the option must equal the known price of the replicating portfolio.

Say that an investor wants to price an option on a five-year bond by a direct application of this logic. The investor would have to begin by making an assumption about how the price of the five-year bond evolves over time. But this is considerably more complicated than making assumptions about how the price of a stock evolves over time. First, the price of a bond must converge to its face value at maturity while the random process describing the stock price need not be constrained in any similar way. Second, because of the maturity constraint, the volatility of a bond's price must eventually get smaller as the bond approaches maturity. The simpler assumption that the volatility of a stock is constant is not so appropriate for bonds. Third, since stock volatility is very large relative to short-term rate volatility, it may be relatively harmless to assume that the short-term rate is constant. By contrast, it can be difficult to defend the assumption that a bond price follows some random process while the short-term interest rate is constant.⁵

These objections led researchers to make assumptions about the random evolution of the interest rate rather than of the bond price. In that way bond prices would naturally approach par, price volatilities would naturally approach zero, and the interest rate would not be assumed to be constant. But this approach raises another set of questions. Which interest rate is assumed to evolve in a particular way? Making assumptions about the 5-year rate over time is not particularly helpful for two reasons. First, 5-year coupon bond prices depend on shorter-term rates as well. Second, pricing an option on a 5-year bond requires assumptions about the bond's future possible prices. But knowing the 5-year rate over time is insufficient because, in a very short time, the option's underlying security will no longer be a 5-year bond. Therefore, one must often make assumptions about the evolution of the entire term structure of interest rates to price bond options and other derivatives. In the one-factor case described in this chapter it has been shown that modeling the evolution of the short-term rate is sufficient,

⁵Because these three objections are less important in the case of short-term options on long-term bonds, practitioners do use stock-like models in this fixed income context. Also, it is often sufficient to assume, somewhat more satisfactorily, that the relevant discount factor is uncorrelated with the price of the underlying bond. See Chapter 18.

combined with arbitrage arguments, to build a model of the entire term structure. In short, despite the enormous importance of the Black-Scholes-Merton analysis, the fixed income context does demand special attention.

Having reached the conclusion at the end of the previous paragraph, there are some contexts in which practitioners invoke assumptions so that the Black-Scholes-Merton models can be applied in place of more difficult-to-implement term structure models. These situations are discussed at length in Chapter 18.

The Evolution of Short Rates and the Shape of the Term Structure

This chapter presents a framework for understanding the shape of the term structure. In particular, it is shown how spot or forward rates are determined by expectations of future short-term rates, the volatility of short-term rates, and an interest rate risk premium. To conclude the chapter, this framework is applied to swap curves in the United States and Japan.

INTRODUCTION

From assumptions about the interest rate process for the short-term rate and from an initial term structure implied by market prices, Chapter 7 showed how to derive a risk-neutral process that can be used to price all fixed income securities by arbitrage. Models that follow this approach, i.e., models that take the initial term structure as given, are called *arbitrage-free* models. A different approach, however, is to start with assumptions about the interest rate process and about the risk premium demanded by the market for bearing interest rate risk and then derive the risk-neutral process. Models of this sort do not necessarily match the initial term structure and are called *equilibrium* models.¹ The strengths and weaknesses of each approach are discussed in subsequent chapters of Part Three.

This chapter describes how assumptions about the interest rate process and about the risk premium determine the level and shape of the term structure. For equilibrium models, an understanding of the relationships between the model assumptions and the shape of the term structure is important in order to make reasonable assumptions in the

¹This nomenclature is somewhat misleading. Equilibrium models, in the context of their assumptions, which do not include market prices for the initial term structure, are also arbitrage-free.

first place. For arbitrage-free models, an understanding of these relationships reveals the assumptions implied by the market through the observed term structure.

Many economists might find this chapter remarkably narrow. An economist asked about the shape of the term structure would undoubtedly make reference to such macroeconomic factors as the marginal productivity of capital, the propensity to save, and expected inflation. The more modest goal of this chapter is to connect the dynamics of the short-term rate of interest and the risk premium with the shape of the term structure. While this goal does fall short of answers that an economist might provide, it is more ambitious than the derivation of arbitrage restrictions on bond and derivative prices given underlying bond prices.

EXPECTATIONS

The word *expectations* implies uncertainty. Investors might expect the one-year rate to be 10%, but know there is a good chance it will turn out to be 8% or 12%. For the purposes of this section alone the text assumes away uncertainty so that the statement that investors expect or forecast a rate of 10% means that investors assume that the rate will be 10%. The sections to follow reintroduce uncertainty.

To highlight the role of interest rate forecasts in determining the shape of the term structure, consider the following simple example. The one-year interest rate is currently 10% and all investors forecast that the one-year interest rate next year and the year after will also be 10%. In that case, investors will discount cash flows using forward rates of 10%. In particular, the price of one-, two- and three-year zero coupon bonds per dollar face value (using annual compounding) will be

$$P^1 = \frac{1}{1.10} \quad (8.1)$$

$$P^2 = \frac{1}{(1.10)(1.10)} = \frac{1}{1.10^2} \quad (8.2)$$

$$P^3 = \frac{1}{(1.10)(1.10)(1.10)} = \frac{1}{1.10^3} \quad (8.3)$$

From inspection of equations (8.1) through (8.3), the term structure of spot rates in this example is flat at 10%. Very simply, investors are willing to lock in 10% for two or three years because they assume that the one-year rate will always be 10%.

Now assume that the one-year rate is still 10%, but that all investors forecast the one-year rate next year to be 12% and the one-year rate in

two years to be 14%. In that case, the one-year spot rate is still 10%. The two-year spot rate, $\hat{r}(2)$, is such that

$$P^2 = \frac{1}{(1.10)(1.12)} = \frac{1}{(1 + \hat{r}(2))^2} \quad (8.4)$$

Solving, $\hat{r}(2) = 10.995\%$. Similarly, the three-year spot rate, $\hat{r}(3)$, is such that

$$P^3 = \frac{1}{(1.10)(1.12)(1.14)} = \frac{1}{(1 + \hat{r}(3))^3} \quad (8.5)$$

Solving, $\hat{r}(3) = 11.998\%$. Hence, the evolution of the one-year rate from 10% to 12% to 14% generates an upward-sloping term structure of spot rates: 10%, 10.995%, and 11.988%. In this case investors require rates above 10% when locking up their money for two or three years because they assume one-year rates will be higher than 10%. No investor, for example, would buy a two-year zero at a yield of 10% when it is possible to buy a one-year zero at 10% and, when it matures, buy another one-year zero at 12%.

Finally, assume that the one-year rate is 10%, but that investors forecast that it will fall to 8% in one year and to 6% in two years. In that case, it is easy to show that the term structure of spot rates will be downward-sloping. In particular, $\hat{r}(1) = 10\%$, $\hat{r}(2) = 8.995\%$, and $\hat{r}(3) = 7.988\%$.


These simple examples reveal that expectations can cause the term structure to take on any of a myriad of shapes. Over short horizons, the financial community can have very specific views about future short-term rates. Over longer horizons, however, expectations cannot be so granular. It would be difficult, for example, to defend the position that the expectation for the one-year rate 29 years from now is substantially different from the expectation of the one-year rate 30 years from now. On the other hand, an argument can be made that the long-run expectation of the short-term rate is 5%: 3% due to the long-run real rate of interest and 2% due to long-run inflation. Hence, forecasts can be very useful in describing the shape and level of the term structure over short-term horizons and the level of rates at very long horizons. This conclusion has important implications for extracting expectations from observed interest rates (see the application at the end of this chapter) and for choosing among term structure models.

VOLATILITY AND CONVEXITY

This section drops the assumption that investors believe that their forecasts will be realized and assumes instead that investors understand the volatility around their expectations. To isolate the implications of volatility on the


shape of the term structure, this section assumes that investors are risk-neutral so that they price securities by expected discounted value. The next section drops this assumption.

Assume that the following tree gives the true process for the one-year rate:



Note that the expected interest rate on date 1 is $.5 \times 8\% + .5 \times 12\%$ or 10% and that the expected rate on date 2 is $.25 \times 14\% + .5 \times 10\% + .25 \times 6\%$ or 10%. In the previous section, with no volatility around expectations, flat expectations of 10% imply a flat term structure of spot rates. That is not the case in the presence of volatility.

The price of a one-year zero is, by definition, $\frac{1}{1.10}$ or .909091, implying a one-year spot rate of 10%. Under the assumption of risk-neutrality, the price of a two-year zero may be calculated by discounting the terminal cash flow using the preceding interest rate tree:



Hence, the two-year spot rate is such that $.82672 = (1 + \hat{r}(2))^{-2}$, implying that $\hat{r}(2) = 9.982\%$.

Even though the one-year rate is 10% and the expected one-year rate in one year is 10%, the two-year spot rate is 9.982%. The 1.8-basis point difference between the spot rate that would obtain in the absence of uncertainty, 10%, and the spot rate in the presence of volatility, 9.982%, is the effect of convexity on that spot rate. This convexity effect arises from the mathematical fact, a special case of *Jensen's Inequality*, that

$$E \left[\frac{1}{1+r} \right] > \frac{1}{E[1+r]} = \frac{1}{1+E[r]} \quad (8.6)$$




FIGURE 8.1 An Illustration of Convexity

Figure 8.1 graphically illustrates this equation. There are two possible values of r and, consequently, of the function $\frac{1}{1+r}$ in the figure,² shown as points A and D. The height or vertical-axis coordinate of point B is the average of these two function values. Under the assumption that the two possible values of r occur with equal probability, this average can be thought of as $E\left[\frac{1}{1+r}\right]$ in (8.6). And under the same assumption, the horizontal-axis coordinates of the points B and C can be thought of as $E[r]$ so that the height of point C can be thought of as $\frac{1}{1+E[r]}$. Clearly, the height of B is greater than that of C, or $E\left[\frac{1}{1+r}\right] > \frac{1}{1+E[r]}$. To summarize, equation (8.6) is true because the pricing function of a zero-coupon bond, $\frac{1}{1+r}$, is convex rather than concave.

Returning to the example of this section, equation (8.6) may be used to show why the one-year spot rate is less than 10%. The spot rate one year from now may be 12% or 8%. According to (8.6),

$$.5 \times \frac{1}{1.12} + .5 \times \frac{1}{1.08} > \frac{1}{.5 \times 1.12 + .5 \times 1.08} = \frac{1}{1.10} \quad (8.7)$$


Dividing both sides by 1.10,

$$\frac{1}{1.10} \left[.5 \times \frac{1}{1.12} + .5 \times \frac{1}{1.08} \right] > \frac{1}{1.10^2} \quad (8.8)$$

²The curve shown is actually a power of $\frac{1}{1+r}$; i.e., the price of a longer-term zero-coupon bond, so that the curvature is more visible.

The left-hand side of (8.8) is the price of the two-year zero-coupon bond today. In words then, equation (8.8) says that the price of the two-year zero is greater than the result of discounting the terminal cash flow by 10% over the first period and by the expected rate of 10% over the second period. It follows immediately that the yield of the two-year zero, or the two-year spot rate, is less than 10%.

The tree presented at the start of this section may also be used to price a three-year zero. The resulting price tree is



The three-year spot rate, such that $.752309 = (1 + \hat{r}(3))^{-3}$, is 9.952%. Therefore, the value of convexity in this spot rate is $10\% - 9.952\%$ or 4.8 basis points, whereas the value of convexity in the two-year spot rate was only 1.8 basis points.

It is generally true that, all else equal, the value of convexity increases with maturity. This will become evident shortly. For now, suffice it to say that the convexity of the pricing function of a zero maturing in N years, $(1 + r)^{-N}$, increases with N . In terms of Figure 8.1, the longer the maturity of the illustrated pricing function, the more convex the curve.

Chapter 4 showed that securities with greater convexity perform better when yields change a lot and claimed that they perform worse when yields do not change by much. The discussion in this section shows that convexity does, in fact, lower bond yields. The mathematical development in a later section ties these observations together by showing exactly how the advantages of convexity are offset by lower yields.

The previous section assumes no interest rate volatility and, consequently, yields are completely determined by forecasts. In this section, with the introduction of volatility, yield is reduced by the value of convexity. So it may be said that the value of convexity arises from volatility. Furthermore, the value of convexity increases with volatility. In the tree introduced at the start of the section, the standard deviation of rates is





FIGURE 8.2 Volatility and the Shape of the Term Structure in Three-Date Binomial Models

200 basis points a year.³ Now consider a tree with a standard deviation of 400 basis points:



The expected one-year rate in one year and in two years is still 10%. Spot rates and convexity values for this case may be derived along the same lines as before. Figure 8.2 graphs three term structures of spot rates: one with no volatility around the expectation of 10%; one with a volatility of 200 basis points a year (the tree of the first example); and one with a volatility of 400 basis points per year (the tree preceding this paragraph). Note that the value of convexity, measured by the distance between the rates assuming no volatility and the rates assuming volatility, increases with

³Chapter 9 describes the computation of the standard deviation of rates implied by an interest rate tree.

volatility. Figure 8.2 also illustrates that the value of convexity increases with maturity.

For very short terms and realistic levels of volatility, the value of convexity is quite small. But since simple examples must rely on short terms, convexity effects would hardly be discernible without raising volatility to unrealistic levels. Therefore, this section had to make use of unrealistically high volatility. The application at the end of this chapter uses realistic volatility levels to present typical convexity values.

RISK PREMIUM

To illustrate the effect of risk premium on the term structure, consider again the second interest rate tree presented in the preceding section, with a volatility of 400 basis points per year. Risk-neutral investors would price a two-year zero by the following calculation:

$$\begin{aligned} .827541 &= \frac{.5 \left[\frac{1}{1.14} + \frac{1}{1.06} \right]}{1.10} \\ &= \frac{.5 [.877193 + .943396]}{1.10} \end{aligned} \quad (8.9)$$

By discounting the expected future price by 10%, equation (8.9) implies that the expected return from owning the two-year zero over the next year is 10%. To verify this statement, calculate this expected return directly:

$$\begin{aligned} .5 \frac{.877193 - .827541}{.827541} + .5 \frac{.943396 - .827541}{.827541} &= .5 \times 6\% + .5 \times 14\% \\ &= 10\% \end{aligned} \quad (8.10)$$

Would investors really invest in this two-year zero offering an expected return of 10% over the next year? The return will, in fact, be either 6% or 14%. While these two returns do average to 10%, an investor could, instead, buy a one-year zero with a certain return of 10%. Presented with this choice, any risk-averse investor would prefer an investment with a certain return of 10% to an investment with a risky return that averages 10%. In other words, investors require compensation for bearing interest rate risk.⁴

Risk-averse investors demand a return higher than 10% for the two-year zero over the next year. This return can be effected by pricing the zero

⁴This is an oversimplification. See the discussion at the end of the section.

coupon bond one year from now at less than the prices of $\frac{1}{1.14}$ and $\frac{1}{1.06}$. Equivalently, future cash flows could be discounted at rates higher than the possible rates of 14% and 6%. The next section shows that adding, for example, 20 basis points to each of these rates is equivalent to assuming that investors demand an extra 20 basis points for each year of duration risk. Assuming this is indeed the fair market *risk premium*, the price of the two-year zero would be computed as follows:

$$.826035 = \frac{.5 \left[\frac{1}{1.142} + \frac{1}{1.062} \right]}{1.10} \quad (8.11)$$

The price in (8.11) is below the value obtained in (8.9) which assumes that investors are risk-neutral. Put another way, the increase in the discounting rates has increased the expected return of the two-year zero. In one year, if the interest rate is 14%, then the price of a one-year zero will be $\frac{1}{1.14}$ or .877193. If the rate is 6%, then the price will be $\frac{1}{1.06}$ or .943396. Therefore, the expected return of the two-year zero priced at .826035 is

$$\frac{.5 [.877193 + .943396] - .826035}{.826035} = 10.20\% \quad (8.12)$$



Hence, recalling that the one-year zero has a certain return of 10%, the risk-averse investors in this example demand 20 basis points in expected return to compensate them for the one year of duration risk inherent in the two-year zero.⁵

Continuing with the assumption that investors require 20 basis points for each year of duration risk, the three-year zero, with its approximately two years of duration risk,⁶ needs to offer an expected return of 40 basis points. The next section shows that this return can be effected by pricing the three-year zero as if rates next year are 20 basis points above their true values and as if rates the year after next are 40 basis points above their true values. To summarize, consider trees (a) and (b) below. If tree (a) depicts the actual or true interest rate process, then pricing with tree (b) provides investors with a risk premium of 20 basis points for each year of duration

⁵The reader should keep in mind that a two-year zero has one year of interest rate risk only in this stylized example: it has been assumed that rates can move only once a year. In reality rates can move at any time so a two-year zero has two full years of interest rate risk.


⁶A three-year zero has two years of interest rate risk only in this stylized example. See the previous footnote.

risk. If this risk premium is, in fact, embedded in market prices, then, by definition, tree (b) is the risk-neutral interest rate process.



The text now verifies that pricing the three-year zero with the risk-neutral process does offer an expected return of 10.4%, assuming that rates actually move according to the true process.

The price of the three-year zero can be computed by discounting using the risk-neutral tree:



To find the expected return of the three-year zero over the next year, proceed as follows. Two years from now the three-year zero will be a one-year zero with no interest rate risk.⁷ Therefore, its price will be determined by discounting at the actual interest rate at that time: $\frac{1}{1.18}$ or .847458, $\frac{1}{1.10}$ or .909091, and $\frac{1}{1.02}$ or .980392. One year from now, however, the three-year zero will be a two-year zero with one year of duration risk. Therefore, its price at that time will be determined by using the risk-neutral rates of

⁷Once again, this is an artifact of this example in which rates change only once a year.

14.20% and 6.20%. In particular, the two possible prices of the three-year zero in one year are

$$.769067 = \frac{.5 (.847458 + .909091)}{1.142} \quad (8.13)$$

and

$$.889587 = \frac{.5 (.909091 + .980392)}{1.062} \quad (8.14)$$

Finally, then, the expected return of the three-year zero over the next year is

$$\frac{.5 (.769067 + .889587) - .751184}{.751184} = 10.40\% \quad (8.15)$$

To summarize, in order to compensate investors for two years of duration risk, the return on the three-year zero is 40 basis points above a one-year zero's certain return of 10%.

Continuing with the assumption of 400-basis-point volatility, Figure 8.3 graphs the term structure of spot rates for three cases: no risk premium; a risk premium of 20 basis points per year of duration risk; and a risk premium of 40 basis points. In the case of no risk premium, the term structure of spot rates is downward-sloping due to convexity. A risk premium of




FIGURE 8.3 Volatility, Risk Premium, and the Shape of the Term Structure in Three-Date Binomial Models

20 basis points pushes up spot rates while convexity pulls them down. In the short end the risk premium effect dominates and the term structure is mildly upward-sloping. In the long end the convexity effect dominates and the term structure is mildly downward-sloping. The next section clarifies why risk premium tends to dominate in the short end while convexity tends to dominate in the long end. Finally, a risk premium as large as 40 basis points dominates the convexity effect and the term structure of spot rates is upward-sloping. The convexity effect is still evident, however, from the fact that the curve increases more rapidly from one to two years than from two to three years.

Just as the section on volatility uses unrealistically high levels of volatility to illustrate its effects, this section uses unrealistically high levels of the risk premium to illustrate its effects. The application at the end of this chapter focuses on reasonable magnitudes for the various effects in the context of the USD and JPY swap markets.

Before closing this section, a few remarks on the sources of an interest rate risk premium are in order. Asset pricing theory (e.g., the Capital Asset Pricing Model, or CAPM) teaches that assets whose returns are positively correlated with aggregate wealth or consumption will earn a risk premium. Consider, for example, a traded stock index. That asset will almost certainly do well if the economy is doing well and poorly if the economy is doing poorly. But investors, as a group, already have a lot of exposure to the economy. To entice them to hold a little more of the economy in the form of a traded stock index requires the payment of a risk premium; i.e., the index must offer an expected return greater than the risk-free rate of return. On the other hand, say that there exists an asset that is negatively correlated with the economy. Holdings in that asset allow investors to reduce their exposure to the economy. As a result, investors would accept an expected return on that asset below the risk-free rate of return. That asset, in other words, would have a negative risk premium.

This section assumes that bonds with interest rate risk earn a risk premium. In terms of asset pricing theory, this is equivalent to assuming that bond returns are positively correlated with the economy or, equivalently, that falling interest rates are associated with good times. One argument supporting this assumption is that interest rates fall when inflation and expected inflation fall and that low inflation is correlated with good times.

The concept of a risk premium in fixed income markets has probably gained favor more for its empirical usefulness than for its theoretical solidity. On average, over the past 75 years, the term structure of interest rates has sloped upward.⁸ While the market may from time to time expect that interest

⁸See, for example, Homer, S., and Richard Sylla, *A History of Interest Rates*, 3rd Edition, Revised, Rutgers University Press, 1996, pp. 394–409.

rates will rise, it is hard to believe that the market expects interest rates to rise on average. Therefore, expectations cannot explain a term structure of interest rates that, on average, slopes upward. Convexity, of course, leads to a downward-sloping term structure. Hence, of the three effects described in this chapter, only a positive risk premium can explain a term structure that, on average, slopes upward.

An uncomfortable fact, however, is that over earlier time periods the term structure has, on average, been flat.⁹ Whether this means that an interest rate risk premium is a relatively recent phenomenon that is here to stay or that the experience of persistently upward-sloping curves is only partially due to a risk premium is a question beyond the scope of this book. In short, the theoretical and empirical questions with respect to the existence of an interest rate risk premium have not been settled.

A MATHEMATICAL DESCRIPTION OF EXPECTATIONS, CONVEXITY, AND RISK PREMIUM

This section presents a decomposition of return and of rates in fixed income markets. The level of mathematics of this section is higher than that of most of this book, but the discussion still aims at intuition.

Assume that all bond prices are determined by a single interest rate factor, namely the instantaneous rate r taking on the value r_t at time t . Then, let $P_t(r_t, T)$ be the price of a T -year zero-coupon bond at time t . By *Ito's Lemma*, a discussion of which is beyond the scope of this book,

$$dP = \frac{\partial P}{\partial r} dr + \frac{\partial P}{\partial t} dt + \frac{1}{2} \frac{\partial^2 P}{\partial r^2} \sigma^2 dt \quad (8.16)$$

where dP , dr , and dt are the changes in price, rate, and time over the next instant, respectively, and σ is the volatility of the instantaneous rate measured in basis points per year. The two first-order partial derivatives $\frac{\partial P}{\partial r}$ and $\frac{\partial P}{\partial t}$ denote the change in the bond price for a unit change in the rate (with time unchanged) and the change in the bond price for a unit change in time (with yield unchanged), respectively, over the next instant. Finally, the second order partial derivative, $\frac{1}{2} \frac{\partial^2 P}{\partial r^2}$, gives the change in $\frac{\partial P}{\partial r}$ for a unit change in rate (with time unchanged) over the next instant. Dividing both sides of (8.16) by price,

$$\frac{dP}{P} = \frac{1}{P} \frac{\partial P}{\partial r} dr + \frac{1}{P} \frac{\partial P}{\partial t} dt + \frac{1}{2} \frac{1}{P} \frac{\partial^2 P}{\partial r^2} \sigma^2 dt \quad (8.17)$$

⁹Ibid.

Thus, equation (8.17) breaks down the return from the zero coupon bond's price changes over the next instant, $\frac{dP}{P}$, into three components. But this equation can be written in a more intuitive form by invoking several facts from earlier chapters.

First, the time t price of a T -year zero-coupon bond can be written in terms of the then T -year continuously compounded spot rate. This spot rate is a function of both the short-term rate factor r_t and T , but will be written more simply as $\hat{r}^c(T)$:

$$P_t = e^{-\hat{r}_t^c(T)T} \quad (8.18)$$

Then, differentiating both sides of (8.18) with respect to t , recognizing that increasing t decreases T one-for-one,

$$\frac{\partial P}{\partial t} = -\frac{\partial P}{\partial T} = \frac{\partial}{\partial T} [\hat{r}_t^c(T) T] \times P_t \quad (8.19)$$

But, combining equations (2.37) and (2.41) of Appendix B in Chapter 2, it can be seen that $\frac{\partial}{\partial T} [\hat{r}_t^c(T) T]$ is just the instantaneous forward rate of term T , written as $f^c(T)$. Hence,

$$\frac{\partial P}{\partial t} = f^c(T) P_t \quad (8.20)$$

Second, by the definitions of duration and convexity, D and C ,

$$D \equiv -\frac{1}{P} \frac{\partial P}{\partial r} \quad (8.21)$$

$$C \equiv \frac{1}{P} \frac{\partial^2 P}{\partial r^2} \quad (8.22)$$

Now substitute equations (8.20) through (8.22) into the return breakdown of (8.17), to see that

$$\frac{dP}{P} = f^c(T) dt - Ddr + \frac{1}{2}C\sigma^2 dt \quad (8.23)$$

The left-hand side of equation (8.23) is the return of the zero-coupon bond. The right-hand side gives the three components of return. The first component equals the return due to the passage of time, which, in this case, is the forward rate of term T . The second and third components equal the return due to changes in the rate. The second term says that increases in rate reduce bond return and that the greater the duration of the bond, the

greater this effect. This term is perfectly consistent with the discussion of interest rate sensitivity in Part Two of the book.

The third term on the right-hand side of equation (8.23) is consistent with the related discussions in Chapter 4. It was shown there that bond return increases with convexity multiplied by the change in rate squared. Here, C is multiplied by the volatility of the rate instead of the rate squared. By the definition of volatility and variance, of course, these quantities are very closely related: variance equals the expected value of the rate squared minus the square of the expected rate.

Chapter 4 showed that, holding duration constant, positive convexity increases the return of a security as rates change, whether they rise or fall. Equation (8.23) has the same implication: the greater the convexity and the greater the volatility of the rate, the greater the return. The text turns in a moment to the cost of this convexity-induced return.

To draw conclusions about the expected returns of bonds with different duration and convexity characteristics, it will prove useful to take the expectation of each side of (8.23), obtaining

$$E \left[\frac{dP}{P} \right] = f^c(T) dt - DE[dr] + \frac{1}{2} C \sigma^2 dt \quad (8.24)$$

Equation (8.24) divides expected return into its mathematical components. These components are analogous to those in equation (8.23): a return due to the passage of time, a return due to expected changes in rate, and a return due to volatility and convexity. To develop equation (8.23) further, the analysis must incorporate the economics of expected return.

Risk-neutral investors demand that each bond offer an expected return equal to the short-term rate of interest. The interest rate risk of one bond relative to another would not affect the required expected returns. Mathematically,

$$E \left[\frac{dP}{P} \right] = r_t dt \quad (8.25)$$

Risk-averse investors demand higher expected returns for bonds with more interest rate risk. The Appendix in this chapter shows that the interest rate risk of a bond over the next instant may be measured by its duration with respect to the interest rate factor and that risk-averse investors demand a risk premium proportional to duration. Letting the risk premium parameter be λ , the expected return equation for risk-averse investors becomes

$$E \left[\frac{dP}{P} \right] = r_t dt + \lambda D dt \quad (8.26)$$

Say, for example, that the short-term rate is 1%, that the duration of a particular bond is five, and that the risk premium is 10 basis points per year of duration risk. Then, according to equation (8.26), the expected return of the bond equals $1\% + 5 \times .1\%$ or 1.5% per year.

Another useful way to think of the risk premium is in terms of the *Sharpe Ratio* of a security, defined as its expected excess return (i.e., its expected return above the short-term interest rate), divided by the standard deviation of the return. Since the random part of a bond's return comes from its duration times the change in rate, see (8.23), the standard deviation of the return equals the duration times the standard deviation of the rate. Therefore, the Sharpe ratio of a bond, S , may be written as

$$S = \frac{E\left[\frac{dP}{P}\right] - r_t dt}{\sigma D dt} \quad (8.27)$$

Comparing equations (8.26) and (8.27), one can see that $S = \frac{\lambda}{\sigma}$. So, continuing with the numerical example, if the risk premium is 10 basis points per year and if the standard deviation is 100 basis points per year, then the Sharpe ratio of a bond investment is 10%.

Equipped with the economic description of expected returns in (8.26), the text can now draw conclusions about the determination of forward rates. Equate the right-hand sides of the expected return in (8.26) and the breakdown of expected return in (8.24) to see that

$$f^c(T) = \left\{ r_t + E\left[\frac{dr}{dt}\right] D \right\} + \lambda D - \frac{1}{2} C \sigma^2 \quad (8.28)$$

Equation (8.28) mathematically describes the determinants of forward rates. The three terms on the right-hand side represent the effects of expectations, risk premium, and convexity, respectively. The first term says that the forward rate is composed of the instantaneous interest rate plus the expected change in that instantaneous rate times the duration of the zero-coupon bond corresponding to the term of the forward rate. In other words, the higher the instantaneous rate, the higher the forward rate; the more rates are expected to increase, the higher the forward rate; and the greater the corresponding duration, i.e., the greater the term of the forward rate, the greater the effect of expected instantaneous rate changes on the forward rate.

The second term on the right-hand side of (8.28) says that the forward rate increases with the risk premium in proportion to the corresponding zero coupon bond duration. In other words, the greater the corresponding interest rate risk and the greater the risk premium, the greater the forward rate.

Chapter 7 noted that pricing bonds as if the short-term rate drifted up by a certain amount each year has the same effect as a risk premium per year of that amount. Inspection of equation (8.28) reveals this equivalence more formally. Increasing the risk premium by a fixed number of basis points is empirically indistinguishable from increasing the expected short-term rate, through $E\left[\frac{dr}{dt}\right]$, by the same number of basis points. This means that the market term structure at any given time cannot be used to distinguish between market expectations of rate changes and risk premium. From a modeling perspective this means that only the risk-neutral process is relevant for pricing. Dividing the drift into expectations and risk premium might be very useful in determining whether the model seems economically reasonable, but this division has no pricing implications.

Finally, the third term of (8.28) shows that the forward rate falls with interest rate volatility and the corresponding zero-coupon convexity. Equation (8.24) shows that the expected return of a bond is enhanced by its convexity in the quantity $\frac{1}{2}C\sigma^2 dt$. But the convexity term in (8.28) shows that the forward rate and, therefore, the expected return due to the passage of time in (8.28), are reduced by exactly that amount. Hence, as claimed in Chapter 7 and as mentioned earlier in this chapter, a bond priced by arbitrage offers no advantage in expected return due to its convexity. In fact, the expected return condition (8.26) ensures that this is so.

APPLICATION: EXPECTATIONS, CONVEXITY, AND RISK PREMIUM IN USD AND JPY SWAP MARKETS

This section uses the simple framework of the previous section to illustrate the magnitudes of the effects described in this chapter in the context of USD and JPY swap markets. EUR and GBP swap markets will be mentioned at the end of the section.

Figure 8.4 is an example of a decomposition of the term structure of forward rates in the USD swap market as of May 28, 2010, into expectations, risk premium, and convexity. This decomposition was achieved by making assumptions sufficient to apply equation (8.28).

The first challenge in applying equation (8.28) is the convexity effect on forward rates. This effect equals $-\frac{1}{2}C\sigma^2$, where C is the convexity of the matching-maturity zero coupon bond with respect to the short-term interest rate and σ is the volatility of that rate. This is difficult to apply precisely without invoking a more specific model, like those presented later in Part Three, for two reasons. First, the convexity formulas from Chapter 4 are with respect to spot rates or zero coupon yields. Second, in reality there does not exist one interest rate volatility but rather a term structure of volatilities. The exact handling of these issues will become clear over the next few chapters. For the purposes of this application, however, the




FIGURE 8.4 A Decomposition of the USD Swap Curve as of May 28, 2010, with a Risk Premium of 9 Basis Points per Year

TABLE 8.1 Spot Rate Volatilities of Selected Terms from the Swaptions Market as of May 28, 2010

Term in Years	Spot Rate Volatility	
	USD Swaps	JPY Swaps
2	105.46	17.23
5	114.45	30.01
10	112.91	45.73
25	91.79	67.02

convexity term in (8.28) is calculated using the convexity with respect to the spot rate—equation (4.51)—and the variance of that spot rate from the term structure of volatilities in the swaption market,¹⁰ which are shown in Table 8.1. Hence, the convexity term in (8.28) for a six-month rate 10 years forward, with a 10-year spot rate of 3.512%, is taken to be

$$-\frac{1}{2} \frac{10 \times 10.5}{\left(1 + \frac{3.512\%}{2}\right)^2} \times \left(\frac{112.91}{10,000}\right)^2 = -.646\% \tag{8.29}$$

¹⁰ More precisely, the average of the caplet (forward) volatilities over a particular term is taken to be the spot rate volatility of that term.

The second challenge in applying equation (8.28) is setting the risk premium. As mentioned earlier, there is no way of separating expectations and risk premium from fixed income security prices alone. For the purposes of this illustration, therefore, several strong assumptions were made. First, consistent with the notation of the previous section, it is assumed that the risk premium is constant although the risk premium may, in theory, depend on calendar time and the level of rates. Second, rate expectations are relatively flat at longer maturities. As mentioned earlier, while the market might expect a particular path of rates in the short term, it is hard to defend any such expectations from 20 to 30 years in the future. Third, the long-run expectation of the short-term rate is about 5%, corresponding to a long-run real rate of 3% and a long-run rate of inflation of 2%. Fourth, the Sharpe ratio of investments in bonds is consistent with historical norms.

As it turns out, a risk premium of 9 basis points per year satisfies the assumptions of the previous paragraph relatively well. The resulting expectations curve in Figure 8.4 is relatively flat at 5% at long maturities. Also, at an interest rate volatility of about 100 basis points, a risk premium of 9 basis points per year gives a Sharpe ratio of 9%, which is a historically plausible order of magnitude. In any case, using a risk premium of 9 basis points per year in (8.28), the risk premium part of a forward rate is just λD , where D is the duration of the appropriate zero coupon bond. For example, the risk premium component of the six-month rate 10 years forward, with a 10-year spot rate of 3.512%, is about 88 basis points:

$$\frac{10}{1 + \frac{3.512\%}{2}} \times \frac{9}{10,000} = .884\% \quad (8.30)$$

Taken as a whole, Figure 8.4 gives an idea of the orders of magnitude of expectations, convexity, and risk premium on forward rates of different terms. It so happens that, over the relevant range, the risk premium and convexity effects to a large extent cancel, leaving expected rates approximately equal to forward rates.

Figure 8.5 performs a similar exercise for JPY swaps. Because the volatilities are lower than for the USD curve, the convexity effect is lower as well. The risk premium is kept at nine basis points per year, which does result in relatively flat rate expectations, although the implied Sharpe ratios are much higher because of the lower volatilities.

Interestingly, a similar decomposition in EUR and GBP does not produce results that are as reasonable. The significantly downward-sloping forward rate curves in the long end in those currencies, displayed in Chapter 2, require a zero if not slightly negative risk premium for long-rate expectations to be flat. This is understandable in light of the discussion in the Overview that




FIGURE 8.5 A Decomposition of the JPY Swap Curve as of May 28, 2010, with a Risk Premium of 9 Basis Points per Year

pension fund and insurance company demand for long-end duration has distorted the long-end of the term structure in these currencies.

APPENDIX: PROOF OF EQUATION (8.26)

This proof follows that of Ingersoll (1987)¹¹ and assumes some knowledge of stochastic processes and their associated notation. This notation is described in Chapters 9 and 10.

Assume that r , the single, instantaneous interest rate factor, follows the process

$$dr = \mu dt + \sigma dw \quad (8.31)$$

Let P be the full price of some security that depends on r and time. Then, by Ito's Lemma,

$$dP = P_r dr + P_t dt + \frac{1}{2} P_{rr} \sigma^2 dt \quad (8.32)$$

¹¹ Ingersoll, J., *Theory of Financial Decision Making*, Rowman & Littlefield, 1987.

where P_r , P_t , and P_{rr} denote the partial first derivatives with respect to r and t and the second partial derivative with respect to r . Dividing both sides of (8.32) by P , taking expectations, and defining α_P to be the expected return of the security,

$$\alpha_P dt \equiv E \left[\frac{dP}{P} \right] = \frac{P_r}{P} \mu dt + \frac{P_t}{P} dt + \frac{1}{2} P_{rr} \sigma^2 dt \quad (8.33)$$

Combining (8.31), (8.32) and (8.33),

$$\frac{dP}{P} - \alpha_P dt = \frac{P_r}{P} \sigma dw \quad (8.34)$$

Since equation (8.34) applies to any security, it also applies to some other security Q :

$$\frac{dQ}{Q} - \alpha_Q dt = \frac{Q_r}{Q} \sigma dw \quad (8.35)$$

Now consider the strategy of investing \$1 in security P and

$$- \frac{P_r Q}{P Q_r} \quad (8.36)$$

dollars in security Q . Using equations (8.34) and (8.35), the return on this portfolio is

$$\frac{dP}{P} - \frac{P_r Q}{P Q_r} \frac{dQ}{Q} = \alpha_P dt - \frac{P_r Q}{P Q_r} \alpha_Q dt \quad (8.37)$$

Notice that terms with the random variable dw have fallen out of equation (8.37). This particular portfolio was, in fact, chosen so as to hedge completely the risk of P with Q . In any case, since the portfolio has no risk it must earn the instantaneous rate r :

$$\alpha_P dt - \frac{P_r Q}{P Q_r} \alpha_Q dt = \left(1 - \frac{P_r Q}{P Q_r} \right) r dt \quad (8.38)$$

Rearranging the terms of (8.38),

$$\frac{\alpha_P - r}{-\frac{P_r}{P}} = \frac{\alpha_Q - r}{-\frac{Q_r}{Q}} \equiv \lambda(r, t) \quad (8.39)$$

Equation (8.39) says that the expected return of any security above the instantaneous rate divided by its duration with respect to that rate must equal some function λ . This function cannot depend on any characteristic of the security because (8.39) is true for all securities. The function may depend on the interest rate factor and time although, this book, for simplicity, assumes that λ is constant. Rewriting (8.39), for each security it must be true that

$$E \left[\frac{dP}{P} \right] \equiv \alpha_P dt = r dt + \lambda D dt \quad (8.40)$$

The Art of Term Structure Models: Drift

Chapters 7 and 8 show that assumptions about the true and risk-neutral short-term rate processes determine the term structure of interest rates and the prices of fixed income derivatives. The goal of this chapter is to describe the most common building blocks of short-term rate models. Selecting and rearranging these building blocks to create suitable models for the purpose at hand is the art of term structure modeling.

This chapter begins with an extremely simple model with no drift and normally distributed rates. The next sections add and discuss the implications of alternate specifications of the drift: a constant drift, a time-deterministic shift, and a mean-reverting drift.

MODEL 1: NORMALLY DISTRIBUTED RATES AND NO DRIFT

The particularly simple model of this section will be called Model 1. The continuously compounded, instantaneous rate r_t is assumed to evolve according to the following equation:

$$dr = \sigma dw \quad (9.1)$$

The quantity dr denotes the change in the rate over a small time interval, dt , measured in years; σ denotes the annual *basis-point volatility* of rate changes; and dw denotes a normally distributed random variable with a mean of zero and a standard deviation of $\sqrt{dt}$.¹

Say, for example, that the current value of the short-term rate is 6.18%, that volatility equals 113 basis points per year, and that the time interval

¹It is beyond the mathematical scope of the text to explain why the random variable dw is denoted as a change. But the text uses this notation since it is the convention of the field.


under consideration is one month or $\frac{1}{12}$ years. Mathematically, $r_0 = 6.18\%$; $\sigma = 1.13\%$; and $dt = \frac{1}{12}$. A month passes and the random variable dw , with its zero mean and its standard deviation of $\sqrt{\frac{1}{12}}$ or .2887, happens to take on a value of .15. With these values, the change in the short-term rate given by (9.1) is

$$dr = 1.13\% \times .15 = .17\% \quad (9.2)$$


or 17 basis points. Since the short-term rate started at 6.18%, the short-term rate after a month is 6.35%.

Since the expected value of dw is zero, (9.1) says that the expected change in the rate, or the drift, is zero. Since the standard deviation of dw is $\sqrt{dt}$, the standard deviation of the change in the rate is $\sigma\sqrt{dt}$. For the sake of brevity, the standard deviation of the change in the rate will be referred to as simply the standard deviation of the rate. Continuing with the numerical example, the process (9.1) says that the drift is zero and that the standard deviation of the rate is $\sigma\sqrt{dt}$, which is $1.13\% \times \sqrt{\frac{1}{12}} = .326\%$ or 32.6 basis points per month.

A rate tree may be used to approximate the process (9.1). A tree over dates 0 to 2 takes the following form:



In the case of the numerical example, substituting the sample values into the tree gives the following:



To understand why these trees are representations of the process (9.1), consider the transition from date 0 to date 1. The change in the interest rate in the up state is $\sigma\sqrt{dt}$ and the change in the down state is $-\sigma\sqrt{dt}$. Therefore, with the probabilities given in the tree, the expected change in the rate, often denoted $E[dr]$, is

$$E[dr] = .5 \times \sigma\sqrt{dt} + .5 \times -\sigma\sqrt{dt} = 0 \quad (9.3)$$

The variance of the rate, often denoted $V[dr]$, from date 0 to date 1 is computed as follows:

$$\begin{aligned} V[dr] &= E[dr^2] - \{E[dr]\}^2 \\ &= .5 \times (\sigma\sqrt{dt})^2 + .5 \times (-\sigma\sqrt{dt})^2 - 0 \\ &= \sigma^2 dt \end{aligned} \quad (9.4)$$

Note that the first line of (9.4) follows from the definition of variance. Since the variance is $\sigma^2 dt$, the standard deviation, which is the square root of the variance, is $\sigma\sqrt{dt}$.

Equations (9.3) and (9.4) show that the drift and volatility implied by the tree match the drift and volatility of the interest rate process (9.1). The process and the tree are not identical because the random variable in the process, having a normal distribution, can take on any value while a single step in the tree leads to only two possible values. In the example, when dw takes on a value of .15, the short rate changes from 6.18% to 6.35%. In the tree, however, the only two possible rates are 6.506% and 5.854%. Nevertheless, as shown in Chapter 7, after a sufficient number of time steps the branches of the tree used to approximate the process (9.1) will be numerous enough to approximate a normal distribution. Figure 9.1 shows the distribution of short rates after one year, or the *terminal distribution* after one year, in Model 1 with $r_0 = 6.18\%$ and $\sigma = 1.13\%$. The tick marks on the horizontal axis are one standard deviation apart from one another.

Models in which the terminal distribution of interest rates has a normal distribution, like Model 1, are called *normal* or *Gaussian* models. One problem with these models is that the short-term rate can become negative. A negative short-term rate does not make much economic sense because people would never lend money at a negative rate when they can hold cash and earn a zero rate instead.² The distribution in Figure 9.1, drawn to encompass three standard deviations above and below the mean, shows that

²Actually, the interest rate can be slightly negative if a security or bank account were safer or more convenient than holding cash.




FIGURE 9.1 Distribution of Short Rates After One Year, Model 1

over a horizon of one year the interest rate process will almost certainly not exhibit negative interest rates. The probability that the short-term rate in the process (9.1) becomes negative, however, increases with the horizon. Over 10 years, for example, the standard deviation of the terminal distribution in the numerical example is $1.13\% \times \sqrt{10}$ or 3.573%. Starting with a short-term rate of 6.18%, a random negative shock of only $\frac{6.18\%}{3.573\%}$ or 1.73 standard deviations would push rates below zero.

The extent to which the possibility of negative rates makes a model unusable depends on the application. For securities whose value depends mostly on the average path of the interest rate, like coupon bonds, the possibility of negative rates typically does not rule out an otherwise desirable model. For securities that are asymmetrically sensitive to the probability of low interest rates, however, using a normal model could be dangerous. Consider the extreme example of a 10-year call option to buy a long-term coupon bond at a yield of 0%. The model of this section would assign that option much too high a value because the model assigns too much probability to negative rates.

The challenge of negative rates for term structure models is much more acute, of course, when the current level of rates is low, as it is at the time of this writing. Changing the distribution of interest rates is one solution. To take but one of many examples, lognormally distributed rates, as will be seen in Chapter 10, cannot become negative. As will become clear later in that chapter, however, building a model around a probability distribution that rules out negative rates or makes them less likely may result in volatilities that are unacceptable for the purpose at hand.




FIGURE 9.2 Rate Curves from Model 1 and Selected Market Swap Rates, February 16, 2001

Another popular method of ruling out negative rates is to construct rate trees with whatever distribution is desired, as done in this section, and then simply set all negative rates to zero.³ In this methodology, rates in the original tree are called the shadow rates of interest while the rates in the adjusted tree could be called the observed rates of interest. When the observed rate hits zero, it can remain there for a while until the shadow rate crosses back to a positive rate. The economic justification for this framework is that the observed interest rate should be constrained to be positive only because investors have the alternative of investing in cash. But the shadow rate, the result of aggregate savings, investment, and consumption decisions, may very well be negative. Of course, the probability distribution of the observed interest rate is not the same as that of the originally postulated shadow rate. The change, however, is localized around zero and negative rates. By contrast, changing the form of the probability distribution changes dynamics across the entire range of rates.

Returning now to Model 1, the techniques of Chapter 7 may be used to price fixed coupon bonds. Figure 9.2 graphs the semiannually compounded par, spot, and forward rate curves for the numerical example along with data from U.S. dollar swap par rates. The initial value of the short-term rate in the example, 6.18%, is set so that the model and market 10-year, semiannually compounded par rates are equal at 6.086%. All of the other data points shown are quite different from their model values. The desirability of fitting

³Fischer Black, "Interest Rates as Options," *Journal of Finance*, Vol. 50, 1995, pp. 1371–1376.

TABLE 9.1 Convexity Effects on Par Rates in a Parameterization of Model 1

Term (years)	Convexity (bps)
2	−0.8
5	−5.1
10	−18.8
30	−135.3

market data exactly is discussed in its own section, but Figure 9.2 clearly demonstrates that the simple model of this section does not have enough flexibility to capture the simplest of term structure shapes.

The model term structure is downward-sloping. As the model has no drift, rates decline with term solely because of convexity. Table 9.1 shows the magnitudes of the convexity effect on par rates of selected terms.⁴ The numbers are realistic in the sense that a volatility of 113 basis points a year is reasonable. In fact, the volatility of the 10-year swap rate on the data date, as implied by options markets, was 113 basis points. The convexity numbers are not necessarily realistic, however, because, as this chapter will demonstrate, the magnitude of the convexity effect depends on the model and Model 1 is almost certainly not the best model of interest rate behavior.

The term structure of volatility in Model 1 is constant at 113 basis points per year. In other words, the standard deviation of changes in the par rate of any maturity is 113 basis points per year. As shown in Figure 9.3, this implication fails to capture the implied volatility structure in the market. The volatility data in Figure 9.3 show that the term structure of volatility is humped, i.e., that volatility initially rises with term but eventually declines. As this shape is a feature of fixed income markets, it will be revisited again in this chapter and in Chapters 10 and 11.

The last aspect of this model to be analyzed is its factor structure. The model's only factor is the short-term rate. If this rate increases by 10 semiannually compounded basis points, how would the term structure change? In this simple model the answer is that all rates would increase by 10 basis points. (See the closed-form solution for spot rates in Model 1 in the Appendix in Chapter 10). Therefore, Model 1 is a model of parallel shifts.

⁴The convexity effect is the difference between the par yield in the model with its assumed volatility and the par yield in the same structural model but with a volatility of zero.




FIGURE 9.3 Par Rate Volatility from Model 1 and Selected Implied Volatilities, February 16, 2001

MODEL 2: DRIFT AND RISK PREMIUM

The term structures implied by Model 1 always look like Figure 9.2: relatively flat for early terms and then downward sloping. Chapter 8 pointed out that the term structure tends to slope upward and that this behavior might be explained by the existence of a risk premium. The model of this section, to be called Model 2, adds a drift to Model 1, interpreted as a risk premium, in order to obtain a richer model in an economically coherent way.

The dynamics of the risk-neutral process in Model 2 are written as

$$dr = \lambda dt + \sigma dw \quad (9.5)$$

The process (9.5) differs from that of Model 1 by adding a drift to the short-term rate equal to λdt . For this section, consider the values $r_0 = 5.138\%$, $\lambda = .229\%$, and $\sigma = 1.10\%$. If the realization of the random variable dw is again .15 over a month, then the change in rate is


$$dr = .229\% \times \frac{1}{12} + 1.10\% \times .15 = .1841\% \quad (9.6)$$

Starting from 5.138%, the new rate is 5.322%.

The drift of the rate is $.229\% \times \frac{1}{12}$ or 1.9 basis points per month, and the standard deviation is $1.10\% \times \sqrt{\frac{1}{12}}$ or 31.75 basis points per month. As discussed in Chapter 8, the drift in the risk-neutral process is a

combination of the true expected change in the interest rate and of a risk premium. A drift of 1.9 basis points per month may arise because the market expects the short-term rate to increase by 1.9 basis points a month, because the short-term rate is expected to increase by one basis point with a risk premium of .9 basis points, or because the short-term rate is expected to fall by .1 basis points with a risk premium of two basis points.

The tree approximating this model is



It is easy to verify that the drift and standard deviation of the tree match those of the process (9.5).

The terminal distribution of the numerical example of this process after one year is normal with a mean of $5.138\% + 1 \times .229\%$ or 5.367% and a standard deviation of 110 basis points. After 10 years, the terminal distribution is normal with a mean of $5.138\% + 10 \times .229\%$ or 7.428% and a standard deviation of $1.10\% \times \sqrt{10}$ or 347.9 basis points. Note that the constant drift, by raising the mean of the terminal distribution, makes it less likely that the risk-neutral process will exhibit negative rates.

Figure 9.4 shows the rate curves in this example along with par swap rate data. The values of r_0 and λ are calibrated to match the 2- and 10-year par swap rates, while the value of σ is chosen to be the average implied volatility of the 2- and 10-year par rates. The results are satisfying in that the resulting curve can match the data much more closely than did the curve of Model 1 shown in Figure 9.2. Slightly unsatisfying is the relatively high value of λ required. Interpreted as a risk premium alone, a value of .229% with a volatility of 110 basis points implies a relatively high Sharpe ratio of about .21. On the other hand, interpreting λ as a combination of true drift and risk premium is difficult in the long end where, as argued in Chapter 8, it is difficult to make a case for rising expected rates. These interpretive difficulties arise because Model 2 is still not flexible enough to explain the shape of the term structure in an economically meaningful way. In fact, the use of r_0 and λ to match the 2- and 10-year rates in this relatively inflexible model may explain why the model curve overshoots the 30-year par rate by about 25 basis points.




FIGURE 9.4 Rate Curves from Model 2 and Selected Market Swap Rates, February 16, 2001

Moving from Model 1 with zero drift to Model 2 with a constant drift does not qualitatively change the term structure of volatility, the magnitude of convexity effects, or the parallel-shift nature of the model.

Models 1 and 2 would be called equilibrium models because no effort has been made to match the initial term structure closely. The next section presents a generalization of Model 2 that is in the class of arbitrage-free models.


THE HO-LEE MODEL: TIME-DEPENDENT DRIFT

The dynamics of the risk-neutral process in the Ho-Lee model are written as

$$dr = \lambda_t dt + \sigma dw \quad (9.7)$$

In contrast to Model 2, the drift here depends on time. In other words, the drift of the process may change from date to date. It might be an annualized drift of -20 basis points over the first month, of 20 basis points over the second month, and so on. A drift that varies with time is called a *time-dependent* drift. Just as with a constant drift, the time-dependent drift over each time period represents some combination of the risk premium and of expected changes in the short-term rate.

The flexibility of the Ho-Lee model is easily seen from its corresponding tree:



The free parameters λ_1 and λ_2 may be used to match the prices of securities with fixed cash flows. The procedure may be described as follows. With $dt = \frac{1}{12}$, set r_0 equal to the one-month rate. Then find λ_1 such that the model produces a two-month spot rate equal to that in the market. Then find λ_2 such that the model produces a three-month spot rate equal to that in the market. Continue in this fashion until the tree ends. The procedure is very much like that used to construct the trees in Chapter 7. The only difference is that Chapter 7 adjusts the probabilities to match the spot rate curve while this section adjusts the rates. As it turns out, the two procedures are equivalent so long as the step size is small enough.

The rate curves resulting from this model match all the rates that are input into the model. Just as adding a constant drift to Model 1 to obtain Model 2 does not affect the shape of the term structure of volatility nor the parallel-shift characteristic of the model, adding a time-dependent drift does not change these features either.

DESIRABILITY OF FITTING TO THE TERM STRUCTURE

The desirability of matching market prices is the central issue in deciding between arbitrage-free and equilibrium models. Not surprisingly, the choice depends on the purpose of building the model in the first place.

One important use of arbitrage-free models is for quoting the prices of securities that are not actively traded based on the prices of more liquid securities. A customer might ask a swap desk to quote a rate on a swap to a particular date, say three years and four months away, while liquid market prices might be observed only for three- and four-year swaps, or sometimes only for two- and five-year swaps. In this situation the swap desk may price the odd-maturity swap using an arbitrage-free model essentially as a means of interpolating between observed market prices.

Interpolating by means of arbitrage-free models may very well be superior to other curve-fitting methods, from linear interpolation to more sophisticated approaches. The potential superiority of arbitrage-free models arises from their being based on economic and financial reasoning. In an arbitrage-free model the expectations and risk premium built into neighboring swap rates and the convexity implied by the model's volatility assumptions are used to compute, for example, the three-year and four-month swap rate. In a purely mathematical curve fitting technique, by contrast, the chosen functional form heavily determines the intermediate swap rate. Selecting linear or quadratic interpolation, for example, results in intermediate swap rates with no obvious economic or financial justification. This potential superiority of arbitrage-free models depends crucially on the validity of the assumptions built into the models. A poor volatility assumption, for example, resulting in a poor estimate of the effect of convexity, might make an arbitrage-free model perform worse than a less financially sophisticated technique.

Another important use of arbitrage-free models is to value and hedge derivative securities for the purpose of making markets or for proprietary trading. For these purposes many practitioners wish to assume that some set of underlying securities is priced fairly. For example, when trading an option on a 10-year bond, many practitioners assume that the 10-year bond is itself priced fairly. (An analysis of the fairness of the bond can always be done separately.) Since arbitrage-free models match the prices of many traded securities by construction, these models are ideal for the purpose of pricing derivatives given the prices of underlying securities.

That a model matches market prices does not necessarily imply that it provides fair values and accurate hedges for derivative securities. The argument for fitting models to market prices is that a good deal of information about the future behavior of interest rates is incorporated into market prices, and, therefore, a model fitted to those prices captures that interest rate behavior. While this is a perfectly reasonable argument, two warnings are appropriate. First, a mediocre or bad model cannot be rescued by calibrating it to match market prices. If, for example, the parallel shift assumption is not a good enough description of reality for the application at hand, adding a time-dependent drift to a parallel shift model so as to match a set of market prices will not make the model any more suitable for that application. Second, the argument for fitting to market prices assumes that those market prices are fair in the context of the model. In many situations, however, particular securities, particular classes of securities, or particular maturity ranges of securities have been distorted due to supply and demand imbalances, taxes, liquidity differences, and other factors unrelated to interest rate models. In these cases, fitting to market prices will make a model worse by attributing these outside factors to the interest rate process. If, for example, a large bank liquidates its portfolio of bonds or swaps with approximately seven years to maturity and, in the process, depresses prices and raises rates

around that maturity, it would be incorrect to assume that expectations of rates seven years in the future have risen. Being careful with the word *fair*, the seven-year securities in this example are fair in the sense that liquidity considerations at a particular time require their prices to be relatively low. The seven-year securities are not fair, however, with respect to the expected evolution of interest rates and the market risk premium. For this reason, in fact, investors and traders might buy these relatively cheap bonds or swaps and hold them past the liquidity event in the hope of selling at a profit.

Another way to express the problem of fitting the drift to the term structure is to recognize that the drift of a risk-neutral process arises only from expectations and risk premium. A model that assumes one drift from years 15 to 16 and another drift from years 16 to 17 implicitly assumes one of two things. First the expectation today of the one-year rate in 15 years differs from the expectation today of the one-year rate in 16 years. Second, the risk premium in 15 years differs in a particular way from the risk premium in 16 years. Since neither of these assumptions is particularly plausible, a fitted drift that changes dramatically from one year to the next is likely to be erroneously attributing non-interest rate effects to the interest rate process.

If the purpose of a model is to value bonds or swaps relative to one another, then taking a large number of bond or swap prices as given is clearly inappropriate: arbitrage-free models, by construction, conclude that all of these bond or swap prices are fair relative to one another. Investors wanting to choose among securities, market makers looking to pick up value by strategically selecting hedging securities, or traders looking to profit from temporary mispricings must, therefore, rely on equilibrium models.

Having starkly contrasted arbitrage-free and equilibrium models, it should be noted that, in practice, there need not be a clear line between the two approaches. A model might posit a deterministic drift for a few years to reflect relatively short-term interest rate forecasts and posit a constant drift from then on. Another model might take the prices of 2-, 5-, 10- and 30-year bond or swap rates as given, thus assuming that the most liquid securities are fair while allowing the model to value other securities. The proper blending of the arbitrage-free and equilibrium approaches is an important part of the art of term structure modeling.

THE VASICEK MODEL: MEAN REVERSION

Assuming that the economy tends toward some equilibrium based on such fundamental factors as the productivity of capital, long-term monetary policy, and so on, short-term rates will be characterized by *mean reversion*. When the short-term rate is above its long-run equilibrium value, the drift is negative, driving the rate down toward this long-run value. When the rate is below its equilibrium value, the drift is positive, driving the rate up toward

this value. In addition to being a reasonable assumption about short rates,⁵ mean reversion enables a model to capture several features of term structure behavior in an economically intuitive way.

The risk-neutral dynamics of the Vasicek model⁶ are written as

$$dr = k(\theta - r)dt + \sigma dw \quad (9.8)$$

The constant θ denotes the long-run value or central tendency of the short-term rate in the risk-neutral process and the positive constant k denotes the speed of mean reversion. Note that in this specification the greater the difference between r and θ , the greater the expected change in the short-term rate toward θ .

Because the process (9.8) is the risk-neutral process, the drift combines both interest rate expectations and risk premium. Furthermore, market prices do not depend on how the risk-neutral drift is divided between the two. Nevertheless, in order to understand whether or not the parameters of a model make sense, it is useful to make assumptions sufficient to separate the drift and the risk premium. Assuming, for example, that the true interest rate process exhibits mean reversion to a long-term value r_∞ and, as assumed previously, that the risk premium enters into the risk-neutral process as a constant drift, the Vasicek model takes the following form:

$$\begin{aligned} dr &= k(r_\infty - r)dt + \lambda dt + \sigma dw \\ &= k\left(\left[r_\infty + \frac{\lambda}{k}\right] - r\right)dt + \sigma dw \end{aligned} \quad (9.9)$$

The process in (9.8) is identical to that in (9.9) so long as

$$\theta \equiv r_\infty + \frac{\lambda}{k} \quad (9.10)$$

⁵While reasonable, mean reversion is a strong assumption. Long time series of interest rates from relatively stable markets might display mean reversion because there happened to be no catastrophe over the time period, that is, precisely because a long time series exists. Hyperinflation, for example, is not consistent with mean reversion and results in the destruction of a currency and its associated interest rates. When mean reversion ends, the time series ends. In short, the most severe critics of mean reversion would say that interest rates mean revert until they don't.

⁶O. Vasicek, "An Equilibrium Characterization of the Term Structure, *Journal of Financial Economics*, 5, 1977, pp. 177–188. It is appropriate to add that this paper started the literature on short-term rate models. The particular dynamics of the model described in this section, which is commonly known as the Vasicek model, is a very small part of the contribution of that paper.

Note that very many combinations of r_∞ and λ give the same θ and, through the risk-neutral process (9.8), the same market prices.

For the purposes of this section, let $k = .025$, $\sigma = 126$ basis points per year, $r_0 = 6.179\%$, and $\lambda = .229\%$. According to (9.10), then, $\theta = 15.339\%$. With these parameters, the process (9.8) says that over the next month the expected change in the short rate is

$$.025 \times (15.339\% - 5.121\%) \frac{1}{12} = .0213\% \quad (9.11)$$

or 2.13 basis points. The volatility over the next month is $126 \times \sqrt{\frac{1}{12}}$ or 36.4 basis points.

Representing this process with a tree is not quite so straightforward as the simpler processes described previously because the most obvious representation leads to a nonrecombining tree. Over the first time step,

$$\begin{array}{lcl} & \nearrow \frac{1}{2} & 5.121\% + \frac{.025(15.339\% - 5.121\%)}{12} + \frac{.0126}{\sqrt{12}} = 5.5060\% \\ 5.121\% & & \\ & \searrow \frac{1}{2} & 5.121\% + \frac{.025(15.339\% - 5.121\%)}{12} - \frac{.0126}{\sqrt{12}} = 4.7786\% \end{array}$$


To extend the tree from date 1 to date 2, start from the up state of 5.5060%. The tree branching from there is

$$\begin{array}{lcl} & \nearrow \frac{1}{2} & 5.506\% + \frac{.025(15.339\% - 5.5060\%)}{12} + \frac{.0126}{\sqrt{12}} = 5.8902\% \\ 5.5060\% & & \\ & \searrow \frac{1}{2} & 5.506\% + \frac{.025(15.339\% - 5.5060\%)}{12} - \frac{.0126}{\sqrt{12}} = 5.1628\% \end{array}$$

while the tree branching from the date 1 down state of 4.7786% is

$$\begin{array}{lcl} & \nearrow \frac{1}{2} & 4.7786\% + \frac{.025(15.339\% - 4.7786\%)}{12} + \frac{.0126}{\sqrt{12}} = 5.1643\% \\ 4.7786\% & & \\ & \searrow \frac{1}{2} & 4.7786\% + \frac{.025(15.339\% - 4.7786\%)}{12} - \frac{.0126}{\sqrt{12}} = 4.4369\% \end{array}$$


To summarize, the most straightforward tree representation of (9.8) takes the following form:



This tree does not recombine since the drift increases with the difference between the short rate and θ . Since 4.7786% is further from θ than 5.5060%, the drift from 4.7786% is greater than the drift from 5.5060%. In this model, the volatility component of an up move followed by a down move does perfectly cancel the volatility component of a down move followed by an up move. But since the drift from 4.7786% is greater, the move up from 4.7786% produces a larger short-term rate than a move down from 5.5060%.

There are many ways to represent the Vasicek model with a recombining tree. One method is presented here, but it is beyond the scope of this book to discuss the numerical efficiency of the various possibilities.

The first time step of the tree may be taken as shown previously:




Next, fix the center node of the tree on date 2. Since the expected perturbation due to volatility over each time step is zero, the drift alone determines the expected value of the process after each time step. After the first time step the expected value is

$$5.121\% + .025(15.339\% - 5.121\%) \frac{1}{12} = 5.1423\% \quad (9.12)$$

After the second time step the expected value is

$$5.1423\% + .025(15.339\% - 5.1423\%) \frac{1}{12} = 5.1635\% \quad (9.13)$$

Take this value as the center node on date 2 of the recombining tree:



The parts of the tree to be solved for, namely, the missing probabilities and interest rate values, are given variable names.

According to the process (9.8) and the parameter values set in this section, the expected rate and standard deviation of the rate from 5.5060% are, respectively,

$$5.5060\% + .025 (15.339\% - 5.5060\%) \frac{1}{12} = 5.5265\% \quad (9.14)$$

and

$$1.26\% \sqrt{\frac{1}{12}} = .3637\% \quad (9.15)$$

For the recombining tree to match this expectation and standard deviation, it must be the case that

$$p \times r^{uu} + (1 - p) \times 5.1635\% = 5.5265\% \quad (9.16)$$

and, by the definition of standard deviation,

$$\sqrt{p (r^{uu} - 5.5265\%)^2 + (1 - p) (5.1635\% - 5.5265\%)^2} = .3637\% \quad (9.17)$$

Solving equations (9.16) and (9.17), $r^{uu} = 5.8909\%$ and $p = .4990$.

The same procedure may be followed to compute r^{dd} and q . The expected rate from 4.7786% is

$$4.7786\% + .025 (15.339\% - 4.7786\%) \frac{1}{12} = 4.8006\% \quad (9.18)$$

and the standard deviation is again 36.37 basis points. Starting from 4.7786%, then, it must be the case that


$$q \times 5.1635\% + (1 - q) \times r^{dd} = 4.8006\% \quad (9.19)$$

and

$$\sqrt{q(5.1635\% - 4.8006\%)^2 + (1 - q)(r^{dd} - 4.8006\%)^2} = .3637\% \quad (9.20)$$

Solving equations (9.19) and (9.20), $r^{dd} = 4.4361\%$ and $q = .5011$.


Putting the results from the up and down states together, a recombining tree approximating the process (9.8) with the parameters of this section is



To extend the tree to the next date, begin again at the center. From the center node of date 2, the expected rate of the process is

$$5.1635\% + .025 \times (15.339\% - 5.1635\%) \frac{1}{12} = 5.1847\% \quad (9.21)$$

As in constructing the tree for date 1, adding and subtracting the standard deviation of .3637% to the average value 5.1847% (obtaining 5.5484% and 4.8210%) and using probabilities of 50% for up and down movements satisfy the requirements of the process at the center of the tree:






FIGURE 9.5 Mean Reversion and the Terminal Distribution of Short Rates

The unknown parameters can be solved for in the same manner as described in building the tree on date 2.

The text now turns to the effects of mean reversion on the term structure. Figure 9.5 illustrates the impact of mean reversion on the terminal, risk-neutral distributions of the short rate at different horizons. The expectation or mean of the short-term rate as a function of horizon gradually rises from its current value of 5.121% toward its limiting value of $\theta = 15.339\%$. Because the mean-reverting parameter $k = .025$ is relatively small, the horizon expectation rises very slowly toward 15.339%. While mathematically beyond the scope of this book, it can be shown that the distance between the current value of a factor and its goal decays exponentially at the mean-reverting rate. Since the interest rate is currently $15.339\% - 5.121\%$ or 10.218% away from its goal, the distance between the expected rate at a 10-year horizon and the goal is

$$10.2180\% \times e^{-.025 \times 10} = 7.9578\% \quad (9.22)$$

Therefore, the expectation of the rate in 10 years is $15.3390\% - 7.9578\%$ or 7.3812%.

For completeness, the expectation of the rate in the Vasicek model after T years is

$$r_0 e^{-kT} + \theta (1 - e^{-kT}) \quad (9.23)$$

In words, the expectation is a weighted average of the current short rate and its long-run value, where the weight on the current short rate decays exponentially at a speed determined by the mean-reverting parameter.

The mean-reverting parameter is not a particularly intuitive way of describing how long it takes a factor to revert to its long-term goal. A more intuitive quantity is the factor's *half-life*, defined as the time it takes the factor to progress half the distance toward its goal. In the example of this section, the half-life of the interest rate, τ , is given by the following equation:

$$(15.339\% - 5.121\%)e^{-.025\tau} = \frac{1}{2}(15.339\% - 5.121\%) \quad (9.24)$$

Solving,

$$\begin{aligned} e^{-.025\tau} &= \frac{1}{2} \\ \tau &= \frac{\ln(2)}{.025} \\ \tau &= 27.73 \end{aligned} \quad (9.25)$$

where $\ln(\cdot)$ is the natural logarithm function. In words, the interest rate factor takes 27.73 years to cover half the distance between its starting value and its goal. This can be seen visually in Figure 9.5 where the expected rate 30 years from now is about halfway between its current value and θ . Larger mean-reverting parameters produce shorter half lives.

Figure 9.5 also shows one-standard deviation intervals around expectations both for the mean-reverting process of this section and for a process with the same expectation and the same σ but without mean reversion ("No MR"). The standard deviation of the terminal distribution of the short rate after T years in the Vasicek model is

$$\sqrt{\frac{\sigma^2}{2k}(1 - e^{-2kT})} \quad (9.26)$$

In the numerical example, with a mean-reverting parameter of .025 and a volatility of 126 basis points, the short rate in 10 years is normally distributed with an expected value of 7.3812%, derived earlier, and a standard deviation of

$$\sqrt{\frac{.0126^2}{2 \times .025}(1 - e^{-2 \times .025 \times 10})} \quad (9.27)$$

or 353 basis points. Using the same expected value and σ but no mean reversion the standard deviation is $\sigma\sqrt{T} = 1.26\%\sqrt{10}$ or 398 basis points.




FIGURE 9.6 Rate Curves from the Vasicek Model and Selected Market Swap Rates, February 16, 2001

Pulling the interest rate toward a long-term goal dampens volatility relative to processes without mean reversion, particularly at long horizons.

To avoid confusion in terminology, note that the mean-reverting model in this section sets volatility equal to 125 basis points “per year.” Because of mean reversion, however, this does not mean that the standard deviation of the terminal distribution after T years increases with the square root of time. Without mean reversion this is the case, as mentioned in the previous paragraph. With mean reversion, the standard deviation increases with horizon more slowly than that, producing a standard deviation of only 353 basis points after 10 years.

Figure 9.6 graphs the rate curves in this parameterization of the Vasicek model. The values of r_0 and θ were calibrated to match the 2- and 10-year par rates in the market. As a result, Figure 9.6 qualitatively resembles Figure 9.4. The mean reversion parameter might have been used to make the model fit the observed term structure more closely, but, as discussed in the next paragraph, this parameter was used to produce a particular term structure of volatility. In conclusion, Figure 9.6 shows that the model as calibrated in this section is probably not flexible enough to produce the range of term structures observed in practice.

A model with mean reversion and a model without mean reversion result in dramatically different term structures of volatility. Figure 9.7 shows that the volatilities of par rates decline with term in the Vasicek model. In this example the mean reversion and volatility parameters are chosen to fit the implied 10- and 30-year volatilities. As a result, the model matches the market at those two terms but overstates the volatility for shorter terms.




FIGURE 9.7 Par Rate Volatility from the Vasicek Model and Selected Implied Volatilities, February 16, 2001

While Figure 9.7 certainly shows an improvement relative to the flat term structure of volatility shown in Figure 9.3, mean reversion in this model generates a term structure of volatility that slopes downward everywhere. Chapter 11 shows that a second factor can produce the humped volatility structure evident in the market.

Since mean reversion lowers the volatility of longer-term par rates, it must also lower the impact of convexity on these rates. Table 9.2 reports the convexity effect at several terms. Recall that the convexity effects listed in Table 9.1 are generated from a model with no mean reversion and a volatility of 113 basis points per year. Since this section sets volatility equal to 126 basis points per year and since mean reversion is relatively slow, the convexity effects for terms up to 10 years are slightly larger in Table 9.2 than in Table 9.1. But by a term of 30 years the dampening effect of mean reversion on volatility manifests itself, and the convexity effect in the Vasicek

TABLE 9.2 Convexity Effects on Par Rates in a Parameterization of the Vasicek Model

Term (years)	Convexity (bps)
2	−1.0
5	−5.8
10	−19.1
30	−74.7




FIGURE 9.8 Sensitivity of Spot Rates in the Vasicek Model to a 10-Basis-Point Change in the Factor

model of about 75 basis points is substantially below the 135 basis point in the model without mean reversion.

Figure 9.8 shows the shape of the interest rate factor in a mean-reverting model, that is, how the spot rate curve is affected by a 10-basis point increase in the short-term rate. By definition, short-term rates rise by about 10 basis points but longer term rates are impacted less. The 30-year spot rate, for example, rises by only 7 basis points. Hence a model with mean reversion is not a parallel shift model.

The implications of mean reversion for the term structure of volatility and factor shape may be better understood by reinterpreting the assumption that short rates tend toward a long-term goal. Assuming that short rates move as a result of some news or shock to the economic system, mean reversion implies that the effect of this shock eventually dissipates. After all, regardless of the shock, the short rate is assumed to arrive ultimately at the same long-term goal.

Economic news is said to be *long-lived* if it changes the market's view of the economy many years in the future. For example, news of a technological innovation that raises productivity would be a relatively long-lived shock to the system. Economic news is said to be *short-lived* if it changes the market's view of the economy in the near but not far future. An example of this kind of shock might be news that retail sales were lower than expected due to excessively cold weather over the holiday season. In this interpretation, mean reversion measures the length of economic news in a term structure model. A very low mean reversion parameter, i.e., a very long half-life, implies that news is long-lived and that it will affect the short rate for many years to

come. On the other hand, a very high mean reversion parameter, i.e., a very short half-life, implies that news is short-lived and that it affects the short rate for a relatively short period of time. In reality, of course, some news is short-lived while other news is long-lived, a feature captured by the multi-factor Gauss+ model presented in Chapter 11.

Interpreting mean reversion as the length of economic news explains the factor structure and the downward-sloping term structure of volatility in the Vasicek model. Rates of every term are combinations of current economic conditions, as measured by the short-term rate, and of long-term economic conditions, as measured by the long-term value of the short rate (i.e., θ). In a model with no mean reversion, rates are determined exclusively by current economic conditions. Shocks to the short-term rate affect all rates equally, giving rise to parallel shifts and a flat term structure of volatility. In a model with mean reversion, short-term rates are determined mostly by current economic conditions while longer-term rates are determined mostly by long-term economic conditions. As a result, shocks to the short rate affect short-term rates more than longer-term rates and give rise to a downward-sloping term structure of volatility and a downward-sloping factor structure.

The Art of Term Structure Models: Volatility and Distribution

This chapter continues the presentation of the elements of term structure modeling, focusing on the volatility of interest rates and on models in which rates are not normally distributed.

TIME-DEPENDENT VOLATILITY: MODEL 3

Just as a time-dependent drift may be used to fit many bond or swap rates, a time-dependent volatility function may be used to fit many option prices. A particularly simple model with a time-dependent volatility function might be written as follows:

$$dr = \lambda(t) dt + \sigma(t) dw \quad (10.1)$$

Unlike the models presented in Chapter 9, the volatility of the short rate in equation (10.1) depends on time. If, for example, the function $\sigma(t)$ were such that $\sigma(1) = 1.26\%$ and $\sigma(2) = 1.20\%$, then the volatility of the short rate in one year is 126 basis points per year while the volatility of the short rate in two years is 120 basis points per year.

To illustrate the features of time-dependent volatility, consider the following special case of (10.1) that will be called Model 3:

$$dr = \lambda(t) dt + \sigma e^{-\alpha t} dw \quad (10.2)$$

In (10.2) the volatility of the short rate starts at the constant σ and then exponentially declines to zero. Volatility could have easily been designed to decline to another constant instead of zero, but Model 3 serves its pedagogical purpose well enough.




FIGURE 10.1 Standard Deviation of Terminal Distributions of Short Rates in Model 3

Setting $\sigma = 126$ basis points and $\alpha = .025$, Figure 10.1 graphs the standard deviation of the terminal distribution of the short rate at various horizons.¹ Note that the standard deviation rises rapidly with horizon at first but then rises more slowly. The particular shape of the curve depends, of course, on the volatility function chosen for (10.2), but very many shapes are possible with the more general volatility specification in (10.1).

Deterministic volatility functions are popular, particularly among market makers in interest rate options. Consider the example of *caplets*. At expiration, a caplet pays the difference between the short rate and a strike, if positive, on some notional amount. Furthermore, the value of a caplet depends on the distribution of the short rate at the caplet's expiration. Therefore, the flexibility of the deterministic functions $\lambda(t)$ and $\sigma(t)$ may be used to match the market prices of caplets expiring on many different dates.²

The behavior of standard deviation as a function of horizon in Figure 10.1 resembles the impact of mean reversion on horizon standard deviation in Figure 9.5. In fact, setting the initial volatility and decay rate in Model 3 equal to the volatility and mean reversion rate of the numerical example of the Vasicek model, the standard deviations of the terminal distributions from the two models turn out to be identical. Furthermore, if

¹This result is presented without derivation.

²For a fuller discussion of caplets see Chapter 18.

the time-dependent drift in Model 3 matches the average path of rates in the numerical example of the Vasicek model, then the two models produce exactly the same terminal distributions.

While these parameterizations of the two models give equivalent terminal distributions, the models remain very different in other ways. As is the case for any model without mean reversion, Model 3 is a parallel shift model. Also, the term structure of volatility in Model 3 is flat. Since the volatility in Model 3 changes over time, the term structure of volatility is flat at levels that change over time, but it is still always flat.

The arguments for and against using time-dependent volatility resemble those for and against using a time-dependent drift. If the purpose of the model is to quote fixed income options prices that are not easily observable, then a model with time-dependent volatility provides a means of interpolating from known to unknown option prices. If, however, the purpose of the model is to value and hedge fixed income securities, including options, then a model with mean reversion might be preferred for two reasons.

First, while mean reversion is based on the economic intuitions outlined earlier, time-dependent volatility relies on the difficult argument that the market has a forecast of short-term volatility in the distant future. A modification of the model that addresses this objection, by the way, is to assume that volatility depends on time in the near future and then settles at a constant.

Second, the downward-sloping factor structure and term structure of volatility in mean-reverting models capture the behavior of interest rate movements better than parallel shifts and a flat term structure of volatility. (Recall the empirical PCA results in Chapter 6). It may very well be that the Vasicek model does not capture the behavior of interest rates sufficiently well to be used for a particular valuation or hedging purpose. But in that case it is unlikely that a parallel shift model calibrated to match caplet prices will be better suited for that purpose.

THE COX-INGERSOLL-ROSS AND LOGNORMAL MODELS: VOLATILITY AS A FUNCTION OF THE SHORT RATE

The models presented so far assume that the basis-point volatility of the short rate is independent of the level of the short rate. This is almost certainly not true at extreme levels of the short rate. Periods of high inflation and high short-term interest rates are inherently unstable and, as a result, the basis-point volatility of the short rate tends to be high. Also, when the short-term rate is very low, its basis-point volatility is limited by the fact that interest rates cannot decline much below zero.

Economic arguments of this sort have led to specifying the basis-point volatility of the short rate as an increasing function of the short rate. The risk-neutral dynamics of the Cox-Ingersoll-Ross (CIR) model are

$$dr = k(\theta - r)dt + \sigma\sqrt{r}dw \quad (10.3)$$

Since the first term on the right-hand side of (10.3) is not a random variable and since the standard deviation of dw equals $\sqrt{dt}$ by definition, the annualized standard deviation of dr (i.e., the basis-point volatility) is proportional to the square root of the rate. Put another way, in the CIR model the parameter σ is constant, but basis-point volatility is not: annualized basis-point volatility equals $\sigma\sqrt{r}$ and increases with the level of the short rate.

Another popular specification is that the basis-point volatility is proportional to rate. In this case the parameter σ is often called *yield volatility*. Two examples of this volatility specification are the Courtadon model,

$$dr = k(\theta - r)dt + \sigma r dw \quad (10.4)$$

and the simplest *lognormal model*, to be called Model 4, a variation of which will be discussed in the next section:

$$dr = ardt + \sigma r dw \quad (10.5)$$

In these two specifications, yield volatility is constant but basis-point volatility equals σr and increases with the level of the rate.

Figure 10.2 graphs the basis-point volatility as a function of rate for the cases of the constant, square root, and proportional specifications. For comparison purposes, σ is set in all three cases such that basis-point volatility equals 100 at a short rate of 8%. Mathematically,

$$\sigma^{bp} = .01 \quad (10.6)$$

$$\sigma^{CIR} \times \sqrt{8\%} = 1\% \implies \sigma^{CIR} = .0354 \quad (10.7)$$

$$\sigma^y \times 8\% = 1\% \implies \sigma^y = 12.5\% \quad (10.8)$$

Note that the units of these volatility measures are somewhat different. Basis-point volatility is in the units of an interest rate (e.g., 100 basis points), while yield volatility is expressed as a percentage of the short rate (e.g., 12.5%).

As shown in Figure 10.2, the CIR and proportional volatility specifications have basis-point volatility increasing with rate but at different speeds. Both models have the basis-point volatility equal to zero at a rate of zero.




FIGURE 10.2 Three Volatility Specifications

The property that basis-point volatility equals zero when the short rate is zero, combined with the condition that the drift is positive when the rate is zero, guarantees that the short rate cannot become negative. In some respects this is an improvement over models with constant basis-point volatility that allow interest rates to become negative. It should be noted again, however, that choosing a model depends on the purpose at hand. Consider a trader who believes the following. One, the assumption of constant volatility is best in the current economic environment. Two, the possibility of negative rates has a small impact on the pricing of the securities under consideration. And three, the computational simplicity of constant volatility models has great value. This trader might very well opt for a model that allows some probability of negative rates.

Figure 10.3 graphs terminal distributions of the short rate after 10 years under the CIR, normal, and lognormal volatility specifications. In order to emphasize the difference in the shape of the three distributions, the parameters have been chosen so that all of the distributions have an expected value of 5% and a standard deviation of 2.32%. The figure illustrates the advantage of the CIR and lognormal models with respect to not allowing negative rates. The figure also indicates that out-of-the-money option prices could differ significantly under the three models. Even if, as in this case, the mean and volatility of the three distributions are the same, the probability of outcomes away from the means are different enough to generate significantly different options prices. (See Chapter 18 for more on these issues.) More generally, the shape of the distribution used in an interest rate model is an important determinant of that model's performance.




FIGURE 10.3 Terminal Distributions of the Short Rate After Ten Years in CIR, Normal, and Lognormal Models

TREE FOR THE ORIGINAL SALOMON BROTHERS MODEL

This section shows how to construct a binomial tree to approximate the dynamics for a lognormal model with a deterministic drift, a model attributed here to researchers at Salomon Brothers in the '80s. The dynamics of the model are as follows:

$$dr = \tilde{a}(t) r dt + \sigma r dw \quad (10.9)$$

By Ito's Lemma, which is beyond the mathematical scope of this book,

$$d[\ln(r)] = \frac{dr}{r} - \frac{1}{2}\sigma^2 dt \quad (10.10)$$

Substituting (10.9) into (10.10),


$$d[\ln(r)] = \left[\tilde{a}(t) - \frac{1}{2}\sigma^2 \right] dt + \sigma dw \quad (10.11)$$

Redefining the notation of the time-dependent drift so that $a(t) = \tilde{a}(t) - \frac{1}{2}\sigma^2$, equation (10.11) becomes


$$d[\ln(r)] = a(t) dt + \sigma dw \quad (10.12)$$

Equation (10.12) says that the natural logarithm of the short rate is normally distributed. Furthermore, by definition, a random variable has a lognormal distribution if its natural logarithm has a normal distribution. Therefore, (10.12) implies that the short rate has a lognormal distribution.

Equation (10.12) may be described as the Ho-Lee model based on the natural logarithm of the short rate instead of on the short rate itself. Adapting the tree for the Ho-Lee model accordingly, the tree for the first three dates is



To express this tree in rate, as opposed to the natural logarithm of the rate, exponentiate each node:



This tree shows that the perturbations to the short rate in a lognormal model are multiplicative as opposed to the additive perturbations in normal models. This observation, in turn, reveals why the short rate in this model cannot become negative. Since e^x is positive for any value of x , so long as r_0 is positive every node of the lognormal tree results in a positive rate.

The tree also reveals why volatility in a lognormal model is expressed as a percentage of the rate. Recall the mathematical fact that, for small values of x , $e^x \approx 1 + x$. Setting $a_1 = 0$ and $dt = 1$, for example, the top node of date 1 may be approximated as

$$r_0 e^\sigma \approx r_0 (1 + \sigma) \quad (10.13)$$

Volatility is clearly a percentage of the rate in equation (10.13). If, for example, $\sigma = 12.5\%$, then the short rate in the up state is 12.5% above the initial short rate.

As in the Ho-Lee model, the constants that determine the drift (i.e., a_1 and a_2) may be used to match market bond prices.

THE BLACK-KARASINSKI MODEL: A LOGNORMAL MODEL WITH MEAN REVERSION

The final model to be presented in this chapter is a lognormal model with mean reversion called the Black-Karasinski model. The model allows volatility, mean reversion, and the central tendency of the short rate to depend on time, firmly placing the model in the arbitrage-free class. A user may, of course, use or remove as much time dependence as desired.

The dynamics of the model are written as

$$dr = k(t) (\ln \tilde{\theta}(t) - \ln r) dt + \sigma(t) r dw \quad (10.14)$$

or, equivalently,³ as

$$d[\ln r] = k(t) (\ln \theta(t) - \ln r) dt + \sigma(t) dw \quad (10.15)$$

In words, equation (10.15) says that the natural logarithm of the short rate is normally distributed. It reverts to $\ln \theta(t)$ at a speed of $k(t)$ with a volatility of $\sigma(t)$. Viewed another way, the natural logarithm of the short rate follows a time-dependent version of the Vasicek model.

As in the previous section, the corresponding tree may be written in terms of the rate or the natural logarithm of the rate. Choosing the former, the process over the first date is

$$\begin{array}{l} \nearrow \frac{1}{2} \quad r_0 e^{k(1)(\ln \theta(1) - \ln r_0)dt + \sigma(1)\sqrt{dt}} = r_1 e^{\sigma(1)\sqrt{dt}} \\ r_0 \\ \searrow \frac{1}{2} \quad r_0 e^{k(1)(\ln \theta(1) - \ln r_0)dt - \sigma(1)\sqrt{dt}} = r_1 e^{-\sigma(1)\sqrt{dt}} \end{array}$$

The variable r_1 is introduced for readability. The natural logarithms of the rates in the up and down states are

$$\ln r_1 + \sigma(1)\sqrt{dt} \quad (10.16)$$

³This derivation is similar to that of moving from equation (10.9) to equation (10.12).

and

$$\ln r_1 - \sigma(1) \sqrt{dt} \quad (10.17)$$

respectively. It follows that the step down from the up state requires a rate of

$$r_1 e^{\sigma(1)\sqrt{dt}} e^{k(2)[\ln \theta(2) - \{\ln r_1 + \sigma(1)\sqrt{dt}\}]dt - \sigma(2)\sqrt{dt}} \quad (10.18)$$

while the step up from the down state requires a rate of

$$r_1 e^{-\sigma(1)\sqrt{dt}} e^{k(2)[\ln \theta(2) - \{\ln r_1 - \sigma(1)\sqrt{dt}\}]dt + \sigma(2)\sqrt{dt}} \quad (10.19)$$

A little algebra shows that the tree recombines only if

$$k(2) = \frac{\sigma(1) - \sigma(2)}{\sigma(1) dt} \quad (10.20)$$

Imposing the restriction (10.20) would require that the mean reversion speed be completely determined by the time-dependent volatility function. But these elements of a term structure model serve two distinct purposes. As demonstrated in this chapter, mean reversion controls the term structure of volatility while time-dependent volatility controls the future volatility of the short-term rate (and the prices of options that expire at different times). To create a model flexible enough to control mean reversion and time-dependent volatility separately, the model has to construct a recombining tree without imposing (10.20). To do so it allows the length of the time step, dt , to change over time.

Rewriting equations (10.18) and (10.19) with the time steps labeled dt_1 and dt_2 gives the following values for the up-down and down-up rates:

$$r_1 e^{\sigma(1)\sqrt{dt_1}} e^{k(2)[\ln \theta(2) - \{\ln r_1 + \sigma(1)\sqrt{dt_1}\}]dt_2 - \sigma(2)\sqrt{dt_2}} \quad (10.21)$$

$$r_1 e^{-\sigma(1)\sqrt{dt_1}} e^{k(2)[\ln \theta(2) - \{\ln r_1 - \sigma(1)\sqrt{dt_1}\}]dt_2 + \sigma(2)\sqrt{dt_2}} \quad (10.22)$$

A little algebra now shows that the tree recombines if

$$k(2) = \frac{1}{dt_2} \left[1 - \frac{\sigma(2)\sqrt{dt_2}}{\sigma(1)\sqrt{dt_1}} \right] \quad (10.23)$$

The length of the first time step can be set arbitrarily. The length of the second time step is set to satisfy (10.23), allowing the user freedom in choosing the mean reversion and volatility functions independently.

APPENDIX: CLOSED-FORM SOLUTIONS FOR SPOT RATES

This appendix lists formulas for spot rates, without derivation, in various models mentioned in the text. These can be useful for some applications and also to gain intuition about applying term structure models. The spot rates of term T , $\hat{r}(T)$, are continuously compounded rates. The discount factors and forward rates can be derived by the formulas developed in Part One.

Model 1

$$\hat{r}(T) = r_0 - \frac{\sigma^2 T^2}{6} \quad (10.24)$$

Model 2

$$\hat{r}(T) = r_0 + \frac{\lambda T}{2} - \frac{\sigma^2 T^2}{6} \quad (10.25)$$

Vasicek

$$\begin{aligned} \hat{r}(T) = & \theta + \frac{1 - e^{-kT}}{kT} (r_0 - \theta) \\ & - \frac{\sigma^2}{2k^2} \left(1 + \frac{1 - e^{-2kT}}{2kT} - 2 \frac{1 - e^{-kT}}{kT} \right) \end{aligned} \quad (10.26)$$

Model 3 with $\lambda(t) = \lambda$

$$\hat{r}(T) = r_0 + \frac{\lambda T}{2} - \sigma^2 \frac{2\alpha^2 T^2 - 2\alpha T + 1 - e^{-2\alpha T}}{8\alpha^3 T} \quad (10.27)$$

Cox-Ingersoll-Ross

Let $P(T)$ be the price of a zero coupon bond maturing at time T (from which the spot rate can be easily calculated). Then,

$$P(T) = A(T) e^{-B(T)r_0} \quad (10.28)$$

where

$$A(T) = \left[\frac{2he^{(k+h)T/2}}{2h + (k+h)(e^{hT} - 1)} \right]^{2k\theta/\sigma^2} \quad (10.29)$$

$$B(T) = \frac{2(e^{hT} - 1)}{2h + (k+h)(e^{hT} - 1)} \quad (10.30)$$

$$h = \sqrt{k^2 + 2\sigma^2} \quad (10.31)$$

The Gauss+ and LIBOR Market Models

The previous chapters in Part Three are extremely useful for learning about term structure models and for basic applications. Most models used in practice, however, are much more complex. This chapter presents two models that can be and that are used in practice. The first, the Gauss+ model, is a multi-factor generalization of the short-rate models presented in previous chapters. The second, the LIBOR market model (LMM), is in a different family of models: its factors are market observable forward rates.

The mathematical sophistication required to understand and apply this chapter is higher than that required to understand most of the rest of the book. In exchange, however, the reader will understand the key concepts behind many state-of-the-art models, be able to implement the Gauss+ model and a simple version of the LMM model, and certainly be well prepared for further study in this field.

THE GAUSS+ MODEL

The Gauss+ model is well-known among practitioners for use in relative value trading and hedging. The assumptions of the model are intuitively appealing and lead to a reasonable balance between tractability and capturing the empirical complexity of term structure dynamics. The goal of this chapter is to enable a determined reader to implement the model and estimate its parameters. Sample results from the USD and EUR swap markets are provided as well.

The dynamics of the *cascade form* of the model are given in equations (11.1) through (11.4). The rates $r(t)$, $m(t)$, and $l(t)$ denote the short-term rate of interest, a medium-term factor, and a long-term factor,

respectively.

$$dr = -\alpha_r (r - m) dt \quad (11.1)$$

$$dm = -\alpha_m (m - l) dt + \sigma_m (\rho dw_1 + \sqrt{1 - \rho^2} dw_2) \quad (11.2)$$

$$dl = -\alpha_l (l - \theta) dt + \sigma_l dw_1 \quad (11.3)$$

$$E [dw_1 dw_2] = 0 \quad (11.4)$$

The long-term factor $l(t)$ is meant to reflect long-term trends in demographics, production technology, etc. It is assumed to mean revert at a speed α_l to some very long-term constant, θ , but to fluctuate around that trend with a volatility $\sigma_l dw_1$. Consistent with the notation of previous chapters, σ_l is a constant and dw_1 is a normally distributed random variable with mean 0 and standard deviation $\sqrt{dt}$. Recalling the introduction of mean reversion in the context of the Vasicek model (see Chapter 9), the parameter θ can be thought of as including both long-term expectations and a risk premium. Also, as in that model, the parameter σ_l is in units of annual basis-point volatility even though the annual volatility of the factor, due to mean reversion, is less than σ_l . Finally, the speed of mean reversion of the long-term factor to θ is expected to be relatively slow. In the USD sample results to follow, θ is 8%, the speed of mean reversion corresponds to a half-life of 67 years, and σ_l is 105 basis points.

The medium-term factor $m(t)$ mean reverts to the long-term factor and is meant to capture monetary or business cycles around the long-term economic trend. The mean-reversion and volatility parameters are α_m and σ_m , respectively. Given the interpretation of m relative to that of l , m reverts to its mean more quickly than l reverts to its mean, implying that $\alpha_m > \alpha_l$.

Like dw_1 , dw_2 is normally distributed with mean zero and standard deviation $\sqrt{dt}$. Furthermore, by (11.4), dw_1 and dw_2 are uncorrelated. The specification of the volatility of dm in equation (11.2) is a convenient way to set its volatility to $\sigma_m \sqrt{dt}$ and its correlation with dl to ρ . To see this, under the assumptions made, the standard deviation of dm is

$$\sqrt{\sigma_m^2 (\rho^2 dt + [1 - \rho^2] dt)} = \sigma_m \sqrt{dt} \quad (11.5)$$

Also, the covariance of dm and dl is

$$\text{Cov} \left[\sigma_m (\rho dw_1 + \sqrt{1 - \rho^2} dw_2), \sigma_l dw_1 \right] = \rho \sigma_m \sigma_l dt \quad (11.6)$$

so the correlation of dm and dl is

$$\frac{\rho\sigma_m\sigma_l dt}{\sigma_m\sqrt{dt} \times \sigma_l\sqrt{dt}} = \rho \quad (11.7)$$

Returning to the structure of the Gauss+ model, the short-term rate $r(t)$ is meant to reflect the activity of a central bank that pegs the short-term rate at its current level and adjusts it, relatively rapidly, toward a level appropriate for the state of the business-cycle. The specification (11.1) captures this in two ways. First, the short-term rate mean reverts to the model's representation of the state of the business cycle, namely, $m(t)$. Second, the instantaneous volatility of the short-term is zero since, as a matter of central bank policy, the short rate is pegged at some level. Of course, the short-term rate will exhibit volatility over time through its direct tracking of $m(t)$ and its indirect tracking of $l(t)$.

The lack of a volatility term in (11.1) is an important feature of the Gauss+ model. Chapter 9 pointed out that mean-reverting models generate a downward-sloping term structure of volatility. Largely because of the activity of central banks, however, empirical and implied term structures of volatility tend to have a hump, i.e., volatility is low for very short-term rates, increases to a peak at intermediate-term rates, and then declines.¹ The term structure of volatility in the Gauss+ model matches this observed market behavior. The lack of an independent volatility term in the dynamics of $r(t)$ keeps short-term rate volatility low. The volatility of $m(t)$ and $l(t)$ significantly impact the distribution of the short-term rate at longer horizons and, therefore, the volatility of longer-term rates. The mean reversion of both $m(t)$ and $l(t)$, however, eventually result in a downward-sloping term structure of volatility, just as in the Vasicek model. The humped-shape term structure of volatility in the Gauss+ model will be illustrated in the sample results to follow.

In passing, the Gauss+ model gets its name from the lack of a volatility term in (11.1). The “Gauss” part of the name indicates that interest rates have a normal or Gaussian distribution. But while most one-, two-, or three-factor normal models have one, two, or three sources of risk, respectively, the Gauss+ model, strictly speaking, has three *state variables* and two sources of risk. It has three state variables in the sense that the state of the world in the model is described by the levels of $r(t)$, $m(t)$, and $l(t)$ or, equivalently, that bond prices or spot rates depend on the values of all three of these rates. (This point is obvious from the expression for the spot rate in equation (11.16) below.) There are, however, only two sources of risk in the model: dw_1 and

¹See the empirical evidence in Chapter 6 and representative data on implied volatilities in Table 18.4.

dw_2 . The “+” in the name indicates the somewhat unusual presence of a factor that is not also a source of risk.

As a final comment on the structure of the cascade form, there is no constant or drift term in the dynamics of $r(t)$ or $m(t)$ that can represent an implicit risk premium in these risk-neutral processes as is the case, for example, in the mean reverting process of the Vasicek model in Chapter 9. The idea is that the horizons of investors are long enough so that no risk premium is earned on a central bank’s deciding to raise or lower rates somewhat later or sooner nor on other short-lived news. This assumption, that a risk premium is earned only on the long-term factor in a multi-factor model, has characterized industry practice and has been supported by recent empirical work.²

The cascade form of the model and its parameters are very useful for intuition about the workings of the model. To solve for prices and rates, however, it is easier to work with a *reduced form* of the model. Appendix A in this chapter shows that the parameters of the cascade form can be recovered from the parameters of the reduced form in equations (11.8) through (11.12). In other words, the parameters θ , α_r , α_m , α_l , σ_m , σ_l , and ρ of the cascade form can be recovered from the parameters θ , α_r , α_m , α_l , σ_{11} , σ_{12} , and σ_{21} of the reduced form. The first four parameters in each of the models are the same, and the relationships between σ_m , σ_l , and ρ of the cascade form and σ_{11} , σ_{12} , and σ_{21} of the reduced form are given in equations (11.13) through (11.15) below.

$$r = \theta + x_1 + x_2 + x_3 \quad (11.8)$$

$$dx_1 = -\alpha_r x_1 dt + \sigma_{11} dw_1 + \sigma_{12} dw_2 \quad (11.9)$$

$$dx_2 = -\alpha_m x_2 dt + \sigma_{21} dw_1 - \sigma_{12} dw_2 \quad (11.10)$$

$$dx_3 = -\alpha_l x_3 dt - (\sigma_{11} + \sigma_{21}) dw_1 \quad (11.11)$$

$$E(dw_1 dw_2) = 0 \quad (11.12)$$

$$\sigma_{11} = \frac{\alpha_r \alpha_m}{(\alpha_r - \alpha_m)(\alpha_r - \alpha_l)} \sigma_l - \frac{\alpha_r}{(\alpha_r - \alpha_m)} \rho \sigma_m \quad (11.13)$$

$$\sigma_{21} = \frac{\alpha_r}{(\alpha_r - \alpha_m)} \rho \sigma_m - \frac{\alpha_r \alpha_m}{(\alpha_r - \alpha_m)(\alpha_m - \alpha_l)} \sigma_l \quad (11.14)$$

$$\sigma_{12} = -\frac{\alpha_r}{(\alpha_r - \alpha_m)} \sigma_m \sqrt{1 - \rho^2} \quad (11.15)$$

²See, for example, J. Cochrane and M. Piazzesi, “Decomposing the Yield Curve,” Working paper, March 13, 2008.

SOLUTION AND ESTIMATION

The reduced form model in equations (11.8) through (11.12) can be solved for zero coupon bond prices or spot rates. Letting $\widehat{r}(T)$ be the T -year spot rate,

$$\widehat{r}(T) = \theta + \bar{x}_1 \frac{1 - e^{-\alpha_r T}}{\alpha_r T} + \bar{x}_2 \frac{1 - e^{-\alpha_m T}}{\alpha_m T} + \bar{x}_3 \frac{1 - e^{-\alpha_l T}}{\alpha_l T} + \Upsilon(T) \quad (11.16)$$

where $\bar{x}_1$, $\bar{x}_2$, and $\bar{x}_3$ are the initial values of the factors and $\Upsilon(T)$ is a function, given in Appendix B in this chapter, that depends on all the parameters of the reduced form of the model and on T but not on the initial values of the factors.

Given parameters of the reduced form of the Gauss+ model, equation (11.16) can be used to price any security with fixed cash flows, along the lines of Chapters 1 and 2 of this book. Furthermore, as mentioned later, some contingent claims can be priced directly in this framework as well, i.e., without having first to capture the dynamics of the Gauss+ model in a tree or some other numerical method. But how should the parameters of the model be set? One popular starting point is to choose parameters so that the two principal components (PCs) implied by the model—there are only two sources of risk in the Gauss+ model—approximate the first two PCs from interest rate data.

Appendix C in this chapter shows that there are two functions of maturity T and of the parameters α_r , α_m , α_l , σ_m , σ_l , and ρ of the cascade form, namely $\Psi_1(T)$ and $\Psi_2(T)$, such that

$$\frac{\partial \widehat{r}(T)}{\partial z_1} = \Psi_1(T) \quad (11.17)$$

$$\frac{\partial \widehat{r}(T)}{\partial z_2} = \Psi_2(T) \quad (11.18)$$

where z_1 and z_2 are uncorrelated normally distributed random variables with unit variance. This result implies that the functions $\Psi_1(T)$ and $\Psi_2(T)$ give the two PCs of spot rates in the Gauss+ model. A change in the factor z_1 changes the one-year spot rate by $\Psi_1(1)$, the two-year spot rate by $\Psi_1(2)$, the three-year spot rate by $\Psi_1(3)$, etc., while a change in the factor z_2 changes the one-year spot rate by $\Psi_2(1)$, etc. An estimation procedure for the parameters α_r , α_m , α_l , σ_m , σ_l , and ρ , therefore, is to find the parameters that minimize the distance, e.g., sum of squared errors, between the functions $\Psi_1(T)$ and $\Psi_2(T)$ and the first two empirical principal components. Sample results from this estimation procedure are given in the next section.

The initial values of the factors, $\bar{x}_1$, $\bar{x}_2$, and $\bar{x}_3$ are typically implied each day from market prices. Just as the initial value of the single factor in the one-factor models of Chapters 9 and 10 might be chosen to match the market 10-year par rate, the initial factors in the Gauss+ model might be chosen to match the market 2-year, 10-year, and 30-year par rates. Alternatively, the three initial values might be used to minimize the sum of squared errors between model and market prices across the curve. The daily fitting of initial factor values is illustrated in the next section as well.

With the parameters α_r , α_m , α_l , σ_m , σ_l , and ρ fixed from the PCs, and with the initial factor values set each day based on prevailing market prices, the parameter θ is still free. It can be chosen based on *a priori* views about the long-term rate of interest and the risk premium or it can be used so that model and market prices over a particular sample period match as closely as possible. Some practitioners change θ every day so as to be able to fit an additional market price each day or so as to improve the match between model and market prices each day, even though this is an internally inconsistent use of the model since, in the model, θ is a constant.

USD AND EUR SAMPLE RESULTS

The mean reversion, volatility, and correlation parameters of the Gauss+ model were estimated, along the lines of the previous section, from principal components of USD and EUR swap rates over the period August 1, 2001, to December 9, 2010.³ The parameter θ was chosen to improve the overall in-sample fit of market to model rates, one of the choices mentioned in the previous section. Table 11.1 reports the estimated parameters of the cascade form of the model. The half-lives corresponding to the long-term factors are, as expected, considerably longer than those corresponding to the short-term rates and the intermediate-term factors. The difference in half-life between the latter is less impressive. In fact, it is often difficult in the Gauss+ model to pin down these coefficients very accurately. The orders of magnitudes of the volatility parameters are quite reasonable, ranging from 105 to 258 basis points. These volatilities, dampened by mean reversion, generate the term structure of volatility, discussed below. For the present, however, note that the relatively high volatilities of the medium-term factors are offset by relatively high mean reversion parameters.

Figures 11.1 and 11.2 show the USD and EUR first and second PCs, respectively, both estimated from the data and from the fitted model. The Gauss+ model clearly has enough flexibility to match these shapes.

³The empirical PCs were estimated with an exponential decay rate of .5. This means that data from one year ago are weighted by about half as much as current data.

TABLE 11.1 Parameter Estimates of the Gauss+ Model for USD and EUR Swaps from August 1, 2001, to December 9, 2010

Parameter	USD		EUR	
	Value	Mean Reversion Half-Life (years)	Value	Mean Reversion Half-Life (years)
α_r	.729	1.0	.719	1.0
α_m	.629	1.1	.619	1.1
α_l	.038	18.4	.01	66.7
σ_m	2.58%		1.78%	
σ_l	1.79%		1.05%	
ρ	.15		-.05	
θ	9.04%		8.00%	

To the extent that the shapes of the empirical PCs are expected to persist, this is a powerful argument in support of this term structure model and the associated approach to estimation.

Figure 11.3 shows the term structure of volatility of swap rates in the estimated Gauss+ models for both USD and EUR. The Gauss+ model is able to capture the feature that volatility is particularly low at short terms, increases with term, and then eventually decays, although in this estimation the term structure of volatility in EUR increases over the whole of the range. In any case, the shapes in Figure 11.3 are much more realistic than those produced by the one-factor models without the “+” described earlier;




FIGURE 11.1 Historical PCs from the USD Swap Curve from August 1, 2001, to December 9, 2010, along with PCs from a Calibrated Gauss+ Model




FIGURE 11.2 Historical PCs from the EUR Swap Curve from August 1, 2001, to December 9, 2010, along with PCs from a Calibrated Gauss+ Model

recall, for example, the monotonically downward-sloping term structure of volatility in Figure 9.7 produced by the Vasicek model. Despite this success, however, for some applications, like market making in derivatives, much more detailed specifications of volatility are required. This and a related extension of the model is discussed in the next section.

The final graph of this section, Figure 11.4, illustrates the fitting of the initial factors on a particular day, namely, February 11, 2011. The graph




FIGURE 11.3 Term Structure of Swap Rate Volatility from Gauss+ Models of USD and EUR Swap Rates Estimated from August 1, 2001, to December 9, 2010




FIGURE 11.4 Fitting to the Term Structure of USD Spot Rates with the Gauss+ Model as of February 11, 2011

shows three USD spot rate curves, corresponding to the left axis, and two error curves, corresponding to the right axis. The three rate curves depict spot rates from the market; from the model with initial factors chosen to fit the 2-, 10-, and 30-year benchmark rates exactly; and from the model with initial factors chosen by least squares to fit the entire curve reasonably well. The two error curves show the difference, in basis points, between market rates and the curves from the two fitting approaches. Both model curves are relatively close to the market, with almost all of the absolute errors less than 20 basis points, but neither approach dominates the other. The benchmark curve matches the most liquid parts of the curve exactly,⁴ but the least-squares approach gives the better overall fit. In any case, fit either way, the model indicates that the 5-year sector is cheap (market rates above model rates) and the 20-year sector rich (market rates below model rates). Presenting this as a possible trading or investment opportunity is, of course, one of the reasons for building term structure models.

MODEL EXTENSIONS

Chapters 9 and 10 made the point that, in certain applications, it is desirable for the model to match the entire initial market term structure of rates. This can be done in the Gauss+ model in the same way as in the Ho-Lee and the

⁴Trading desks typically fit benchmark par swap rates, which are by far the most traded, rather than spot rates as in this illustration.

original Salomon Brothers models, i.e., by making the drift time dependent. Mathematically, equation (11.8) of the reduced form becomes

$$r^* = \theta(t) + x_1 + x_2 + x_3 \quad (11.19)$$

Appendix D in this chapter shows that, given an initial estimate of the Gauss+ model with the constant θ set to zero, finding the function $\theta(t)$ that fits the initial term structure is very simple: set $\theta(t)$ equal to the difference between the continuously compounded forward rates in the market and in the initial model. Put another way, the forward rates in the initial model plus $\theta(t)$ equal the forward rates in the model defined by (11.19) which, by construction, match the forward rates in the market.

Another extension of the model is to match implied market volatilities in order to capture convexity properly when interpolating between observed market rates or to price interest rate derivatives. The estimation methodology presented in the previous section is designed to capture historical dynamics, including volatilities and correlations, but this in no way means that implied volatilities of options, e.g., swaptions, will be priced correctly by the model.⁵ The Gauss+ model can be modified to fit traded volatilities in the same way as Model 3 or the Black-Karasinski model of Chapter 10, namely, by changing volatility constants to be functions of time. A simple and popular way to do this for the Gauss+ model is to multiply each of the volatility parameters in (11.9) through (11.11) by a function of time, denoted here as $\eta(t)$. The special case of $\eta(t) = 1$ for all t is the original Gauss+ model.

One use of $\eta(t)$ is to match swaption prices with a fixed underlying swap maturity at various expiration dates. A practitioner might decide, for example, to match the market prices of options on a 10-year swap at the well-traded expirations of three months, six months, one year, three years, five years, and 10 years. In that case $\eta(t)$ would be a step function, with a value between 0 and .25 set to match the three-month option price, a value between .25 and .5 that, together with the value of $\eta(t)$ from 0 to .25, matches the six-month option, etc. It should be emphasized once again that the desirability and usefulness of matching market quantities depends on the application at hand. Such calibration is certainly appropriate, for example, when pricing and making markets in exotic derivatives or when making convexity corrections to Eurodollar futures rates (see Chapter 15). By contrast, for relative value trading in swaps themselves, precise calibration to many volatility products is normally not worth the effort.

Table 11.2 illustrates a function $\eta(t)$ fitted on February 11, 2011, and appended to the Gauss+ model estimated in the previous section. Swaption

⁵A swaption is the option to receive or pay fixed in a swap at a preset rate at some time in the future. See Chapter 18.

TABLE 11.2 A Time-Dependent Volatility Factor in the Gauss+ Model Fitted to Swaption Prices on 10-Year Swaps at Various Expiries as of February 11, 2011

Option Expiry (years)	Implied Volatility (bps per day)	$\eta(t)$
.25	6.9	.87
.50	7.0	.86
1	7.1	.82
2	7.1	.75
3	7.0	.68
4	6.9	.63
5	6.7	.61

prices are for at-the-money options and are given in terms of implied volatility in basis points per day. The relatively low values of $\eta(t)$ indicate that the implied volatilities of swaptions are low relative to the historical volatility built into the initial calibration of the Gauss+ model. Or, taking the time-dependent model literally, swaption prices indicate that the volatility of the 10-year swap rate will fall in the future.

Figure 11.5 shows the term structure of swap rate volatility at various points in the future in the time-dependent USD model. The fall in volatility from one curve to the next is not so dramatic as the decline of $\eta(t)$ itself since the volatility of a term rate is, qualitatively, an average of volatilities over the relevant time horizon.

Strictly speaking, pricing European swaptions in the Gauss+ model requires some numerical implementation of the model, in sharp contrast with the closed-form solutions available for discount factors, spot rates, and forward rates. In practice, however, there exist good closed-form approximations that make use of Black-Scholes style option pricing models.⁶

⁶Using equation (13.11), write the forward swap rate from t to T in terms of discount factors and the annuity factor as $(d(t) - d(t+T))/A(t, t+T)$. Since discount factors are exponential functions of the factors in the Gauss+ model, this expression links swap rates to the underlying factors. Take derivatives, then, with respect to the factors to approximate the change in the swap rate as its partial derivatives with respect to the factors times the change in the factors. As it turns out, these partial derivatives are typically close to constant so that the change in the swap rate can be approximated as a linear function of the changes in the factors. Since the factors are normally distributed with a known variance-covariance matrix, the swap rate is approximately normal and its term volatility can be easily computed. The techniques of Chapter 18 can then be used to price swaptions in a Black-Scholes context.




FIGURE 11.5 Term Structures of USD Swap Rate Volatility Over Time with a Time-Dependent Volatility Factor in the Gauss+ Model Calibrated to Swaption Prices as of February 11, 2011

THE LIBOR MARKET MODEL

The LIBOR Market Model⁷ (LMM) differs from the models of Chapters 9 and 10 and from the Gauss+ model of this chapter. Instead of positing the form of one or several somewhat abstract factors, the factors of LMM can be understood as observable forward rates, which together describe the entire term structure. To take a simple example, the factors of an implementation of LMM might be 30 one-year forward rates, starting with the forward rate from today to the end of the year and ending with the forward rate from year 29 to year 30. Also, in another contrast with short-rate models, LMM does not require the specification of drift functions; its key inputs turn out to be the volatilities and correlations of the chosen forward rates.

The LMM framework is particularly appealing when the objective is to price and hedge derivatives, particularly “exotics,” given the prices of other, more liquidly traded derivatives. First, since the initial values of the factors are the forward rates themselves, LMM automatically ensures that the initial model and market term structures are the same. Second, since there are many factors, the model is flexible enough to capture volatilities and correlations of rates across the term structure. Third, since the volatilities of forward rates are relatively easily connected to the prices of the most liquid derivatives, the calibration of LMM to derivatives prices is relatively straightforward.

⁷A. Brace, D. Gatarek and M. Musiela, “The Market Model of Interest Rate Dynamics,” *Mathematical Finance* 1997, 7(2) pp. 127–154.

Presentations of the LMM model tend to be highly mathematical. The discussion here aims to convey the essential ideas with a minimum amount of mathematics. As a result, the interested reader should consult other sources for more general and theoretically rigorous presentations.⁸

Notation and Introduction

To introduce LMM, consider the simple context of three factors that are adjacent τ -year forward rates. Most commonly, $\tau = .25$ for three-month forwards, although τ might also be $\frac{1}{2}$ or $\frac{1}{12}$ for semiannual or monthly forwards. In any case, the simple setup here requires only four fixed points in time: T_0 ; $T_1 = T_0 + \tau$; $T_2 = T_1 + \tau$; and $T_3 = T_2 + \tau$. Then, denote the forward rates and discount factors as follows:

- $f_t(T_0, T_1) = f_t(T_0, T_0 + \tau)$: the forward rate from T_0 to $T_1 = T_0 + \tau$ as of $t \leq T_0$
- $f_t(T_1, T_2) = f_t(T_1, T_1 + \tau)$: the forward rate from T_1 to $T_2 = T_1 + \tau$ as of $t \leq T_1$
- $f_t(T_2, T_3) = f_t(T_2, T_2 + \tau)$: the forward rate from T_2 to $T_3 = T_2 + \tau$ as of $t \leq T_2$
- $d_t(T)$: the discount factor to time T as of $t \leq T$

A contingent claim makes payments on any or all of the dates T_0 , T_1 , T_2 , and T_3 , where each payment can depend on any or all of the forward rates at the time the payment is made. Since the term structure of interest rates to a particular maturity is defined by the set of adjacent forward rates to that maturity, this description of a claim is very general. A simple example would be a derivative on a bond or swap, since bonds and swaps depend on all forward rates to their maturity dates. A more exotic example, a special case of which will be used as an illustration in a later subsection, would be a derivative making a payment that is a function of current and past forward rates. Note that an exotic derivative's payments can easily depend on past values of forward rates. For instance, for $t \geq T_1$ the forward rate $f_t(T_1, T_2)$ is fixed at $f_{T_1}(T_1, T_2)$, but the payment of a contingent claim at times T_1 , T_2 , or T_3 can depend on the value of that realized forward rate.

As was the case for the short-rate models, the pricing of contingent claims in the LMM framework begins with assumptions about the probability distributions of the factors. For present purposes, to present ideas

⁸See L. Andersen and V. Piterbarg, *Interest Rate Modeling, Volume II*, Atlantic Financial Press, 2010; D. Brigo and F. Mercurio, *Interest Rate Models: Theory and Practice*, Springer, 2001; R. Rebonato, *Modern Pricing of Interest Rate Derivatives: The LMM and Beyond*, Princeton University Press, 2002.

simply and without the complexities of continuous time notation, assume that local changes in the three forward rates, i.e., changes over short time intervals, follow a multivariate normal distribution.⁹ Notationally, let the local change of the forward rate $f_t(T_{i-1}, T_i)$ be normally distributed with mean $\mu_i(t)$ and volatility $\sigma_i(t)$ and let the local correlation of the forward rates $f_t(T_{i-1}, T_i)$ and $f_t(T_{j-1}, T_j)$ be denoted $\rho_{i,j}(t)$. Borrowing terminology from continuous time models, these quantities will be referred to here as the instantaneous drifts, volatilities, and correlations of the forward rates. In any case, the next step in building the short-rate models was to transform the drift of the factors so as to obtain the risk-neutral process for the short rate, which, in turn, was used to price contingent claims. LMM also transforms the drifts of the factors so that contingent claims can be priced, but the transformations are different.

Developing the appropriate transformations in LMM requires the following result, which is proved in Chapter 18. There exist probabilities, called the “ $T + \tau$ forward measure,” such that two conditions hold. First, with $E_t^{T+\tau}[\cdot]$ denoting the time- t expectation under these probabilities and with s denoting some time after t ,

$$f_t(T, T + \tau) = E_t^{T+\tau} [f_s(T, T + \tau)] \quad (11.20)$$

In words, the forward rate from T to $T + \tau$ is a *martingale* in that its time- t value equals the expectation of its future realization at time s .

Second, let V_t denote the arbitrage-free time- t value of a contingent claim that makes a payout at time $T + \tau$. Then, V_t and V_s are related through the expectation $E_t^{T+\tau}[\cdot]$ as follows:

$$\frac{V_t}{d_t(T + \tau)} = E_t^{T+\tau} \left[\frac{V_s}{d_s(T + \tau)} \right] \quad (11.21)$$

And, in the special case $s = T + \tau$, equation (11.21) reduces to

$$V_t = d_t(T + \tau) \times E_t^{T+\tau} [V_{T+\tau}] \quad (11.22)$$

so that the value of the contingent claim at time $T + \tau$ equals its expected payout discounted back to time t .

The probabilities implicit in the $T + \tau$ forward measure are analogous to the risk-neutral probabilities used in the context of short-rate models. Equation (11.22) says that the $T + \tau$ forward measure probabilities can be

⁹The probability distribution of a set of random variables is multivariate normal if each random variable, given the values of all the others, is normally distributed.

used to price contingent claims by taking expectations under those probabilities and then discounting. Risk-neutral probabilities are used to price contingent claims by taking expectations of discounted payoffs under those probabilities. A more formal connection between the $T + \tau$ forward measure and the risk-neutral measure will be made in Chapter 18.

Pricing Contingent Claims and the Choice of Measure

The exposition of LMM to this point can be summarized as follows. First, changes in the three forward rate factors follow a multivariate normal distribution, but the drifts or means of the factors to be used for pricing are as yet to be determined. Second, for a particular $T + \tau$ forward measure, the drift of the forward rate $f_t(T, T + \tau)$ is zero and contingent claims can be priced using equation (11.21).

Consistent with typical applications of LMM, fix the forward measure to match the last possible payment date, which is called the *terminal measure* and, in this case, is the T_3 forward measure. Under this measure contingent claims can be priced and $f_t(T_2, T_3)$ has no drift. Therefore, for contingent claims such that the expectation in (11.21) can be computed knowing only $f_t(T_2, T_3)$, the model is complete. In the more general and usual case, however, when the valuation of the contingent claim depends on all of the forward rates, (11.21) cannot be applied without knowing the drift of these other forward rates under the T_3 measure. Note, of course, that if another measure, like the T_2 measure, had been fixed instead, then (11.21) could be applied to price claims that require knowing only $f_t(T_1, T_2)$, which, under the T_2 measure, has no drift. Claims with values that depended on the dynamics of $f_t(T_2, T_3)$ and $f_t(T_0, T_1)$, however, could not be priced without knowing their drifts under the T_2 measure.

Returning to the T_3 forward measure, the rest of this subsection develops the methodology for pricing contingent claims. First, the text expands on the pricing of claims that can be computed knowing only $f_t(T_2, T_3)$. Second, the drift of $f_t(T_1, T_2)$ under the T_3 forward measure is derived. Third, the drift of $f_t(T_0, T_1)$ under the T_3 forward measure is derived. Together, these results completely describe the pricing of contingent claims in this three-factor version of LMM. Fourth, for use beyond the exposition in this chapter, general expressions for forward rates drifts under various forward measures are presented.

Pricing Claims Requiring the Dynamics of $f_t(T_2, T_3)$ Alone Say that a derivative makes a single payment at time T_3 that is a function of the forward rate $f_{T_2}(T_2, T_3)$. Pricing this contingent claim is straightforward given the assumptions made so far. According to equation (11.20), the probabilities of the T_3 forward measure are such that $f_t(T_2, T_3)$ has no drift.

Also, according to (11.22), the value of a contingent claim on $f_{T_2}(T_2, T_3)$ making a payment at time T_3 is equal to the discounted value of the expected value of that payoff under the T_3 forward measure. Hence, to price the claim, invoke the assumption made earlier that local changes in $f_t(T_2, T_3)$ are normally distributed with volatility $\sigma_3(t)$, impose a mean of zero as the local transformation to the T_3 forward measure, and then calculate the discounted expected value of the payment.

While closed-form solutions and very good approximations exist for the prices of some contingent claims in the LMM framework, valuation usually relies on *Monte Carlo simulation*.¹⁰ To simulate one *path* of the forward rate $f_t(T_2, T_3)$, start at time 0 with the initial forward rate $f_0(T_2, T_3)$. Then, for a step of size Δt , set the change in the forward rate, $\Delta f_0(T_2, T_3)$, to the draw of a normal random variable with mean zero and standard deviation $\sigma_3(0)\sqrt{\Delta t}$. Add this realized change to the initial forward rate to get the forward rate at time Δt , i.e., $f_{\Delta t}(T_2, T_3)$. Then compute the next change in the forward rate as a normal random variable with mean zero and standard deviation $\sigma_3(\Delta t)\sqrt{\Delta t}$, etc. Continue in this manner to obtain eventually the value of $f_{T_2}(T_2, T_3)$ for this particular path from which the payment of the contingent claim at time T_3 can be calculated. Other paths can be simulated in the same way, with each generating an observation of the payment of the contingent claim. Finally, when it has been judged that sufficiently many realizations have been observed, average all the realized contingent-claim payments across paths as an estimate of $E_t^{T_3}[V_{T_3}]$ in (11.22) and discount to time t as an estimate of the time- t value of the contingent claim.

The approach just described for pricing a single payment at time T_3 that is contingent on the value of $f_{T_2}(T_2, T_3)$ also works for a payment made at time T_2 . To see this, note that by equation (11.21),

$$\begin{aligned}\frac{V_t}{d_t(T_3)} &= E_t^{T_3} \left[\frac{V_{T_2}}{d_{T_2}(T_3)} \right] \\ V_t &= d_t(T_3) \times \\ &\quad E_t^{T_3} [V_{T_2} (1 + \tau f_{T_2}(T_2, T_3))] \end{aligned} \quad (11.23)$$

where the second line of (11.23) follows from the definitions of discount factors and forward rates. To value this time- T_2 claim, therefore, simulate paths as before to obtain values for $f_{T_2}(T_2, T_3)$ at time T_2 . From these values compute the future value of the time- T_2 claim to time T_3 , i.e., the term $V_{T_2} (1 + \tau f_{T_2}(T_2, T_3))$, which is inside the expectation in (11.23). Finally, average these future values across paths and discount the result to time t .

¹⁰ For a book-length treatment, see Paul Glasserman, *Monte Carlo Methods in Financial Engineering*, 2003. Also, Chapter 20 discusses Monte Carlo simulation in more detail.

In contrast to claims on $f_{T_2}(T_2, T_3)$ paid at time T_2 or T_3 , a claim on $f_t(T_2, T_3)$ that pays at time T_1 or T_0 cannot be valued along the lines just described. Consider a claim on $f_{T_1}(T_2, T_3)$ paying at time T_1 . According to (11.21),

$$\frac{V_t}{d_t(T_3)} = E_t^{T_3} \left[\frac{V_{T_1}}{d_{T_1}(T_3)} \right]$$

$$V_t = d_t(T_3) E_t^{T_3} [V_{T_1} (1 + \tau f_{T_1}(T_1, T_2)) (1 + \tau f_{T_1}(T_2, T_3))] \quad (11.24)$$

where the second line of (11.24) again follows from the definitions of discount factors and forward rates. Since (11.24) depends on $f_{T_1}(T_1, T_2)$, this claim cannot be priced without knowing the drift of $f_{T_1}(T_1, T_2)$ under the T_3 measure. The text now turns to the calculation of that drift, which, of course, is also required to price any claim that depends on $f_t(T_1, T_2)$ and that is to be valued under the T_3 measure.

The Drift of $f_t(T_1, T_2)$ Under the T_3 Forward Measure Consider an agreement as of time t to pay $\bar{f}$ and to receive $f_s(T_1, T_2)$ at time T_2 for some $t \leq s \leq T_1$. The fair value of $\bar{f}$ can be quickly computed using the T_2 forward measure. Since the value of the agreement today is 0, it follows from (11.22) that

$$0 = d_t(T_2) \times E_t^{T_2} [f_s(T_1, T_2) - \bar{f}] \quad (11.25)$$

Invoking (11.20), the fact that $f_t(T_1, T_2)$ is a martingale under the T_2 measure, (11.25) can be solved for $\bar{f}$:

$$\begin{aligned} \bar{f} &= E_t^{T_2} [f_s(T_1, T_2)] \\ &= f_t(T_1, T_2) \end{aligned} \quad (11.26)$$

The drift of $f_t(T_1, T_2)$ under the T_3 measure will now be calculated by finding the fair value of $\bar{f}$ using the T_3 measure and equating the result with that of (11.26). Applying (11.21) here,

$$\begin{aligned} 0 &= E_t^{T_3} \left[\frac{f_s(T_1, T_2) - \bar{f}}{d_{T_2}(T_3)} \right] \\ &= E_t^{T_3} [(f_s(T_1, T_2) - \bar{f}) (1 + \tau f_{T_2}(T_2, T_3))] \\ \bar{f} E_t^{T_3} [(1 + \tau f_{T_2}(T_2, T_3))] &= E_t^{T_3} [f_s(T_1, T_2) (1 + \tau f_{T_2}(T_2, T_3))] \end{aligned} \quad (11.27)$$

Making use of a fact about the covariance of two random variables X_1 and X_2 , i.e., that $\text{Cov}(X_1, X_2) = E(X_1 X_2) - E(X_1)E(X_2)$, the right-hand side of (11.27) can be rewritten so that

$$\begin{aligned} \bar{f} E_t^{T_3} [(1 + \tau f_{T_2}(T_2, T_3))] &= \text{Cov}_t [f_s(T_1, T_2), (1 + \tau f_{T_2}(T_2, T_3))] \\ &\quad + E_t^{T_3} [f_s(T_1, T_2)] E_t^{T_3} [1 + \tau f_{T_2}(T_2, T_3)] \\ \bar{f} &= \frac{\tau \text{Cov}_t [f_s(T_1, T_2), f_{T_2}(T_2, T_3)]}{E_t^{T_3} [(1 + \tau f_{T_2}(T_2, T_3))]} \\ &\quad + E_t^{T_3} [f_s(T_1, T_2)] \end{aligned} \quad (11.28)$$

Since $f_s(T_1, T_2)$ does not change from s to T_2 , the covariance of $f_s(T_1, T_2)$ and $f_{T_2}(T_2, T_3)$ is equal to the covariance of $f_s(T_1, T_2)$ and $f_s(T_2, T_3)$. Also, using the value for $\bar{f}$ computed in (11.26) and the fact that $f_t(T_2, T_3)$ is a martingale under the T_3 measure, (11.28) becomes

$$E_t^{T_3} [f_s(T_1, T_2)] - f_t(T_1, T_2) = -\frac{\tau \text{Cov}_t [f_s(T_1, T_2), f_s(T_2, T_3)]}{1 + \tau f_t(T_2, T_3)} \quad (11.29)$$

To take the last step, use the (annualized) instantaneous volatility and correlation notation given earlier to write the numerator of the right-hand side of (11.29) as $\tau \rho_{2,3}(t) \sigma_2(t) \sigma_3(t) \times (s - t)$. Then, divide both sides of (11.29) by $s - t$ and let s approach t so that the left-hand side of (11.29) becomes the instantaneous drift of $f_t(T_1, T_2)$:

$$\mu_2(t) \equiv -\frac{\tau \rho_{2,3}(t) \sigma_2(t) \sigma_3(t)}{1 + \tau f_t(T_2, T_3)} \quad (11.30)$$

Equation (11.30) defines the drift of $f_t(T_1, T_2)$ under the T_3 measure, as desired. To review its use, consider applying (11.21) to a claim that pays some function of both $f_{T_1}(T_1, T_2)$ and $f_{T_1}(T_2, T_3)$ at time T_1 . Starting at the initial values of the forward rates, $f_t(T_1, T_2)$ and $f_t(T_2, T_3)$, simulate paths for the two rates under the T_3 measure from time t to time T_1 . More specifically, taking a small step from time t to time $t + \Delta t$, the changes in the two forward rates are found by drawing two random variables from a bivariate normal distribution with means $\mu_2(t) \Delta t$ and zero, volatilities $\sigma_2(t) \sqrt{\Delta t}$ and $\sigma_3(t) \sqrt{\Delta t}$, and correlation $\rho_{2,3}(t)$. Then, from the terminal values of $f_{T_1}(T_1, T_2)$ and $f_{T_1}(T_2, T_3)$ along each path compute the payoff of the claim at time T_1 divided by the discount factor $d_{T_1}(T_3)$, i.e., multiplied by $(1 + \tau f_{T_1}(T_1, T_2))(1 + \tau f_{T_1}(T_2, T_3))$. Finally, compute the value of the claim by averaging these results across paths and multiplying by $d_t(T_3)$.

To summarize the exposition to this point, the three forward rates were assumed to follow a multivariate normal distribution with given instantaneous volatilities and correlations. Drifts of the forward rates that can be used for pricing, i.e., those under the T_3 measure, are now known for $f_t(T_1, T_2)$ and $f_t(T_2, T_3)$. Once the drift for $f_t(T_0, T_1)$ under the T_3 measure is found, the model will be complete in that it can be used to price any contingent claim making payments that depend on any of the three forward rates.

The Drift of $f_t(T_0, T_1)$ Under the T_3 Forward Measure Consider an agreement as of time t to pay f^* and to receive $f_s(T_0, T_1)$ at time T_2 for some $t \leq s \leq T_0$. Since the value of the agreement today is 0, the fair value of f^* under the T_2 forward measure is, by (11.22),

$$\begin{aligned} 0 &= d_t(T_2) \times E_t^{T_2} [f_s(T_0, T_1) - f^*] \\ f^* &= E_t^{T_2} [f_s(T_0, T_1)] \end{aligned} \quad (11.31)$$

Since $f_t(T_0, T_1)$ is not a martingale under the T_2 measure, equation (11.31) cannot be simplified immediately, but will be expressed more usefully in a moment. For now, find the fair value of f^* under the T_3 measure and equate the result to the right-hand side of (11.31). Applying (11.21) to the claim now under consideration,

$$0 = E_t^{T_3} \left[\frac{f_s(T_0, T_1) - f^*}{d_{T_2}(T_3)} \right] \quad (11.32)$$

Proceeding just as in (11.27) through (11.29), but using f^* from (11.31), equation (11.32) becomes

$$E_t^{T_3} [f_s(T_0, T_1)] - E_t^{T_2} [f_s(T_0, T_1)] = - \frac{\tau \text{Cov} [f_s(T_0, T_1), f_s(T_2, T_3)]}{1 + \tau f_t(T_2, T_3)} \quad (11.33)$$

To express $E_t^{T_3} [f_s(T_0, T_1)]$ more simply, the text now returns to simplifying the right-hand side of (11.31), namely, $E_t^{T_2} [f_s(T_0, T_1)]$. Consider a claim that pays f^{**} to receive $f_s(T_0, T_1)$ at time T_1 . By (11.22), combined with the facts that the value of the claim today is zero and that $f_t(T_0, T_1)$ is a martingale under the T_1 measure,

$$\begin{aligned} 0 &= d_t(T_1) \times E_t^{T_1} [f_s(T_0, T_1) - f^{**}] \\ f^{**} &= f_t(T_0, T_1) \end{aligned} \quad (11.34)$$

Pricing the claim under the T_2 measure, however, using (11.21), gives that

$$0 = E_t^{T_2} \left[\frac{f_s(T_0, T_1) - f^{**}}{d_{T_1}(T_2)} \right] \quad (11.35)$$

Proceeding once again as in (11.27) through (11.29), but using f^{**} from (11.34), equation (11.35) becomes

$$E_t^{T_2} [f_s(T_0, T_1)] - f_t(T_0, T_1) = - \frac{\tau \text{Cov} [f_s(T_0, T_1), f_s(T_1, T_2)]}{1 + \tau f_t(T_1, T_2)} \quad (11.36)$$

Then, substituting $E_t^{T_2} [f_s(T_0, T_1)]$ from (11.36) into (11.33),

$$\begin{aligned} E_t^{T_3} [f_s(T_0, T_1)] - f_t(T_0, T_1) = & - \frac{\tau \text{Cov} [f_s(T_0, T_1), f_s(T_1, T_2)]}{1 + \tau f_t(T_1, T_2)} \\ & - \frac{\tau \text{Cov} [f_s(T_0, T_1), f_s(T_2, T_3)]}{1 + \tau f_t(T_2, T_3)} \end{aligned} \quad (11.37)$$

Finally, using the established notation and letting s approach t ,

$$\mu_1(t) = - \frac{\tau \rho_{1,2}(t) \sigma_1(t) \sigma_2(t)}{1 + \tau f_t(T_1, T_2)} - \frac{\tau \rho_{1,3}(t) \sigma_1(t) \sigma_3(t)}{1 + \tau f_t(T_2, T_3)} \quad (11.38)$$

With this drift, this simple version of LMM is complete. To summarize, to price any contingent claim that depends on the three forwards and makes all of its payments at T_0 , T_1 , T_2 , and T_3 , follow these three steps:

1. Simulate paths for the three forward rates. For each step of size Δt from time t , draw three random variables from a multivariate normal distribution with means $\mu_1(t) \Delta t$, $\mu_2(t) \Delta t$, and 0; with volatilities $\sigma_1(t) \sqrt{\Delta t}$, $\sigma_2(t) \sqrt{\Delta t}$, and $\sigma_3(t) \sqrt{\Delta t}$; and with correlations $\rho_{1,2}(t)$, $\rho_{1,3}(t)$, and $\rho_{2,3}(t)$.
2. Along each path, calculate the future value to time T_3 of the payments from the claim.
3. Average the results across paths and discount to the present.

The General Expression for Drift Changes The simple version of LMM developed here, with three forward rate factors and with the T_3 forward measure fixed as the pricing measure, required the computation of the drifts of $f_t(T_0, T_1)$ and $f_t(T_1, T_2)$ under the T_3 measure. In general, in a model with many more forward rates, use of LMM might require the drift of the forward rate $f_t(T_{k-1}, T_k)$ under the T_i forward measure. The general expression for

these drifts can be derived along the lines of this subsection, but is presented here without proof:¹¹

$$i > k : \mu_k(t) = - \sum_{j=k+1}^i \frac{\tau \rho_{k,j}(t) \sigma_k(t) \sigma_j(t)}{1 + \tau f_t(T_{j-1}, T_j)} \quad (11.39)$$

$$i = k : \mu_k(t) = 0 \quad (11.40)$$

$$i < k : \mu_k(t) = \sum_{j=i+1}^k \frac{\tau \rho_{k,j}(t) \sigma_k(t) \sigma_j(t)}{1 + \tau f_t(T_{j-1}, T_j)} \quad (11.41)$$

If, as in the simple model here, the LMM pricing measure is taken as the terminal measure, then the drifts that have to be computed, like those derived in (11.30) and (11.38), all fall into the category $i > k$.

Calibrating the Instantaneous Volatility and Correlation Functions

The previous subsection took the instantaneous volatility and correlation functions of the forward rates, i.e., $\sigma_i(t)$ and $\rho_{i,j}(t)$, as given. This subsection discusses how practitioners typically set these functions. In short, after assuming some functional form, parameters are set by a combination of calibrating to the market prices of volatility products and estimating the historical behavior of forward rates. With respect to volatility products,¹² short-term rate futures options and caplets are essentially options on forward rates, and caps are portfolios of caplets. As such, these derivatives provide direct information on the volatilities of forward rates. Swaptions are options on swaps. Since the value of a given swap depends on many forward rates simultaneously, swaptions depend not only on the volatilities of forward rates but also on the correlations across forward rates.

It might seem at first that the prices of options on forward rates should be used to calibrate the instantaneous volatility functions while swaption prices should be used to calibrate the correlation functions. In practice, however, this turns out to be too cumbersome a procedure. Basically, the pricing of swaptions depends on the term correlations, i.e., on the correlations of forward rates at option expiration, not on the instantaneous correlations of changes in forward rates, which are the building blocks of

¹¹ For a more rigorous treatment, see, for example, L. Andersen and V. Piterbarg, *Interest Rate Modeling, Volume II*, 2010; and D. Brigo and F. Mercurio, *Interest Rate Models: Theory and Practice*, Springer, 2001, Chapter 6.

¹² All of the products mentioned here are described in detail in Chapter 18.

LMM. Furthermore, because both volatilities and correlations are time dependent, the relationship between correlations of forward rates at expiration and instantaneous correlations is particularly complex.¹³

In light of the discussion of the previous paragraph, a popular methodology for calibrating the LMM model is to use historical data to calibrate the correlation functions and to use some subset of volatility products to calibrate the volatility functions. This procedure also happens to have the advantage of providing relative value indicators across volatility products. For example, if the cap market is used to calibrate volatilities, then the prices of swaptions using that calibrated model are indicators of the relative value of swaptions *versus* caps. In any case, the text continues by describing popular specifications and calibrations of the correlation functions and then moves on to the same for the volatility functions.

Specification and Calibration of the Instantaneous Correlation Functions

An appealing specification of the instantaneous correlation of $f_t(T_{i-1}, T_i)$ and $f_t(T_{j-1}, T_j)$, i.e., $\rho_{i,j}(t)$, is the following, where $T_i > T_j$:¹⁴

$$\rho_{i,j}(t) = \rho_\infty + (1 - \rho_\infty) e^{-\kappa[T_j-t] \times (T_i-T_j)} \quad (11.42)$$

where

$$\kappa[T-t] = a_\infty + (a_0 - a_\infty) e^{-\beta(T-t)} \quad (11.43)$$

The specification of (11.42) and (11.43) captures two important features of the correlation structure of forward rates so long as $1 - \rho_\infty > \rho_\infty$ and $a_0 > a_\infty$. First, the instantaneous correlation between two forward rates should decrease with the distance between the forward rates, i.e., with $T_i - T_j$. For example, the three-month rates 2 and 5 years forward should be more correlated than the three-month rates 2 and 10 years forward. Second, for a given difference between the maturities of forward rates, the instantaneous correlation between two forward rates should increase with the maturity of the earlier forward, i.e., with $T_j - t$ for a fixed $T_i - T_j$. For example, the three-month rates 2 and 3 years forward should be less correlated than the three-month rates 9 and 10 years forward. Economically, while the market may have views of the future that distinguish between short-term

¹³ Spread options, which make payments based on the difference between rates at expiration, are more suitable for extracting correlations across forward rates. These options, however, are not traded with enough liquidity across terms to drive term structure model calibrations.

¹⁴ See L. Andersen and V. Piterbarg, *Interest Rate Modeling, Volume II*, by Atlantic Financial Press, 2010.

TABLE 11.3 Selected Model Correlations for USD
Three-Month Forward Rates of Different Forward Times, in
Years, Estimated from February 6, 2002, to March 27, 2011

	.25	2	5	10	30
.25	1.00	.63	.37	.27	.25
2	.63	1.00	.91	.78	.47
5	.37	.91	1.00	.90	.61
10	.27	.78	.90	1.00	.66
30	.25	.47	.61	.66	1.00

rates two and three years forward, it is much less likely to have views that can distinguish between short-term rates nine and 10 years forward.

Calibrating the functions (11.42) and (11.43) to changes in USD three-month forward rates from February 6, 2002, to March 27, 2011, gives the following parameters: $\rho_\infty = .25$, $a_0 = .610029$, $a_\infty = .029683$, and $\beta = 1.861247$. Table 11.3 gives a selection of the estimated model correlations; the empirical correlations very closely match these model quantities. Note that, as supposed, $1 - \rho_\infty > \rho_\infty$ and $a_0 > a_\infty$, so that the correlations do exhibit the desired properties mentioned earlier. For example, the correlation between the three-month rates 2 and 5 years forward, at 91%, exceeds that of the more separated three-month rates 2 and 10 years forward, at 78%. Also, the correlation of the three-month rates .25 and 5 years forward, at 37%, is less than that of the more distant but almost equally separate three-month rates 5 and 10 years forward, at 90%.

Specification and Calibration of the Instantaneous Volatility Functions A reasonable but relatively simple specification of the instantaneous volatility of the forward rate $f_t(T_{i-1}, T_i)$ is the following:

$$\sigma_i(t) = l_\infty + (l_0 - l_\infty + a(T_{i-1} - t)) e^{-b(T_{i-1} - t)} \quad (11.44)$$

The specification of $\sigma_i(t)$ in (11.44) is relatively simple for a few reasons. First, the volatility function is *time homogeneous* in that volatility depends only on the difference between the forward time T_{i-1} and the calendar time t , i.e., only on the term of the forward rate. This means, for example, that the volatility of the three-month rate three months forward now is assumed to be the same as the volatility of the three-month rate three months forward in five years. This is not necessarily realistic, however, particularly in the short end: when, for example, central banks are keeping rates low for an extended period, it is likely that volatilities of fixed-term forwards will increase in the medium term.

In any case, that volatility in (11.44) depending only on the term of the forward rate makes the functional form particularly easy to interpret. At a term of zero, volatility is l_0 ; at a term of infinity, volatility is l_∞ ; and, it is easy to show that volatility reaches a peak at a term of $\frac{1}{b} + \frac{l_\infty - l_0}{a}$.

A second reason that the functional form of (11.44) is relatively simple is that the volatility of the forward rate does not depend on the level of the forward rate. As discussed in Chapter 10, this sort of dependence is particularly useful for establishing the shape of the probability distribution of the interest rate. Third, volatility is deterministic in equation (11.44); even apart from the possibility of a dependence on rates, there is no randomness to the level of volatility. This assumption is particularly restrictive when modeling derivatives prices because implied volatility levels are constantly changing. In short, the specifications of volatility functions in the version of LMM used in practice are substantially more complex in order to model the volatility skew, discussed in Chapter 18, in a tractable manner. For pedagogical purposes, however, (11.44) will be used throughout this chapter.

As mentioned previously, the parameters of a specification such as (11.44) are calibrated so as to fit the market prices of volatility products. Consider first using caplets. Since a caplet maturing at time T_{i-1} is essentially an option on $f_{T_{i-1}}(T_{i-1}, T_i)$, it turns out that, if the forward rate is normally distributed, a caplet can be priced using the normal Black-Scholes formula (see Chapter 18). Furthermore, since the instantaneous volatility of the forward rate changes over time as, for example, in (11.44), the appropriate implied volatility to use in the Black-Scholes formula is the average of that instantaneous volatility from the present time t to time T_{i-1} . In particular, letting $\hat{\sigma}_i$ denote the implied volatility of the caplet with the underlying forward $f_t(T_{i-1}, T_i)$,

$$\hat{\sigma}_i = \sqrt{\frac{1}{T_{i-1} - t} \int_t^{T_{i-1}} \sigma_i^2(s) ds} \quad (11.45)$$

Fitting the volatility specification (11.44) to USD caplet prices as of July 6, 2006, results in the following parameters values: $l_\infty = .508\%$, $l_0 = .602\%$, $a = .00398$, and $b = .29773$. Figure 11.6 shows the resulting instantaneous volatility function along with the resulting model caplet volatilities. The humped shape of these volatility functions is typical of market-implied volatilities and is easily captured by the functional form (11.44). By the way, note that, using the expression given earlier, the instantaneous volatility function peaks at a term of 3.12 years.

Table 11.4 explores the difference between market and model caplet volatilities. Column (3) of the table shows selected model caplet volatilities, which match those in Figure 11.6. Column (2) gives the respective market volatilities and Column (4) gives the difference between the two. The




FIGURE 11.6 Model Instantaneous and Caplet Volatilities Fit to USD Caplet Prices as of July 6, 2006

volatility function is able to fit the market relatively well, although there are oscillations of richness and cheapness across terms.

If a closer fit is required, the volatility function (11.44) can be generalized, e.g., by multiplying $\sigma_i(t)$ by a different constant for each forward rate, $f_t(T_{i-1}, T_i)$, and using these constants to match caplet volatilities exactly. Note that, since each constant depends on T_{i-1} but not on t , their inclusion breaks the time homogeneity of the volatility specification. As mentioned in the earlier discussion of the time homogeneity property, however, it may very well be desirable to introduce calendar effects in volatility, particularly for near-dated forwards.

Before turning to calibrating the volatility functions with swaptions rather than caplets, it is worth noting that the caplet calibration is, in

TABLE 11.4 Selected Market vs. Model Caplet Volatilities for Calibrations Using Caplet and Swaption Prices, as of July 6, 2006

(1)	(2)	(3)	(4)	(5)	(6)
Expiration	Caplet Vol	Caplet Model	Calibration Mkt-Model	Swaption Model	Calibration Mkt-Model
.5	59.48	60.91	−1.42	68.69	−9.21
2	86.26	84.13	2.13	85.66	.61
5	92.09	96.46	−4.37	95.68	−3.59
10	87.35	86.15	1.20	90.27	−2.91
20	69.84	64.74	5.10	76.45	−6.60
30	45.61	53.52	−7.91	69.16	−23.55

practice, not so straightforward as presented to this point. Caps, i.e., portfolios of caplets, are much more liquid than caplets themselves. But, as discussed in Chapter 18, extracting at-the-money caplet volatilities from at-the-money caps is a challenging exercise in its own right. As a result, Eurodollar or Euribor futures options, also discussed in Chapter 18, are sometimes substituted for caplets. In fact, using options from these markets actually presents an additional opportunity for pinning down the volatility functions. The *quarterly* rate futures options are like caplets in that they are options, that mature at time T_{i-1} , on a forward from T_{i-1} to T_i . Put another way, both caplets and these options provide information about the average volatility of the forward rate from the present to T_{i-1} . *Mid-curve* futures options, or *mid-curves*, however, are options on a forward from T_{i-1} to T_i that mature at some earlier time, before T_{i-1} . Therefore, mid-curves provide information about the average volatility of the forward rate to that earlier date, which is information that is not available from caplets.

The discussion now turns to calibrating the instantaneous volatility functions from swaptions. In fact, swaptions are more liquid than caps, especially at longer maturities. The difficulty is that, also as mentioned previously, calibrating to swaptions is not so straightforward as calibrating to caplets because swaption prices are complicated functions of forward rate volatilities and correlations. As it turns out, however, there is an approximation that facilitates the calibration of (11.44) to swaption prices. Swap rates can be reasonably approximated as a linear combination of forward rates with fixed weights. But linear combinations of normally distributed random variables are themselves normal. Hence, when forward rates are assumed to be jointly normal, as they have been here, swap rates are approximately normal with volatilities that can be calculated from the volatilities and correlations of the forward rates. These swap rate volatilities can then be compared with market implied volatilities of swaptions. In short, specifications like (11.44) can and are commonly calibrated to the swaption market.

Table 11.5 shows the results of calibrating LMM to the USD swaption market as of July 6, 2006, using the volatility specification (11.44) and the correlation specification (11.42) and (11.43) calibrated to historical data, as described previously. Swaptions will be discussed in detail in Chapter 18, but, for now, Table 11.5 is read as saying that the swaption to enter into a five-year swap (swap tenor) in two years (option maturity) trades at a volatility of 87.8 basis points. According to the panel giving the difference between market and model volatilities, LMM does fit swaption volatilities rather well, although there are regions of over- and under-pricing. By the way, as with the caplet calibration, constants can be added to the specification of the volatility functions so that selected swaption volatilities can be fitted exactly.

To conclude the discussion of the calibration of volatilities, return to Table 11.4. As discussed previously, columns (2), (3), and (4) of this table

TABLE 11.5 Market and Fitted USD Swaption Volatilities, in Basis Points, from the LMM Model Calibrated to Swaption Prices as of July 6, 2006

Swap Tenor	1	2	3	4	5	10	30
Option Maturity	Market Implied Volatility						
.5	68.4	74.5	76.2	77.0	77.8	76.8	74.3
2	87.8	88.6	88.3	88.1	87.8	86.0	79.2
5	93.6	93.2	92.2	91.5	90.8	87.6	76.6
10	88.1	87.1	86.2	85.4	84.5	79.7	
30	64.0	62.8	62.2	61.7	61.2	57.5	
	Model Implied Volatility						
.5	73.6	77.3	79.6	80.9	81.6	81.4	77.7
2	88.0	88.6	88.7	88.4	87.8	84.6	79.2
5	95.8	94.1	92.4	90.9	89.6	84.6	78.8
10	89.9	88.0	86.3	84.8	83.6	79.5	
30	69.6	69.0	68.5	68.0	67.7	66.4	
	Market-Model Volatility						
.5	-5.3	-2.7	-3.4	-3.9	-3.8	-4.6	-3.4
2	-.2	0	-.4	-.3	0	1.4	0
5	-2.2	-.9	-.2	.6	1.2	3.0	-2.2
10	-1.8	-.9	-.1	.5	.9	.2	
30	-5.6	-6.2	-6.2	-6.3	-6.5	-8.9	

compare market and model caplet volatilities in an LMM model calibrated to caplet prices. Columns (2), (5), and (6) compare market and model caplet volatilities in an LMM model calibrated to swaption prices, in particular, the calibration used in Table 11.5. Column (6), therefore, is a report on the volatilities of caplets relative to swaptions. As of July 6, 2006, market caplet volatilities are significantly lower than would be consistent with volatilities in the swaption market. In fact, swaptions and caplets, or equivalently caps, do usually trade rich or cheap relative to one another. Caps tend to be used by floating-rate borrowers to hedge the risks of rising rates while swaptions are used by participants in the mortgage-backed securities market to hedge prepayments that accelerate when rates fall. (See Chapter 20.) Hence, when fears of rising rates dominate, caps tend to trade rich while, when fears of falling rates dominate, swaptions tend to trade rich. At the time of the calibrations reported here, in July 2006, it was widely believed that the Board of Governors of the Federal Reserve System had finished increasing rates in the monetary tightening cycle that had begun in the summer of 2004. As a result, the demand for caps fell and they traded cheap relative to swaptions.

Pricing an Interest Rate Exotic

This subsection illustrates the use of LMM in pricing an illustrative exotic derivative as of July 6, 2006. This exotic makes one payment on March 21, 2007, which is determined as a function of the history of both the three-month rate forward to December 20, 2006, and the three-month rate forward to March 21, 2007. In particular, calculate the geometric average of the three-month rate forward to December 20, 2006, as observed on July 6, 2006, September 20, 2006, and December 20, 2006. Then calculate the geometric average of the three-month rate forward to March 21, 2007, as observed on the same three dates. Finally, subtract the second geometric average from the first. That is the payoff of the exotic per unit face value. As of July 6, 2006, the three-month rate forward to December 20, 2006, was 5.68% while the three-month rate forward to March 21, 2007, was 5.65%.

To place this problem in the LMM framework, let time 0 be July 6, 2006, and define other times and forward rates as in Table 11.6. Using the notation in this table, the payoff of the exotic on May 21, 2007, is equal to

$$\begin{aligned} & \left[f_0(T_1, T_2) \times f_{T_0}(T_1, T_2) \times f_{T_1}(T_1, T_2) \right]^{1/3} \\ & - \left[f_0(T_2, T_3) \times f_{T_0}(T_2, T_3) \times f_{T_1}(T_2, T_3) \right]^{1/3} \end{aligned} \tag{11.46}$$

The strategy for valuing the exotic will be to simulate the forward rates $f_t(T_1, T_2)$ and $f_t(T_2, T_3)$. Since the last and only cash flow of the exotic is on March 21, 2007, i.e., at time T_2 , set the pricing measure as the T_2 forward measure. This means that $f_t(T_1, T_2)$ has no drift under the pricing measure. The drift of the other required forward, namely $f_t(T_2, T_3)$, has an instantaneous drift under the T_2 measure given by (11.41) with $i = 2$ and $k = 3$, that is,

$$\begin{aligned} \mu_3(t) &= \frac{\tau \rho_{3,3}(t) \sigma_3(t) \sigma_3(t)}{1 + \tau f_t(T_2, T_3)} \\ &= \frac{.25 \sigma_3^2(t)}{1 + .25 f_t(T_2, T_3)} \end{aligned} \tag{11.47}$$

TABLE 11.6 Notation for the Exotic Derivative Pricing Example

Start Date	Time	End Date	Time	Forward Rate
9/20/06	T_0	12/20/06	T_1	$f_t(T_0, T_1)$
12/20/06	T_1	3/21/07	T_2	$f_t(T_1, T_2)$
3/21/07	T_2	6/21/07	T_3	$f_t(T_2, T_3)$

To simulate the first path, start at the given initial value of the forwards, $f_0(T_1, T_2) = 5.68\%$ and $f_0(T_2, T_3) = 5.65\%$. Then take a step from time 0 to time T_0 , which, in years, is a step size equal to the number of days between July 6, 2006, and September 20, 2006, i.e., 76, divided by 365, or .208.

To take the step from time 0 to time T_0 , begin by drawing the changes in the two forward rates, $\Delta f_0(T_1, T_2)$ and $\Delta f_0(T_2, T_3)$, as two random variables that are bivariate normal with means 0 and $\mu_3(0) \times .208$, with volatilities $\sigma_2(0)\sqrt{.208}$ and $\sigma_3(0)\sqrt{.208}$, and with correlation $\rho_{2,3}(0)$. According to Appendix E in this chapter this is accomplished by drawing and transforming two independent standard normal variables, say $z_2 = -.506$ and $z_3 = -.923$, so that

$$\Delta f_0(T_1, T_2) = \sigma_2(0)\sqrt{.208} \times z_2 \quad (11.48)$$

$$\begin{aligned} \Delta f_0(T_2, T_3) = & \mu_3(0) \times .208 \\ & + \sigma_3(0)\sqrt{.208} \left(z_1 \rho_{2,3}(0) + z_2 \sqrt{1 - \rho_{2,3}(0)^2} \right) \end{aligned} \quad (11.49)$$

To compute the numerical values of the changes in (11.48) and (11.49), use the specifications and parameters given in the text to compute that $\sigma_2(0) = .7489\%$, $\sigma_3(0) = .8121\%$, and $\rho_{2,3}(0) = .9499$. Then compute that $\mu_3(0) = .034$ from (11.47) using the value of $\sigma_3(0)$ just given and the initial value $f_0(T_2, T_3) = 5.65\%$. Finally, return to (11.48) and (11.49) with all of these values, including the given draws of the two independent standard normal random variables, to find that $\Delta f_0(T_1, T_2) = -.173\%$ and $\Delta f_0(T_2, T_3) = -.285\%$.

Adding the increments just computed to the initial value of the two forward rates gives the value of these forwards at time T_0 : $f_{T_0}(T_1, T_2) = 5.507\%$ and $f_{T_0}(T_2, T_3) = 5.365\%$.

The simulation continues with a step of .249 years, from T_0 or September 20, 2006, to T_1 or December 20, 2006. This is accomplished as was the step from 0 to T_0 . Use the functional forms and parameters given in the text to compute that $\sigma_2(T_0) = .6874\%$, $\sigma_3(T_0) = .7601\%$, and $\rho_{2,3}(T_0) = .9297$. Then use $\sigma_3(T_0)$, $f_{T_0}(T_2, T_3)$, and (11.47) to compute that $\mu_3(T_0) = .036$. Draw another two independent standard normal variables, say $z_2 = -.605$ and $z_3 = .783$, and write down the changes in the forward rates from T_0 to T_1 :

$$\Delta f_{T_0}(T_1, T_2) = \sigma_2(T_0)\sqrt{.249} \times z_2 \quad (11.50)$$

$$\begin{aligned} \Delta f_{T_0}(T_2, T_3) = & \mu_3(T_0) \times .249 + \sigma_3(T_0)\sqrt{.249} \\ & \times \left(z_1 \rho_{2,3}(T_0) + z_2 \sqrt{1 - \rho_{2,3}(T_0)^2} \right) \end{aligned} \quad (11.51)$$

Filling in the values now available, $\Delta f_{T_0}(T_1, T_2) = -.208\%$ and $\Delta f_{T_0}(T_2, T_3) = -.104\%$. Adding these increments to the values of the forward rates at time T_0 gives the two forward rates at time T_1 : $f_{T_1}(T_1, T_2) = 5.299\%$ and $f_{T_1}(T_2, T_3) = 5.262\%$.

All of the forward rates are now known to compute the payoff of the claim on March 21, 2007, from (11.46):

$$[5.680\% \times 5.507\% \times 5.299\%]^{1/3} - [5.650\% \times 5.365\% \times 5.262\%]^{1/3} \\ = 5.493\% - 5.423\% \quad (11.52)$$

$$= .070\% \quad (11.53)$$

Dividing this by $d_{T_2}(T_2) = 1$, as required by (11.21), gives the same seven basis points.

To value the exotic derivative in this example, repeat this process for many paths, average the result, and multiply by the discount factor $d_0(T_2)$.

APPENDIX A: EQUIVALENCE OF THE CASCADE AND REDUCED FORMS OF THE GAUSS+ MODEL

Begin with equations (11.9), (11.10), (11.11) of the reduced form of the model. In matrix form:

$$\begin{pmatrix} dx_1 \\ dx_2 \\ dx_3 \end{pmatrix} = - \begin{pmatrix} \alpha_r & 0 & 0 \\ 0 & \alpha_m & 0 \\ 0 & 0 & \alpha_l \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} dt \\ + \begin{pmatrix} \sigma_{11} & \sigma_{12} & 0 \\ \sigma_{21} & -\sigma_{12} & 0 \\ -(\sigma_{11} + \sigma_{21}) & 0 & 0 \end{pmatrix} \begin{pmatrix} dw_1 \\ dw_2 \\ 0 \end{pmatrix} \quad (11.54)$$

where $E(dw_1 dw_2) = 0$ as in (11.12). More compactly,

$$d\mathbf{x} = -\alpha \mathbf{x} dt + \Sigma d\mathbf{w} \quad (11.55)$$

Now define the matrix $\mathbf{A}$ to be

$$\mathbf{A} = \begin{pmatrix} 1 & 1 & 1 \\ 0 & \frac{\alpha_r - \alpha_m}{\alpha_r} & \frac{\alpha_r - \alpha_l}{\alpha_r} \\ 0 & 0 & \frac{(\alpha_r - \alpha_l)(\alpha_m - \alpha_l)}{\alpha_r \alpha_m} \end{pmatrix} \quad (11.56)$$

Let $\mathbf{u} = (u_1, u_2, u_3)'$ be a transformation of the variables $\mathbf{x}$ such that

$$\mathbf{u} = \mathbf{A}\mathbf{x} + \boldsymbol{\theta} \quad (11.57)$$

where $\boldsymbol{\theta} = (\theta, \theta, \theta)'$. Premultiply both sides of (11.57) by the inverse of $\mathbf{A}$ and solve for $\mathbf{x}$:

$$\mathbf{x} = \mathbf{A}^{-1}(\mathbf{u} - \boldsymbol{\theta}) \quad (11.58)$$

Next, take the differential of both sides of (11.57):

$$d\mathbf{u} = \mathbf{A}d\mathbf{x} \quad (11.59)$$

Putting the pieces together, substitute (11.55) into (11.59)

$$d\mathbf{u} = \mathbf{A}[-\alpha x dt + \Sigma d\mathbf{w}] \quad (11.60)$$

and then substitute (11.58) into (11.60):

$$\begin{aligned} d\mathbf{u} &= \mathbf{A}[-\alpha \{\mathbf{A}^{-1}(\mathbf{u} - \boldsymbol{\theta}) dt\} + \Sigma d\mathbf{w}] \\ d\mathbf{u} &= -\mathbf{A}\alpha\mathbf{A}^{-1}(\mathbf{u} - \boldsymbol{\theta}) dt + \mathbf{A}\Sigma d\mathbf{w} \end{aligned} \quad (11.61)$$

Given the definitions of $\mathbf{A}$ and α , it is straightforward, although tedious, to show that

$$\mathbf{A}^{-1} = \begin{pmatrix} 1 & -\frac{\alpha_r}{\alpha_r - \alpha_m} & \frac{\alpha_r \alpha_m}{(\alpha_r - \alpha_m)(\alpha_r - \alpha_l)} \\ 0 & \frac{\alpha_r}{\alpha_r - \alpha_m} & -\frac{\alpha_r \alpha_m}{(\alpha_r - \alpha_m)(\alpha_m - \alpha_l)} \\ 0 & 0 & \frac{\alpha_r \alpha_m}{(\alpha_r - \alpha_l)(\alpha_m - \alpha_l)} \end{pmatrix} \quad (11.62)$$

$$-\mathbf{A}\alpha\mathbf{A}^{-1} = \begin{pmatrix} -\alpha_r & \alpha_r & 0 \\ 0 & -\alpha_m & \alpha_m \\ & & -\alpha_l \end{pmatrix} \quad (11.63)$$

$$\begin{aligned} -\mathbf{A}\alpha\mathbf{A}^{-1}(\mathbf{u} - \boldsymbol{\theta}) dt &= \begin{pmatrix} -\alpha_r & \alpha_r & 0 \\ 0 & -\alpha_m & \alpha_m \\ & & -\alpha_l \end{pmatrix} \begin{pmatrix} u_1 - \theta \\ u_2 - \theta \\ u_3 - \theta \end{pmatrix} dt \\ &= \begin{pmatrix} -\alpha_r (u_1 - u_2) dt \\ -\alpha_m (u_2 - u_3) dt \\ -\alpha_l (u_3 - \theta) dt \end{pmatrix} \end{aligned} \quad (11.64)$$

Also, given the definitions of $\mathbf{A}$ and Σ it can be shown that

$$\mathbf{A}\Sigma = \begin{pmatrix} 0 & 0 & 0 \\ \rho\sigma_m & \sqrt{1-\rho^2}\sigma_m & 0 \\ \sigma_l & 0 & 0 \end{pmatrix} \quad (11.65)$$

Having performed these computations, substitute (11.64) and (11.65) into (11.61) to see that

$$du_1 = -\alpha_r (u_1 - u_2) dt \quad (11.66)$$

$$du_2 = -\alpha_m (u_2 - u_3) dt + \rho\sigma_m dw_1 + \sqrt{1-\rho^2}\sigma_m dw_2 \quad (11.67)$$

$$du_3 = -\alpha_l (u_3 - \theta) dt + \sigma_l dw_1 \quad (11.68)$$

But with $u_1 = r$, $u_2 = m$, and $u_3 = l$, equations (11.66) through (11.68) are identical to (11.1) through (11.3). Hence the transformations (11.57) and (11.58) can be used to move from the reduced form to the cascade form and back again.

APPENDIX B: THE FUNCTION $\Upsilon(T)$ FOR THE GAUSS+ MODEL

This appendix presents the function $\Upsilon(T)$ without derivation. The spot rate in the model is given by (11.16), which is reproduced here as (11.69):

$$\hat{r}(T) = \theta + \bar{x}_1 \frac{1 - e^{-\alpha_r T}}{\alpha_r T} + \bar{x}_2 \frac{1 - e^{-\alpha_m T}}{\alpha_m T} + \bar{x}_3 \frac{1 - e^{-\alpha_l T}}{\alpha_l T} + \Upsilon(T) \quad (11.69)$$

where

$$\begin{aligned} \Upsilon(T) = & -\frac{1}{2} (G_1^2 + G_2^2) + (\sigma_{11} G_1 + \sigma_{12} G_2) \frac{1 - e^{-\alpha_r T}}{\alpha_r^2 T} \\ & + (\sigma_{21} G_1 - \sigma_{12} G_2) \frac{1 - e^{-\alpha_m T}}{\alpha_m^2 T} - (\sigma_{11} + \sigma_{21}) G_1 \frac{1 - e^{-\alpha_l T}}{\alpha_l^2 T} \\ & - (\sigma_{11}\sigma_{21} - \sigma_{12}^2) \frac{1 - e^{-(\alpha_r + \alpha_m)T}}{\alpha_r \alpha_m (\alpha_r + \alpha_m) T} + \sigma_{11} (\sigma_{11} + \sigma_{21}) \frac{1 - e^{-(\alpha_r + \alpha_l)T}}{\alpha_r \alpha_l (\alpha_r + \alpha_l) T} \\ & + \sigma_{21} (\sigma_{11} + \sigma_{21}) \frac{1 - e^{-(\alpha_m + \alpha_l)T}}{\alpha_m \alpha_l (\alpha_m + \alpha_l) T} - (\sigma_{11}^2 + \sigma_{12}^2) \frac{1 - e^{-2\alpha_r T}}{4\alpha_r^3 T} \\ & - (\sigma_{12}^2 + \sigma_{21}^2) \frac{1 - e^{-2\alpha_m T}}{4\alpha_m^3 T} - (\sigma_{11} + \sigma_{21})^2 \frac{1 - e^{-2\alpha_l T}}{4\alpha_l^3 T} \end{aligned} \quad (11.70)$$

and

$$G_1 = \frac{\sigma_{11}}{\alpha_r} + \frac{\sigma_{21}}{\alpha_m} - \frac{\sigma_{11} + \sigma_{21}}{\alpha_l} \quad (11.71)$$

$$G_2 = \frac{\sigma_{12}}{\alpha_r} - \frac{\sigma_{12}}{\alpha_m} \quad (11.72)$$

APPENDIX C: ESTIMATING THE PARAMETERS OF THE GAUSS+ MODEL

This appendix describes how to define $\Psi_1(T)$ and $\Psi_2(T)$ in equations (11.17) and (11.18) of the text. The notation of Appendix A is continued here.

Define the 3x1 vector $\mathbf{b}$ such that

$$\mathbf{b} = \begin{pmatrix} \frac{1 - e^{-\alpha_r T}}{\alpha_r T} \\ \frac{1 - e^{-\alpha_m T}}{\alpha_m T} \\ \frac{1 - e^{-\alpha_l T}}{\alpha_l T} \end{pmatrix} \quad (11.73)$$

Then, equation (11.16) can be rewritten as

$$\hat{r}(T) = \theta + \Upsilon(T) + \mathbf{b}'\mathbf{x} \quad (11.74)$$

Substituting for $\mathbf{x}$ from equation (11.58),

$$\begin{aligned} \hat{r}(T) &= \theta + \Upsilon(T) + \mathbf{b}'\mathbf{A}^{-1}(\mathbf{u} - \boldsymbol{\theta}) \\ &= \theta + \Upsilon(T) - \mathbf{b}'\mathbf{A}^{-1}\boldsymbol{\theta} + \mathbf{b}'\mathbf{A}^{-1}\mathbf{u} \end{aligned} \quad (11.75)$$

As shown toward the end of Appendix A in this chapter, $u_1 = r$, $u_2 = m$, and $u_3 = l$. So, partitioning the 1x3 vector $\mathbf{b}'\mathbf{A}^{-1}$ into its first element, b_r , and a 1x2 vector, $\mathbf{b}'_{ml}$, of its other two elements, (11.75) can be further rewritten as

$$\hat{r}(T) = \theta + \Upsilon(T) - \mathbf{b}'\mathbf{A}^{-1}\boldsymbol{\theta} + b_r r + \mathbf{b}'_{ml} \begin{pmatrix} m \\ l \end{pmatrix} \quad (11.76)$$

The partials of $\hat{r}(T)$ with respect to m and l are given by (11.76), but, since m and l are correlated, these partials cannot be equated to principal

components. The final step, therefore, is to transform the variables m and l into uncorrelated variables. Define the matrix S such that

$$S = \begin{pmatrix} \rho\sigma_m & \sigma_m\sqrt{1-\rho^2} \\ \sigma_l & 0 \end{pmatrix} \quad (11.77)$$

It follows that

$$S^{-1} = \frac{1}{\sigma_m\sigma_l\sqrt{1-\rho^2}} \begin{pmatrix} 0 & \sigma_m\sqrt{1-\rho^2} \\ \sigma_l & \rho\sigma_m \end{pmatrix} \quad (11.78)$$

and that

$$\begin{aligned} S^{-1} \begin{pmatrix} m \\ l \end{pmatrix} &= \frac{1}{\sigma_m\sigma_l\sqrt{1-\rho^2}} \begin{pmatrix} 0 & \sigma_m\sqrt{1-\rho^2} \\ \sigma_l & \rho\sigma_m \end{pmatrix} \begin{pmatrix} m \\ l \end{pmatrix} \\ &= \begin{pmatrix} \frac{l}{\sigma_l} \\ \frac{1}{\sqrt{1-\rho^2}} \left[\frac{m}{\sigma_m} - \frac{\rho l}{\sigma_l} \right] \end{pmatrix} \end{aligned} \quad (11.79)$$

It is easily verified that the variance-covariance matrix of (11.79) is the identity matrix.

Returning now to equation (11.76), insert a multiplication by the identity matrix written as SS^{-1} :

$$\hat{r}(T) = \theta + \Upsilon(T) - \mathbf{b}'\mathbf{A}^{-1}\boldsymbol{\theta} + b_r r + \mathbf{b}'_{ml} SS^{-1} \begin{pmatrix} m \\ l \end{pmatrix} \quad (11.80)$$

and then define the random variables z_1 and z_2 such that

$$\begin{pmatrix} z_1 \\ z_2 \end{pmatrix} = S^{-1} \begin{pmatrix} m \\ l \end{pmatrix} \quad (11.81)$$

But, as just discussed, the variance-covariance of the right-hand side of (11.81) and, therefore, of z_1 and z_2 is the identity matrix. Hence, z_1 and z_2 are uncorrelated with unit variance. They are, of course, normally distributed, since m and l are normally distributed.

Finally, substituting (11.81) into (11.80),

$$\hat{r}(T) = \theta + \Upsilon(T) - \mathbf{b}'\mathbf{A}^{-1}\boldsymbol{\theta} + b_r r + \mathbf{b}'_{ml} S \begin{pmatrix} z_1 \\ z_2 \end{pmatrix} \quad (11.82)$$

which means that the functions $\Psi_1(T)$ and $\Psi_2(T)$ of the text are defined such that

$$[\Psi_1(T), \Psi_2(T)] = \mathbf{b}'_{ml} \mathbf{S} \quad (11.83)$$

Note that $\mathbf{b}$ and $\mathbf{A}^{-1}$, which determine $\mathbf{b}'_{ml}$, depend only on the mean reversion coefficients of the cascade form of the Gauss+ model and $\mathbf{S}$ depends on the two volatilities and the correlation. Hence, by (11.83), $\Psi_1(T)$ and $\Psi_2(T)$ depend only on these parameters as well, i.e., not on θ nor on the initial values of the factors.

APPENDIX D: FITTING THE INITIAL TERM STRUCTURE IN THE GAUSS+ MODEL

Zero coupon bond prices in the model defined by (11.19) are given by

$$\begin{aligned} P^*(T) &= E \left[\exp \left(- \int_0^T r^*(t) dt \right) \right] \\ &= \exp \left(- \int_0^T \theta(t) dt \right) E \left[\exp \left(- \int_0^T r(t) dt \right) \right] \\ &= \exp \left(- \int_0^T \theta(t) dt \right) P(T) \end{aligned} \quad (11.84)$$

where P and r are the zero coupon bond prices and the short-term rate in the original formulation of the Gauss+ model (11.8) with the constant θ set equal to zero. Taking the derivative of both sides of (11.84) with respect to T ,

$$\begin{aligned} \frac{dP^*}{dT} &= \exp \left(- \int_0^T \theta(t) dt \right) \frac{dP}{dT} - \theta(T) \exp \left(- \int_0^T \theta(t) dt \right) P(T) \\ \frac{dP^*}{dT} &= \frac{P^*(T)}{P(T)} \frac{dP}{dT} - \theta(T) P^*(T) \\ \theta(T) &= \frac{1}{P^*(T)} \frac{dP^*}{dT} - \frac{1}{P(T)} \frac{dP}{dT} \end{aligned} \quad (11.85)$$

By the definition of forward rates, see equation (2.41), the two terms on the right-hand side of (11.85) are the forward rates in these respective models. Let $f^*(t)$ denote the continuously compounded forward rates in the market, or equivalently, in the adjusted Gauss+ model, and let $f(t)$

denote the initial Gauss+ model with θ equal to zero. Equation (11.85) then becomes

$$\theta(T) = f^*(T) - f(T) \quad (11.86)$$

as was to be shown.

APPENDIX E: DRAWING RANDOM NUMBERS FROM A MULTIVARIATE NORMAL DISTRIBUTION

The problem is to draw N random variables from a multivariate normal distribution with means given by the $N \times 1$ vector μ and variances and covariances given by the $N \times N$ matrix Σ .

Software is widely available for drawing random variables from a standard normal distribution, i.e., with mean 0 and volatility 1, it is assumed here that the starting point is an $N \times 1$ vector z with each element a draw from a standard normal distribution.

Software is also widely available to perform a Choleski decomposition, i.e., to decompose a positive definite matrix, like the variance-covariance matrix Σ , into the product of a (lower triangular) matrix, A and its transpose A' :

$$AA' = \Sigma \quad (11.87)$$

Then, using the draws from the standard normal distribution in the vector z , define the $N \times 1$ vector x such that

$$x = \mu + Az \quad (11.88)$$

This vector x is a draw from the desired multivariate normal distribution. To see this, note that

$$\begin{aligned} E[x] &= \mu + AE[z] \\ &= \mu \end{aligned} \quad (11.89)$$

and that

$$\begin{aligned} V[x] &= V[Az] \\ &= AV[zz']A' \end{aligned} \quad (11.90)$$

$$= AA' \quad (11.91)$$

$$= \Sigma \quad (11.92)$$

where $V[zz']$ is the identity matrix since z is a vector of standard normal variables.

For convenience, it is easily verified that, in the two-dimensional case,

$$\Sigma = \begin{pmatrix} \sigma_1^2 & \rho\sigma_1\sigma_2 \\ \rho\sigma_1\sigma_2 & \sigma_2^2 \end{pmatrix} \quad (11.93)$$

and

$$A = \begin{pmatrix} \sigma_1 & 0 \\ \rho\sigma_2 & \sigma_2\sqrt{1-\rho^2} \end{pmatrix} \quad (11.94)$$

Selected Securities and Topics

Parts One, Two, and Three provided the basic tools with which to price and measure the risk of fixed income securities and portfolios. This part can now apply these tools to the details of various products and markets. The chapter titles are for the most part explanation enough for the purposes of introduction. A few additional comments are warranted, however, and are made here.

Chapter 12 covers repo markets and financing. This is an important subject in its own right, both under normal conditions and particularly under conditions of financial stress. But financing is a crucial part of many other valuation and risk problems. Forward and futures prices in bond markets, for example, depend directly on financing rates. And even more fundamentally, as argued in Chapter 17, the foundations of arbitrage pricing depend on financing arrangements and on the relationships of financing rates across securities.

Chapter 13 presents preliminaries on forward and futures contracts. The preliminary subject matter necessary to understand these contracts is substantial enough to distract from the narrative flow of chapters about individual products. Hence, this material is collected here and applied in the rest of Part Four, particularly in the chapters on note and bond futures and on short-term rate derivatives.

Chapter 17 presents material that may be completely new to many readers. First, it revisits the fundamentals of arbitrage pricing under realistic financing arrangements, deriving more realistic conditions under which portfolios of bonds or swaps can be found to replicate other bonds or swaps. Second, the chapter revisits the pricing of swaps when the riskless investable

rate and the rate earned on posted collateral (e.g., federal (fed) funds) does not equal the floating rate index (e.g., LIBOR). The result is a two-curve swap pricing methodology that, since the financial crisis of 2007–2009, has rapidly become the industry standard.

Most of Chapter 18 describes fixed income option products and associated valuation methodologies. The final part of the chapter, however, in justifying the use of the Black-Scholes model in certain fixed income contexts, introduces concepts and techniques that have become staples in finance, i.e., numeraires and martingale pricing. The presentation is done with a minimum of mathematics so that a broader audience has access to these important topics.

Repurchase Agreements and Financing

This chapter is about repurchase agreements or repos, which were introduced in the Overview. Repos are short-term contracts that are used to lend money on the security of usually high-grade collateral, to finance the purchase of bonds, and to borrow bonds to be sold short.

Financial institutions have traditionally relied on repos to finance some portion of fixed income inventory. Repo financing, as secured, short-term borrowing, is typically a relatively inexpensive way to borrow money. The practice can leave firms in a perilous situation, however, should lenders of cash through repos, in times of trouble, fail to renew their loans. This turned out to be an issue in the financial crisis of 2007–2009, which is illustrated in this chapter by two cases, one about liquidity management at Bear Stearns and one about the financing relationship between Lehman Brothers and JPMorgan Chase.

The last part of the chapter focuses on repo rates and, in particular, on the *specials* market in the United States, where market participants lend money at relatively low rates predominantly in order to borrow the most-recently issued and most liquid U.S. Treasury bonds. The behavior of these rates is examined in some detail and linked empirically to the auction cycle of U.S. government bonds.

REPURCHASE AGREEMENTS: STRUCTURE AND USES

A *repurchase agreement* or *repo* is a contract in which a security is traded at some initial price with the understanding that the trade will be reversed at some future date at some fixed price. Repos are used by several different types of market participants for different purposes; figures 12.1 and 12.2 begin the discussion by illustrating a simplified trade between generic counterparties.

At initiation of the repo, depicted in Figure 12.1, counterparty A sells €100 million face amount of the DBR 4s of January 4, 2037, to counterparty



FIGURE 12.1 The Initiation of a Repo Trade

B, for settlement on May 31, 2010, at an invoice price of €111.772 million. At the same time, counterparty A agrees to repurchase that €100 million face amount three months later, for settlement on August 31, 2010, at a purchase price equal to the original invoice price plus interest at a *repo rate* of .23%. Using the actual/360 convention of most money market instruments, and noting that there are 92 days between May 31, 2010, and August 31, 2010 the repurchase price is

$$€111,772,000 \left(1 + \frac{.23\% \times 92}{360} \right) = €111,837,697.10 \qquad (12.1)$$

Hence, at the termination or unwind of the repo, depicted in Figure 12.2, counterparty A repurchases the €100 million face amount of the bund from counterparty B for about €111.838 million. (Bund is another name for a DBR.)



FIGURE 12.2 The Unwinding of a Repo Trade

The next three subsections describe the three reasons to do repo: to lend funds short-term on a secured basis, to finance a long position in a security, and to borrow a security in order to sell it short.

Repos and Cash Management

Investors holding cash for liquidity or safekeeping purposes often find *investing in repo* to be an ideal solution. The most significant example of this is the money market mutual fund industry, which invests on behalf of investors willing to accept relatively low returns in exchange for liquidity and safety. In terms of Figures 12.1 and 12.2, a money market fund would be in the position of counterparty B, lending money while taking collateral and then, at maturity, collecting the loan plus interest and returning the collateral. Holding collateral makes the lender less vulnerable to the creditworthiness of a counterparty because, in the event of a default by counterparty A, counterparty B, in this case the money market fund, can sell the repo collateral to recover any amounts owed. In summary, relative to super-safe and liquid non-interest-bearing bank deposits, repo investments pay a short-term rate without sacrificing much liquidity or incurring significant default risk.

Municipalities constitute another significant category of repo investors. As the timing of tax receipts has little to do with the schedule of public expenditures, municipalities tend to run cash surpluses from tax receipts so as to have money on hand to meet expenditures. These tax revenues cannot be invested in risky securities, but neither should the cash collected lie idle. Short-term loans backed by collateral, like repos, again satisfy both revenue and safety considerations. Other institutions with similar cash management issues that choose to invest in repo are mutual funds, insurance companies, pension funds, and even some nonfinancial corporations. It is worth noting, however, that many lenders in the repo market during the recent financial crisis realized that they were not well positioned, either in expertise or operational ability, to take possession of and liquidate repo collateral.

Since repo investors place a premium on liquidity, they tend to lend *overnight*, rather than for *term*, which refers to any maturity longer than one day. Many investors planning to lend cash through the repo market for an extended period of time will, rather than lend for term, engage in an *open repo*, i.e., a one-day repo that renews itself day-to-day until cancelled by either party. Nevertheless, investors willing to take on some additional liquidity and counterparty risk, in addition to interest rate risk, do lend through term repos. These are available at various maturities, out to several months, although demand declines rapidly with term.

Since safety is the other key consideration of investors in repo, only securities of the highest credit quality are typically accepted as collateral. The most common choices are government securities, debt issues of government-sponsored entities (GSEs), and mortgage-backed securities guaranteed by

the government or the GSEs. (See the Overview for institutional descriptions.) Even taking high-quality securities as collateral, however, a lender of cash faces the risk that a borrower defaults at the same time those securities decline in value.¹ In that eventuality, selling the collateral might not fully cover the loss of the loan amount. Therefore, repo agreements often provide *haircuts* through which investors require borrowers to deliver securities worth more than the amount of the loan. In the example of this section, counterparty B might lend only €106 million against the €111.772 million of securities and, of course, collect only €106 million plus interest at maturity. In addition, repo agreements are normally subject to *margin calls*, through which the borrower of cash supplies extra collateral in declining markets but may withdraw collateral in advancing markets. Again using the example of this section, should the value of the bund collateral decline from its initial value of €111.772 million to €110 million, the borrower would have to put up the €1.772 million difference in additional collateral to protect the investor's loan. Combining the haircut and repricing features, after the drop in bund value the investor would still have €111.772 million of collateral against the outstanding loan of €106 million.

While repo investors care about the quality of the collateral they accept, they do not usually care about which particular bond they accept. Hence, while repo investors can be very particular about which classes of securities they will take as collateral, e.g., Euro-area government bonds with less than five years to maturity, they will not insist on receiving any particular security within that delineated class. For this reason these investors are said to accept *general collateral*, which trades at *general collateral repo rates*. The types and determination of repo rates are discussed later in this chapter.

Repos and Long Financing

Financial institutions are the typical borrowers of cash in repo markets. Say that a client wants to sell €100 million face amount of the DBR 4s of January 4, 2037, to the trading desk of a financial institution. The desk will buy the bonds and eventually sell them to another client. Until that buyer is found, however, the trading desk needs to raise money to pay the client. Put another way, it needs to *finance* the purchase of the bonds. Rather than draw on the scarce capital of the financial institution for this purpose, the trading desk will *repo* or *repo out* the securities, or *sell the repo*. This means that it will

¹In risk management parlance, however, this is called a *right-way risk*. If a repo borrower, typically a well-established financial institution, were to default, the market response would probably be a “flight-to-quality trade” in which securities of the most creditworthy governments increase in value.




FIGURE 12.3 Back-to-Back Repo Trades

borrow the purchase amount from someone, like a money market fund, and use the DBR 4s, which it just bought, as collateral. Thus, the trading desk acts as counterparty A in Figure 12.1. Of course, any haircuts applied will require the trading desk to use some of its capital to make up the difference between the purchase price of the securities and the amount borrowed from the repo counterparty.

When the bunds are ultimately sold to some buyer, the desk will, still as counterparty A but now in Figure 12.2, unwind the repo, using the proceeds from the sale of the bunds to repay the repo loan and using the returned collateral to make delivery of those bunds to that buyer. If no buyer is found before the expiration of the repo, the trading desk will have to *roll* or renew the repo for another period with the same counterparty or unwind that repo and find a different counterparty to finance the bond. This latter option is illustrated in Figure 12.3. The trading desk, still as counterparty A, will repay counterparty B the approximately €111.838 million due and take back the bunds; then borrow funds from counterparty C and deliver the bunds as collateral. Note that since the cash obtainable from counterparty C depends on the price of the bond at the time of the roll, while the cash due to counterparty B depends on the amount owed from the previously agreed-upon transaction, this renewal of the repo may leave counterparty A, the trading desk, with a cash surplus or deficit. Of course, had the trading desk hedged the price risk of its inventory, the profit and loss from the hedge would offset this cash surplus or deficit.

In the example of this subsection, the financial institution used the repo market to finance its inventory for the purpose of making markets. Other uses include financing its proprietary positions² and positions for customers. Repo for proprietary positions can be described by Figures 12.1 and 12.2,

²Currently, the “Volcker rule” is envisioned as limiting the magnitude of proprietary positions held by financial institutions.

with the relevant trading desk again as counterparty A, but with internal rather than customer motivations for purchasing and then selling the bunds. Repo for financing customer positions, at initiation, can be described in terms of Figure 12.3. This time a customer, e.g., a hedge fund, is counterparty B, who wants to finance the purchase of the DBR 4s. The trading desk of the financial institution, counterparty A, does a repo with the customer, lending cash and taking the DBR 4s as collateral. The trading desk then does a back-to-back repo with counterparty C, who provides the cash and takes the collateral originally supplied by the hedge fund. Without haircuts the cash amounts shown would be the same, but, in practice, the haircuts charged on each leg of the trade depend on the creditworthiness and negotiating power of the relevant counterparties.

The issues surrounding financial institutions' use of repo to finance their businesses are discussed later in this chapter.

Reverse Repos and Short Positions

Professional investors often want to short a bond, either as an outright bet that interest rates will rise, as a hedge, or as part of a relative value bet that the price of another security will rise relative to the price of the security being sold short. Say that a hedge fund wants to short the DBR 4s of January 4, 2037. It sells the bund, but then needs to borrow it from somewhere in order to make delivery. In terms of Figure 12.1, the hedge fund is counterparty B, initiating the transaction not because it wants to lend cash but because it wants to borrow the bund. From the point of view of the hedge fund, it will do a *reverse repurchase agreement*,³ will *reverse* or *reverse in* the securities, or will *buy the repo*.

After initiating the reverse, the hedge fund will, at some point in time, be ready to *cover* its short, i.e., to neutralize its economic exposure to the bund by buying that bund back. At that time the hedge fund will buy the bund and then unwind its reverse as in Figure 12.2. Specifically, the hedge fund, as counterparty B, will buy the bund at market and then deliver it to counterparty A, who, in exchange, will return the hedge fund's cash with interest. If the return on the bund has been less than the repo rate of interest, the hedge fund will have made money on an outright short position, while if the return on the bund has been greater than the repo rate of interest, the hedge fund will have lost money on an outright short. Of course, the short might very well have been part of a larger trade in which case the profit and loss (P&L) has more components.

³The practice of calling the trade a reverse repo is particularly confusing because the same trade that is a reverse repo for the borrower of a security is a repo for the lender of that security.

While Figures 12.1 and 12.2 were used to explain both repo investment and reverse repos, it is important to keep in mind that the former are initiated in order to invest cash while the latter are initiated to borrow a bond. So while repo investors are willing to accept general collateral, reverses require the delivery of a particular bond. Repo transactions that require the delivery of a particular bond are called *special trades* and they take place at *special collateral* rates. The specials market is discussed further later in this chapter.

REPO, LIQUIDITY MANAGEMENT, AND THE FINANCIAL CRISIS OF 2007–2009

The Overview introduced liquidity risk and noted that broker-dealers rely less on repo financing currently than they did before the 2007–2009 crisis. A financial institution can borrow funds in many ways, some of which are more stable than others, i.e., some of which can be easily maintained under conditions of financial stress and some of which cannot be so easily maintained. The most stable source of funds is equity capital because equity holders do not have to be paid according to any particular schedule and because they cannot compel a redemption of their shares. Slightly less stable is long-term debt because bondholders have to be paid interest and principal as set out in bond indentures. At the other extreme of funding stability is short-term unsecured funding, like commercial paper: these borrowings have to be repaid in a matter of weeks or months, as they mature, when the institution, under adverse conditions, might not be able to borrow money elsewhere. Not surprisingly, the more stable sources of funds are usually more expensive in terms of the expected return required by the providers of funds. Through *liquidity management*, firms balance the costs of funding against the risks of being caught without the financing necessary to survive.

In the spectrum of financing choices, repo markets are relatively liquid and repo borrowing rates relatively low. Though, by nature of its short maturities, repo is on the less stable side of the funding spectrum, although more stable than short-term, unsecured borrowing. After all, repo collateral should prevent repo lenders from bolting too quickly in response to unfavorable rumors or news. Nevertheless, if repo investors do lose confidence in a financial institution, that institution's repo financing can disappear as fast as the repos mature, which is mostly overnight. The beleaguered institution would no longer be able to facilitate customer trades by holding inventory, would not be able to facilitate customer financing, and might not remain an acceptable counterparty for derivative and even spot security transactions. Furthermore, the institution would have to sell inventory and proprietary positions to repay repo lenders, which sales, given their size and public nature, would likely turn into fire sales and result in significant losses.

Essentially then, while significant business losses rather than financing are the usual cause of a financial institution's difficulties, the loss of financing is often the killing blow. The same argument, of course, applies to all leveraged investors, like part of the hedge fund world. To the extent that a firm borrows money to finance positions, losing the confidence of repo and other secured financing counterparties can result in fire sales, substantial losses, and possible bankruptcy.

The risks of repo funding juxtaposed with those of repo investing create tensions between borrowers and lenders of cash as well as difficulties for regulators. Borrowers want to extend the term of their repo borrowings,⁴ sometimes at the encouragement of their regulators, so as to have more time, should conditions for refinancing deteriorate, to arrange alternate financing, to raise capital, or even to sell corporate entities. Lenders, on the other hand, want to shorten the term of their repo lendings, sometimes at the encouragement of their regulators, so as to minimize exposure to borrower defaults. Prices, in this case repo rates, allocate repos of various terms across borrowers and lenders, but the financial system as a whole cannot both extend the maturities of secured financing and contract the maturities of secured lending.

In the run-up to the financial crisis of 2007–2009, borrowers financed lower-quality collateral, like lower-quality corporate bonds and lower-quality mortgage-backed securities, at the relatively low rates and haircuts available in the repo market. Lenders, for their part, accepted this collateral in exchange for rates somewhat higher than those available when lending on higher-quality collateral. The resulting expansion of collateral accepted for repo did not work out well during the crisis, particularly for borrowers who were unable to meet margin calls caused by declining security values, who were unable to post sufficient collateral in response to lenders' raising haircuts, or who were unable to replace lost financing arrangements when lower-quality collateral found fewer and fewer takers. The worst-hit borrowers suffered collateral liquidations, losses, capital depletion, and business failure.

Case Study: Repo Financing and the Collapse of Bear Stearns

This is an excerpt from the testimony of Paul Friedman before the Financial Crisis Inquiry Commission on May 5, 2010.⁵ Mr. Friedman, a Senior

⁴As an operational aside, the term repo financing usually includes *rights of substitution* that enable the borrower of cash who needs to sell a particular bond being financed to replace that bond as collateral with other bonds of comparable value and quality.

⁵Source: <http://fcic.gov/hearings/pdfs/2010-0505-Friedman.pdf>

Managing Director at Bear Stearns, was responsible for its fixed income repo desk at the time of the firm's demise in March 2008.

Bear Stearns generally financed its business by borrowing funds on a secured and unsecured basis and through the use of equity capital. During 2006, Bear Stearns decided to reduce the amount of short-term unsecured funding, primarily commercial paper, that it borrowed. The firm made this decision primarily based on its belief, which I shared, that commercial paper tended to be confidence-sensitive, and could become unavailable at a time of market stress, while secured borrowing based on high-quality collateral is generally less credit sensitive and therefore more stable.

Bear Stearns implemented this strategy in late 2006 and 2007, and succeeded in reducing its short-term unsecured financing from \$25.8 billion at the end of fiscal 2006 to \$11.6 billion at the end of fiscal 2007, and specifically reduced its commercial paper borrowing from \$20.7 billion to \$3.9 billion. That funding was replaced by secured funding, principally repo borrowing. . .

As part of the firm's transition away from unsecured borrowing, Bear Stearns also substantially increased the average term of its secured funding during the first half of 2007. Bear Stearns was able to obtain longer term repo facilities of six months or more to finance assets such as whole loans and non-agency mortgage-backed securities, and generally limit its use of short-term secured funding to finance Treasury or agency securities. By increasing the amount of its long-term secured funding, the firm believed that it could better withstand a liquidity event.

From approximately August 2007 to the beginning of 2008, however, the fixed income repo markets started experiencing instability, in which fixed income repo lenders began shortening the duration of their loans and asking all borrowers to post higher quality collateral to support those loans. Although the firm was successful in obtaining some long-term fixed income repo facilities, by late 2007 many lenders, both traditional and nontraditional, were showing a diminished willingness to enter into such facilities.

During the week of March 10, 2008, Bear Stearns suffered from a run on the bank that resulted, in my view, from an unwarranted loss of confidence in the firm by certain of its customers, lenders, and counterparties. In part, this loss of confidence was prompted by market rumors, which I believe were unsubstantiated and untrue, about Bear Stearns' liquidity position. Nevertheless, the loss of confidence had three related consequences: prime brokerage clients withdrew their cash and unencumbered securities at

a rapid and increasing rate;⁶ repo market lenders declined to roll over or renew repo loans, even when the loans were supported by high-quality collateral such as agency securities; and counterparties to non-simultaneous settlements of foreign exchange trades refused to pay until Bear Stearns paid first... [T]his loss of confidence in Bear Stearns... resulted in a rapid flight of capital from the firm that could not be survived.

Case Study: JPMorgan Chase's Repo Exposure to Lehman Brothers

The counterparty risk of lending money through a repo is that the borrower defaults and the value of the collateral turns out to be insufficient to cover the loan amount. Any sterile discussion of the topic cannot do justice to the bare-knuckle fighting over collateral that takes place when a counterparty is at risk of default. A striking and well-publicized example of this through 2008 was the repo exposure of JPMorgan Chase (JPM) to Lehman Brothers.

JPM was Lehman's tri-party repo clearing agent. When repo investors lend money to a financial institution through the tri-party repo system,⁷ taking collateral as security, their loans are, literally, overnight.⁸ During the day, however, the tri-party repo agent is lending this money to the financial institution on a secured basis.⁹ Furthermore, given the operational structure of the industry, a broker-dealer could not stay in business without its tri-party agent performing this function. Returning to JPM and Lehman, before Lehman's final week, JPM's tri-party lending to Lehman typically exceeded \$100 billion.¹⁰ Furthermore, JPM had historically not taken any haircuts on its tri-party, intraday advances, but began to do so in early 2008. In Lehman's case, JPM phased in haircuts so as to match, by mid-August 2008, the haircuts posted to overnight repo investors.

⁶Unencumbered securities are securities that have not been posted as collateral or otherwise committed. This part of the testimony seems to imply that Bear Stearns relied on customer cash and customer securities (the latter could be posted as collateral to raise funds) to finance other businesses of the firm. In the spectrum of financing stability, customer cash and unencumbered securities are extremely unstable sources of funding as they can be withdrawn without notice at any time.

⁷For a more complete description, see Bruce Tuckman, "Systemic Risk and the Tri-Party Repo Clearing Banks," *Center for Financial Stability Policy Paper*, February 2010. www.centerforfinancialstability.org/research/Tri-Party-Repo20100203.pdf

⁸This is the reason that overnight repo trades are called that and not one-day trades.

⁹As of the time of this writing, an industry task force is working to eliminate this transfer of intraday risk from repo lenders to the tri-party repo agents.

¹⁰See "Written Statement of Barry Zubrow Before the Financial Crisis Inquiry Commission," September 1, 2010, p. 2. fcic-static.law.stanford.edu/cdn_media/fcic-testimony/2010-0901-Zubrow.pdf

Against this background, the following two excerpts describe differing viewpoints of the events of September 2008, the first from a lawsuit filed by the estate of Lehman Brothers Holdings Inc. (LBHI) against JPM¹¹ and the second from the testimony of Barry Zubrow, Chief Risk Officer of JPM, before the Financial Crisis Inquiry Commission.

First, from Lehman's estate:

On the brink of LBHI's bankruptcy, JPMorgan leveraged its life and death power as the brokerage firm's primary clearing bank to force LBHI into a series of one-sided agreements and to siphon billions of dollars in critically-needed assets. The purpose of these last-minute maneuvers was to leapfrog JPMorgan over other creditors by putting itself in the position of an overcollateralized creditor, not just for clearing obligations, but for any and all possible obligations of LBHI or any of its subsidiaries that JPMorgan believed could result from an LBHI bankruptcy. The effect of JPMorgan's actions—taken with the benefit of unparalleled inside knowledge—was devastating. JPMorgan not only took billions of dollars more than it needed from LBHI, but it also accelerated LBHI's freefall into bankruptcy by denying it an opportunity for a more orderly wind-down, costing the LBHI estate tens of billions of dollars in lost value...

In the weeks preceding LBHI's bankruptcy filing, JPMorgan's top management were the ultimate insiders to the evolving crisis, enjoying real-time access to the key decision-makers at the United States Treasury and the Federal Reserve Bank of New York. JPMorgan's investment bankers were also attempting to assist Lehman's primary potential bidder, the Korea Development Bank, and consequently had first-hand knowledge of its intentions regarding a potential acquisition. JPMorgan also had direct access to internal financial information about Lehman, including an opportunity to review and comment on Lehman's presentation to the rating agencies. At one crucial point, JPMorgan was invited to a meeting with Lehman to consider rescue financing proposals, but instead used it as an opportunity to probe Lehman's financial condition and business plans from a risk management perspective. With all of the bank's tentacles encircling the financial crisis at Lehman, JPMorgan was uniquely positioned to capitalize on the opportunities that crisis presented...

¹¹ Lehman Brothers Holdings Inc., and Official Committee of Unsecured Creditors of Lehman Brothers Holdings Inc., against JPMorgan Chase Bank, N.A.

JPMorgan . . . drained LBHI of desperately needed cash by making repeated demands that LBHI increase the amount of collateral payments it posted. In the last four business days before LBHI's Chapter 11 filing, JPMorgan seized \$8.6 billion of cash collateral, including over \$5 billion in cash on the final business day. All the while that JPMorgan was aggressively leveraging its position to grab increasingly more collateral, JPMorgan knew that it was already overcollateralized by billions of dollars.

JPMorgan's . . . unjustified demands for billions in additional collateral, and its refusal to return that collateral in the critical days before LBHI's bankruptcy filing, severely constrained LBHI's liquidity and impeded its ability to pursue and implement alternatives and initiatives that would have resulted in the preservation of billions in value. . . .

Next, from Barry Zubrow of JPM:

Increasing margin requirements [during the course of 2008] . . . did not protect JPMorgan fully from the risks it faced in extending tens of billions of dollars of credit to broker-dealers each morning . . . JPMorgan, unlike any single tri-party investor, took on a broker-dealer's entire tri-party repo book each day. This meant it would face far greater risks in a liquidation scenario. Furthermore, JPMorgan had no assurance that investors would return to fund the broker-dealer in the evening . . . with the cash necessary to repay JPMorgan's intraday advances. Moreover, the haircuts negotiated between investors and the broker-dealers did not, in many cases, fully reflect the liquidation risk for the increasingly large amount of structured, difficult-to-value securities that were being financed through the tri-party repo program. . . .

By late August and early September 2008, Lehman's deteriorating financing condition was becoming increasingly apparent. . . . In addition, it came to light that many of the securities Lehman had pledged to JPMorgan in June were illiquid, structured debt instruments that appeared to have been assigned overstated values. Nevertheless, JPMorgan . . . continued to . . . act on a business-as-usual basis.

But JPMorgan's exposure to Lehman was growing. This included exposure in areas unrelated to tri-party repo clearing. . . . JPMorgan searched for a way to protect itself without triggering a run on Lehman. . . . JPMorgan determined that it could continue to face Lehman in the market if it had \$5 billion in additional collateral . . . [This] was far from sufficient to cover all of JPMorgan's potential exposures to Lehman . . . but JPMorgan believed that it

was an amount that Lehman could reasonably provide . . . Lehman executives agreed to pledge additional collateral, and . . . did not indicate that JPMorgan's request was putting undue pressure on Lehman. . . .

During the second week of September 2008 . . . a broad review of Lehman's collateral securities . . . indicated that some of the largest pieces of collateral pledged to JPMorgan were illiquid, could not reasonably be valued, and were supported largely by Lehman's own credit. . . . When the true nature of Lehman's collateral came to light on September 11, 2008, it became apparent that JPMorgan . . . would need additional collateral if it were to continue supporting Lehman. JPMorgan decided that \$5 billion in cash was . . . appropriate . . . even though its potential collateral shortfall was greater, as it was a number that JPMorgan believed Lehman could handle. . . . Later that night, JPMorgan sent Lehman a letter stating that, if Lehman did not post the collateral by the open of business the next day, JPMorgan would exercise its right to decline to extend credit to Lehman. . . . Lehman delivered \$5 billion of cash collateral during the morning and early afternoon [of September 12]. . . .

Throughout all of this . . . JPMorgan continued to make enormous—discretionary—extensions of credit to the ailing bank, and it continued to trade with Lehman. . . .

GENERAL AND SPECIAL REPO RATES

As mentioned earlier in this chapter, repo trades can be divided into those using general collateral (GC) and those using special collateral or *specials*. In the former, the lender of cash is willing to take any particular security, although the broad categories of acceptable securities might be specified with some precision. In specials trading, the lender of cash initiates the repo in order to take possession of a particular security. For these trades, therefore, it makes more sense to say that “counterparty A is lending a security to counterparty B and taking cash as collateral” as opposed to saying that “counterparty B is lending cash and taking a security as collateral,” although the two statements are economically equivalent. For this reason, by the way, when using the words “borrow” or “lend” in the repo context, it is best to specify whether cash or securities are being borrowed or lent. Also, as another note on market terminology, bonds most in demand to be borrowed are said to be *trading special*, although any request for specific collateral is a specials trade.

Each day there is a GC rate for each bucket of collateral and each repo term. The most commonly cited GC rates are for repos where any U.S.

Treasury collateral is acceptable, and “the” GC rate refers to the overnight rate for U.S. Treasury collateral. With respect to special rates, there can be one for each security for each term, e.g., the 3 $\frac{5}{8}$ s of August 15, 2019, to September 30, 2010. But every special rate is typically less than the GC rate: being able to borrow cash at a relatively low rate induces holders of securities that are in great demand to lend those securities, while being forced to lend cash at a relatively low rate allocates securities that are in great demand to potential borrowers of that security. Differences between the GC rate and the specials rates for particular securities and terms are called *special spreads*.

Relating GC and specials trades to the market participants discussed earlier in this chapter, GC trades suit repo investors: they obtain the highest rate for the collateral they are willing to accept. Traders intending to short particular securities have to do specials trades and must decide whether they are willing to lend money at rates below GC rates in order to borrow those securities. Funding trades are predominantly GC. Should an institution find itself wanting to borrow money against a security that is trading special, however, it will lend that security in the specials market and borrow cash at a rate below GC, rather than financing that security as part of a GC trade.

In the United States the GC rate is typically close to, but below, the federal funds rate. The latter, discussed further in Chapter 15, is the unsecured rate for overnight loans between banks in the Federal Reserve system. By contrast, repo loans secured by U.S. Treasury collateral are safer and should trade at a lower rate of interest. From October 23 through July 1, 2010, for example, the GC rate was, on average, about 16 basis points below the fed funds rate. This fed funds–GC spread can vary, however, with the demand for Treasury collateral. When the U.S. government was running surpluses and paying down debt in the late 1990s and early 2000s, so that U.S. Treasuries were becoming scarcer and expected to become scarcer still, the fed funds–GC spread widened to reflect the decreasing supply of Treasury collateral.

The fed funds–GC spread also widens during times of financial stress. At such times the demand to hold Treasury bonds and to lend cash on Treasury collateral increases as part of flight-to-quality trades. In addition, willingness to lend Treasury bonds in repo declines as market participants fear that collateral may not be returned, either because a counterparty will fail to return collateral or because a counterparty’s counterparty will fail to do so. Shortly after the collapse of Bear Stearns, for example, after the Board of Governors of the Federal Reserve System had hurriedly lowered its target for the fed funds rate to 2.25%, GC traded at below .50%. Similarly, extremely wide spreads prevailed in the months after the failure of Lehman Brothers.

Special rates for a particular issue to a particular date are determined by the demand of borrowing that issue to that date relative to the supply available. This statment is obvious in some ways, but the important point is that the demand and supply to borrow and lend issues is not the same as the

demand and supply to buy and sell issues. In fact, because some owners of U.S. Treasuries, for institutional reasons, do not lend bonds in repo markets, the amount of a particular issue available for borrow might be somewhat less or very much less than the amount outstanding, depending on the distribution of ownership of that issue across various types of institutions. Put another way, a bond that trades rich relative to neighboring bonds, in the sense of the metrics of Part One, implies a high demand to own that bond relative to the outstanding supply. But the bond may or may not trade very special in repo depending on the extent traders want to short it relative to the supply available for borrow. Given this reasoning, predicting the special spreads of individual bonds is quite difficult. Having said that, there is one predominant explanatory factor for special spreads in the United States, namely, the auction cycle: the most recently issued bonds of each maturity trade special. This is the topic of the next subsection.

Special Spreads in the United States and the Auction Cycle

As mentioned in Chapter 1, the U.S. government sells bonds of different maturities according to a fixed schedule. As of this writing, for example, a new 10-year issue is sold every three months. The most recently issued bond of a given maturity is called the *on-the-run* (OTR) or *current* issue while all other issues are called *off-the-run* (OFR). However, the second most recently issued bond of a given maturity does have its own designation as the *old* issue; the third most recent as the *double-old* issue; etc. As a general rule, at each maturity, the OTR trades the most special, followed by the old, followed by the double-old, etc. Table 12.1 lists the more recent 10- and 30-year U.S. Treasury bonds along with representative overnight repo rates and spreads as of May 28, 2010. The special spreads equal the GC rate minus the respective bond repo rates.

Table 12.1 illustrates how the more recently issued bonds at each maturity trade more special. The table also shows that the OTR 10-year trades more special than the OTR 30-year, a regularity that has been true for some time. The discussion now turns to why special rates are related to the auction cycle.

Current issues tend to be the most liquid.¹² This means that their bid-ask spreads are particularly low and that trades of large size can be conducted

¹² This effect is particularly pronounced in the United States. In Germany, the deliverability of a bond into highly liquid futures contracts is the best determinant of liquidity. See Jacob W. Ejsing and Jukka Sihvonen, "Liquidity Premia in German Government Bonds," European Central Bank Working Paper Series, no. 1081, August 2009. In Japan, liquidity characteristics develop from a mix of the auction cycle and futures contract deliverability.

TABLE 12.1 Special Repo Spreads for Selected, Recently Issued U.S. Treasury Bonds as of May 28, 2010

General Collateral Rate		.23%			
Coupon	Maturity	Issue Date	Designation	Repo Rate	Spread
$3\frac{1}{2}\%$	5/15/20	5/17/10	OTR 10yr	.09%	.14%
$3\frac{5}{8}\%$	2/15/20	2/16/10	Old 10yr	.21%	.02%
$3\frac{3}{8}\%$	11/15/19	11/16/09	Dbl-Old 10yr	.21%	.02%
$4\frac{3}{8}\%$	5/15/40	5/17/10	OTR 30yr	.18%	.05%
$4\frac{5}{8}\%$	2/15/40	2/16/10	Old 30yr	.20%	.03%
$4\frac{3}{8}\%$	11/15/39	11/16/09	Dbl-Old 30yr	.21%	.02%

relatively quickly. This phenomenon is partly self-fulfilling. Since everyone expects a recent issue to be liquid, investors and traders who require liquidity flock to that issue and thus endow it with the anticipated liquidity. Also, the dealer community, which trades as part of its business, tends to own a lot of a new issue until it *seasons* and is distributed to buy-and-hold investors. As a matter of historical interest, the OTR 30-year bond had been such a dominant issue in terms of liquidity that traders called it “the bond.” This nickname persists to this day despite the decline of the bond’s importance relative to that of shorter maturities, in particular of the 10-year.

The extra liquidity of newly issued Treasuries makes them ideal candidates not only for long positions but for shorts as well. Most shorts in Treasuries are for relatively brief holding periods: a trading desk hedging the interest rate risk of its current position; a corporation or its underwriter hedging an upcoming sale of its own bonds; or an investor betting that interest rates will rise. All else being equal, holders of these relatively brief short positions prefer to sell particularly liquid Treasuries so that, when necessary, they can cover their short positions quickly and at relatively low transaction costs.

Investors and traders who are long an OTR bond for liquidity reasons require compensation if they are to sacrifice that liquidity by lending that bond in the repo market. At the same time, investors and traders wanting to short the OTR securities are willing to pay for the liquidity of shorting these particular bonds when borrowing them in the repo market. As a result, OTR securities tend to trade special.

The auction cycle is an important determinant not only of which bonds trade special, but also of how special individual bonds trade over the course of the auction cycle. This will be illustrated first by examining the history of special spreads for the OTR 10-year Treasury and then by examining the




FIGURE 12.4 OTR 10-Year Special Spread, July 1997 to July 2010, Part I




FIGURE 12.5 OTR 10-Year Special Spread, July 1997 to July 2010, Part II

term structure of special spreads for the OTR 10-year Treasury as of May 28, 2010.

Figures 12.4 through 12.6 show the history of the OTR 10-year Treasury special spread from July 1997 to July 2010; the 13-year history is broken up into three graphs for better readability.¹³ The vertical lines indicate

¹³ Note that data from the aftermath of the Lehman bankruptcy, from November 2008 to February 2009, is missing from Figure 12.5. Events at that time will be discussed in the next subsection.




FIGURE 12.6 OTR 10-Year Special Spread, July 1997 to July 2010, Part III

10-year Treasury auctions. These are either auctions of new OTR securities, in which case the OTR security changes over the vertical line, or *re-openings* of existing OTR securities (i.e., auctions that increase the size of an already existing issue), in which case the same security is featured on both sides of the vertical line.

Several lessons may be drawn from these graphs. First, special spreads are quite volatile on a daily basis, reflecting supply and demand for special collateral each day. Second, special spreads can be quite large: spreads of hundreds of basis points are quite common. Third, special spreads do attain higher levels over some periods rather than others, a feature that will be discussed in the next subsection. Fourth, and the main theme of this subsection, while the cycle of OTR special spreads is far from regular, these spreads tend to be small immediately after auctions and to peak before auctions. It takes some time for a short base to develop. Immediately after an auction of a new OTR security, shorts can stay in the previous OTR security or shift to the new OTR. This substitutability tends to depress special spreads. Also, the extra supply of the OTR security immediately following a re-opening auction tends to depress special spreads. In fact, a more detailed examination of special spreads indicates that re-opened issues do not get as special as do new issues. In any case, as time passes after an auction, shorts tend to migrate toward the OTR security, and its special spread tends to rise. Furthermore, as many market participants short the OTR to hedge purchases of the to-be-issued next OTR, the demand to short the OTR and, therefore, its special spread, can increase dramatically or spike going into the subsequent auction.




FIGURE 12.7 Average OTR 10-Year Special Spread as a Function of the Auction Cycle, July 1997 to July 2010

While not shown in these figures, the special spread of the 5-year OTR behaves quite similarly to the 10-year. The pattern of spreads of shorter-maturity OTRs is similar although these spreads tend not to be nearly so wide. This difference is primarily due to the more frequent issuance of shorter-maturity Treasuries that prevents a particular issue from becoming far and away the most liquid bond or most-favored short. Finally, the 30-year had historically been as liquid and its special spreads as large and volatile as that of the 10-year, but this has not been the case since the years leading up to a discontinuation of 30-year bond issuance in 2001. Subsequently, apart from some very active specials trading surrounding the announcement of the re-introduction of “the bond” in 2005 and its sale in 2006, the specials spread of the OTR 30-year has been quite muted relative to those of shorter maturities.

Because special spreads in Figures 12.4 through 12.6 are so volatile, Figure 12.7 reports the average special spread as a function of the auction cycle. The horizontal axis represents time into an auction cycle, measured as the days since the issue of the 10-year OTR divided by the total number of days between issue dates. The curve gives the average of the special spread across cycles of the 13-year history depicted in Figures 12.4 through 12.6. As expected, the average special spread increases over the cycle, spiking as the subsequent auction approaches.

The auction-driven pattern of special spreads can be seen not only from historical data but also from the term structure of special spreads of an individual issue. Table 12.2 gives the spot and forward term structure of special spreads for the OTR 10-year Treasury as of May 28, 2010. The terms

TABLE 12.2 Term Structure of Special Rates and Spreads for the 10-Year, OTR U.S. Treasury as of May 28, 2010. TYM0, TYU0, and TYZ0 Are the Tickers of the Relevant 10-Year Futures Contracts

Term	Term Date	Term Days	Term Rate	Term GC	Term Spread	Forward Spread
Overnight	6/1/10	4	.09%	.23%	.14%	.14%
1 Week	6/4/10	7	.00%	.23%	.23%	.35%
TYM0	6/30/10	33	−.10%	.22%	.32%	.34%
2 Months	7/28/10	61	−.05%	.23%	.28%	.23%
3 Months	8/30/10	94	.00%	.25%	.25%	.20%
TYU0	9/30/10	125	.04%	.26%	.22%	.13%
TYZ0	12/31/10	217	.17%	.30%	.13%	.01%

listed are representative of commonly-traded terms for OTR issues. These typically include fixed terms from the pricing date (e.g., one month) and expiration dates of the relevant futures contracts. The latter trade because many market participants are interested in OTR *basis trades*, i.e., trading the OTR against the futures contract into which it is deliverable. (See Chapter 14.) Note that the overnight rate is the business day following the pricing date.

The term spreads in Table 12.2 are simply the differences between the term GC rates and the respective term special rates. For this table the forward spreads are computed from the term spreads, but forward repo trades do exist. In any case, to illustrate the calculation, the forward special spread from June 30, 2010, to July 28, 2010, is such that investing at the spread to June 30 and then at forward spread from June 30 to July 28 is equivalent to investing to July 28. Let s^{fwd} be this forward spread. Then, using the numbers supplied in the table, s^{fwd} is approximated (see “Characteristics of Spot, Forward, and Par Rates” in Chapter 2) by

$$33 \times .32\% + (61 - 33) \times s^{fwd} \approx 61 \times .28\% \\ s^{fwd} \approx .23\% \quad (12.2)$$

The projected (and realized) 10-year auction schedule as of May 28, 2010, was a re-opening of the current $3\frac{1}{2}$ s of May 15, 2020, both in the middle of June and July, to be followed by the issue of a new OTR in the middle of August 2010. In light of the discussion in this subsection and the historical evidence, the special spread would be expected to increase into these auctions. According to the implied forward spreads, the spread is projected to increase into the June re-opening. The $3\frac{1}{2}$ s are projected to stay special into and somewhat past the July and August auctions as well, but, for the period September 30 to December 31, the forward special spread is

only one basis point. In other words, by the time the then-current 10-year has been around for a month, the specialness of the $3\frac{1}{2}$ s is projected to have dissipated.

Special Spreads in the United States and the Level of Rates

By graphing special spreads rather than special rates, Figures 12.4 through 12.6 hide a factor that has historically limited special spreads. Until very recently, there was no explicit penalty for a *fail*, i.e., for failing to deliver a bond that had been sold. This has implied that the special rate could not fall below 0%. Reason is as follows. If a trader had shorted the OTR 10-year and failed to deliver upon settlement, the trader would not receive the cash from the sale and, consequently, would lose one day of interest on that cash. But what if the trader could borrow the bond overnight in the repo market at 0%, i.e., lend money at 0%, so as to be able to make delivery? The economics of that borrow to the trader is the same as failing; in both cases no interest is earned on the proceeds from selling the bond. Therefore, because no trader would borrow the bond if the special rate were 0% or less, the special rate should never be less than 0%. Equivalently, the special spread should not exceed the GC rate.¹⁴

Figure 12.8 superimposes the Fed's target rate for fed funds on the overnight, 10-year special spread. Clearly, over all but the most recent period, the special spread has been limited by the level of rates. The level of rates, therefore, is part of the explanation for the periods of relatively high and relatively low special spreads observed in this figure and in Figures 12.4 through 12.6.

In 2009, however, the treatment of fails changed. In October and November 2008, as part of the reaction to the Lehman bankruptcy that September, fails to deliver the 10-year OTR climbed to record levels, \$5.311 trillion in the week ending October 22, relative to a pre-crisis average of \$165 billion.¹⁵ Regulators were extremely unhappy with the situation as it was viewed as a threat to the liquidity and efficiency of the U.S. Treasury market. With their prodding, an industry group called the Treasury Market Practices Group adopted a penalty rate for fails, which took effect on May

¹⁴ This is not strictly true because there are such non-monetary costs of fails as regulatory capital requirements. For a case study on negative OTR 10-year special rates in the second half of 2003, see "Repurchase Agreements with Negative Interest Rates," by Michael Fleming and Kenneth Garbade, *Current Issues in Economics and Finance*, Volume 10, Number 5, April 2004. www.newyorkfed.org/research/current_issues/ci10-5/ci10-5.html

¹⁵ Liz Capo McCormick, "Treasury Traders Paid to Borrow as Fed Examines Repos," Bloomberg, November 24, 2008.




FIGURE 12.8 OTR 10-Year Special Spread and the Fed Funds Target Rate, July 1997 to July 2010

1, 2009, equal to the greater of 3% minus the fed funds target rate or zero. Essentially, when the fed funds rate is near zero, the penalty is near 3%, i.e., failing to deliver \$100 million of a bond costs $\$100\text{mm} \times \frac{3\%}{360}$ or \$8,333 per day. As the fed funds rate increases, the penalty falls. The logic there is that since higher interest rates are typically associated with higher opportunity costs of failing, high penalties are not necessary in high-rate environments to prevent episodes of system-wide fails.

In light of the imposition of a penalty for failing to deliver, the new upper limit for the special spread is the penalty rate rather than the GC rate. In fact, soon after the imposition of the penalty, demand to short the OTR Treasury in June 2009 drove the special spread up to this limit. This episode can be seen to the far right of Figure 12.8.

Valuing the Financing Advantage of a Bond Trading Special in Repo

Chapter 1 showed that the OTR 10-year, the $3\frac{1}{2}$ s of May 15, 2020, was trading extremely rich on May 28, 2010, slightly more than 2 per 100 face value, relative to the C-STRIPS curve. OTR bonds often trade at a premium that is in part due to their liquidity advantages, i.e., the ability to turn positions in these bonds back into cash with minimum effort, even in a crisis, and in part due to their financing advantages, i.e., the ability to lend these bonds and borrow cash at a relatively low rate. It is obvious from Table 12.2 that, in the case of the $3\frac{1}{2}$ s, almost all of the premium is due to

liquidity. Nevertheless, it is useful here and even more so in other situations to translate special spreads into price or yield premiums.

The financing value of a bond is the value, over the entire life of the bond, of lending it in repo, borrowing cash at its special rate, and investing that cash at the higher GC rate. The key assumption then, is how special the bond will trade and for how long. Professional repo traders have an opinion about how the special spreads of particular issues will evolve over time that can be used in this analysis. Another approach is to accept the market's view as expressed in the term structure of special spreads. According to Table 12.2, it is reasonable to assume that the $3\frac{1}{2}$ s will trade as GC past September 30, 2010: the forward special rate from then to December 31 is only one basis point. Also, there is no reason to expect that the issue will ever in its life trade special again. Hence, the financing value of the bond is its financing value over the 125 days from the pricing date, May 28, 2010, to September 30, 2010. But, as will now be shown, this financing value can be easily calculated from the term special spread of .22% to September 30, 2010.

The value of lending 100 of cash at a spread of .22% for 125 days is simply

$$100 \times \frac{125 \times .22\%}{360} = .076 \quad (12.3)$$

or 7.6 cents per 100 market value of the bond. At a price of 101.90, therefore, assuming no haircuts, the financing advantage of the $3\frac{1}{2}$ s is worth only about 7.7 cents per 100 face amount, a very small part of its total premium of over 2 dollars. Finally, to translate the dollar value of specialness into a yield value, simply divide by the *DV01*. In this case, with the *DV01* of the $3\frac{1}{2}$ s approximately equal to .085, the value of the special spread is $\frac{.077}{.085}$ or only about .9 basis points.

Forwards and Futures: Preliminaries

Chapter 1 described the determination of a price for spot settlement, that is, for the immediate exchange of an asset for cash. This chapter begins by describing the determination of a price for forward settlement, that is, for the exchange of an asset for cash at some future date. In swap markets, explicit forward agreements are common. In bond markets—with one caveat to be made in a moment—explicit forward contracts are rare but understanding them is important nonetheless. First, spot and repo positions are very commonly combined to create the economic equivalent of a forward contract on a bond. Second, futures contracts on bonds and rates, which are enormously important in fixed income markets (and the subjects of Chapters 14 and 15, respectively), are best understood as variants of forward contracts. Third, it is often useful for technical reasons to think of a fixed income option as an option on an appropriately defined forward position rather than as an option on a spot position.

Having said that explicit forward contracts are rare in bond markets, it must be noted that in one sense this is not true at all. Almost all bond trades that are thought of as spot transactions are actually for settlement one or two days after the trade date. Strictly speaking then, almost all bond trades are forward trades! In practice (and in Chapter 1), however, discount factors extracted from the market prices of normally settled trades are simply defined as spot discount factors.

After discussing forward prices, forward swap rates, and forward yields, this chapter continues the preparation for Chapters 14 and 15 by describing the daily resettlement feature of futures contracts and by showing why and how daily resettlement differentiates futures contracts and futures prices from forward contracts and forward prices.

FORWARD CONTRACTS AND FORWARD PRICES

Forward contracts are agreements to settle the trade of an asset at some date in the future so that the purchase price and the asset change hands at that future date. Importantly, however, the price at which future settlement takes place is fixed at the initiation of the forward contract. Looking ahead to the example of this section, on May 28, 2010, a trader might commit to purchase the U.S. Treasury $3\frac{5}{8}$ s due August 15, 2019, on September 30, 2010, at a price of 101.71. In this example, the initiation or trade date of the forward contract is May 28, 2010; the *underlying security* is the $3\frac{5}{8}$ % Treasury bond; the *forward date*, *expiration date*, *delivery date*, or *contract maturity date* is September 30, 2010; and the *forward price* is 101.71. A trader committing to purchase the bond on the forward date is the *buyer* of the contract and is *long* the forward, while the trader on the other side of the transaction is the *seller* of the contract and is *short* the forward.

By definition, the forward price is such that the buyer and seller are willing to enter into the forward agreement without any initial exchange of cash. This implies that the initial value of the forward contract is zero. Over time, however, the value of the forward position may rise or fall. Any increase in the price of the underlying security benefits the buyer of the forward contract, who has committed to purchase the security at a price that, in retrospect, seems relatively low. In this situation, the value of the forward contract to the buyer would be positive, meaning that the seller would have to pay the buyer to exit the contract. Conversely, the seller of the forward contract would benefit from any decrease in the price of the underlying security, after which the value of the forward would be negative and the buyer would have to pay the seller to exit the contract. It is important to emphasize that the value of a forward contract is conceptually different from the forward price: the former is the current market value of an existing contract while the latter is the forward settlement price at which a new contract can be struck without any initial exchange of cash.

The forward price of a bond to a particular settlement date, relative to its spot price and its repo rate to that settlement date, can be determined by an arbitrage argument. Specifically, the following two strategies both result in the purchase of the bond on the forward settlement date with no other exchange of cash: 1) buying a forward; 2) purchasing the bond for spot settlement and selling its repo to the forward settlement date, i.e., borrowing its purchase price through that date. By arbitrage, then, the forward price must be such that the cost of purchasing the bond on the forward settlement date is the same in strategies 1) and 2).

This arbitrage argument will be illustrated with a forward agreement to purchase the U.S. Treasury $3\frac{5}{8}$ s of August 15, 2019, as of May 28, 2010, for delivery on September 30, 2010. This particular example appears again in Chapter 14 since the $3\frac{5}{8}$ s are deliverable into the ten-year note




FIGURE 13.1 Forward Agreement on May 28, 2010, to Purchase the U.S. Treasury $3\frac{5}{8}$ s of August 15, 2019, for Delivery on September 30, 2010

futures contract expiring on September 30, 2010. In any case, the necessary background information is supplied in Figure 13.1 and Table 13.1. (Note that the accrued interest calculation of this bond as of spot settlement was worked out in Chapter 1.)

Consider the following set of trades which, as will soon be apparent, are equivalent to a forward agreement:

On May 28, 2010 (for settlement on June 1, 2010):

- Buy the $3\frac{5}{8}$ s of August 15, 2019.
- Sell the repo:
 - Borrow the full price of the bond, $102.8125 + 1.0615$ or 103.8740 .
 - Deliver the bond as collateral against the loan.
- No cash is generated or required.

On August 15, 2010:

- Repo loan balance has grown to $103.8740 \left(1 + \frac{.3\% \times 75}{360}\right)$ or 103.9389 .
- Apply the bond's coupon payment of $\frac{1}{2} \times 3\frac{5}{8}$ or 1.8125 to reduce this loan balance to 102.1264 .
- No cash is generated or required.

TABLE 13.1 Selected Data for Calculating the Forward Price of the $3\frac{5}{8}$ s of August 15, 2019, as of May 28, 2010, for Delivery on September 30, 2010

Bond	$3\frac{5}{8}$ s of 8/15/19
Spot Price:	102.8125
Repo Rate to Fwd Settlement:	.3%
Accrued Interest as of Spot Settlement:	1.0615
Accrued Interest as of Fwd Settlement:	.4531

On September 30, 2010:

- Pay off the loan balance of $102.1264 \left(1 + \frac{.3\% \times 46}{360}\right)$ or 102.1655.
- Take back the bond given as collateral.
- Effectively, buy the bond on the forward settlement date for a full price of 102.1655 or, equivalently, for a flat price of 101.7124 plus accrued interest of .4531.

Note that all of the quantities in these trades are known as of the pricing and trade date, May 28, 2010. In summary then, by buying the $3\frac{5}{8}$ s for spot settlement and selling the repo to the forward settlement date, a trader locks in a price of 101.7124 for buying the bond on the forward settlement of September 30, 2010. In other words, the arbitrage-free forward price is 101.7124.

Let $p_0(d)$ denote the (flat) forward price of the bond where spot transactions settle on date 0 and forward settlement is d days later. In the example, date 0 is June 1, 2010, and delivery, on September 30, 2010, is 121 days later. Algebraically, then, the forward price in the example can be written as follows:

$$p_0(121) = \left[(102.8125 + 1.0615) \left(1 + \frac{.3\% \times 75}{360} \right) - \frac{3\frac{5}{8}}{2} \right] \left(1 + \frac{.3\% \times 46}{360} \right) - .4531 \quad (13.1)$$

The derivation of the arbitrage forward price can, of course, be made more general. To this end, define the following variables:

- p_0 : price for 100 face amount for spot settlement on date 0
- c : coupon rate, so $100c$ is the coupon payment per 100 face amount
- d_1 : the number of days from spot settlement to the coupon date
- d_2 : the number of days from the coupon date to forward settlement
- $d = d_1 + d_2$: the number of days from spot to forward settlement
- $AI(\cdot)$: accrued interest for 100 face amount as a function of days from spot settlement.¹
- r : the repo rate from the spot to forward settlement dates

¹Defined this way, implicit in the accrued interest function, $AI(\cdot)$, is the coupon schedule of the bond. For illustration, in the example of the text, $AI(0) = 1.0615$ and $AI(121) = .4531$.

Then, the forward price is

$$p_0(d) = \left[\{p_0 + AI(0)\} \left(1 + \frac{rd_1}{360}\right) - \frac{100c}{2} \right] \left(1 + \frac{rd_2}{360}\right) - AI(d) \quad (13.2)$$

The interpretation of the forward price is a bit easier to see in terms of full prices. Denote the full forward and spot prices as $P_0(d) = p_0(d) + AI(d)$ and $P_0 = p_0 + AI(0)$, respectively, to rewrite equation (13.2) as

$$\begin{aligned} P_0(d) &= \left[P_0 - \frac{\frac{100c}{2}}{1 + \frac{rd_1}{360}} \right] \left(1 + \frac{rd_1}{360}\right) \left(1 + \frac{rd_2}{360}\right) \\ &\approx \left[P_0 - \frac{\frac{100c}{2}}{1 + \frac{rd_1}{360}} \right] \left(1 + \frac{rd}{360}\right) \end{aligned} \quad (13.3)$$

where the second line of (13.3) neglects the very small interest on interest term. In words, the full forward price is the future value, to the delivery date, of the full spot price less the present value of the interim coupon payment.

If there are several coupon payments between the spot and forward settlement dates, the sum of the present value of these coupon payments has to be deducted from the full spot price in equation (13.3). If there is no coupon payment between spot and forward settlement, the full forward price is just the future value of the full spot price. The arbitrage proof of these cases is left as an exercise. For convenience, though, the full forward price with no intermediate coupon is given here:

$$P_0(d) = P_0 \left(1 + \frac{rd}{360}\right) \quad (13.4)$$

A final special case worth mentioning is that, on the delivery date, spot and forward delivery mean the same thing so the forward price equals the spot price. Mathematically, in (13.4), when $d = 0$, $P_0(0) = P_0$.

Turning from the determination of a forward price to the calculation of the value of a forward contract, suppose that the spot price of the $3\frac{5}{8}$ s jumped from 102.8125 to 103 immediately after a long bought the bond forward at the forward price just computed, 101.7124. Since nothing but the spot price has changed, replace the old with the new spot price in equation (13.1) to get a new market forward price of 101.9002. This means that the value of the long's existing contract, with its forward contract price of 101.7124, is worth $101.9002 - 101.7124$ or .1878 as of the delivery date. To see this, note that as of the forward delivery date new longs in the forward contract have to pay 101.9002 while the existing long has to pay only 101.7124. Alternatively, imagine that the existing long sells a forward contract after

the price jumps to its new level. Then, on the forward settlement date, that long will have to sell one bond at a price of 101.9002 and buy one bond at a price of 101.7124, for a profit on that date of the .1878 difference. In any case, given the value of the existing forward contract as of the forward delivery date, simply discount to the present to get the value of the existing contract as of spot settlement. In the example at hand it makes sense to discount the forward delivery date contract value of .1878 by the .3% repo rate for 121 days to get a present value of .1876.

More generally, denoting the forward price of an existing contract by $\bar{P}$, and preserving the notation $P_0(d)$ for the current market forward price for delivery in d days, the value of the existing contract to the long as of the spot settlement date is

$$\frac{P_0(d) - \bar{P}}{1 + \frac{rd}{360}} \quad (13.5)$$

THE FORWARD DROP AND CASH CARRY

In the example of the previous section, the forward price of 101.7124 is less than the spot price of 102.8125. As it turns out, forward prices are usually less than spot prices and the phenomenon is commonly known as the *forward drop*.

To understand the intuition behind the forward drop, imagine that a trader has funds equal to the spot price and wants to own the bond as of the forward settlement date. There are two possible strategies, which, by arbitrage, have to be equally appealing: 1) use the funds to buy the bond spot and accrue coupon interest to the forward date; 2) enter into an agreement to buy the bond forward and invest the funds at the repo rate until forward settlement. If the coupon interest from strategy 1), the spot-purchase strategy, exceeds the repo interest from strategy 2), the forward-purchase strategy, then the two strategies will be equally appealing only if the forward price is less than the spot price. Conversely, if the repo interest from strategy 2) exceeds the coupon interest from strategy 1), then the two strategies will be equally appealing only if the forward price is greater than the spot price.

In theory, then, there is a forward drop when the coupon is high relative to the repo rate but not when the coupon is low relative to the repo rate. Since the term structure of interest rates is usually upward-sloping, however, coupon rates are usually significantly above short-term repo rates. Therefore, in practice, there is usually a forward drop. Put another way, it is usually the case that the spot-purchase strategy has the advantage of earning relatively high coupon income, while the forward-purchase strategy has the advantage of a relatively low purchase price.

This section concludes with a note about cash carry, which, in bond forward and futures markets, is called carry. As introduced in Chapter 3, cash carry in the context of coupon bonds denotes the direct cash flows from holding a bond, namely, coupon interest minus the cost of financing. But rearranging (13.2), and neglecting higher-order terms,

$$p_0 - p_0(d) \approx \frac{100c}{2} + AI(d) - AI(0) - \{p_0 + AI(0)\} \frac{rd}{360} \quad (13.6)$$

In words, the forward drop is equal to coupon interest minus financing costs and, therefore, equal to cash carry. See expression (3.36). Hence, denoting carry for the d days between spot and forward settlements by $\kappa(d)$,

$$\kappa(d) \equiv p_0 - p_0(d) \quad (13.7)$$

FORWARD BOND YIELDS

A bond's forward yield is defined as the single rate such that discounting the bond's cash flows from the forward settlement date to maturity gives that bond's (full) forward price. To illustrate, continue with the $3\frac{5}{8}$ s of August 15, 2019, and the accompanying data from the first section of this chapter. With a full forward price of 102.1655, the forward yield, y , for delivery on September 30, 2010, is²

$$102.1655 = \left(1 + \frac{y}{2}\right)^{1 - \frac{138}{184}} \left[\frac{3.625}{y} \left(1 - \frac{1}{\left(1 + \frac{y}{2}\right)^{18}}\right) + \frac{100}{\left(1 + \frac{y}{2}\right)^{18}} \right] \quad (13.8)$$

Solving, $y = 3.399\%$. The intuition relating forward bond yields to spot yields is omitted as it is analogous to that relating forward swap rates to par spot rates, which is discussed in the next section.

FORWARD SWAP RATES

As mentioned in the introduction, much of the material about forward contracts in this chapter is preliminary to the material to follow on futures. This section stands on its own, however, because forward contracts, rather than futures contracts, trade in the swap market.

Chapter 2 introduced swap contracts for spot delivery, meaning that interest on both legs starts to accrue at the spot settlement date. By contrast,

²The reader may find it useful to refer to Appendix A in Chapter 3.

forward-starting swaps begin accruing interest at some time further in the future. Consider a two-year swap that starts in one year, called a “1-year 2-year swap” or a “1- into 2-year swap” and written as a “1yr2yr swap” or as a “1x2 swap.” The fixed leg of this swap, including its fictional notional payment (see Chapter 2), is analogous to a forward agreement to trade a three-year bond in one year, when it will have become a two-year bond. But in the bond market, where a specific bond is bought or sold for forward settlement, traders speak of trading a three-year bond for settlement in one year. In the swap market, by contrast, where forward starting swaps are created whenever two counterparties elect to do so, there is no need to refer to a spot starting swap and traders speak directly of a two-year swap starting in one year.

As argued in Part One, swaps—unlike bonds—are extremely well priced by arbitrage: because swaps can be created in unlimited supply, there is nothing unique about any particular swap and, therefore, only cash flows matter for pricing. To price a forward swap, then, simply discount the appropriate cash flows, as shown below, and rely on the equivalence of arbitrage and discounting demonstrated in Chapter 1.

Denote the forward price of a T -year swap starting in t years with fixed rate c by $P(t, t + T; c)$, which notation emphasizes that the swap starts at time t and matures at time $t + T$. The forward price is such that both counterparties will agree today to enter into the forward swap at time t without any other exchange of cash. This will be the case only if the present value of the transaction is zero. Mathematically,

$$-P(t, t + T; c) d(t) + \frac{c}{2} [d(t + .5) + \cdots + d(t + T)] + d(t + T) = 0 \quad (13.9)$$

In words, there is no net present value to paying the price $P(t, t + T; c)$ on date t and receiving the swap cash flows, which start at time $t + .5$ and end at time $t + T$. Solving,

$$P(t, t + T; c) = \frac{1}{d(t)} \left\{ \frac{c}{2} [d(t + .5) + \cdots + d(t + T)] + d(t + T) \right\} \quad (13.10)$$

In principle, a forward swap with any rate c could trade in the market so long as the forward price is set as in (13.10). In practice, however, forward swap rates almost always mean *forward par swap rates*, i.e., forward swap rates that correspond to a forward swap price of par. Denoting the txT forward par swap rate by $C(t, t + T)$ and setting $P(t, t + T; c) = 1$, it

follows from (13.10), that

$$1 = \frac{1}{d(t)} \left\{ \frac{C(t, t+T)}{2} [d(t+.5) + \dots + d(t+T)] + d(t+T) \right\} \quad (13.11)$$

For a sample calculation, consider a USD 1x1.5 forward par swap rate, $C(1, 2.5)$, as of May 28, 2010. With $t = 1$ and $T = 1.5$, equation (13.11) becomes

$$1 = \frac{1}{d(1)} \left\{ \frac{1}{2} C(1, 2.5) [d(1.5) + d(2) + d(2.5)] + d(2.5) \right\} \quad (13.12)$$

Then, substituting the discount factors given in Table 2.1 into equation (13.12), $C(1, 2.5) = 1.832\%$.

INTEREST RATE SENSITIVITY OF FORWARDS

What is the interest rate sensitivity or $DV01$ of a forward contract on a bond? Focusing for the moment on yield-based $DV01$, a question practitioners often pose is whether the $DV01$ should be computed with respect to a change in the spot yield or with respect to a change in the forward yield. Figure 13.2 describes this question in terms of the example of this chapter, a forward on the $3\frac{5}{8}\%$ s of August 15, 2019, for delivery on September 30, 2010, with spot settlement on June 1, 2010. Shifting the spot yield shifts the interest rate used for discounting over the period from June 1, 2010, to August 15, 2019. Shifting the forward yield shifts the interest rate used for discounting over the period September 30, 2010, to August 15, 2019. The difference between the two shifts is that the spot yield shift essentially shifts the repo rate by the same amount as the rest of the curve while the forward yield shift essentially leaves the repo rate unchanged.

Before discussing the choice between the two shifts, it will be useful to calculate the two $DV01$ alternatives in this example. Bumping the forward yield on the right-hand side of equation (13.8) down by one basis point




FIGURE 13.2 Various Rate Shifts for the $3\frac{5}{8}\%$ s of August 15, 2019, for Spot and Forward Delivery

gives a new forward price of 102.2424 and, therefore, a forward $DV01$ of $102.2424 - 102.1655$ or .0769.³

To compute the spot $DV01$, first use the information in Figure 13.1 and Table 13.1 to write the spot price-yield equation for this example:

$$103.8740 = \left(1 + \frac{y}{2}\right)^{1 - \frac{75}{181}} \left[\frac{3.625}{y} \left(1 - \frac{1}{\left(1 + \frac{y}{2}\right)^{19}}\right) + \frac{100}{\left(1 + \frac{y}{2}\right)^{19}} \right] \quad (13.13)$$

Solving, $y = 3.268\%$. Then, bump the yield on the right-hand side of 13.13 down by one basis point to get a new spot price of 103.9543 and a $DV01$ of $103.9543 - 103.8740$ or .0803.

The spot $DV01$ at .0803 is larger than the forward $DV01$ at .0769 because, as just explained, the spot $DV01$ essentially shifts the repo rate as well. In fact, the difference between the spot $DV01$ and the forward $DV01$, $.0803 - .0769$ or .0034, is exactly the $DV01$ of the repo: from equation (4.42), the $DV01$ of a zero coupon bond of short maturity is approximately equal to its maturity, in years, divided by 100, which, in the case of this repo, is $\frac{1}{100} \times \frac{121}{360}$ or .0034.

Returning to the question of which $DV01$ should be used, the determinant of the choice is now clear. If one believes that the repo rate moves one-for-one with the rest of the curve, then a shift of the spot yield makes sense; if one believes that the repo rate stays constant, a shift of the forward yield makes sense. While it is certainly true that the repo rate is less volatile than longer term rates, it is an exaggeration to say that it remains constant. As a result, some practitioners would hedge a forward by assuming that the repo rate changes by a fixed fraction of the yield change. For example, assuming that the repo rate moves 40% as much as the forward yield, this approach would give a $DV01$ estimate for the forward contract equal to the $DV01$ of the forward (i.e., .0769) plus 40% of the $DV01$ of the repo (i.e., $40\% \times .0034 = .0014$) or .0783.

The weakness of any of these solutions, that is, using the spot $DV01$, the forward $DV01$, or a blend of the forward and repo $DV01$ s, is that they all assume a fixed relationship between changes in long-term yields and the repo rate. The situation, however, really calls for a two-factor approach: allow the long-term rates to move independently of the repo rate. To see that this is a better approach, consider the practicalities of hedging with one of the three approaches. Assume, for example, that a practitioner has decided to use the blended $DV01$ of .0783 to hedge a long forward on

³Note that computing $DV01$ s with respect to full prices gives the same result as with respect to spot prices since accrued interest does not change as yield changes.

the $3\frac{5}{8}$ s due August 15, 2019, a bond with slightly more than nine years remaining to maturity. The practitioner would then probably choose to sell an appropriate face amount of liquid 10-year bonds. But this means that the exposure to extremely short-term repo rates is being hedged by 10-year rates. Much more sensible would be to hedge the forward yield with the liquid 10-year bond and to hedge the repo risk separately with short-term futures contracts. Chapter 15 describes the latter hedge in detail.

DAILY SETTLEMENT OF FUTURES CONTRACTS

Chapters 14 and 15 will present and analyze bond and note futures and short-term interest rate futures, respectively. The purpose of this and the next two sections is to describe the difference between forward and futures contracts in a more theoretical way, as preparation for the material in those chapters.

The first section in this chapter explained the value of a forward contract over time, using the example of buying a forward on May 28, 2010, to purchase the $3\frac{5}{8}$ s of August 15, 2019, on September 30, 2010, at a price of 101.7124. It was then pointed out that should the forward price jump to 101.9002 immediately after that transaction, the profit to the buyer, as of September 30, 2010, is $101.9002 - 101.7124$ or .1878. And should the forward price rise, over time, to a price of 105, the profit would be $105 - 101.7124$ or 3.2876, again, as of the delivery date.

A futures contract is similar to a forward agreement at initiation in that the futures price is such that the buyer and seller are willing to enter into the futures contract without any exchange of cash. But futures and forwards differ with respect to subsequent cash flows in that futures contracts are subject to *daily settlement*. This means two things. First, at the end of each day, the losing counterparty pays the change in the value of the futures contract to the winning counterparty. Second, the counterparties tear up their prior contract, at the prior futures price, and enter into a new contract at that day's closing futures price. Returning to the example, assume for the moment (though this will be corrected in the next two sections) that forward and futures prices are the same. At the end of the day in which the futures price jumps from 101.7124 to 101.9002, the seller of the contract pays the buyer the difference of .1878. Also, the two counterparties roll their contract into a new contract with a futures price of 101.9002.

The essential difference between forwards and futures contracts then, from a cash flow perspective, is that the profit or loss from a forward contract is realized at the delivery date while the profit or loss from a futures contract is realized over time. If, in the example, the price for forward delivery eventually rises from 101.7124 to 105, the buyer of a forward contract will

buy the bond for 101.7124 on September 30, 2010, and sell it in the market for 105 at a profit, on that date, of 3.2876. The buyer of a futures contract, by contrast, will have collected that 3.2786 over time, as the futures price rose. On the delivery date, after the last daily settlement payment, the buyer of the futures contract has a contract to buy the bond at 105. Doing so, and selling the bond at market for 105, generates no extra profit or loss.

The flip-side of the cash flow perspective is the perspective with respect to contract value. A forward contract accumulates or loses value over time while the value of a futures contract, after each daily settlement, is zero. After the first day in the example, with the forward price increasing from 101.7124 to 101.9002, a long position in the forward contract is, as of the delivery date, worth the difference of .1878. After the same change in a futures price, however, the buyer collects .1878 through the daily settlement payment and is put into a new contract at the prevailing market futures price of 101.9002, which, as a contract at the prevailing market price, is worth zero.

While this and the next two sections describe the cash flow implications of daily settlement, there are also implications for counterparty risk. These implications, however, are confounded by the fact that futures normally trade through a central counterparty while forward contracts normally trade over-the-counter. In any case, since the counterparty risk of derivative contracts is a broader issue than the difference between forwards and futures, the relevant issues are collected and discussed in Chapter 16, in the context of interest rate swaps.

This section closes with some comments on terminology. Two terms are often used interchangeably, although somewhat confusingly, with the term daily settlement. The first such term is *mark-to-market*. Strictly speaking, mark-to-market is the process of adjusting security prices in some accounting framework to match market values. For example, securities in the trading books of a bank have to be marked-to-market when reported on its balance sheet while securities in its “held-to-account” ledgers might be carried at cost. But the term mark-to-market, strictly speaking, does not necessarily refer to any transaction or to any exchange of cash flows between parties to a transaction.

The second term used interchangeably and somewhat confusingly with daily settlement is *variation margin*. Margin refers most generally to cash or security collateral that is posted in order to secure obligations under a contract, meaning that the collateral can be seized in the event of a default. Variation margin, then, refers most generally to collateral given or returned as an adjustment to the collateral already pledged, most usually in response to increased or decreased exposures under a contract. The margin calls described in Chapter 12, which require repo borrowers to post more collateral as the value of existing collateral declines, is an example of variation margin. In any case, in this general meaning of margin, counterparties taking cash

collateral have to pay interest on the amount posted and, when all contract obligations have been fulfilled, have to return collateral to its original owner. By contrast, daily settlement payments of futures are the gains and losses of the economic contract. They are paid irrevocably and no interest is ever paid by a receiving counterparty. Nevertheless, daily settlement payments are often referred to as variation margin. The usage probably arises from the fact that paying futures losses as they occur reduces exposure to trading counterparties, as does the posting of margin.

FORWARD AND FUTURES PRICES IN A TERM STRUCTURE MODEL

For ease of exposition in this section, focus on the case of a forward or futures on a bond with no intermediate coupon payments so that equation (13.4) obtains for the forward price. Also, to allow for a more general presentation, replace the number of days in earlier sections with n periods of unspecified length. Finally, write the discount factor over those n periods as $d(n)$. Then, the forward price in (13.4) becomes

$$P_0(n) = \frac{P_0}{d(n)} \quad (13.14)$$

Begin the analysis of pricing forwards and futures in a term structure model in the context of a binomial tree, along the lines of Part Three. Let there be three dates, labeled 0, 1, and 2, and let the forward and futures delivery dates be on date 2 so that $n = 2$ in (13.14). As in the notation of Part Three, let r_0 denote the initial short-term rate and let r_i^s denote the short-term rate on date i , state s . Finally, let the probability of an up move from date 0 to date 1 be .6 and the probability of an up move from date 1 to date 2 (from any date 1 state) be .5.

Take as given that the prices of a particular security in the three states of date 2 have been computed using the later dates of a risk-neutral process, not described here. These three prices depend, of course, on the different values of the short-term rate on date 2 and are denoted P_2^{uu} , P_2^{ud} , and P_2^{dd} . Using the methods of Part Three, the prices of the security today, P_0 , may be computed backward along the three-date tree, starting from the given prices on date 2. Algebraically,

$$P_0 = \frac{1}{1 + r_0} \left[.6 \times \frac{.5 P_2^{uu} + .5 P_2^{ud}}{1 + r_1^u} + .4 \times \frac{.5 P_2^{ud} + .5 P_2^{dd}}{1 + r_1^d} \right] \quad (13.15)$$

or, rearranging terms,

$$P_0 = .3 \frac{P_2^{uu}}{(1+r_0)(1+r_1^u)} + .3 \frac{P_2^{ud}}{(1+r_0)(1+r_1^u)} \\ + .2 \frac{P_2^{ud}}{(1+r_0)(1+r_1^d)} + .2 \frac{P_2^{dd}}{(1+r_0)(1+r_1^d)} \quad (13.16)$$

Each term of equation (13.16) is the product of a price times the probability of reaching that price all discounted along the path to that price. In the first term, for example, the probability of moving up and then up again to the price of P_2^{uu} is $.6 \times .5$ or $.3$. Discounting P_2^{uu} along that path means discounting using r_0 and r_1^u . More generally, consider a security worth P_n in period n when short-term rates take on the values $r_0, r_1, \dots, r_{n-1}$ from date 0 to date $n-1$. By the logic of (13.16), the price of the security today is given by the following expectation:

$$P_0 = E \left[\frac{P_n}{\prod_{i=0}^{n-1} (1+r_i)} \right] \quad (13.17)$$

where Π is the standard product notation so that $\prod_{i=0}^{n-1} (1+r_i)$ is equal to the product of the terms $(1+r_i)$ from $i=0$ to $n-1$.

In words, equation (13.17) says that the price today equals the expected discounted value of its future value—in particular, of its value on date n . This equation also reveals the reason for using the term expected discounted value rather than discounted expected value.

The discount factor to date 2 implied by the tree is the same, of course, as the price of a zero-coupon bond maturing on date 2. But a zero-coupon bond maturing on date 2 is worth 1 in every state, or, making this a special case of (13.17), $P_n = 1$ for all states. Therefore,

$$d(n) = E \left[\frac{1}{\prod_{i=0}^{n-1} (1+r_i)} \right] \quad (13.18)$$

This completes the derivation of forward rates in a term structure model because, by (13.14), the forward price is the right-hand side of (13.17) divided by the right-hand side of (13.18).

Turning to the derivation of a futures price in a term structure model, continue with the same three-date tree described at the start of this section and with the same bond prices on date 2. Let F_i^s denote the futures price on date i , state s , immediately after the daily settlement of date i . Let F_0

denote the futures price today. Then, to begin, recall that the futures price for immediate delivery is, by definition, the same as the spot price of the security. Hence, at expiration of a futures contract, the futures price equals the spot price at that time and state. Hence, $F_2^{uu} = P_2^{uu}$, $F_2^{ud} = P_2^{ud}$, and $F_2^{dd} = P_2^{dd}$.

As of the up state on date 1, the futures price is F_1^u . If the price of the underlying bond moves to P_2^{uu} on date 2, then that will be the date 2 futures price, and the daily settlement payment on a long position of one contract will be $P_2^{uu} - F_1^u$. Similarly, if the price moves to P_2^{ud} on date 2, then the settlement payment will be $P_2^{ud} - F_1^u$. Since the tree has been assumed to be the risk-neutral pricing tree, the value of the contract in the up state of date 1 must equal the expected discounted value of its cash flows. But, by the definition of futures contracts, the value of a futures contract after its daily settlement payment must equal zero. Putting these two facts together,

$$\frac{.5 \times (P_2^{uu} - F_1^u) + .5 \times (P_2^{ud} - F_1^u)}{1 + r_1^u} = 0 \quad (13.19)$$

Then, solving for the unknown futures price,

$$F_1^u = .5 \times P_2^{uu} + .5 \times P_2^{ud} \quad (13.20)$$

Applying the same logic to the down state of date 1 gives

$$F_1^d = .5 \times P_2^{ud} + .5 \times P_2^{dd} \quad (13.21)$$

Moving to date 0, setting the expected discounted settlement payment equal to zero implies that

$$\frac{.6 \times (F_1^u - F_0) + .4 \times (F_1^d - F_0)}{1 + r_0} = 0 \quad (13.22)$$

Or,

$$F_0 = .6 \times F_1^u + .4 \times F_1^d \quad (13.23)$$

Finally, substituting (13.20) and (13.21) into (13.23),

$$\begin{aligned} F_0 &= .3 \times P_2^{uu} + .3 \times P_2^{ud} + .2 \times P_2^{ud} + .2 \times P_2^{dd} \\ &= .3 \times P_2^{uu} + .5 \times P_2^{ud} + .2 \times P_2^{dd} \end{aligned} \quad (13.24)$$

In words, under the risk-neutral process the futures price equals the expected price of the underlying security as of the delivery date. More generally,

$$F_0 = E[P_n] \quad (13.25)$$

THE FUTURES-FORWARD DIFFERENCE

This section derives a formula for the difference between a forward price and a futures price to deliver the same bond on the same settlement date. Recall that for any random variables, X and W , their covariance can be written as follows:

$$\text{Cov}(X, W) = E[XW] - E[X]E[W] \quad (13.26)$$

Letting $X = P_n$, be the price of the underlying security on the delivery date, and let $W = [\prod_{i=0}^{n-1} (1 + r_i)]^{-1}$, represent the realized discount factor from spot to forward settlement, then (13.26) becomes

$$\text{Cov}\left(P_n, \frac{1}{\prod_{i=0}^{n-1} (1 + r_i)}\right) = E\left[\frac{P_n}{\prod_{i=0}^{n-1} (1 + r_i)}\right] - E[P_n]E\left[\frac{1}{\prod_{i=0}^{n-1} (1 + r_i)}\right] \quad (13.27)$$

In words, this covariance equals the expected discounted value of the price of the underlying minus the discounted expected value of that price. These two quantities are not the same! In any case, substituting the definitions of P_0 , $d(n)$, and F_0 , from (13.17), (13.18), and (13.25), respectively, into (13.27),

$$\text{Cov}\left(P_n, \frac{1}{\prod_{i=0}^{n-1} (1 + r_i)}\right) = P_0 - F_0 d(n) \quad (13.28)$$

Rearranging terms and then using the definition of the forward price from (13.14),

$$\begin{aligned} F_0 &= \frac{P_0}{d(n)} - \frac{1}{d(n)} \text{Cov} \left(P_n, \frac{1}{\prod_{i=0}^{n-1} (1 + r_i)} \right) \\ &= P_0(n) - \frac{1}{d(n)} \text{Cov} \left(P_n, \frac{1}{\prod_{i=0}^{n-1} (1 + r_i)} \right) \end{aligned} \quad (13.29)$$

Since the price of the underlying security on the delivery date is likely to be relatively low if rates from now to then are relatively high and *vice versa*, the covariance term in (13.29) is likely to be positive.⁴ And, if this is indeed the case, it follows that

$$F_0 < P_0(n) \quad (13.30)$$

The intuition behind bond futures prices being lower than their forward prices is as follows. Assume for a moment that futures and forward prices were the same. Daily changes in the value of the forward contract generate no cash flows while daily changes in the value of the futures contract generate daily settlement payments. Daily settlement gains can be invested and losses must be financed, but, on average these effects do not cancel. Bond prices tend to be high when short rates are low, and *vice versa*, so daily settlement gains are invested at low rates while daily settlement losses are financed at high rates. But this means that, at equal prices, the futures contract would be worth less than the forward contract. Therefore, priced properly relative to one another, the futures price is less than the forward price, as in (13.30).

Having completed this analysis, it can now be argued that for commonly traded bond futures contracts, like those discussed in Chapter 14, the size of the futures-forward difference is quite small. The reason for this is that the covariance in (13.29) applies to the period from spot to forward settlement. But as this time period is quite short for the most actively traded contracts, usually from zero to four months, the covariance is small as well. For a very rough order of magnitude calculation, consider the following parameters: an underlying bond with a DV01 of .08, like that of a 10-year bond; an annual interest rate volatility of 100 basis points across the term structure (though this overstates the likely volatility of the short-term rate); a contract maturity date of three months, so that the DV01 of the discount factor to the

⁴This discussion does not necessarily apply to forwards and futures on securities outside the fixed income context. Consider, for example, a forward and a futures on oil. In this case it is more difficult to determine the covariance between the discounting factor and the underlying security.

delivery date is approximately .0025%; and a correlation of one between the price of the bond and the discount factor (though this overstates the likely correlation). With these parameters, the volatility of the bond price to the delivery date is $.08 \times 100 \times \sqrt{.25}$ or 4. The volatility of the discount factor is $.0025\% \times 100 \times \sqrt{.25}$ or .125%. Finally, since the correlation is assumed to be one, the relevant covariance is just the product of the standard deviations, $4 \times .125\%$ or .005. Dividing this covariance by the three-month discount factor gives a futures-forward effect of about half a cent on a bond price of 100 or, in terms of a bond with a DV01 of .08, about $\frac{1}{16}$ of a basis point.

While the futures-forward effect is small for commonly traded bond futures, it is not so small for commonly traded rate futures, since the maturities of the latter can be significantly longer. The next section, therefore, applies the analysis of this section to rate futures.

FORWARD RATES *VERSUS* FUTURES ON RATES

As will become clear in Chapter 15, money market futures are really futures on rates, not on bond prices. Consider Eurodollar (ED) futures contracts, which are discussed in detail in that chapter. Let L_n be the three-month LIBOR, the interest rate referenced by the ED futures contracts as of the delivery date n . Then, the terminal settlement price of the futures, F_0^R , is defined as

$$F_n^R = 100 - 100 \times L_n \quad (13.31)$$

Note that this futures price is not a price in the sense of discounting some payment at some interest rate; it is simply a way of expressing an interest rate.

Following the development of the previous section, the ED futures price as of date 0 is

$$F_0^R = E[100 - 100 \times L_n] = 100 - 100 \times E[L_n] \quad (13.32)$$

Since the 100 terms exist just to make the futures price look like a price, it is more natural to focus on the *futures rate*, r^{fut} , defined as

$$r^{fut} \equiv \frac{100 - F_0^R}{100} = E[L_n] \quad (13.33)$$

where the second equality follows from (13.32). It is also more natural, then, to express the difference between futures and forward contracts not in terms of prices but in terms of the difference between the futures rate, r^{fut} , and the forward rate, r^{fwd} , to be defined presently.

Continuing with the ED futures contract, the underlying security is a 90-day deposit or, equivalently, a zero-coupon bond, under the actual/360 convention. Hence, by definition, the forward price of that zero is related to the forward rate, r^{fwd} , by the pricing equation

$$P_0(n) = \frac{1}{1 + \frac{90r^{fwd}}{360}} \quad (13.34)$$

The ED futures rate in terms of the underlying interest rate dynamics is given by (13.33). To compare the futures and forward rates, note that the futures on the zero coupon price, F_0 , by the logic of the previous section, is

$$F_0 = E[P_n] = E\left[\frac{1}{1 + \frac{90L_n}{360}}\right] \quad (13.35)$$

By an application of Jensen's inequality,

$$E\left[\frac{1}{1 + \frac{90L_n}{360}}\right] > \frac{1}{1 + \frac{90E[L_n]}{360}} \quad (13.36)$$

Finally, combining (13.30) and (13.33) through (13.36),

$$P_0(n) = \frac{1}{1 + \frac{90r^{fwd}}{360}} > F_0 > \frac{1}{1 + \frac{90r^{fut}}{360}} \quad (13.37)$$

This equation shows that the difference between forwards and futures on rates has two separate effects. The first inequality represents the difference between the forward price and the futures on a price. This difference is properly called the futures-forward effect since it arises from the daily settlement of futures contracts. The second inequality represents the difference between a futures on a price and a futures on a rate which, as evident from (13.36), is a convexity effect. The combination of the two effects, expressed as the difference between the observed forward rates on deposits and Eurodollar futures rates, will be referred to as the total futures-forward effect.

It follows immediately from (13.37) that

$$r^{fut} > r^{fwd} \quad (13.38)$$

Accordingly, the futures rate exceeds the forward rate or, equivalently, the total futures-forward difference is positive.

To illustrate the magnitude of the futures-forward effect in rate futures, the text now presents, without proof, a closed-form expression for futures and forward rates from a term structure model. The model used is a normal model with initial rate r_0 , drift λ , and short-rate annual basis-point volatility σ . The rates r^{fut} and r^{fwd} are taken to be the continuously compounded futures and forward rates, respectively, of a β -year zero-coupon bond for delivery in t years. Then, in this special case,

$$r^{fut} = r_0 + \lambda \frac{(t + \beta)^2 - t^2}{2\beta} - \frac{\sigma^2 \beta^2}{6} \quad (13.39)$$

$$r^{fwd} = r_0 + \lambda \frac{(t + \beta)^2 - t^2}{2\beta} - \frac{\sigma^2}{6} \left[\frac{(t + \beta)^3 - t^3}{\beta} \right] \quad (13.40)$$

Subtracting (13.40) from (13.39),

$$r^{fut} - r^{fwd} = \frac{\sigma^2 t^2}{2} + \frac{\sigma^2 \beta t}{2} \quad (13.41)$$

Equation (13.41) supports the comment at the end of the previous section that the futures-forward difference can be large for commonly-traded rate futures. Continuing with the example of ED contracts, the deposit is 90 days, so that β is approximately .25 years, and the impact of the second term is limited. Unlike note and bond futures contracts, however, ED contracts have maturities that extend to 10 years, although only the first two or so years are particularly liquid. Table 13.2 uses $\beta = .25$ and an annual volatility of 100 basis points to compute numerical values of (13.41) at selected maturities. The resulting differences can be quite large, in contrast to

TABLE 13.2 Futures-Forward Difference for Rate Futures on 3-Month Deposits at Selected Maturities calculated under a Normal Term Structure Model with Constant Drift and an Annual Volatility of 100 Basis Points

Expiration Years	Futures-Forward Difference Basis Points
.25	.0625
.5	.1875
1	.6250
2	2.2500
5	13.1250
10	51.2500

the order of magnitude of commonly-traded note and bond futures derived in the previous section.

Equation (13.41) also shows explicitly that the total futures-forward effect increases with interest rate volatility. The pure futures-forward effect arises because daily settlement gains are invested at low rates while daily settlement losses are financed at high rates. With no interest rate volatility there are no daily settlement cash flows and no investment or financing of those flows. The convexity effect depends on volatility as well, as demonstrated in Chapter 8.

TAILS

In several important contexts market participants need to hedge forward contracts with futures contracts. In Chapter 14, this problem surfaces in the construction of basis trades; in Chapter 15 it surfaces when a liquid futures contract is chosen as the hedge for a forward loan. The purpose of this section is to present an approximation, which is commonly used in the industry, for calculating the number of futures contracts required to hedge an otherwise identical forward contract.

Consider a forward and futures contract on the same underlying security for delivery in d days with a term repo rate of r . To develop the approximate hedge of this section, assume that the forward and futures prices change, but that short-term rates, to the term of the delivery date, do not change. Then, by equation (13.29), the forward and futures prices change by the same amount, say Δ . The futures contract, with its daily settlement feature, pays Δ immediately to the long. The change benefits a long in the forward contract only as of the delivery date, so, by (13.5), the value of the forward contract changes by

$$\frac{\Delta}{1 + \frac{rd}{360}} \quad (13.42)$$

Hence, less than one futures contract is needed to hedge the change in value of one forward contract. More precisely, to hedge N^{fwd} forward contracts, with N^{fut} futures contracts, set

$$N^{fut} = \frac{N^{fwd}}{1 + \frac{rd}{360}} \quad (13.43)$$

The difference between N^{fwd} and N^{fut} is known as the *tail* of the hedge, and taking account of this difference when hedging is known as *tailing the hedge*.

Note and Bond Futures

Futures contracts on government bonds are liquid and require relatively little capital to establish sizable positions. Consequently, these contracts are often the instruments of choice for hedging longer-term interest rate risk and for speculating on the direction of these rates.¹

With the theoretical preliminaries of futures contracts and the differences between futures and forward contracts established in Chapter 13, this chapter focuses on the mechanics of U.S. Treasury futures and, in particular, on how the options embedded in these contracts affect valuation and risk. The chapter concludes with a case study that critiques a once-popular trade of a futures contract against an underlying Treasury note.

Futures contracts that trade in Europe and Japan embed only one of the options present in U.S. futures contracts. Therefore, while focused on U.S. Treasury futures, the treatment of this chapter can easily be applied to these international markets.

MECHANICS

This section and the next describe the workings of U.S. note and bond futures contracts. The motivations behind the design of these contracts are explained later in this chapter.

Futures contracts on U.S. government bonds do not have one underlying security. Instead, there is a *basket* of underlying securities defined by some set of rules. The 10-year note contract expiring in September, 2010, for example, with the ticker TYU0, includes as an underlying security any U.S. Treasury note that matures in 6.5 to 10 years as of September 1, 2010. As of May 28, 2010, this rule included all of the securities listed in Table 14.1. The conversion factors listed in the table, as well as the chosen order of the securities, are discussed presently.

¹For a book-length treatment of the subject, see Burghardt and Belton, *The Treasury Bond Basis*, Third Edition, McGraw-Hill, 2005.

TABLE 14.1 The Deliverable Basket into TYU0

Rate	Maturity	Conversion Factor
$3\frac{1}{4}$	3/31/17	.8538
$4\frac{1}{2}$	5/15/17	.9202
$3\frac{1}{8}$	4/30/17	.8471
$2\frac{3}{4}$	5/31/17	.8272
$4\frac{3}{4}$	8/15/17	.9314
$4\frac{1}{4}$	11/15/17	.9012
$3\frac{7}{8}$	5/15/18	.8732
4	8/15/18	.8774
$3\frac{1}{2}$	2/15/18	.8547
$3\frac{3}{4}$	11/15/18	.8587
$3\frac{5}{8}$	8/15/19	.8401
$3\frac{1}{8}$	5/15/19	.8107
$2\frac{3}{4}$	2/15/19	.7909
$3\frac{3}{8}$	11/15/19	.8195
$3\frac{5}{8}$	2/15/20	.8332
$3\frac{1}{2}$	5/15/20	.8210

The seller of a futures contract, or the short, commits to *deliver* a set quantity of a bond in that contract's basket during the *delivery month*. The seller may choose which bond to deliver and when to deliver during the delivery month. These options are called the *quality option* and the *timing option*, respectively. The buyer of the futures contract, or the long, commits to buy or *take delivery* of the bonds chosen by the seller at the time chosen by the seller. In the case of TYU0, with a *contract size* of \$100,000, each contract requires the delivery of \$100,000 face amount of bonds during the delivery month of September 2010. Delivery must take place between the *first delivery date* of September 1, 2010, and the *last delivery date* of September 30, 2010.

Throughout the trading day, market forces determine futures prices. And at the end of each day, the exchange on which the futures trade, determines a *settlement price* that is designed to reflect the price of the last trade of the day. The daily settlement process, described in Chapter 13, is based on daily changes in this settlement price. Table 14.2 lists the settlement prices of TYU0 from May 17 to May 28, 2010, along with the daily settlement payments arising from a long position of one contract. To illustrate, the settlement price falls from May 21 to May 24 by 4 ticks or 32nds per 100

TABLE 14.2 Settlement Prices of TYU0 and Daily Settlement Payments from a Long Position of One Contract

Date	Price	Change (ticks)	Daily Settlement
5/17/10	118-16+		
5/18/10	119-05	20.5	\$641
5/19/10	119-08	3	\$ 94
5/20/10	119-31+	23.5	\$734
5/21/10	120-14+	15	\$469
5/24/10	120-10+	-4	-\$125
5/25/10	120-28+	18	\$563
5/26/10	120-10+	-18	-\$563
5/27/10	119-16	-26.5	-\$828
5/28/10	119-28	12	\$375

face amount. Therefore, on a contract's \$100,000 face amount, the loss to a long position is $\$100,000 \times \frac{4}{32} \%$ or \$125.

The price at which a seller delivers a particular bond to a buyer is determined by the settlement price of the futures contract and by the *conversion factor* of that particular bond. Let the settlement price of the futures contract at time t be F_t and let the conversion factor of bond i be cf^i . Then the delivery price is $cf^i \times F_t$ and the invoice price for delivery is this delivery price plus accrued interest: $cf^i \times F_t + AI_t^i$. The conversion factors for TYU0 are listed in Table 14.1. Since the conversion factor for the $4\frac{1}{2}$ s of May 15, 2017 is .9202, if the futures settlement price is 100, any delivery of the $4\frac{1}{2}$ s occurs at a flat price of $.9202 \times 100$ or 92.02.

Each contract trades until the *last trade date*. The settlement price at the end of that day is the *final settlement price*. This final settlement price is used for the daily settlement payment and for any deliveries that have not yet been made. The last trade date of TYU0 is September 21, 2010. Any delivery from then on, through the last delivery date of September 30, 2010, is based on the final settlement price determined on September 21, 2010. This feature of U.S. futures contracts gives rise to the *end-of-month* option, which is discussed later in this chapter.

The quality option is the most significant embedded option in futures contracts. To simplify the presentation, the timing and end-of-month options are ignored until discussed explicitly. Ignoring these two options is equivalent to assuming that the first delivery date, the last trade date, and the last delivery date are all the same. In fact, this simplification does describe the government bond futures contracts that trade in Europe and Japan.

COST OF DELIVERY AND THE DETERMINATION OF THE FINAL SETTLEMENT PRICE

The *cost of delivery* measures how much it costs a short to fulfill the commitment to deliver a bond through a futures contract. Having decided to deliver bond i , the short has to buy the bond at its market price and then deliver it at the futures price. If the price of bond i at time t is p_t^i , then the cost of delivery is

$$p_t^i + AI_t^i - (cf^i \times F_t + AI_t^i) = p_t^i - cf^i \times F_t \quad (14.1)$$

The short will minimize the cost of delivery by choosing which bond to deliver from among the bonds in the delivery basket. The bond that minimizes the cost of delivery is called the *cheapest-to-deliver* or the CTD.

To illustrate the pricing and interest rate sensitivity of futures contracts, this chapter uses a Vasicek-style term structure model calibrated to market quantities as of May 28, 2010. For the purposes of this section, a particular realization of the model as of TYU0's delivery date, September 30, 2010, has been selected. At that realization, the seven-year par rate is 2.77%, the futures price is 121.2039, and bond prices are as given in Table 14.3. With this data and the contract's conversion factors, the table calculates the costs of delivery across bonds.

As an example of the calculations in the table, the cost of delivering the 3 $\frac{5}{8}$ s of August 15, 2019, is found by applying (14.1) as follows:

$$103.1007 - .8401 \times 121.2039 = 1.277 \quad (14.2)$$

According to Table 14.3, the 4 $\frac{1}{2}$ s of May 15, 2017, is the CTD since it has the lowest cost of delivery, which in this case is zero. The next to CTD is the 3 $\frac{1}{4}$ s of March 31, 2017, with a cost of delivery of .013. Mathematically, since the CTD is chosen so as to minimize the cost of delivery (14.1), it follows that, for any bond i ,

$$p_t^i - cf^i \times F_t \geq p_t^{CTD} - cf^{CTD} \times F_t \quad (14.3)$$

where equality holds only for any jointly-CTD bond.

Since this section ignores the timing and end-of-month options, the determination of the final settlement price is quite simple:

$$F_T = \frac{p_T^{CTD}}{cf^{CTD}} \quad (14.4)$$

where T denotes the last delivery date.

TABLE 14.3 Cost of Delivery Calculations Using a September 30, 2010, Realization of Prices from a Vasicek-Style Model Calibrated as of May 28, 2010

Futures Price: 121.2039

Coupon	Maturity	Price	Conv. Factor	Cost of Delivery	Price ÷ Factor
3¼	3/31/17	103.4967	.8538	.013	121.2189
4½	5/15/17	111.5318	.9202	.000	121.2039
3⅛	4/30/17	102.7042	.8471	.032	121.2421
2¾	5/31/17	100.3229	.8272	.063	121.2802
4¾	8/15/17	113.0395	.9314	.150	121.3651
4¼	11/15/17	109.7547	.9012	.526	121.7873
3⅞	5/15/18	106.5626	.8732	.727	122.0369
4	8/15/18	107.1652	.8774	.821	122.1395
3½	2/15/18	104.3789	.8547	.786	122.1234
3¾	11/15/18	105.0173	.8587	.940	122.2981
3⅝	8/15/19	103.1007	.8401	1.277	122.7243
3⅛	5/15/19	99.5649	.8107	1.305	122.8135
2¾	2/15/19	97.2761	.7909	1.416	122.9941
3⅜	11/15/19	101.0901	.8195	1.763	123.3558
3⅝	2/15/20	102.9906	.8332	2.003	123.6084
3½	5/15/20	102.0509	.8210	2.543	123.3008

Equation (14.4) is proved by showing that there is an arbitrage opportunity if that equation does not hold. First assume that $F_T > \frac{p_T^{CTD}}{cf^{CTD}}$. In this case a trader could buy the CTD, sell the contract, and deliver the CTD for a profit of

$$cf^{CTD} \times F_T - p_T^{CTD}$$

(14.5)

But the quantity (14.5) is positive by assumption, implying that the trade described constitutes a riskless arbitrage opportunity. Hence it cannot be the case that $F_T > \frac{p_T^{CTD}}{cf^{CTD}}$.

Next assume the reverse inequality, that $F_T < \frac{p_T^{CTD}}{cf^{CTD}}$. In this case a trader could sell the CTD, buy the contract, and take delivery of the bond delivered by the short. If the short delivers the CTD, the profit from this strategy is

$$p_T^{CTD} - cf^{CTD} \times F_T$$

(14.6)

which, by the current assumption, is positive. If the short delivers some other bond j , the trader would buy back the CTD just sold and sell bond j instead, for a total profit of

$$p_T^j - cf^j \times F_T \geq p_T^{CTD} - cf^{CTD} \times F_T \quad (14.7)$$

where the inequality follows from (14.3). But by the current assumption the right-hand side of (14.7) is positive, so the left-hand side is as well. Therefore, whatever bond is delivered by the short, there exists a riskless arbitrage opportunity. Hence it cannot be the case that $F_T < \frac{p_T^{CTD}}{cf^{CTD}}$ either. And ruling out both of these strict inequalities establishes the validity of (14.4), as desired.

Having determined the final settlement price in (14.4), the relationships among all the bonds in the basket, the CTD, and the futures price as of the last delivery date can be summarized neatly. First, it follows immediately from (14.4) that

$$p_T^{CTD} - cf^{CTD} \times F_T = 0 \quad (14.8)$$

In words, the cost of delivering the CTD on the last delivery date is zero.

Second, combining (14.4) with the CTD condition (14.3) and rearranging terms,

$$\frac{p_T^i}{cf^i} \geq \frac{p_T^{CTD}}{cf^{CTD}} = F_T \quad (14.9)$$

for any bond i , where equality holds only for any jointly-CTD bond. In words, the CTD is the bond with the smallest ratio of price to conversion factor and the futures price equals this minimum ratio. The last column of Table 14.3 illustrates the workings of equation (14.9) in the example of this section.

MOTIVATIONS FOR A DELIVERY BASKET AND CONVERSION FACTORS

The design of bond futures contracts purposely avoids a single underlying security. One reason is to ensure that the liquidity of the futures contract does not depend on the liquidity of a single, underlying bond, which might lose its liquidity for idiosyncratic reasons, e.g., being accumulated by a few large, buy-and-hold investors. Another reason for avoiding a single underlying bond is to avoid losing liquidity to the threat of a *squeeze*. A trader squeezes

a contract by simultaneously purchasing many contracts and a large fraction of a deliverable bond issue, hoping to sell the position at a profit as traders who had sold the contract scramble to buy the bond to make delivery or, failing that, to buy back the contracts they had sold.² The threat of a squeeze can prevent a contract from attracting volume and liquidity by making shorts hesitant to take positions.

The existence of a basket of securities effectively avoids the problems of a single deliverable if the cost of delivering the next to CTD is not that much higher than the cost of delivering the CTD. In Table 14.3, for example, the difference between the cost of delivering the next to CTD and the cost of delivering the CTD (i.e., zero) is only 1.3 cents per 100 face amount. Therefore, if the $4\frac{1}{2}$ s of May 15, 2017, cannot be economically purchased, because they have lost liquidity or because they are the target of a squeeze, the harm to shorts is limited to 1.3 cents: for any larger cost of acquiring the $4\frac{1}{2}$ s, shorts would purchase and deliver the $3\frac{1}{4}$ s of March 31, 2017, instead.

The difference between the cost of delivering the CTD and the cost of delivering the next to CTD is as small as it is because of the conversion factors. If the contract did not provide for conversion factors, or, equivalently, if all conversion factors were one, the CTD in Table 14.3 would be the one with the lowest price, namely, the $2\frac{3}{4}$ s of February 15, 2019, with a price of 97.28, and the futures price would settle at that same 97.28. But then the next to CTD would be the $3\frac{1}{8}$ s of May 15, 2019, with a price of 99.56, implying a cost of delivery of 2.28. This cost of delivering the next to CTD is very large compared with the actual cost of delivering the next to CTD, that is, 1.3 cents. The problem here with not having conversion factors is that delivery of \$100,000 face amount of the $2\frac{3}{4}$ s of February 15, 2019, with its relatively low coupon of $2\frac{3}{4}\%$, is considered just as good as delivering \$100,000 of any of the other, higher-coupon bonds.

Conversion factors in futures contracts reduce the difference in delivery costs across bonds by adjusting delivery prices for coupon rates. For TYU0 the *notional coupon* of the contract is 6%. The precise role of this coupon is discussed shortly, but the basic idea is to set the conversion factor of a bond with a coupon rate of 6% equal to 1 so that its delivery price, i.e., conversion factor times the futures price, equals the futures price. Bonds with a coupon rate below 6%, typically worth less than a bond with a coupon equal to 6%, are assigned conversion factors less than 1 so that their delivery prices are below the futures price. Finally, bonds with a coupon rate above 6%, typically worth more than bonds with a coupon rate of 6%, are assigned conversion factors greater than 1. The conversion factors in Table 14.3 clearly increase with coupon, although there are no

²Futures exchanges can assess severe penalties for failing to deliver on the obligations of a contract.

conversion factors greater than 1 because, in the low-rate environment of the preceding several years, no U.S. Treasuries have been issued with coupons at or above 6%.

Conversion factors are computed by the futures exchanges and are easily available. The precise rule for computing conversion factors is not particularly intuitive, but there is a fairly good approximation that is useful for intuition about futures contracts: the conversion factor of a bond is approximately equal to its price per dollar face amount as of the last delivery date with a yield equal to the notional coupon rate. An easy illustration is the 3 $\frac{1}{4}$ s of March 31, 2017, which mature in exactly 6.5 years from the delivery date of September 30, 2010. In this particular case the conversion factor of .8538 is very precisely estimated using the price-yield relationship of equation (3.14):

$$\frac{3.25\%}{6\%} \left(1 - \frac{1}{\left(1 + \frac{6\%}{2}\right)^{2 \times 6.5}} \right) + \frac{1}{\left(1 + \frac{6\%}{2}\right)^{2 \times 6.5}} = .8538 \quad (14.10)$$

To understand why conversion factors set according to this rule reduce the differences in delivery costs across bonds, assume that the approximation just discussed holds exactly and that term structure is flat at the notional coupon rate. In this case, the price of each bond is the value of 100 face amount at a yield of 6% while the conversion factor is the value of a unit face amount at a yield of 6%. Hence, the ratio of price to conversion factor is 100 for every bond. Furthermore, by the logic of the previous section, the futures price is 100; the cost of delivery of each bond is zero; and all bonds are jointly CTD. To summarize, if the conversion-factor approximation holds exactly and if the term structure is flat at the notional coupon rate, then conversion factors perfectly adjust delivery prices. No bond is preferable to any other with respect to delivery.

IMPERFECTION OF CONVERSION FACTORS AND THE DELIVERY OPTION AT EXPIRATION

Most of the time—that is, whenever the term structure is not flat at the notional coupon rate—conversion factors used in futures contracts do not adjust delivery prices nearly so well. Figure 14.1 illustrates this point by graphing the price divided by conversion factor for three bonds in the TYU0 basket against yield, assuming a flat term structure. As discussed in the previous section, when conversion factors are the prices of unit face amounts of bonds at a yield of 6% and the term structure is flat at 6%, conversion factors adjust prices perfectly and the ratio of price to conversion factor is near 100 for all bonds. But since actual conversion factors are not exactly




FIGURE 14.1 CTD Analysis for TYU0 at Delivery: A Flat Term Structure of Yields

equal to prices of unit face amounts, it turns out that, at a flat term structure of 6%, the $2\frac{3}{4}$ s of May 31, 2017, has a very slightly smaller ratio of price to conversion factor than the other bonds, making it the CTD.

Figure 14.1 shows that the ratios of price to conversion factor differ across bonds to a greater extent as yields move away from 6%. From yields above about 6.15%, the $3\frac{1}{2}$ s of May 15, 2020, becomes markedly CTD. For rates lower than about 5%, the $4\frac{1}{2}$ s of May 15, 2017, is only marginally CTD relative to the $2\frac{3}{4}$ s, but is pronouncedly CTD relative to the $3\frac{1}{2}$ s. The CTD changes noticeably with yield because, as yield moves away from the notional coupon, conversion factors adjust delivery prices less and less well. To understand why this is so, consider the slope of the converted price-yield curves in Figure 14.1. The slope for bond i is

$$\frac{1}{cf^i} \frac{dP^i}{dy} \quad (14.11)$$

But at a yield of 6% the conversion factor bond i is approximately equal to its price per dollar face value. Hence, apart from a missing 100 in the denominator, the slope of the converted price-yield curve in expression (14.11) at a yield of 6% equals the negative of the duration of bond i .

As yield increases in Figure 14.1, the prices of all bonds fall, but the price of the bond with the largest negative slope, which was just shown to be the bond with the largest duration, falls the most. In this case, the $3\frac{1}{2}$ s have the highest duration and fall the most in price. But because conversion factors

are fixed, the delivery price of the $3\frac{1}{2}$ s stays the same relative to those of the other bonds. In other words, as yields increase above the notional coupon rate, the cost of delivering the $3\frac{1}{2}$ s falls more than that of any other bond. Therefore, while all bonds are about equally attractive to deliver at a yield of 6%, as yield increases the $3\frac{1}{2}$ s become CTD. Graphically, the ratio of the price to conversion factor of the $3\frac{1}{2}$ s falls below that of the other bonds.

As yield falls below the notional coupon rate, the prices of all the bonds increase but the price of the bond with the lowest duration, namely the $4\frac{1}{2}$ s, increases the least.³ But, since conversion factors are fixed, the delivery price of the $4\frac{1}{2}$ s stays the same relative to other bonds. Therefore, while all the bonds are about equally attractive to deliver at a yield of 6%, as yield decreases the $4\frac{1}{2}$ s become CTD.

Figure 14.1 is a stylized example in that it assumes a flat term structure, which makes the CTD analysis relatively straightforward. In reality, the term structure can take on a wide variety of shapes that will affect the determination of the CTD. In general, anything that cheapens a bond relative to others makes it more likely to be CTD. If the curve steepens, for example, long-maturity bonds cheapen on a relative basis and are more likely to be CTD. If the curve flattens, shorter-maturity bonds are more likely to become CTD. And if any bond-idiosyncratic factors cause a particular bond to cheapen or to richen, that bond is more likely, or less likely, respectively, to be CTD.

According to (14.9), the futures price is equal to the smallest ratio of price to conversion factor across bonds. Graphically, the futures price is the lower envelope of all the converted price-yield curves. Figure 14.2 graphs the futures price as the lower envelope of the converted price-yield curves of Figure 14.1. Note that the futures price at delivery is negatively convex in that its slope falls as yield falls: as yield decreases the slope of the futures price moves from resembling that of the relatively high duration $3\frac{1}{2}$ s of May 15, 2020, to the intermediate duration $2\frac{3}{4}$ s of May 31, 2017, and then, ultimately, to the lower duration $4\frac{1}{2}$ s of May 15, 2017.

One way to think about the quality option at expiration is as the value of the option to switch from delivering one particular bond in the basket to delivering another. Consider a trader who owned and was planning to deliver the $3\frac{1}{2}$ s of May 15, 2020. In Figure 14.2, the value of switching deliverables at a given yield is the difference between the converted price of the $3\frac{1}{2}$ s and the futures price at that yield. When yield is relatively high and the $3\frac{1}{2}$ s are CTD, this difference is zero and the option to switch has no value. The value of the option is higher when yield is in the intermediate

³The $4\frac{1}{2}$ s of May 15, 2017, have a lower duration than the $2\frac{3}{4}$ s of May 31, 2017, because of a higher coupon rate and a slightly earlier maturity date.




FIGURE 14.2 Futures Price for TYU0 at Delivery: A Flat Term Structure of Yields

range and the CTD is the $2\frac{3}{4}s$. Finally, the value of the option is highest when yield is relatively low and the CTD has moved all the way to the $4\frac{1}{2}s$.

The quality option at expiration may also be expressed using another bond as benchmark, like the $2\frac{3}{4}s$. Then, at expiration, the option is worth nothing in the intermediate range of yield but has value for relatively high or relatively low levels of yield.

GROSS AND NET BASIS

Transactions in futures are usually either outright (i.e., buying or selling futures alone) or against bond or forward bond positions in the form of *basis trades*. Basis trades take a view on the cheapness or richness of the futures contract relative to the prices of the bonds in the delivery basket. These trades are important to arbitrageurs who profit from these trades but, from a market perspective, the activity of these arbitrageurs keeps the price of a futures contract near its fair value relative to cash bonds. This section defines basis trades, defines the terms gross and net basis, and then relates the profit and loss (P&L) from basis trades to changes in the net basis.

Before delving into details, it is useful to make the following point about hedging bonds with futures. By equation (14.4), the change in the price of the futures contract at expiration is the change in the price of the CTD divided by its conversion factor. Therefore, the change in the price of cf^{CTD} contracts equals the change in the price of the bond or, equivalently, for a

TABLE 14.4 Definition of Long and Short Basis Trades

	Long Basis of Bond i
In terms of repo	Buy G^i face amount of bond i ; Sell repo of bond i to the last delivery date; Sell $cf^i \times G^i$ face amount of futures contracts.
In terms of forwards	Buy G^i of bond i forward to the last delivery date; Sell $cf^i \times G^i$ face amount of futures contracts.
	Short Basis of Bond i
In terms of repo	Sell G^i face amount of bond i ; Buy repo of bond i to the last delivery date; Buy $cf^i \times G^i$ face amount of futures contracts.
In terms of forwards	Sell G^i of bond i forward to the last delivery date; Buy $cf^i \times G^i$ face amount of futures contracts.

fixed CTD, a long position in a contract-sized face amount of the bond is hedged by selling not one contract but cf^{CTD} contracts. With this in mind, basis trades are as defined in Table 14.4.⁴ Note that buying or selling the basis as described here involves no cash outlay: repo finances the purchase of a bond or invests the proceeds of its sale, and the forward and futures trades, by definition, require no cash.⁵

Let p_t^i be the spot price of bond i at time t , $p_t^i(T)$ be its forward price to the last delivery date T at time t , and let F_t be the futures price at time t . Then the *gross basis* and *net basis* of bond i at time t , GB_t^i and NB_t^i , respectively, are defined as

$$GB_t^i \equiv p_t^i - cf^i \times F_t \quad (14.12)$$

$$NB_t^i \equiv p_t^i(T) - cf^i \times F_t \quad (14.13)$$

As discussed in Chapter 13, the forward drop is synonymous with cash carry or, more simply, carry. Generalizing the notation of equation (13.7) for any deliverable bond, let the carry of bond i to the last delivery date T be $\kappa^i(T)$. Then

$$p_t^i(T) = p_t^i - \kappa(T) \quad (14.14)$$

⁴This text defines basis trades as futures *versus* forward bond positions. Practitioners also use the term for futures *versus* spot bond positions.

⁵This discussion, of course, abstracts from futures margin requirements and repo haircuts.

and the definition of net basis in (14.13) can be rewritten as

$$NB_t^i \equiv p_t^i - \kappa(T) - cf_t^i \times F_t = GB_t^i - \kappa(T) \tag{14.15}$$

The right-hand side of (14.15) explains the terminology of net basis: it is the gross basis net of carry.

As pointed out in Chapter 13, the forward price equals the spot price at delivery, or, equivalently, carry equals zero. This has two implications for the basis quantities at delivery. First, from (14.15), gross basis equals net basis. Second, from (14.1), both measures also equal the cost of delivery.

Table 14.5 calculates the gross and net basis for all of the bonds in the TYU0 basket as of May 28, 2010. Note that the spot price, carry, gross basis, and net basis are, by common practice, quoted in ticks or 32nds. As an example, consider the 4½s of May 15, 2017. According to (14.12), the

TABLE 14.5 TYU0 and Its Deliverable Basket as of May 28, 2010

Futures Price:		119–28					
Repo to Delivery:		.30%, except .04% for the 3½s of 5/15/20					
Rate	Maturity	Conv. Factor	Spot Price	Gross Basis	Fwd Price	Carry	Net Basis
3¼	3/31/17	.8538	103-10+	31.3	102.3584	31.0	.3
4½	5/15/17	.9202	111-22	44.1	110.3207	43.7	.4
3⅛	4/30/17	.8471	102-15+	30.0	101.5605	29.6	.5
2¾	5/31/17	.8272	100-00	26.9	99.1917	25.9	1.0
4¾	8/15/17	.9314	113-08	51.1	111.7868	46.8	4.3
4¼	11/15/17	.9012	109-25	56.0	108.4947	41.2	14.8
3⅞	5/15/18	.8732	106-14	56.4	105.2709	37.3	19.1
4	8/15/18	.8774	107-02	60.3	105.8421	39.1	21.2
3½	2/15/18	.8547	104-06	55.4	103.1303	33.8	21.5
3¾	11/15/18	.8587	104-26	60.0	103.6853	36.1	24.0
3⅝	8/15/19	.8401	102-26	67.4	101.7124	35.2	32.2
3⅛	5/15/19	.8107	99-05	63.2	98.2289	29.7	33.5
2¾	2/15/19	.7909	96-25+	63.6	95.9813	26.1	37.5
3⅜	11/15/19	.8195	100-22+	78.9	99.6951	32.3	46.6
3⅝	2/15/20	.8332	102-21	88.8	101.5560	35.2	53.6
3½	5/15/20	.8210	101-23+	106.1	100.5973	36.4	69.8

gross basis is $(111 + \frac{22}{32}) - .9202 \times (119 + \frac{28}{32})$ which equals about 1.379 or $1.379 \times 32 = 44.1$ ticks. The forward price is calculated along the lines described in Chapter 13. By (14.14), carry equals $(111 + \frac{22}{32} - 110.3207)$, which is 1.367 or 43.7 ticks. And finally, by (14.15) the net basis is $44.1 - 43.7$ or .4 ticks.

Chapter 13 showed that the P&L of a futures position can be considered as realized at delivery, like a forward position, after adjusting the futures position for the tail. For ease of exposition, it is now assumed that all basis positions are properly tailed so that the text can treat a futures position as if its profit were realized at expiration. In other words, in the background of the discussion is an unmentioned tail adjustment. The case study at the end of the chapter explicitly describes this adjustment.

Given the tail adjustment, the P&L on the delivery date from a long basis trade in bond i initiated at time t and taken off at time s is the profit of the long forward position in bond i to the delivery date T minus the profit of the long futures position. Mathematically, this P&L is

$$G^i \times [P_s^i(T) - P_t^i(T)] - G^i \times cf^i \times [F_s - F_t] \quad (14.16)$$

And using the definition of net basis in (14.13), this P&L becomes

$$G^i \times [NB_s^i - NB_t^i] \quad (14.17)$$

In words, expression (14.17) says that the delivery-date P&L from a long basis position in a bond equals the size of the bond position times the change in that bond's net basis.

Before concluding this section, it is noted that net basis is far more useful than gross basis in analyzing basis trades. However, because it is particularly easy to observe, gross basis is very commonly used in practice, especially to quote bond prices relative to futures prices. In fact, traders buy and sell packages of a bond and its conversion factor-weighted number of futures contracts at that bond's quoted gross basis.

THE QUALITY OPTION BEFORE DELIVERY

This section describes the quality option before the delivery date and relates the value of this option to net basis, both algebraically and graphically.

Continuing to assume that futures positions are properly tailed, selling futures can be viewed as selling a bond forward and buying the option to switch and deliver another bond. Therefore, combining a futures contract with buying a particular bond forward, i.e., buying that bond's basis, can be viewed as a pure purchase of the option to switch away from that bond.

Similarly, selling a bond's basis can be viewed as the pure sale of the option to switch away from that bond. Finally, because the net basis of a bond is the price of buying or selling its basis, the net basis can be viewed as the price of the quality option with respect to that bond.

One immediate implication of this reasoning is that if the net basis of any bond is zero, then the quality option with respect to that bond is worthless and selling that bond forward is the same, again assuming proper tailing, as selling the futures contract. Mathematically, when the net basis of bond i equals zero in (14.13), $F_t = \frac{P_i^i(T)}{c_f^i}$.

The bonds in Table 14.5 are in order of ascending net basis, which order has been preserved throughout the tables in this chapter. The bond with the lowest net basis, in this case the $3\frac{1}{4}$ s of March 31, 2017, is usually called the CTD. Strictly speaking, it is not correct to call any bond the CTD before the first delivery date. The smaller a bond's net basis, however, the lower the value of the option to switch away from it, and the closer it is to being the CTD. In the same sense, then, the $4\frac{1}{2}$ s of May 15, 2017, and the $3\frac{1}{8}$ s of April 30, 2017, with net bases within .1 or .2 ticks of the $3\frac{1}{4}$ s of March 31, 2017, are essentially jointly CTD. In any case, under both the flat term structure scenarios detailed earlier in this chapter and the Vasicek-style model results to follow, the $4\frac{1}{2}$ s of May 15, 2017, turn out to be the most likely CTD at rate levels prevailing as of the pricing date.

Figure 14.3 uses the data in Table 14.5 and a Vasicek-style model to illustrate the value of the quality option and the concept of CTD before




FIGURE 14.3 TYU0 and Two Deliverables as of May 28, 2010

delivery. On the vertical axis is the forward price divided by the conversion factor. On the horizontal axis is the interest rate factor in the model expressed as a seven-year par U.S. Treasury rate. Unlike Figure 14.2, which depicts the futures price on the delivery date, the futures price in Figure 14.3 is not equal to the minimum of the converted prices. In fact, the futures price is strictly less than this price for each bond. Intuitively, before expiration the value of the quality option is positive and the minimum net basis is positive. At the current level of the 7-year par rate, 2.77%, the futures price is closest to the converted price of the $4\frac{1}{2}$ s of May 15, 2017, so that, of the two bonds portrayed, it is the CTD. At sufficiently higher rates, the $3\frac{5}{8}$ s of August 15, 2019, would be CTD in this sense.

Figure 14.4 graphs the net basis for three bonds in the TYU0 basket using the same data and model as those used to derive Figure 14.3. The net basis graphs behave like the quality options they represent. The net basis of the $4\frac{1}{2}$ s of May 15, 2017, increases with rates: this bond is firmly CTD at low rates but moves away from being CTD as rates increase. In option parlance, the net basis of the $4\frac{1}{2}$ s behaves like a call on rates or, equivalently, like a put on bond prices. The $3\frac{5}{8}$ s of August 15, 2019, is or is close to being CTD at high rates but moves away from CTD status as rates fall. This net basis, therefore, behaves like a put on rates or a call on bond prices. Finally, the $4\frac{1}{4}$ s of November 15, 2017, is as close as it will ever be to CTD when rates are near the level of the pricing date, i.e., 2.77%, but moves further from CTD status as rates fall or rise. Thus, the net basis of this bond behaves like a straddle on rates or prices.




FIGURE 14.4 Net Basis of Three TYU0 Deliverables as of May 28, 2010




FIGURE 14.5 DV01 of TYU0 and Two Deliverables as of May 28, 2010

Figure 14.5 graphs the $DV01$ of the futures contract and of two bonds in the basket in the same data and model framework as the other figures. The intuition behind the regions of negative convexity of the futures contract follows from a discussion earlier in this chapter. At high rates the futures contract resembles the relatively high duration $3\frac{5}{8}$ s of August 15, 2019, while at low rates the contract resembles the relatively low duration $4\frac{1}{2}$ s of May 15, 2017. The interest rate behavior of a futures contract is, therefore, quite different from that of a coupon bond and, for this reason, when hedging bonds with futures or *vice versa*, the hedge may very well need rebalancing as rates change.

While excellent for intuition and for understanding the interest rate risk of futures, Figures 14.3 to 14.5 should not be taken too literally. As mentioned earlier in this chapter, the slope of the term structure as well as idiosyncratic changes in bond prices also determine the CTD, the price of the futures contract, and its interest rate sensitivity. Put another way, Figures 14.3 to 14.5 rely heavily on the assumption of one factor. In fact, the case study at the end of this chapter describes a trade that does not go so well as anticipated because of changes in the slope of the term structure. In practice then, a one-factor model may not be sufficient for futures applications. One obvious solution is to use a two-factor model for both pricing and hedging. Another solution is to use a one-factor model for pricing and, for safety, to compute a derivative with respect to some measure of the slope of the term structure. Furthermore, to ensure that a futures position is not too exposed to the idiosyncratic risk of a particular bond, it is prudent to

compute futures price sensitivities with respect to changes in individual bond yields.

Before closing this section it is worth noting that, with rates as low as they are as of the pricing date, the value of optionality in futures contracts is relatively low. This can be seen from the example of this chapter. First, in Table 14.5, the net bases of the bonds close to CTD are very low, implying that the quality option is not worth very much even though the pricing date of May 28, 2010, is a full four months before the last delivery date. Second, the range of rates used in the figures to illustrate the behavior of futures contracts is wider than the range of rates likely to be realized over the four months to delivery. The futures exchange could, as it has done in the past, lower the notional coupon so that conversion factors adjust delivery prices more accurately and, in the process, increase the value of the quality option. But for some time now, the exchange has opted not to do so.

SOME NOTES ON PRICING THE QUALITY OPTION IN TERM STRUCTURE MODELS

Having set up a term structure model in the form of a tree, pricing the quality option is straightforward. Start at the delivery date. At each node compute the ratio of price to conversion factor for each bond. Find the bond with the minimum ratio and set the futures price equal to that ratio. This is the tree equivalent of the lower envelope in Figure 14.3. Then, given these terminal values of the futures price, prices on earlier dates can be computed along the lines described in Chapter 13.

The algorithm described in the previous paragraph assumes that the prices of the bonds are available on the last delivery date. These bond prices can be computed in one of two ways. If a model with a closed-form solution for spot rates is being used, these rates can be used to compute bond prices as of the delivery date. Otherwise, the tree has to be extended to the maturity date of the longest bond in the basket and bond prices have to be computed using the usual tree methodology. Obviously the first solution is faster and less subject to numerical error, but each user must decide if a model with a closed-form solution is suitable for the purpose at hand.

As discussed in Part Three, pricing models usually assume that some set of securities is fairly priced. In the case of futures the standard assumption is that the forward prices of all the bonds in the deliverable basket are fair. This was done in the Vasicek-style model used to generate the analysis of this chapter. Technically this calibration can be accomplished by attaching an option adjusted spread (OAS) to each bond such that its forward price in the model matches the forward price in the market. The assumption that bond prices are fair is popular because many market participants first uncover an investment or trade opportunity in bonds and then determine whether

futures contracts should replace some or all of the bonds in the trade. This is the case because futures can be complex securities and, as a result, many investors and traders use futures only when they offer advantages in value, liquidity, or both. Separating the value of futures relative to bonds from the relative value of the bonds themselves allows for a clean consideration of the costs and benefits of using futures.

Term structure models commonly used for pricing futures contracts fall into two main categories. First are one- or two-factor short-rate models like those described in Part Three. The advantage of these models is that they are relatively easy to implement and, for the most part, flexible enough to capture the yield curve dynamics driving futures prices. With only one or even two factors, however, these models cannot capture the idiosyncratic price movements of any particular bond relative to its near neighbors in the deliverable basket.

The second type of model used in practice allows for a richer set of relative price movements across the deliverable basket. These models essentially allow each bond to follow its own price or yield process. The cost of this flexibility is model complexity of two types. First, ensuring that these models are arbitrage free takes some effort. Second, the user must specify the parameters that describe the stochastic behavior of all bond prices in the basket. For TYU0, for example, a user might have to specify a volatility for each of the 16 bonds in the basket along with their 120 correlation coefficients.

Futures traders often describe their models in terms of the *beta* of each bond in the basket relative to a benchmark bond in the basket. The beta of a bond represents the expected change in the yield of that bond given a one-basis point change in the yield of the benchmark. A bond with a beta of 1.02, for example, implies that the bond's yield is expected to change 2% more than the yield of the benchmark bond. The beta of a particular bond can be thought of as the coefficient from a regression of changes in its yield on changes in the benchmark's yield. Note that in a one-factor model the beta of a bond is simply the ratio of the volatility of that bond yield to the volatility of the benchmark's yield.

THE TIMING OPTION

The party short the futures contract may deliver at any time during the delivery month. The delivery period of TYU0, for example, extends from September 1 to September 30, 2010. To understand whether delivery should be made early or late, consider a trader who is short the futures contract and wants to be short the CTD at the end of the delivery month. In an early delivery strategy, the trader would buy the CTD repo, deliver the CTD early, and stay short the CTD to the futures expiration date. In a late delivery

strategy, the trader would stay short the futures contract until the expiration date. The determinants of the best policy are carry and option value. Under the early delivery strategy, the trader pays carry on the CTD and sacrifices any value left in the quality option. Under the late delivery strategy, the trader pays no carry and can switch bonds if the CTD changes. Clearly, if carry is positive, it is optimal to delay delivery. If carry is negative, however, the carry advantage of delivering early must be weighed against the sacrifice of the quality option.

Having described the timing option in theory, recall from Chapter 13 that the carry of bonds in the deliverable basket is usually positive. Therefore, while the timing option does exist, it is usually optimal to delay delivery.

THE END-OF-MONTH OPTION

At the last trading date, the final settlement price is set. An immediate implication is that the hedging argument made previously in this chapter no longer applies: a position in bond i *versus* a conversion-weighted amount of futures contract is no longer hedged against bond price changes. Denote the final settlement price by F_S . With the futures price fixed at F_S , delivery proceeds do not change as the price of the bond changes. Therefore, the hedged position is to be short one futures contract against a contract-sized amount of the bonds. Regardless of market price changes, this position can always be liquidated by selling the bonds through the contract for the fixed amount $cf^i \times F_S$.⁶

Turning now to the end-of-month option, it has just been established that a position long bond i and short a matching face amount of futures contracts is worth $cf^i \times F_S$. If bond i is the current CTD, then its price can rise or fall, but, so long as it remains CTD, the value of the position can be worth no more or less than $cf^i \times F_S$. If bond prices change such that the CTD changes, however, so that for some new CTD

$$P^{CTD} - cf^{CTD} \times F_S < P^i - cf^i \times F_S \quad (14.18)$$

then a trader holding the long bond-short futures package can sell bond i , buy the new CTD, and deliver the new CTD instead of bond i . This option

⁶To avoid confusion, the reader should note that the difference between the conversion factor-weighted futures hedge and the one-to-one futures hedge as of the last trading date is known by practitioners as a tail. This tail, however, has no connection with the term as defined in Chapter 13, namely, as the correction applied to hedges so that the payoffs of futures positions resemble those of otherwise identical forward positions.

to switch bonds after the last trading date is what is called the end-of-month option. The P&L from the switch is

$$P^i - P^{CTD} + (cf^{CTD} - cf^i) \times F_S \quad (14.19)$$

Before the last trade date, the futures price reflects any cheapening of the CTD. After the last trade date, however, equation (14.19) shows that any cheapening of the CTD leads to greater and greater profits.

Despite this potential for great value, the end-of-month option does not turn out to be worth much in practice. First, since the end-of-month period is short, fewer than seven business days, for example, for TYU0,⁷ bond prices do not have time to move very much. Second, traders who are long bonds and short futures actively seek opportunities to profit by switching bond positions against futures. This attention tends to dominate the trading of bonds in the deliverable basket after the last trading date. As a result, any time a bond begins to cheapen, all the shorts express an interest in switching and the cheapening of the bond comes to an abrupt halt.

TRADING CASE STUDY: NOVEMBER '08 BASIS INTO TYM0 (JUNE 2000)

On February 28, 2000, the 10-year note contract expiring in June 2000, TYM0,⁸ appeared cheap relative to the prices of the bonds in the delivery basket according to most models used by the industry. Table 14.6 gives some background information about the contract and prices at the time.

To take advantage of the perceived cheapness of the contract, many traders sold the 4 $\frac{3}{4}$ s of November 15, 2008, net basis at 7.45 ticks, hoping that the 4 $\frac{3}{4}$ s would remain CTD and, therefore, that its net basis would go to zero. Table 14.7 illustrates why many traders thought this was a good trade. The table is constructed using a horizon date of May 19, 2000. (The reason for this choice will become clear shortly.) The rows of the table give scenarios of parallel shifts in forward yields for delivery on the last delivery date of TYM0 (i.e., June 30, 2000). For example, the scenario of +20 basis points is the scenario in which forward yields of all bonds for delivery on June 30, 2000, increase 20 basis points from the trade date of February 28, 2000, to the horizon date of May 19, 2000. The table also gives the futures price and the net basis of the 4 $\frac{3}{4}$ s in the various scenarios according to a particular

⁷The last trading day is seven business days before the last delivery day, but the short has to give notice of the bond to be delivered before the last delivery day.

⁸The tickers cycle over a period of 10 years, so TYM0 is also the ticker for the contract expiring in June 2010.

TABLE 14.6 TYM0 and Its Deliverable Basket as of February 28, 2000

Pricing Date:		2/28/00						
Last Delivery Date:		6/30/00						
Futures Price:		95-9						
	Bond	Conv. Factor	Price	Gross Basis	Term Repo	Fwd Yield	Carry	Net Basis
4 $\frac{3}{4}$	11/15/08	.9195	87-24 $\frac{5}{8}$	5.1	5.55%	6.667%	-2.4	7.5
6	8/15/09	1.0000	96-3	26.0	4.90%	6.637%	13.2	12.8
5 $\frac{1}{2}$	5/15/09	.9662	92-16 $\frac{7}{8}$	14.9	5.70%	6.632%	1.2	13.7
5 $\frac{1}{2}$	2/15/08	.9702	92-31 $\frac{5}{8}$	17.5	5.75%	6.692%	0.9	16.6
5 $\frac{5}{8}$	5/15/08	.9769	93-20 $\frac{1}{8}$	17.6	5.80%	6.681%	0.8	16.7
6 $\frac{1}{8}$	8/15/07	1.0071	96-26 $\frac{1}{4}$	27.6	5.84%	6.716%	4.2	23.4
6 $\frac{5}{8}$	5/15/07	1.0342	99-19 $\frac{7}{8}$	34.6	5.84%	6.734%	7.2	27.4
6 $\frac{1}{4}$	2/15/07	1.0133	97-20 $\frac{3}{4}$	35.2	5.84%	6.723%	5.0	30.2
5	2/15/10	1.0358	100-16	57.9	3.85%	6.550%	27.7	30.2

pricing model which, in addition to other assumptions, has the cheapness of the futures contract converging to zero by the horizon date. Column (4) of the table gives the predicted P&L from being short \$100,000,000 November '08 net basis from February 28, 2000, to the horizon date. To calculate these numbers, recall from (14.17) that the P&L of a basis trade equals the change in the net basis times the face amount of bonds.⁹ So, for example, in the +20 scenario, the P&L of the short basis position is

$$\$100,000,000 \times \frac{\frac{1}{32}(7.45 - .9)}{100} = \$204,688 \quad (14.20)$$

Many traders thought that selling the November '08 basis was a good trade based on data like that presented in column (4) of Table 14.7. The scenarios cover the most likely outcomes. There are 77 days from the trade date to the horizon date. Assuming an interest rate volatility of 100 basis points per year, the volatility over 77 days is $100 \times \sqrt{\frac{77}{365}}$ or 46 basis points, so a scenario span from -80 to +80 covers -1.74 to 1.74 standard deviations. The trade starts to lose money in a bond price rally of a bit more than 60 basis points, but makes money for any smaller rally and any sell-off. In the context of the table, personal preference determines whether the potential gains are large enough relative to the risks borne. In any case, it should be noted in

⁹The tail held to realize this P&L is discussed at the end of the case study.

TABLE 14.7 Parallel Shift Scenario Analysis of Selling USD 100mm November '08 Basis into TYM0 as of February 28, 2000

Initial Nov. '08 Net Basis: 7.45					
Horizon Date: 5/19/00					
Futures Option Strike: 95					
Futures Option Price: $1 - \frac{16.5}{32}$					
Number of Options: 47					
(1)	(2)	(3)	(4)	(5)	(6)
Shift (bps)	Futures Price	Nov. '08 Net Basis (32nds)	Basis P&L (\$)	Call P&L 100 Face	Total P&L (\$)
-80	100.3666	13.2	-179,688	3.851	1,308
-60	99.2198	7.1	10,938	2.704	138,034
-40	98.0200	3.2	132,813	1.504	203,518
-20	96.7683	1.5	185,938	0.253	197,813
0	95.4937	1.0	201,563	-1.022	153,532
+20	94.2260	0.9	204,688	-1.516	133,453
+40	92.9756	0.9	204,688	-1.516	133,453
+60	91.7426	0.9	204,688	-1.516	133,453
+80	90.5241	1.1	198,438	-1.516	127,203

passing that the attractiveness of the computed P&L profile is very much related to the computed initial cheapness of the contract: had the contract not been so cheap, the computed P&L profile would have been less attractive.

A criticism of making a trading decision based on Table 14.7 is that the table does not really describe all of the risks involved in the basis trade. If the curve flattens, then one of the shortest bonds in the basket will become CTD and the net basis of the $4\frac{3}{4}$ s will rise. If the curve steepens, then one of the longer bonds in the basket will become CTD and the net basis of the $4\frac{3}{4}$ s will rise. Also, if the $4\frac{3}{4}$ s for some reason richens relative to the other bonds in the basket, then its net basis will rise. None of these risks is included in Table 14.7.

Some traders looking at the payoff profile of the basis in Table 14.7 did not like the dramatically falling and ultimately negative P&L after a rally of between 40 and 80 basis points. To make the P&L profile look better, many traders bought 95 strike call options on TYM0 expiring on May 19, 2000, for 1.51 per 100 face amount of futures. The payoff from calls on 100 face of futures is given in column (5) of the table. With a rally of 60 basis points, for example, the payoff is $99.2198 - 95 - 1.516$ or 2.704. Choosing to purchase 47 calls (i.e., calls on 47 contracts covering $47 \times \$100,000$

or \$4.7 million face) evens out the payoff profile nicely, as shown in column (6) of Table 14.7. In a rally of 60 basis points, for example, the 47 calls are worth

$$\$4,700,000 \times \frac{2.704}{100} = \$127,088 \quad (14.21)$$

Adding this to the P&L of \$10,938 from the net basis position alone gives a total P&L of \$138,034. In any case, it is understandable that some traders would choose to sacrifice some upside in a sell-off to limit the loss in a large rally.

The option position also reduces risk in a way not shown in Table 14.7. If volatility were to rise over the trading horizon, the value of the quality option and therefore the net basis would rise, and the short basis position would suffer losses. But the long option position's value would increase as well and at least partially offset these losses. Of course, this protection lasts only to the option expiration date while the short basis position is priced to delivery. Nevertheless, as the case will show, preventing losses over the course of the trade can be as important as the final P&L profile.

It should now be clear that the P&L analysis was done to a May 19, 2000, horizon because the option on TYM0 expires on that date. Options on futures are set to expire before the first delivery date so that they cannot expire after delivery has taken place. This convention often makes basis trading difficult because delivery usually occurs on the last delivery date, more than a month after the relevant option has expired. A trader can try to correct for this mismatch by using less liquid over-the-counter bond options, available at any maturity, or, at the risk of another kind of mismatch, by using options on the next futures contract (in this case, TYU0, which expires in September, 2000).

Tables 14.8, 14.9, and 14.10 show how the trade worked out. Table 14.8 reports the forward yields to June 30, 2000, of each bond in the basket on the initial trade date (February 28, 2000), on two intermediate dates of interest (April 3 and April 10), and on the option expiration date (May 19). In addition, columns (5), (7), and (9) report changes in forward yields over the indicated periods. Table 14.9 reports the futures price, the net basis of each bond, and the option price on these same four dates. Table 14.10 reports the components of the cumulative P&L of the trade.

As evident from column (5) of Table 14.8, from the initiation of the trade to April 3, forward yields fell approximately in parallel by 47 basis points. As a result, as can be seen from the net bases in Table 14.9, the CTD moved toward the shorter end of the basket, to the $6\frac{1}{8}$ s of August 15, 2007, and the $5\frac{1}{2}$ s of February 15, 2008. The P&L implications of the changes in market yields are in Table 14.10. The November '08 basis rose to 11.06 for a loss of \$112,813. The option position, however, gained \$87,391, making the total loss only \$25,422. It should also be noted that, as of April 3, 2000,

TABLE 14.8 Forward Yields to June 30, 2000, Delivery as of Selected Dates

Forward Yields and Changes								
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Bond		2/28	4/3	Δ 4/3 v. 2/28	4/10	Δ 4/10 v. 4/3	5/19	Δ 5/19 v. 2/28
$6\frac{1}{4}$	2/15/07	6.7234	6.2505	-47.3	6.0814	-16.9	6.7656	4.2
$6\frac{5}{8}$	5/15/07	6.7344	6.2637	-47.1	6.1002	-16.4	6.7668	3.2
$6\frac{1}{8}$	8/15/07	6.7155	6.2514	-46.4	6.0915	-16.0	6.7530	3.8
$5\frac{1}{2}$	2/15/08	6.6920	6.2253	-46.7	6.0467	-17.9	6.7175	2.6
$5\frac{5}{8}$	5/15/08	6.6814	6.2035	-47.8	6.0144	-18.9	6.7083	2.7
$4\frac{3}{4}$	11/15/08	6.6670	6.1885	-47.9	5.9924	-19.6	6.6851	1.8
$5\frac{1}{2}$	5/15/09	6.6319	6.1579	-47.4	5.9584	-20.0	6.6354	.3
6	8/15/09	6.6370	6.1464	-49.1	5.9485	-19.8	6.6033	-3.4
$6\frac{1}{2}$	2/15/10	6.5502	6.0468	-50.3	5.8481	-19.9	6.5159	-3.4

the value of the contract lost about 1.5 ticks of its cheapness, i.e., it had converged somewhat closer to a model-based fair value. This implies that the P&L loss to that date would have been about $\frac{\$100\text{mm}}{100} \times \frac{1.5}{32}$ or \$46,875 larger had the value of the futures contract not richened at all.

From April 3 to April 10, 2000, column (7) of Table 14.8 reports that the forward yields continued to fall by between 16 and 20 basis points. Over this period, however, the forward yield curve flattened by about 3 basis points. In addition, the futures contract had cheapened by about 2.5 ticks over these few days. The combination of these effects was disastrous for the trade. As can be seen from Table 14.9, the flattening rally moved all the shorter-term bonds closer to CTD. The net basis of every bond from the $6\frac{1}{4}$ s of February 15, 2007, to the $5\frac{5}{8}$ s of May 15, 2008, fell below that of the $4\frac{3}{4}$ s of November 15, 2008. As for P&L, Table 14.10 shows that the net basis of the $4\frac{3}{4}$ s increased to over 22 for a loss on the net basis position of \$455,000. The option position gained \$127,781, mitigating the damage to a loss of \$327,219.

Note that the loss of \$327,219 is greater than any number in the predicted P&L of Table 14.7. Part of this is due to the steepening of the forward yield curve and part is due to the additional cheapening of the futures contract. In any case, the trading lesson is that intermediate losses can be much greater than horizon losses. In other words, even if the analysis of Table 14.7 turned out to be correct, the losses in the interim could be great and perhaps too great to bear. In particular, a trader showing a loss of \$455,000 or \$327,291 on this trade might have been ordered to reduce or close the position. In that situation the trader would never see the results of Table 14.7. In

TABLE 14.9 TYM0 Futures Price, Futures Option Price, and Net Basis Values as of Selected Dates

		2/28	4/3	4/10	5/19
Futures Price:		95-9	98-8 $\frac{1}{2}$	99-6 $\frac{1}{2}$	95-9 $\frac{1}{2}$
95 Calls Price:		1.516	3.375	4.234	.297
Bond		Net Basis			
6 $\frac{1}{4}$	2/15/07	30.2	13.3	12.1	22.7
6 $\frac{5}{8}$	5/15/07	27.4	11.5	9.8	21.3
6 $\frac{1}{8}$	8/15/07	23.4	9.9	8.8	16.3
5 $\frac{1}{2}$	2/15/08	16.6	9.7	14.2	11.5
5 $\frac{5}{8}$	5/15/08	16.7	13.9	21.3	11.3
4 $\frac{3}{4}$	11/15/08	7.5	11.1	22.0	3.5
5 $\frac{1}{2}$	5/15/09	13.7	19.2	32.6	12.5
6	8/15/09	12.8	22.9	36.7	19.4
6 $\frac{1}{2}$	2/15/10	30.2	47.3	63.5	37.4

fact, one explanation at the time for the cheapening of the futures contract from April 3 to April 10, was that many traders were forced to liquidate short basis positions in this contract. Since such liquidations entail selling futures and buying bonds, enough activity of this sort could certainly cheapen the contract relative to bonds.

Column (9) of Table 14.8 reports that by May 19 the forward yield curve had returned to the levels existing at the start of the trade, on February 28, but had flattened by between 3 and 4 basis points. Table 14.9 reveals that

TABLE 14.10 Cumulative Profit and Loss from the November '08 Basis Trade, With and Without Futures Options

Face amount basis: -100 mm					
Face amount calls: 4.7 mm					
Date	Nov '08 Net Basis	Option Price	P&L from Net Basis	P&L from Options	Total P&L
2/28/00	7.45	1.516			
4/3/00	11.06	3.375	-112,813	87,391	-25,422
4/10/00	22.01	4.234	-455,000	127,781	-327,219
5/19/00	3.51	.297	123,125	-57,281	65,844

this yield curve restored the $4\frac{3}{4}$ s to CTD and reduced its net basis to 3.51. With respect to the P&L, in Table 14.10, the futures contract returned to its original level, but the options lost most of their time value. The total P&L of the trade to its horizon turned out to be \$65,844. Note that this profit is substantially below the predicted P&L of about \$153,532. First, the forward yield curve did flatten, making the shorter-maturity bonds closer to CTD than predicted by the parallel shift scenarios. Second, while the model assumed that the futures contract would be fair relative to the bonds on May 19, it turned out that the contract was still somewhat cheap to bonds on that date. A quick way to quantify these effects is to notice that the net basis of the $4\frac{3}{4}$ s on the horizon date was 3.51 while it had been predicted to be close to 1. This difference of 2.51 ticks is worth $\frac{\$100\text{mm}}{100} \times \frac{2.51}{32}$ or \$78,438 in P&L. Adding this to the actual P&L of \$65,844 would bring the total to \$144,282, much closer to the predicted number. By the way, a trader can, at least in theory, capture any P&L shortfall due to the cheapness of the futures contract on the horizon date by subsequent trading.

By working with the net basis directly this case implicitly assumes that the tail was being managed. So, before concluding the case, some of the relevant calculations are described. The conversion factor of the $4\frac{3}{4}$ s was .9195, so, without the tail, the trade would have purchased about 920 contracts against the sale of \$100mm bonds. On February 28, 2000, there were 122 days to the last delivery date and the repo rate for the $4\frac{3}{4}$ s to that date was 5.55%. Hence, using the rule of Chapter 13, the tail was

$$920 \times \frac{5.55\% \times 122}{360} = 17 \quad (14.22)$$

contracts. In other words, 920 minus 17 or 903 contracts should have been bought against the bond position. On April 3, 2000, the required tail had fallen to 13 contracts, or, equivalently, the futures position should have increased to 920 minus 13 or 907 contracts.

To estimate the order of magnitude of the tail on P&L, note that the futures price rose from 95-9 to 98-8 $\frac{1}{2}$ over that time period, making the tail worth about 2.98 per 100 face amount of contracts. Using an average tail of 15 contracts over the period, i.e., \$1.5 million face, the tail in this trade turned out to be worth

$$\$1.5\text{mm} \times \frac{2.98}{100} = \$44,700 \quad (14.23)$$

In other words, had the the tail not been managed, the difference between the P&L of the basis trade and the bond position times the change in net basis would have been a significant \$44,700.

Short-Term Rates and Their Derivatives

Short-term borrowing and lending, whether through marketable securities or loans, is a large and important segment of financial markets. While these short-term transactions occur at very many different rates, most of these rates are actually spreads off a smaller number of benchmark rates. For example, a bank will charge a customer whatever rate it thinks appropriate for a one-month loan, but will probably start with the current level of one-month *London Interbank Offered Rate (LIBOR)* and add a spread appropriate for the credit risk of a loan to that customer.

This chapter covers two of the most widely used short-term rate benchmarks, namely the set of LIBOR indexes and the federal (fed) funds rate. In addition to defining and discussing the rates themselves, this chapter presents the securities used to hedge exposures to those rates, Forward Rate Agreements (FRAs) and Eurodollar futures in the case of LIBOR and fed funds futures and Overnight Indexed Swap (OIS) in the case of the fed funds rate. In addition to the basic material, this chapter includes an application showing how to extract implied probabilities of the Board of Governors of the Federal Reserve System policy changes from fed fund futures prices, a case study of shorting the Treasury Eurodollar (TED) spread of a particular U.S. Treasury bond, and a discussion of the LIBOR-OIS spread as a leading indicator of system-wide financial stress.

A final point to be made here is that this chapter ignores the counterparty risk of derivative contracts, i.e., the risk that a counterparty, including a futures exchange, fails to fulfill its obligations under a contract. The justification for this treatment is that collateral agreements are often arranged so as to allay these concerns. This will be discussed further in Chapter 16.

LIBOR AND LIBOR-RELATED SECURITIES

LIBOR Fixings

LIBOR, is the generic name for a set of short-term rate indexes. The *British Bankers Association (BBA)* surveys banks each day with the following

TABLE 15.1 BBA Survey of Banks and the Fixing of 3-Month LIBOR as of July 2, 2010

Bank	EUR	USD	GBP	JPY
Abbey National	.700		.770	
Barclays	.750	.530	.720	.240
BNP Paribas			.730	
Bank of America		.490	.730	.250
Bank of Tokyo-Mitsubishi UFJ	.740	.600	.780	.250
Citibank	.740	.500	.750	.230
Credit Suisse	.800	.530		
Deutsche Bank	.700	.500	.740	.210
HSBC	.690	.480	.690	.250
JPMorgan Chase	.700	.510	.700	.240
Lloyds	.740	.530	.760	.250
Mizuho Corporate Bank	.770		.780	.250
Norinchukin Bank		.600		.250
Royal Bank of Scotland	.630	.590	.690	.240
Rabobank	.730	.510	.720	.220
Royal Bank of Canada	.735	.565	.7325	
Sumitomo Mitsui Banking Corp Europe				.250
Societe Generale	.730	.565	.730	.240
UBS AG	.565	.529	.700	.230
WestLB AG	.790	.590		.290
3-Month LIBOR Fixing	.72688	.53363	.73156	.24500

question: “At what rate could you borrow funds, were you to do so by asking for and then accepting inter-bank offers in a reasonable size just prior to 11 A.M. (London time)?” The BBA polls 16 banks for each of several currencies, and includes rates with terms ranging from overnight to 12 months, although most market attention is focused on the three-month rate. Then, for each currency and term, the BBA drops the four highest rates and the four lowest rates and computes the average of the remaining eight rates. This average is the LIBOR *fixing* for that currency, term, and day. Table 15.1 shows the process for the three-month fixing in EUR, USD, GBP, and JPY as of July 2, 2010. The non-bold entries were dropped before averaging the eight entries in bold type to obtain the fixing.¹

¹There is no significance to the table’s dropping one of several equal rates rather than another, e.g. dropping the JPY quotes of .25% by Mizuho, Norinchukin, and Sumitomo and keeping the .25% quotes of Bank of America, Bank of Tokyo, HSBC, and Lloyds.

LIBOR rates are particularly important in financial markets because many other rates are keyed off LIBOR. For example, many borrowing rates are quoted as a spread to LIBOR, so that a company might be given the opportunity to borrow money at LIBOR+150, that is, at 150 bps above LIBOR. Also, Eurodollar futures, discussed in this section, and the floating leg of interest rate swaps, discussed in Chapter 16, are set off LIBOR rates. LIBOR is quoted on an actual/360 basis and assumes $T + 2$ settlement. This means, for example, that a \$1 million three-month LIBOR deposit set at .53363% on July 2, 2010, would pay interest from two business days after the trade, i.e., on July 6, 2010, to three months after that, i.e., on October 6, 2010. The proceeds of this 92-day deposit on October 6, 2010, would be

$$\$1,000,000 \left(1 + \frac{.53363\% \times 92}{360} \right) = \$1,001,364$$

(15.1)

Forward Rate Agreements

FRAs are over-the-counter derivatives that allow market participants to lock in the forward rates of the swap curve, as defined in Chapter 2. To see this, first consider the forward three-month swap agreement depicted in the second column of Table 15.2. On May 28, 2010, party A agrees to pay party B a fixed rate of 2% on \$100 million for the three months settling on March 14, 2012. At the same time, party B agrees to pay party A three-month LIBOR on the \$100mm over those three months. Note from the table that the LIBOR used to compute the floating rate payment is the rate observed on March 12, 2012, two business days before the start of the loan. Denoting LIBOR on that observation date as $\bar{L}$, noting that there are 92 days between March 14, 2012, and June 14, 2012, and applying the money market actual/360 day-count convention for both the fixed and floating legs of the swap, on June 14, 2012, A has to make a new payment

TABLE 15.2 Cash Flows of a Forward 3-Month Swap vs. a FRA

Date	Forward Swap	FRA
5/28/10	trade date	
3/12/12	3-month LIBOR observed to be $\bar{L}$	
3/14/12	interest accrual begins	A pays net of: $\frac{\$100\text{mm} \times \frac{92}{360} (2\% - \bar{L})}{1 + \frac{92}{360} \bar{L}}$
6/14/12	A pays net of: $\$100\text{mm} \times \frac{92}{360} \times (2\% - \bar{L})$	

of $\$100\text{mm} \times \frac{92}{360} \times (2\% - \bar{L})$. (Of course, if $\bar{L} > 2\%$, then A receives a payment on that date.)

A FRA is very much like the forward swap just described, except that instead of the net payment being made on the maturity date, the present value of that net payment is exchanged on the settlement date. Furthermore, this present value is calculated using the same LIBOR fixing. Hence, in the example, as recorded in the third column of Table 15.2, party A pays party B a net of

$$\frac{\$100\text{mm} \times \frac{92}{360} (2\% - \bar{L})}{1 + \frac{92}{360} \bar{L}} \quad (15.2)$$

To see how FRAs can be used to lock in forward rates, consider a corporation that will be receiving \$100 million on March 14, 2012, and, knowing that the money will not be spent until June 14, 2012, would like to lock in a rate of return over that period. To do so, the corporation can agree to receive fixed on the FRA just described. For whatever $\bar{L}$ is realized, on March 14, 2012, the corporation will have its \$100 million plus the net payment from the FRA given in (15.2). If the corporation rolls this total amount, which could be greater or less than \$100 million, at the realized three-month LIBOR rate for three months, its proceeds on June 14, 2012, would be

$$\left(\$100\text{mm} + \frac{\$100\text{mm} \times \frac{92}{360} (2\% - \bar{L})}{1 + \frac{92}{360} \bar{L}} \right) \left(1 + \frac{92}{360} \bar{L} \right) \quad (15.3)$$

$$= \$100\text{mm} \left(1 + \frac{92}{360} \bar{L} + \frac{92}{360} (2\% - \bar{L}) \right) \quad (15.4)$$

$$= \$100\text{mm} \left(1 + \frac{92}{360} \times 2\% \right) \quad (15.5)$$

Hence, the FRA position has enabled the corporation to lock in the FRA rate of 2% on a forward loan. But this implies that the FRA rate is exactly the forward rate defined in Chapter 2.

While not so liquid as Eurodollar futures, to be discussed in the next section, FRA dates can be customized to the needs of the user. While the most liquid FRAs reference three-month LIBOR, so that the maturity date is three months after the settlement date, users can customize that settlement date. By contrast, 90-day Eurodollar futures contracts trade at only four settlement or expiration dates per year. Both FRAs and Eurodollar futures contracts do trade for different forward loan intervals, like one month instead of the three in the example, but at significant sacrifices of liquidity.

TABLE 15.3 Selected Eurodollar Futures Prices as of May 28, 2010

Ticker	Expiration	Price	Rate (%)
EDM0	6/14/10	99.400	.600
EDU0	9/13/10	99.155	.845
EDZ0	12/13/10	99.005	.995
EDH1	3/14/11	98.875	1.125
EDM1	6/13/11	98.705	1.295
EDU1	9/19/11	98.495	1.505
EDZ1	12/19/11	98.245	1.755
EDH2	3/19/12	98.010	1.990

Eurodollar Futures

Eurodollar (ED) futures allow investors and traders to manage their exposures to short-term interest rates. The most liquid of these are quarterly futures designed to hedge \$1 million 90-day LIBOR deposits that mature in March, June, September, and December over the next 10 years. As of May 28, 2010, for example, the first contract matures in June 2010 and the last in March 2020. And, of these contracts, the very most liquid mature in the first few years. The first two columns of Table 15.3 list the tickers and expiration dates of the first eight contracts. The tickers are a concatenation of “ED” for a 90-day Eurodollar contract, a month (H for March, M for June, U for September, and Z for December), and a year. Hence, EDH2 is a 90-day Eurodollar futures contract expiring in March 2012.² The expiration dates of the contracts are two London business days before the third Wednesday of the contract month. The Wednesdays of these contract months, by the way, are known as *IMM (International Money Market) dates*.

The third column of Table 15.3 gives futures prices as of May 28, 2010, while the fourth gives futures rates in percent, defined as 100 minus the corresponding prices. The first four contracts are also known as “fronts.” Subsequent groups of four contracts are traditionally referred to by colors, in the following sequence: reds, greens, blues, golds, purples, oranges, pinks, silvers, and coppers.

To describe how ED futures work, focus on EDH2. On its expiration date of March 19, 2012, the contract price is set at 100 minus 100 times the fixing of three-month LIBOR on that date. So, for example, if the fixing is 1.75%, the final contract price is $100 - 100 \times 1.75\%$ or 98.25. It is important to emphasize that the contract price is only a convention for quoting a 90-day rate: a price of 98.25 means only that the 90-day

²When a contract expires, a new contract with the same ticker is added to the end of the contract list. For example, when EDM0 expired in June 2010, a new EDM0 was listed, which expires in June 2020.

rate is 1.75%. The contract price is not the price of a 90-day zero at the contract rate.

Defining the security underlying ED futures contracts to be a \$1,000,000 90-day deposit is another way of defining the '01 of the contract, i.e., the change in the value of the contract for a one basis-point change in rate, as \$25. To see this, note that the proceeds of a 90-day deposit at EDH2's rate of 1.99% are

$$\$1,000,000 \times \left(1 + \frac{1.99\% \times 90}{360} \right) = \$1,004,975 \quad (15.6)$$

while, if the rate increased by one basis point, to 2%, the proceeds would be

$$\$1,000,000 \times \left(1 + \frac{2.00\% \times 90}{360} \right) = \$1,005,000 \quad (15.7)$$

or \$25 higher.³

Throughout the trading day, market forces determine the prices of futures contracts. At the end of each day, the futures exchange determines a settlement price that is designed to reflect the last trade price of the day. For EDH2 Table 15.4 gives the prices, rates, changes, and daily settlement amounts for a long position in one contract from May 28, 2010, to June 11, 2010. The daily settlement process, in general terms, was discussed in Chapter 13. In the current context, however, it might be added that, with a contract '01 of \$25, the daily settlement amount is simply \$25 times the change in rate in basis points. So, for example, with the price of EDH2 increasing from 97.93 on June 3 to 98.09 on June 4, an increase of 16 basis points, the daily settlement amount for a long of one contract on June 4 is $16 \times \$25$, or \$400. It is also useful to add that this profit comes about because the long, on June 3, had committed to buy a deposit, i.e., to lend money, at a rate of 2.07%. When the rate falls to 1.91% on June 4, this commitment is worth the difference of 16 basis points as of the delivery date. The daily settlement feature of futures contracts, however, pays out this profit immediately rather than on the delivery date.

³Note that there are two small inconsistencies in using three-month LIBOR to settle ED contracts. First, the underlying security of an ED contract is a 90-day deposit even though three-month LIBOR need not be a 90-day rate, as illustrated in the example of the previous section. Second, the LIBOR rate used to settle an ED contract at expiration is really the rate on a deposit for settlement two business days later.

TABLE 15.4 Daily Settlement Payments on One Long EDH2 Contract from June 1 to June 11, 2010

Date	EDH2 Price	EDH2 Rate	Change	Daily Settlement
		(%)	(bps)	(\$)
5/28/2010	98.010	1.990		
6/1/2010	98.010	1.990	0.0	0.00
6/2/2010	97.975	2.025	-3.5	-87.50
6/3/2010	97.930	2.070	-4.5	-112.50
6/4/2010	98.090	1.910	16.0	400.00
6/7/2010	98.105	1.895	1.5	37.50
6/8/2010	98.100	1.900	-0.5	-12.50
6/9/2010	98.130	1.870	3.0	75.00
6/10/2010	98.050	1.950	-8.0	-200.00
6/11/2010	98.155	1.845	10.5	262.50

A final comment to make about the daily settlement process in the context of ED futures is that when these contracts expire and the last daily settlement payments are paid and received, the longs and shorts have no further obligations. In particular, longs do not have to buy a 90-day deposit from shorts at the rates implied by final settlement prices. Futures contracts that do not require delivery of the underlying security at expiration are said to be *cash settled*.⁴ Futures contracts that do require delivery of an underlying security, like the note and bond futures discussed in Chapter 14, are said to be *physically settled*.

Euribor and TIBOR Futures

Euribor futures, EUR-denominated interest rate futures, are extremely liquid contracts and are structured very much like Eurodollar futures. The underlying rate is three-month *Euribor*, which is not the same as three-month Euro LIBOR. Euribor of various maturities come from surveys in the same spirit as LIBOR, but are produced by the European Banking Federation. The survey of over 50 European banks is larger than the LIBOR survey, the top and bottom 15% of responses are discarded, and the results are published at 11 A.M. Central European time.

The yen-denominated interest rate futures contract is based on three-month *TIBOR* (*Tokyo Interbank Offered Rate*). The TIBOR survey, conducted by the Japanese Bankers Association, is an average of rates supplied

⁴Note that the phrase “cash settlement” is sometimes used, in an entirely different context, to mean same-day settlement, e.g., buying a U.S. Treasury for cash settlement means that the trade settles the same day as opposed to normal $T + 1$ settlement.

by 18 banks, mostly Japanese, with the highest and lowest pair discarded. Banks in the Yen LIBOR survey, by contrast, are much more geographically diverse, as can be seen from Table 15.1. This difference between the survey panels has led to positive TIBOR-LIBOR spreads, as TIBOR panel banks have been perceived as riskier and facing more difficult funding conditions than the broader set of banks in the Yen LIBOR panel.

Hedging Lending or Borrowing with Eurodollar Futures

This subsection presents an example to illustrate how ED futures are used to hedge future lending. To make the important points clear, the dates of the example are contrived to match exactly the dates of ED contracts. More realistic settings are dealt with in the next subsection.

Consider the case of a corporation that, as of May 28, 2010, has plans to retain earnings over the near future to finance an expenditure on June 17, 2012. One of its earnings distributions will be \$100,000,000 paid on March 19, 2012. This subsection focuses on the corporation's desire to lock in a return on that \$100,000,000 over the 90 days from March 19, 2012, through June 17, 2012.

If the 90-day rate on March 19, 2012, turns out to be .49%, the corporation will have proceeds on June 17, 2012, of

$$\$100,000,000 \left(1 + \frac{.49\% \times 90}{360} \right) = \$100,122,500 \quad (15.8)$$

On the other hand, if the 90-day rate turns out to be 3.49% on March 19, 2012, the proceeds available would be

$$\$100,000,000 \left(1 + \frac{3.49\% \times 90}{360} \right) = \$100,872,500 \quad (15.9)$$

Rather than face this uncertainty, the corporation might very well prefer to buy a forward contract on May 28, 2010, to lend money at 1.99% on March 19, 2012, for 90 days. In that case, regardless of the 90-day rate that actually obtains on March 19, 2012, the corporation will have locked in proceeds of

$$\$100,000,000 \left(1 + \frac{1.99\% \times 90}{360} \right) = \$100,497,500 \quad (15.10)$$

While this is a perfectly good hedge in theory, it is usually harder and more expensive to find a suitable counterparty to write a customized forward

TABLE 15.5 Hedge Results from Buying 100 EDH2 at a Rate of 1.99% If All EDH2 Profit and Loss were Realized at Expiration

90-Day LIBOR	ED P&L 3/19	ED P&L 6/17	Investment Proceeds 6/17	Total 6/17
.49%	\$375,000	\$375,459	\$100,122,500	\$100,497,959
1.99%	\$0	\$0	\$100,497,500	\$100,497,500
3.49%	-\$375,000	-\$378,272	\$100,872,500	\$100,494,228

contract than to hedge in the ED futures market. Recalling that EDH2 expires on March 19, 2012, and that its underlying security is a \$1 million 90-day deposit, a first pass at a hedge would be for the corporation to buy 100 EDH2 contracts on May 28, 2010, at the rate given in Table 15.3, i.e., 1.99%. If all of the P&L (profit and loss) of EDH2 were realized at its expiration, the hedge would work very well, as shown in Table 15.5. If, at expiration, EDH2 is settled at a 90-day LIBOR of .49%, 150 basis points below 1.99%, the P&L on the 100 contracts would be $100 \times 150 \times \$25$ or \$375,000, which could be invested at .49% through June 17, 2012, for a total of \$375,459. Furthermore, at .49%, the proceeds of investing the \$100 million of retained earnings are \$100,122,500, as given by (15.8). Hence, the total available to the corporation on June 17, 2012, would be \$100,497,959. Similarly, in the case that 90-day LIBOR at expiration is 150 basis points above 1.99%, at 3.49%, the 100 contracts would lose \$375,000, growing to a loss of \$378,272 by June 17, 2012. Combining this with proceeds of \$100,872,500, given by (15.9), results in an available total of \$100,494,228. Hence, for a wide range of values for 90-day LIBOR on March 19, 2012, the corporation has successfully locked in approximately \$100,497,500 on June 17, 2012.⁵

Table 15.5 is not a perfectly accurate representation of the hedge, however, because ED contracts are subject to daily settlement: the \$375,000 profit or loss is paid not on March 19, 2012, but over time, from the trade date on May 28, 2010, to expiration on March 19, 2012. To see the potential impact of daily settlement in this example, consider the extreme example in which 90-day LIBOR jumps or falls immediately, on May 28, 2010, to each of the scenario levels of Table 15.5 and remains at that level until the expiration of EDH2. Table 15.6 shows the results. With 90-day LIBOR falling immediately to and remaining at .49%, the daily settlement payment of \$375,000 would be made immediately and invested for 751 days to June 17, 2012, at .49%, growing to \$378,833.⁶ Adding this to the investment

⁵Note that the total would be exactly \$100,497,500 in each of the scenarios were it not for the 90 days of interest on the ED P&L.

⁶For simplicity, this example assumes simple interest for these scenarios.

TABLE 15.6 Hedge Results from Buying 100 EDH2 at a Rate of 1.99% If, on May 28, 2010, Rates Changed to and Remained at Their Terminal Levels of March 19, 2012

90-Day LIBOR	ED P&L 5/28	ED P&L 6/17	Investment Proceeds 6/17	Total 6/17
.49%	\$375,000	\$378,833	\$100,122,500	\$100,501,333
1.99%	\$0	\$0	\$100,497,500	\$100,497,500
3.49%	-\$375,000	-\$402,302	\$100,872,500	\$100,470,198

proceeds gives a total of \$100,501,333 in this scenario. However, if 90-day LIBOR jumps to 3.49%, the immediate daily settlement loss of \$375,000 has to be financed at 3.49% for 751 days, where financing may be explicit borrowing or the opportunity cost of not being able to invest the foregone funds. Through June 17, 2012, the loss grows to \$402,302. Note, as pointed out in Chapter 13, that the cost of financing the loss of \$375,000 (i.e., \$27,302) exceeds the earnings from investing gains of \$375,000 (i.e., \$3,833) because the loss occurs when rates are relatively high while the gain occurs when rates are relatively low. In any case, the total proceeds from the 3.49% scenario comes to \$100,470,198.

The more accurate depiction of the hedge in Table 15.6 performs relatively well in percentage terms, but the outcome in the 3.49% scenario is noticeably worse than in the .49% scenario. Furthermore, the daily settlement effect would be worse if rates were higher and if the dates in question were further in the future. Therefore, as discussed in Chapter 13, practitioners tend to tail the hedge. Applying the rule developed in that section, accounting for daily settlement reduces the number of contracts by the present value factor to the delivery date. Although it would be more accurate to discount by the rate appropriate from the trade date to the delivery date, for the purposes of this example discounting will be done at 1.99%. Therefore, noting that there are 661 days from the trade date to the delivery date in this example, the recommended hedge is to buy

$$\frac{100}{\left(1 + \frac{1.99\% \times 661}{360}\right)} = 96.47 \quad (15.11)$$

contracts instead of 100. In industry jargon, the tail is 3.53 contracts. Rounding to a whole number of 96 contracts, Table 15.7 displays the results. As planned, this tailed hedge exhibits less variance in terminal proceeds than the unadjusted hedge in Table 15.6.

Before closing the discussion of the tail, it should be noted that the tail has to be adjusted over time. As time to delivery approaches, the present value factor gradually increases to one and the number of contracts in the

TABLE 15.7 Hedge Results from Buying 100 EDH2 at a Rate of 1.99% If, on May 28, 2010, Rates Changed to and Remained at Their Terminal Levels of March 19, 2012

90-Day LIBOR	ED P&L 5/28	ED P&L 6/17	Investment Proceeds 6/17	Total 6/17
.49%	\$360,000	\$363,680	\$100,122,500	\$100,486,180
1.99%	\$0	\$0	\$100,497,500	\$100,497,500
3.49%	-\$360,000	-\$386,210	\$100,872,500	\$100,486,290

tailed hedge approaches the unadjusted hedge. In the example, the tail of four declines to zero, or, equivalently, the tailed hedge increases from 96 to 100.⁷

TED Spreads

The value of a security relative to other securities depends on the choice of a benchmark set of securities that are assumed to be priced fairly. At the short-end of the term structure, ED futures are often thought of as fairly priced for two somewhat related reasons. First, they are quite liquid relative to many other short-term fixed income securities. Second, they are immune to individual security effects that complicate the determination of fair value. This second point is perfectly analogous to the argument, made in the introduction to Part One, that swaps are commodities while bonds have idiosyncratic characteristics.

*TED spreads*⁸ use rates implied by ED futures to assess the value of a short-term security relative to ED futures rates or to assess the value of one security relative to another. The idea is to find the spread such that discounting a security's cash flows at ED futures rates minus that spread gives that security's market price. This is analogous to spread defined in Chapter 3 (with some important caveats to be discussed later), but the TED spread is subtracted from rather than added to the base curve. This sign convention is chosen so that TED spreads of U.S. Treasuries turn out to be positive: since Treasury rates are almost always below LIBOR rates, subtracting a positive spread from ED futures rates can equate discounted Treasury cash flows with Treasury prices.

⁷Note that the numerical example of this subsection happens not to be affected by changing the tail over time since all of the ED P&L was assumed to occur on May 28, 2010.

⁸TED spreads originally compared rates on T-bill futures, which are no longer actively traded, with rates on Eurodollar futures. The name is a concatenation of *T* for Treasury and *ED* for Eurodollar.




FIGURE 15.1 Eurodollar Futures Contracts Used for Computing the TED Spreads of the $1\frac{3}{8}$ s of March 15, 2012, and the $4\frac{1}{2}$ s of March 31, 2012

To illustrate the computation of TED spreads, consider two U.S. Treasury bonds as of May 28, 2010, for settlement on June 1, 2010: the $1\frac{3}{8}$ s due March 15, 2012, and the $4\frac{1}{2}$ s due March 31, 2012. Figure 15.1 illustrates the discounting approach. The top line highlights the coupon dates of the $1\frac{3}{8}$ s while the bottom line highlights the coupon dates of the $4\frac{1}{2}$ s. The center line shows the division of time into periods covered by the various ED contracts. The first period is the *stub*, the time period from the settlement date to the start date of the LIBOR deposit underlying the first ED contract, in this case EDM0. The base curve, from which the TED spread will be subtracted, is made up of the collection of ED futures rates, each applied over the region indicated in Figure 15.1.

Table 15.8 presents the same date information as Figure 15.1, together with the relevant ED rates. The start dates of the periods are the settlement dates of the LIBOR deposits underlying the respective futures contracts. EDM0, for example, expires on June 14, 2010, so its underlying deposit starts two business days later, on June 16. The end dates are simply the start dates of the next ED contract. Defining the dates and rates in this way is somewhat problematic, but discussion of this point is deferred until the end

TABLE 15.8 Relevant Dates and Rates for two TED Spread Calculations as of May 28, 2010

Start Date	End Date	Contract	Rate	Days
6/1/2010	6/16/2010	STUB	.3407%	15
6/16/2010	9/15/2010	EDM0	.6000%	91
9/15/2010	12/15/2010	EDU0	.8450%	91
12/15/2010	3/16/2011	EDZ0	.9950%	91
3/16/2011	6/15/2011	EDH1	1.1250%	91
6/15/2011	9/21/2011	EDM1	1.2950%	98
9/21/2011	12/21/2011	EDU1	1.5050%	91
12/21/2011	3/21/2012	EDZ1	1.7550%	91
3/21/2012	3/31/2012	EDH2	1.9900%	10

of this subsection. Turning to the rates in Table 15.8, since the stub period is 15 days, a rate close to two-week LIBOR is chosen for that rate. The remaining rates are just the respective ED futures rates from Table 15.3.

Having explained the base rate curve to be used, Table 15.9 describes the calculation of the TED spread for $1\frac{3}{8}$ s of March 15, 2012. Let s be the TED spread, which, to repeat, is the single spread such that discounting a bond's cash flows at ED futures rates minus that spread gives the bond's price.⁹ The last row of Table 15.9 shows the bond's flat price, accrued interest, and full price. The main body of rows shows the reciprocal of the discount factor used for the bond's cash flow dates. (The reciprocal is displayed for easier readability.) For example, the discount factor to March 15, 2011, is determined by discounting over the 15 days of the stub period at the stub rate of .3407% minus the spread; over the 91 days from June 16, 2010, to September 15, 2010, at the EDM0 rate of .6% minus the spread; over the 91 days from September 15, 2010, to December 15, 2010, at the EDU0 rate of .845% minus the spread; and over the 90 days from December 15, 2010, to March 15, 2011, at the EDZ0 rate of .995% minus the spread. These particular day counts are shown in Figure 15.1.

Solving for the value of the spread so that the present value of the cash flows of the $1\frac{3}{8}$ s equals its full price gives $s = .500\%$. In words, the $1\frac{3}{8}$ s trade 50 basis points below the LIBOR curve as expressed through the stub and ED contracts in Table 15.8. The TED spread of the $4\frac{1}{2}$ s can be computed in a similar manner, with its full price of 107.809170 as of May 28, 2010, to give a spread of 53.4 basis points. A trader or investor could use these spreads to decide if Treasuries in this maturity range are priced appropriately relative

⁹The definition in this text corresponds to the *Spread Adjusted TED* on the Bloomberg TED screen.

TABLE 15.9 Discount Factors for TED Spread Calculation for the $1\frac{3}{8}$ s of March 15, 2012, as of May 28, 2010
Flat Price: 101-9.5; Accrued Interest: .291440; Full Price: 101.588315

Date	Reciprocal of Discount Factor
9/15/10	$\left(1 + \frac{(.3407\% - s) \times 15}{360}\right) \left(1 + \frac{(.6\% - s) \times 91}{360}\right)$
3/15/11	$\left(1 + \frac{(.3407\% - s) \times 15}{360}\right) \left(1 + \frac{(.6\% - s) \times 91}{360}\right) \left(1 + \frac{(.845\% - s) \times 91}{360}\right) \left(1 + \frac{(.995\% - s) \times 90}{360}\right)$
9/15/11	$\left(1 + \frac{(.3407\% - s) \times 15}{360}\right) \dots \left(1 + \frac{(1.125\% - s) \times 91}{360}\right) \left(1 + \frac{(1.295\% - s) \times 92}{360}\right)$
3/15/12	$\left(1 + \frac{(.3407\% - s) \times 15}{360}\right) \dots \left(1 + \frac{(1.505\% - s) \times 91}{360}\right) \left(1 + \frac{(1.755\% - s) \times 85}{360}\right)$

to LIBOR and if they are priced appropriately relative to one another. The mechanics of such trades are discussed in the next subsection.

The TED spread methodology described in this subsection is commonly used for computing spreads for short-term securities. The methodology, however, does have several conceptual problems. First, cash flows should be discounted at forward rates, not futures rates. This is not too much of a worry in practice since futures-forward differences are relatively small at shorter maturities, when TED spreads are used. Also, when comparing the TED spreads of two bonds, the futures-forward effect is present in both bonds and, therefore, will be even less important. Second, using the futures rates as described in this subsection does not use the right interest rate over partial periods. In computing the TED spread for the $4\frac{1}{2}$ s, for example, as shown in Figure 15.1, the EDH2 rate of 1.99%, which is the futures rate over the three months starting March 21, 2012, is applied to the period from March 21, 2012, to March 31, 2012. But with an upward-sloping term structure, the rate over just those 10 days is less than 1.99%. Third, the rates and dates in Table 15.8 do not really line up with each other. For example, the LIBOR deposit underlying EDM1 is from June 15, 2011, to September 15, 2011, and the market, presumably, prices EDM1 accordingly. Nevertheless, Table 15.8 uses the EDM1 rate of 1.295% from June 15, 2011, all the way through to September 21, 2011, because the underlying deposit for the next contract, EDU1, does not start until September 21, 2011.

TED spreads are widely used despite the problems listed in the previous paragraph. The great transparency and liquidity of futures contracts, combined with the simplicity of the TED spread, outweighs the relatively minor futures-forward differences and the rate-date discrepancies. Practitioners not willing to make this trade-off have to compute spreads by discounting with rates from a rigorous curve-fitting methodology, as described in Chapter 21.

TABLE 15.10 ED Futures Hedges and Forward Bucket '01s for USD 100 Million Face Amount of the $1\frac{3}{8}$ s due March 15, 2012, and the $4\frac{1}{2}$ s due March 31, 2012, as of May 28, 2010

Contract	$1\frac{3}{8}$ s of 3/15/2012		$4\frac{1}{2}$ s of 3/31/2012	
	Bucket '01	Contracts	Bucket '01	Contracts
Ticker	(\$)		(\$)	
STUB	423	17	449	18
EDM0	2,567	103	2,725	109
EDU0	2,548	102	2,676	107
EDZ0	2,547	102	2,665	107
EDH1	2,529	101	2,617	105
EDM1	2,721	109	2,807	112
EDU1	2,510	100	2,554	102
EDZ1	2,343	94	2,547	102
EDH2	0	0	281	11
Total	18,190	728	19,321	773

Hedging Bonds with ED Futures

Suppose a trader believes that a TED spread is too wide, e.g., that the appropriate spread between the $1\frac{3}{8}$ s and the LIBOR curve is less than 50 basis points, or alternatively, that the Treasury bond is too rich in price relative to the LIBOR curve. The trade would then be to short the $1\frac{3}{8}$ s and buy some portfolio of ED futures. What is the right hedging portfolio?

Relying on the pricing methodology used to compute TED spreads, computing hedge portfolios is quite straightforward:

1. Decrease a particular futures (or stub) rate by one basis point.
2. Keeping the TED spread unchanged, calculate the resulting change in the value of the bond, i.e., calculate the bucket '01 with respect to that futures rate.¹⁰
3. Divide the change by \$25, the basis-point value of all ED futures contracts, to get the number of contracts that will hedge that bucket risk.
4. Repeat steps 1 to 3 for all pertinent futures rates.

Table 15.10 shows the hedge calculations for \$100 million face amount of the $1\frac{3}{8}$ s and $4\frac{1}{2}$ s. To describe the order of magnitude of these results, focus first on the $1\frac{3}{8}$ s. Since the full price of the bonds is about \$101.6 million and since each ED contract has \$1 million face amount, the rough order of

¹⁰This step results in the *spread-adjusted hedge* of the Bloomberg TED screen.

magnitude of the hedge should be 101.6 of each of the contracts, or 102 to the nearest contract. The number of contracts assigned to the stub and to EDZ1 will be less than this, of course, because the stub period covers only 15 days of risk and the bond's maturity of March 15, 2012, means that protection is not required for the full EDZ1 period. See Figure 15.1. Note, by the way, that the stub period has to be hedged with securities other than 90-day ED contracts, like fed fund futures, which are discussed later in this chapter. But returning to orders of magnitude, the number of EDM1 contracts in the hedge will be higher than 102 since the period assigned to it in Table 15.8, from June 15, 2011, to September 21, 2011, is 98 days, longer than the 91 days assigned to the other contracts.

It should be pointed out that the hedging results in Table 15.10 automatically account for the tail: setting the change in the present value of the bond equal to the change in the daily settlement payment of the hedge is the whole point of the tail. See Chapter 13. To show this is the case with an example, consider the EDU1 holding against the $1\frac{3}{8}$ s. The tail approximation sets the number of contracts in the hedge equal to the present value of the unadjusted hedge of 101.6 contracts. With 477 days between bond settlement and the expiration of EDU1 on September 21, 2011, and using the weighted average of rates in Table 15.8 over that period, which comes to .9569%, the tailed hedge would be

$$\frac{101.6}{\left(1 + \frac{.9569\% \times 477}{360}\right)} = 100.3 \quad (15.12)$$

or about the 100 contracts given through the more careful methodology of Table 15.10. More generally, the tail effect causes the number of contracts in the hedge to fall with their terms to expiration, an effect clearly seen in Table 15.10.

Rough numbers for the contracts hedging the $4\frac{1}{2}$ s can be understood similarly. With a full price of 107.8, the roughest approximation is a hedge of about 108 of each contract. The actual number will be much lower for the stub, higher for EDM1, and much lower for EDH2, due to the relative number of days covered by each of those contracts. And because of the tail, the general trend is for the number of contracts to decline with the terms of the contracts.

The bucket '01s in Table 15.10 are very similar to the forward bucket '01s presented in Chapter 5, except that the '01s here are with respect to futures rather than forward rates. To take an example, the price of the $1\frac{3}{8}$ s increases by \$2,529 per \$100 million face amount of the $1\frac{3}{8}$ s for every basis-point decline in EDH1's rate.

The last row of Table 15.10 gives the totals of the rows. The \$18,190 and \$19,321 sums of the bucket '01s for the $1\frac{3}{8}$ s and the $4\frac{1}{2}$ s, each per

\$100 million face amount, are DV01-equivalents of .018190 and .019321 per 100 face amount with respect to a parallel shift of ED rates.

The focus of this subsection has been hedging a position in one bond with ED futures. The analysis can easily be extended to hedge the residual risk from trading of one bond *versus* another. For example, were a trader to sell \$100mm of the $4\frac{1}{2}$ s and buy a DV01-neutral face amount of the $1\frac{3}{8}$ s, the differences between the contracts required to hedge each of the individual bond positions would hedge the residual curve risk of the combined bond position.

FED FUNDS AND RELATED SECURITIES

The Fed Funds Rate

Banks are required to keep a minimum level of reserves at the Fed to support their deposit liabilities. At any particular time, however, individual banks may find that they have excess cash that can be invested more profitably elsewhere or that they have a cash deficit and need to borrow to meet minimum reserve requirements. The market in which banks trade funds overnight is called the *federal funds* or *fed funds* market. While only banks can borrow or lend in the fed funds market, the central position of banks in the financial system causes other short-term interest rates to move together with the fed funds rate. To take just one example, repo rates, which were introduced in Chapter 12, are particularly correlated with fed funds rates.

The Fed or, more specifically, the Federal Open Market Committee (FOMC), sets monetary policy in the United States. An important component of this policy is the *targeting* or *pegging* of the fed funds rate at a rate deemed consistent with the goals of price stability and full employment. Since banks trade freely in the fed funds market, the Fed does not directly set the fed funds rate. By using the tools at its disposal, however, including buying and selling securities in its *open market operations*, mostly through repo, the Fed has enormous power to keep the fed funds rate close to the desired target. The bold line in Figure 15.2 shows the time series of the target rate, along with the realized rate, to be discussed in a moment.

In an effort to make its workings more transparent, the Fed began to announce its target level for the fed funds rate in 1995. Furthermore, starting in 2000, the FOMC began issuing statements about its deliberations and decisions. While often taken for granted today, this level of transparency is quite different from that in the past, where “Fed watchers” had to observe and analyze open market operations to figure out what the Fed was trying to accomplish.




FIGURE 15.2 Effective vs. Target Fed Funds Rate

The schedule of FOMC meetings is announced well in advance. And while the Fed may change its target rate between meetings, such changes are being reserved more and more for extraordinary circumstances. As of this writing, the last intra-meeting change to the target rate was on October 8, 2008, when the Fed lowered the rate from 2% to 1.50% in reaction to the events surrounding the bankruptcy of Lehman Brothers. The intra-meeting changes before that were in 2007 (September 17, April 18, and January 3).¹¹

Each business day the Fed calculates and publishes the weighted average rate at which banks borrow and lend money in the fed funds market, a rate called the *fed funds effective rate*. Figure 15.2 shows the time series of the effective rate against the target from May 1997 to November 2010. For the most part, the Fed does keep the fed funds rate close to the target rate. From the start of the sample period to December 15, 2010, the effective rate, on average, was less than two basis points below the target rate. From December 16, 2010, to the end of the sample, with the target rate not a single rate but a range from 0 to 25 basis points, the effective rate averaged about 17 basis points.

The relatively isolated and large deviations of effective from target, visible in Figure 15.2, are due to particular pressures in financial markets.

¹¹In an intra-meeting policy change at the onset of the subprime crisis, on August 17, 2007, the Fed lowered its *discount rate*, which is the rate at which banks can, in case of need, borrow directly from the Fed. The fed funds target rate, however, was not reduced until the subsequent scheduled meeting in September.

Examples include the Russian debt crisis (when the effective rate spiked to over 80 basis points above target at one time and to over 150 basis points at another) and the bankruptcy of Lehman Brothers (when fed funds effective spiked from 60 to 80 basis points above target). The reason for these spikes is that, during times of financial upheaval, the value of liquidity or cash rises dramatically; individuals might rush to withdraw cash from vehicles not considered completely safe while banks, other financial institutions, and corporations might become reluctant to lend cash, even when secured by collateral. (See the cases of Bear Stearns and Lehman Brothers in Chapter 12.)

At times of financial uncertainty, the Fed tries to stabilize conditions by “injecting liquidity into the system,” that is, by lending cash on acceptable collateral. In fact, as a result of these policy responses, most significant deviations of the effective rate from target in recent years have been for the effective rate to be significantly below target. At the earliest signs of crisis, in August 2007, the effective rate fell to 50 to 70 basis points below target; in the weeks after Lehman’s bankruptcy, it fell to 133 below target and, slightly later, fluctuated between 50 and 90 basis points below target; and finally, following the market disruption in the wake of September 11, 2001, the effective rate traded from about 80 to 180 basis points below target.

Another example of fed funds effective differing from target is the *year-end effect*. Historically, significant excess demand by banks to borrow funds at the end of a calendar year, largely for accounting reasons, caused fed funds to trade significantly above target over the “turn” or year-end. One of the last instances of this can be seen in Figure 15.2, as effective traded about 50 basis points above target at year-end 1997. In subsequent years, however, due to a slackening demand for year-end funds or more aggressive compensation by the Fed, the effective rate traded below target over the turn: about 70 basis points below at year-end 1998, 150 below at year-end 1999 (a turn compounded by “Y2K” concerns), and over 100 below at year-end 2000. Since then, however, the turn has become mostly a non-event, with effective trading at about target, although there are occasional surprises, e.g., year-end 2007, when effective traded at about 120 basis points below target.

Fed Fund Futures

Just as ED futures are used to hedge exposure to short-term LIBOR rates, *fed fund futures* are used to hedge exposure to fed funds rates. Table 15.11 lists the first several fed fund contracts as of May 28, 2010. Note that the tickers are a concatenation of “FF” for fed funds, a letter indicating the month of the contract’s expiration, and a single digit for the year of the contract. Each contract expires on the last day of the month, but, as in any futures contract, changes in contract value are settled daily.

The fed funds futures contract is designed as a hedge to a \$5,000,000 30-day deposit in fed funds. First, the daily settlement payment of a contract

TABLE 15.11 Selected Fed Funds Futures Prices and Rates as of May 28, 2010

Ticker	Month	Price	Rate
FFM0	June	99.780	.220
FFN0	July	99.770	.230
FFQ0	August	99.760	.240
FFU0	September	99.750	.250
FFV0	October	99.735	.265
FFX0	November	99.705	.295

is set at \$41.67 per basis point since changing the rate of a \$5,000,000 30-day loan by one basis point changes the interest payment by \$41.67:

$$\$5,000,000 \times \frac{.0001 \times 30}{360} = \$41.67 \quad (15.13)$$

Second, the final settlement price of a fed funds contract in a particular month is set to 100 minus 100 times the average of the daily effective fed funds rates over that month. Table 15.12 illustrates this calculation: in June 2010 average effective was .177% implying a final settlement price for the June contract of $100 - 100 \times .177\%$ or 99.823. Note that nonbusiness days, indicated by italics in the table, are included in the average with rates of the previous business day, consistent with the market convention that an investor earns Friday's rate for Friday, Saturday, and Sunday. Note too that

TABLE 15.12 Calculation of the Settlement Price of FFM0, the 30-Day Fed Fund Futures Contract That Matured in June 2010

Date	Rate	Date	Rate	Date	Rate
6/1/2010	.20	6/11/2010	.18	6/21/2010	.17
6/2/2010	.20	6/12/2010	.18	6/22/2010	.18
6/3/2010	.19	6/13/2010	.18	6/23/2010	.17
6/4/2010	.19	6/14/2010	.18	6/24/2010	.16
6/5/2010	.19	6/15/2010	.19	6/25/2010	.16
6/6/2010	.19	6/16/2010	.19	6/26/2010	.16
6/7/2010	.19	6/17/2010	.19	6/27/2010	.16
6/8/2010	.19	6/18/2010	.18	6/28/2010	.17
6/9/2010	.18	6/19/2010	.18	6/29/2010	.15
6/10/2010	.18	6/20/2010	.18	6/30/2010	.09
Average rate					.176
Average to nearest tenth of a basis point:					.177
Settlement Price					99.823

the average is rounded to the nearest tenth of a basis point before being used to calculate the futures settlement price.

Hedging Borrowing and Lending with Fed Fund Futures over 30-Day Months

To see how fed funds futures contracts works as a hedge, consider the case of a small regional bank that has surplus cash of \$5 million over the month of June 2010. The bank plans to lend this \$5 million overnight in the fed funds market over the month but wants to hedge the risk that a falling effective fed funds rate will reduce the interest it earns. Therefore, the bank buys one June fed fund futures contract for 99.780 at the end of the last trading day in May, corresponding to a rate of .220%. See Table 15.11.

Assume for the moment that the bank earns simple interest on its fed funds lending. Then, given the realized effective rates in Table 15.12, the bank would earn interest in the fed funds market of

$$\$5,000,000 \times \frac{30 \times .17\bar{6}\%}{360} = \$736.11 \quad (15.14)$$

In addition, the sum of the daily settlements from having bought the June contract at 99.780 (or a rate of .22%) and held it to expiration at a final settlement price of 99.823 (or a rate of .177%) is

$$\$5,000,000 \times \frac{30 \times (.22\% - .177\%)}{360} = \$41.67 \times 4.3 \quad (15.15)$$

$$= \$179.18 \quad (15.16)$$

The direct investment proceeds of \$736.11 plus the hedge profit of \$179.18 equals \$915.29. Were it not for the rounding of the rate for the determination of the final settlement price, the sum of the left-hand sides of equation (15.14) and (15.16) would identically equal interest earned for 30 days at .22%, i.e., at the rate to be locked in by hedging with the June futures contract:

$$\$5,000,000 \times \frac{30 \times .22\%}{360} = \$916.67 \quad (15.17)$$

This fed fund futures hedge is less than perfect for two other reasons. First, the bank does not really earn simple interest on its fed funds investments over the month: it can compound the interest it earns daily. Second, since the fed fund futures contract is settled daily, its profit or loss depends on the financing of daily settlement payments. These two effects are typically small given the short maturity of liquid futures contracts, but would become more significant at much higher levels of rates.

A bank is the subject of the hedging example of this subsection because only banks can participate in the fed funds market. But, as mentioned earlier, because many short-term rates are highly correlated with the fed funds rate, other financial institutions, corporations, and investors use fed funds futures to hedge other short-term rate risks. For example, a corporation that needs to borrow money over the month of June might hedge against the risk that rising rates increase the cost of that borrowing by selling June fed funds futures. In this context, however, it is important to note that this hedge will insulate the corporation’s borrowing costs from changes in the general level of interest rates, but not from changes in the spread of general corporate borrowing rates over fed funds nor, of course, from changes in the spread of that particular corporation’s borrowing rate over fed funds. The difference between the actual risk (e.g., changes in a corporation’s borrowing rate) and the risk reduced by the hedge (e.g., changes in the fed funds rate) is known as *basis risk*.

Hedging Borrowing and Lending with Fed Fund Futures over Horizons Other than 30 Days

When fed funds futures are used to hedge interest rate risk over horizons other than full 30-day calendar months, buying or selling one contract per \$5 million face amount is not the correct hedge. This subsection is a detail of trading, but its results will be used in the TED spread trading case study later in this chapter.

To put the hedging issues of this subsection in a context, consider the ED hedges of the $1\frac{3}{8}$ s of March 15, 2012, in Table 15.10. On June 14, 2010, EDM0 expires and the bond position is again exposed to interest rate risk from that date up to September 15, 2010, when protection from the EDU0 position kicks in. (See Table 15.8.)

Figure 15.3 shows the regions to be hedged and their corresponding hedging instruments. June fed fund futures are needed to hedge the risk over the remainder of that month; July and August fed fund futures are needed to hedge the risk over 31-day months, and September fed fund futures are needed to hedge the risk from the start of September up to September 15. The correct hedges in each of these situations are now considered in turn.



FIGURE 15.3 Hedging the $1\frac{3}{8}$ s of March 15, 2012, as of June 14, 2010, with Fed Fund and ED Futures

Hedging from Within a 30-Day Month to the End of the Month In Figure 15.3, June fed fund futures are used to hedge the interest rate risk of the $1\frac{3}{8}$ s from the expiration of EDM0 on June 14 to June 30. The risk of a one basis-point decline in rates over those 17 days, June 14 to June 30, inclusive,¹² is

$$\$101,600,000 \times \frac{17 \times .01\%}{360} \quad (15.18)$$

What is the risk of the June fed funds contract over the same holding period? Since the contract settles based on the average of the effective rate over the month, a one basis-point change in the rate over the last 17 days of June will not change the average or the final settlement rate by a full basis point but only by $\frac{17}{30}$ of a basis point. Hence, the change in the settlement value of the 30-day fed fund contract is

$$\$5,000,000 \times \frac{30 \times \frac{17}{30} \cdot 01\%}{360} = \$5,000,000 \times \frac{17 \times .01\%}{360} \quad (15.19)$$

which, adjusting for notional amount, is exactly the exposure to be hedged in (15.18). Hence, hedging the long position in the $1\frac{3}{8}$ s requires the sale of $\frac{101.6}{5}$ or approximately 20 contracts. More generally, within a 30-day month, the number of required contracts is one per \$5 million face amount. Furthermore, it follows that if a fed fund hedge of one contract per \$5 million of face amount is put on at or before the start of a 30-day month, the hedge rolls off appropriately during the month. In other words, the hedge is valid within the month as well and does not have to be adjusted during the month.

Months with 31, 28, or 29 Days In Figure 15.3, the interest rate risk of July 2010 and August 2010 is to be hedged with July and August fed fund futures, respectively. How many contracts of each should be sold against the \$101.6 million of the $1\frac{3}{8}$ s? Over a month with 31 days, the risk of a one-basis point change in the simple rate of interest over the entire month is

$$\$5,000,000 \times \frac{31 \times .01\%}{360} = \$43.06 \quad (15.20)$$

But, by construction, fed fund futures contracts are designed to hedge 30-day deposits with an exposure of \$41.67 per basis point. Hence, to hedge the interest rate risk of a \$5,000,000 loan over a 31-day month requires $\frac{43.06}{41.67} = \frac{31}{30}$ or 1.033 contracts. Similarly, hedging the interest risk of a

¹²These contracts expire at 11 A.M. London time, so a U.S. trader is exposed to the rate risk of the 14th as well.

28- or 29-day month requires $\frac{28}{30} = .933$ or $\frac{29}{30} = .967$ contracts, respectively, per \$5 million loan amount.

As time passes, hedges for 31-, 29-, or 28-day months do not need to be adjusted, just as argued in the case of 30-day months.

Hedging from the Start of a Month to Sometime During the Month In Figure 15.3, the interest rate risk from September 1 to September 14, inclusive,¹³ is to be hedged with September fed fund futures. At any time before the start of September, there are 14 days of interest rate risk and 30 days of a 30-day month at risk in a 30-day contract. Hence, each \$5 million face amount can be hedged with $\frac{14}{360} \div \left[\left(\frac{30}{30} \right) \times \frac{30}{360} \right]$ or .467 contracts.

During the month of September, however, the problem is a bit more complicated. Consider the close of business of September 4. There are 10 days of risk left to be hedged and 26 days of a 30-day month at risk in a 30-day contract, implying a hedge of $\frac{10}{360} \div \left[\left(\frac{26}{30} \right) \times \frac{30}{360} \right]$ or .385 contracts per \$5 million face amount. But this means that the .467 contracts per \$5 million that were correct before September decline to .385 by September 4. In other words, the correct number of contracts in the hedge declines as time passes. Somewhat more generally, as of the close of business of the m^{th} day of September, $1 \leq m \leq 14$, the correct hedge is $\frac{14-m}{30-m}$ contracts per \$5 million face amount.

Application: Market Expectations of Fed Policy Changes at the Start of the 2004–2006 Tightening Cycle

In deciding whether to buy or sell fed fund futures, market participants certainly consider their expectations of the fed funds effective rate that will prevail in the future and, by implication, the target rates that will be set by the FOMC at its upcoming policy meetings. Making the simplifying assumption that fed funds futures rates are determined solely by consensus expectations, market analysts commonly extract the expectations implied by market rates. Traders and investors then compare these implied expectations with their individual views about future rates to determine how to trade short-term rate products. Examples would include a trader who sells fed fund futures, thinking that market rate expectations are too low (and futures prices too high), and a portfolio manager who decides not to sell futures to hedge an

¹³EDU0 expires on September 13 at 11 A.M. London time. Its underlying three-month LIBOR deposit settles on September 15 at a rate determined on the morning of the 15th. Hence it is reasonable to use EDU0 to cover the risk on the 15th. Most important, however, is to make sure that the total number of days covered by the combination of fed funds and ED contracts is correct.

TABLE 15.13 Fed Fund Futures Rates as of April 15, 2004

Ticker	Month	Rate
FFK4	May	1.020%
FFM4	June	1.030%
FFN4	July	1.105%
FFQ4	August	1.220%
FFU4	September	1.315%
FFV4	October	1.415%
FFX4	November	1.550%
FFZ4	December	1.665%
FFF5	January '05	1.745%

exposure to short-term interest rates, thinking that the cost of such insurance is too high, i.e., market rate expectations are too high (and futures prices too low). The question of whether market expectations are the sole determinants of fed funds futures prices will be discussed at the end of this application.

In late 2010 and early 2011, expectations of FOMC decisions with respect to the target rate were not particularly interesting: the overwhelming consensus was that the FOMC would keep the target rate at 0 to 25 basis points for an extended period of time. A more interesting application is market expectations as of April 15, 2004, when it was commonly believed that the FOMC would start a tightening cycle, gradually increasing the target rate from its then-current level of 1%. That this was indeed the beginning of a tightening cycle is clearly evident from Figure 15.2.

Table 15.13 gives the first 10 fed fund futures contracts and their rates as of April 15, 2004. Table 15.14 lists the relevant FOMC meeting dates,¹⁴ implied market expectations of FOMC rate increases, the calculations of which will be described presently, and, in the last column, the target rates actually chosen at the respective FOMC meetings. Figure 15.4 presents the same information as this table in a picture; note that rates jump in this figure the day after the meeting dates.

As of April 15, 2004, the market assigned a small probability, only 9.2%, that the FOMC would increase the target rate and start the tightening cycle at its May 4 meeting. And, as it turned out, the FOMC did not change the target rate at that meeting. But the market on April 15 did not do so well in predicting the outcome of subsequent meetings. At each of the next five meetings the FOMC increased the target rate by 25 basis points, while the market had expected a significantly slower increase in the target rate, or


¹⁴Some meetings last for two days, in which case the announcement of the target rate decision is made on the second day. Therefore, Table 15.14 lists the dates of the one-day meetings or of the second day of the two-day meetings.

TABLE 15.14 A Set of Expectations of FOMC Target Rates from Fed Fund Futures Prices as of April 15, 2004

FOMC Meeting Date	Market-Implied Probability of a 25bp Tightening	Expected Rate Path	Rate After Meeting
3/16/04			1.00%
5/4/04	9.2%	1.023%	1.00%
6/30/04	32.8%	1.105%	1.25%
8/10/04	67.9%	1.275%	1.50%
9/21/04	56.1%	1.415%	1.75%
11/10/04	81.0%	1.617%	2.00%
12/14/04	51.0%	1.745%	2.25%

equivalently, had assigned probabilities significantly less than 100% to each of these 25-basis point increases.

The discussion now turns to the calculation of the implied probabilities and expected rates reported in Table 15.14. The target rate on April 15, 2004, is 1% and, according to Table 15.13, the May fed funds futures rate is 1.02%. To extract market expectations about the next FOMC meeting, on May 4, make the following five assumptions, which will be discussed further below: 1) the FOMC will change the target rate only on a meeting date; 2) on a meeting date the FOMC will either increase the target rate by 25 basis points, with probability p , or leave the rate unchanged with probability $1 - p$; 3) the change in the target rate moves fed funds

**FIGURE 15.4** An Expected Path for the Fed Funds Target Rate as of April 15, 2004 vs. the Realized Path

effective the day after the meeting; 4) the average difference between fed funds effective and the fed funds target rate is zero over contract months; and 5) futures prices are determined solely by expectations.

Under the assumptions just made, the fed funds target rate from May 1 to May 4, inclusive, is 1%, while the rate from May 5 to May 31, inclusive, is 1% with probability $1 - p$ and 1.25% with probability p . Hence, for the average fed funds rate over the May contract month to equal its market rate of 1.02%, it must be that

$$\frac{4 \times 1\% + 27 \times [(1 - p) \times 1\% + p \times 1.25\%]}{31} = 1.02\%$$

$$1\% + \frac{27}{31} \times p \times .25\% = 1.02\% \quad (15.21)$$

Solving, $p = 9.185\%$, which also means that the expected rate after the meeting is $(1 - p) \times 1\% + p \times 1.25\%$ or 1.023%.

To take one more example, Table 15.14 reports that the expected rate after the June 30 meeting is 1.105%. What probability of a 25 basis-point increase at the August 10 meeting, to an expectation of $1.105\% + .25\%$, or 1.355%, correctly prices the August futures contract? Proceeding as before, solve for p such that

$$\frac{10 \times 1.105\% + 21 [(1 - p) 1.105\% + p \times 1.355\%]}{31} = 1.22\%$$

$$1.105\% + \frac{21}{31} \times p \times .25\% = 1.22\% \quad (15.22)$$

Solving, $p = 67.905\%$.

The rest of the results in Table 15.14 are calculated along the same lines. It is worth pointing out, however, that there are more fed fund futures contracts than there are FOMC meetings. This has two implications. One, an analyst has to select which contracts are to be used to calibrate which probabilities. Two, prices of contracts not used will not necessarily be priced correctly by these calibrated probabilities. For example, with FOMC meetings on May 4 and June 30, and with an expected rate of 1.023% after the May meeting determined by the May futures price, the June futures price should be 1.023% as well. But the market price of that contract, given in Table 15.13, implies a rate of 1.03%.¹⁵ A different approach, by the way, is

¹⁵For completeness, note that the calculations behind Table 15.14 use the October contract to imply the probability of a tightening at the September 21 meeting and the January contract to imply the result of the December 14 meeting. These choices mean that the resulting expected path of rates will not necessarily price the September and December contracts correctly.

to solve for the expected path that minimizes some measure of the discrepancy between the monthly averages of rates along that path and all market futures prices.

This subsection concludes by discussing the assumptions made at the beginning of the application. Support for the assumption that the FOMC will change the target rate only on meeting dates has been presented earlier in this chapter. Except in extraordinary times, the complexity of a model that incorporates the possibility of inter-meeting changes is probably not worth the trouble. The assumption that the FOMC changes rates by zero or 25 basis points was particularly suited to the anticipated tightening cycle as of April 2004. The prevailing view was that the FOMC would increase rates gradually, which was in fact the case: no single meeting increased the target rate by more than 25 basis points over the entire tightening cycle. In other macroeconomic situations, however, changes of more than 25 basis points per meeting might well require consideration. An obvious difficulty from the modeling perspective is that one fed fund futures price can imply only one probability, not a full distribution of probabilities, i.e., the probability of no move, of a 25 basis-point move, of a 50 basis-point move, etc. Possible solutions include making somewhat arbitrary assumptions about this probability distribution or using the prices of other short-term securities, like options on ED futures,¹⁶ to infer a richer set of market probabilities.

The assumption that changes in the fed funds target rate affect fed funds effective the day after the meeting stems from the fact that most fed funds trading takes place early in the day. By the time the FOMC announces its decision in the afternoon of a meeting day, most trading for that day has finished so that fed funds effective has, for the most part, been already determined. The assumption that fed funds effective averages to the target rate is made here purely for convenience. Any additional information that market professionals have about the difference between effective and target, e.g., year-end effects or other calendar regularities, should certainly be incorporated into the analysis.

The last and far from least important assumption made above is that futures prices are determined solely by expectations. Chapter 8 discussed the interest rate risk premium, which tends to cause forward rates to exceed the expectation of future rates,¹⁷ and there is no reason to believe that this risk

¹⁶Using ED futures options for this purpose requires adjustments to account for the spread between LIBOR and fed funds. Making such adjustments is not difficult, however, given the existence of LIBOR-fed funds basis swaps. See Chapter 16.

¹⁷Chapter 8 also showed that convexity tends to lower forward rates relative to expectations but, in the context of this chapter, which focuses on very short-term rates, the convexity effect is very small.

premium is zero at the short end of the curve. It is not difficult to incorporate a known risk premium into the calculations of this application, although, as explained in Chapter 8, it is difficult to estimate the risk premium in the first place.

Overnight Index Swaps

With *overnight index swaps (OIS)*, counterparties swap payments at a fixed rate for payments based on fed funds effective. Just as FRAs are an over-the-counter alternative to ED futures for hedging exposure to LIBOR rates, OIS are an over-the-counter alternative to FF futures for hedging exposure to fed funds effective. OIS have several advantages relative to fed funds. First, as will become clear presently, OIS are more accurate than fed fund futures for hedging the risk of rolling loans at fed funds effective. Second, OIS dates are customizable, although particular maturities and forward structures do trade with much more liquidity than others. Third, OIS are liquid out to two to three years, while fed fund futures are liquid out to four to eight months.

OIS have become particularly well-known since the financial crisis of 2007–2009 through the *LIBOR-OIS spread*, which has been adopted as a premier measure of stress in the financial system. The idea behind using this spread as a measure of stress, along with the historical behavior of that spread, will be presented in the next section.

Consider \$100 million notional amount of a 14-day OIS swap, from June 1 to June 14, 2010, in which party A agrees to pay .2% to party B while party B agrees to pay compounded fed funds effective. The fixed side of the swap is simple, with party A paying $\$100\text{mm} \times \frac{14 \times .2\%}{360}$ or \$7,777.78 at expiration of the swap. Using the fed funds effective rates from Table 15.12, Table 15.15 shows how to calculate the floating rate payment. Starting with a notional amount of \$100 million, the notional accumulates at fed funds effective on an actual/360 basis, where this accumulation includes compounding. This means that interest over the first day, June 1, is based on the original notional amount and fed funds effective for June 1, i.e., $\$100\text{mm} \times \frac{.2\%}{360}$ or \$555.556, while interest over June 10 is based on the accumulated notional as of June 9 and fed funds effective for June 10, i.e., $\$100,004,777.87 \times \frac{.18\%}{360}$ or 500.02. Note too that simple interest is earned over non-business days, consistent with money market convention. Hence, the total interest over June 4, 5, and 6 equals $\$100,001,638.90 \times \frac{3 \times .19\%}{360}$ or \$1,583.36. The floating payment of the OIS at expiration is the sum of the interest payments in Table 15.15 or \$7,278.01.

The important point to take away from the calculations of Table 15.15 is that the resulting floating rate payment replicates the interest a lender would earn from rolling over an initial principal amount of \$100 million in the fed funds market for 14 days. Hence, the combination of the following three trades locks in daily compounded lending at .2% over the life of the

TABLE 15.15 Calculating the Floating Payment of a 14-Day OIS Swap

Date	Rate	Days	Interest	Accumulated Notional
				100,000,000.00
6/1/10	.20%	1	555.56	100,000,555.56
6/2/10	.20%	1	555.56	100,001,111.11
6/3/10	.19%	1	527.78	100,001,638.90
6/4/10	.19%	3	1,583.36	100,003,222.26
6/7/10	.19%	1	527.79	100,003,750.05
6/8/10	.19%	1	527.80	100,004,277.85
6/9/10	.18%	1	500.02	100,004,777.87
6/10/10	.18%	1	500.02	100,005,277.89
6/11/10	.18%	3	1,500.08	100,006,777.97
6/14/10	.18%	1	500.03	100,007,278.01
Total		14	7,278.01	

swap: receive .2% fixed rate from the OIS swap; pay the realized interest amount calculated as in Table 15.15; and earn interest from rolling \$100 million in the fed funds market.

For OIS with maturities greater than one year, payments are annual. To determine the payment dates, start from the maturity date and count back in annual increments. An 18-month OIS, for example, would make its last payment in 18 months and its earlier and first payment in 6 months.

The most obvious uses of OIS are for hedging borrowing and lending in the fed funds market or borrowing and lending in a market where the rate tracks fed funds closely. One important example of the latter is the repo market. As discussed in Chapter 12, the general collateral repo rate is usually somewhat below the fed funds rate. Hence, a market participant that is borrowing or lending overnight in the repo market can hedge changes in the repo rate by overlaying an OIS swap. The hedge is not perfect, of course, because the spread between repo and fed funds can change, as also discussed in Chapter 12.

The ability to customize OIS dates creates additional uses for the product. Two relatively liquid examples of this include the following: one, OIS between two sequential IMM dates, which, in combination with ED futures, can be used to hedge or speculate on future spreads between LIBOR and fed funds rates; two, OIS between two sequential meeting dates of the FOMC, which can be used to hedge or speculate on the policy action taken at the first of those two meetings.

This subsection, as part of a section on fed funds, described USD OIS. OIS in other currencies are quite similar in spirit, although based on different

overnight rates and, in some cases, different day-count conventions. For EUR the OIS reference rate is EONIA; for JPY, TONAR;¹⁸ and for GBP, SONIA.

LIBOR-OIS AS AN INDICATOR OF FINANCIAL STRESS

Consider a scenario in which a bank has decided to fund itself at a rate that varies with fed funds effective and is choosing between two ways of doing so. First, the bank can borrow directly in the fed funds market at the fed funds rate. Second, the bank could borrow for three months at a three-month LIBOR rate of 1%, pay fed funds effective through a three-month OIS, and receive .8% over three months through that same OIS. In this example, the three-month LIBOR-OIS spread is 1%–.8% or 20 basis points, and the second of the bank's funding options results in a total funding cost of fed funds plus 20 basis points. Why would the bank prefer this second option, at fed funds plus 20, when it could borrow overnight directly at fed funds flat?

The answer is that, in the second option, the bank locks up its funding source for three months. Should there be a credit event, which lowers the creditworthiness of the individual bank or the perceived creditworthiness of the banking system, or should there be a liquidity event, which makes investors extremely reluctant to part with cash, the bank will have some time to sort out its funding sources and requirements. And in the special case of a short, intense crisis, the locked-up funding might allow the bank simply to wait until the storm passes. By contrast, had the bank been borrowing directly in the overnight fed funds market as a credit or liquidity event struck, its access to the fed funds market and other funding alternatives might vanish too quickly to be replaced. And were that the case, the bank might have to sell assets or even itself at fire-sale prices or, in the most extreme crises, to declare bankruptcy.

The implications of this discussion are that banks should be willing to pay something to lock in their funding and that a market price for locking up three-month funding relative to overnight funding is given by the LIBOR-OIS spread. Furthermore, the greater the concern about credit and liquidity conditions, the greater that spread would be.

This interpretation of the LIBOR-OIS spread, combined with its behavior over 2007–2009, to be described presently, has led to its adoption as a key indicator of financial stress. Figure 15.5 shows the three-month LIBOR-OIS spread from October 2005 to July 2010, for both USD and EUR. The averages of these spreads from October 2005 to the beginning of July 2007

¹⁸The Overview briefly describes these rates.




FIGURE 15.5 Three-Month LIBOR Minus Three-Month OIS Rates in USD and EUR

were about 8 basis points in USD and about 6 basis points in EUR, suggestive of minimal perceived credit and liquidity risks. As the subprime crisis intensified through the second half of 2007 and into 2008, these spreads widened to new levels of 50 to 100 basis points. Then, from the bankruptcy of Lehman Brothers in mid-September 2008, spreads spiked, with the USD spread reaching a peak of 365 basis points on October 10, 2008. Since then, as shown in the figure, the spread has fallen considerably.

TRADING CASE STUDY: SHORTING THE TED SPREAD OF THE $1\frac{3}{8}$ S OF MARCH 15, 2012

This case illustrates how to short a bond against ED futures, with focus on hedging the stub and the complications of the hedging contracts' maturing over the life of the trade. Based on the analysis earlier in this chapter, a trader decides to short the $1\frac{3}{8}$ s of March 15, 2012, against ED futures at a TED spread of 50 basis points. The ED futures hedge is along the lines of the analysis earlier in this chapter. To hedge the stub, however, the trader finds it most convenient to buy fed fund futures, constructing this part of the hedge along the lines of the hedging analyses earlier in this chapter. Furthermore, when EDM0 matures, the new stub period is hedged with fed fund futures.¹⁹

¹⁹It should be noted that hedging these stub periods with fed fund futures is not really internally consistent, since fed funds is not a LIBOR-based rate. If the LIBOR-fed funds basis is a particular concern, a LIBOR-fed funds basis swap can be added to the trade. See Chapter 16.

TABLE 15.16 Shorting the TED Spread of the $1\frac{3}{8}\%$ of March 15, 2012, from Settlement Dates of June 1, 2012, to July 6, 2012

Trade Date	Security	Hedge Coverage		Principal/ Contracts	Initial Price	Ending Price	P&L
		Start	End				
5/28/10	$1\frac{3}{8}\%$	6/1/10	3/15/12	−100mm	101.588	101.891	−302,650
5/28/10	FFM0	6/1/10	6/16/10	10	99.780	99.803*	950
5/28/10	EDM0	6/16/10	9/15/10	103	99.400	99.463	16,223
6/14/10	FFM0	6/16/10	7/1/10	10	99.810	99.823	542
6/14/10	FFN0	7/1/10	8/1/10	21	99.790	99.810	1,750
6/14/10	FFQ0	8/1/10	9/1/10	21	99.780	99.800	1,750
6/14/10	FFU0	9/1/10	9/15/10	10	99.770	99.795	1,042
5/28/10	EDU0	9/15/10	12/15/10	102	99.155	99.375	56,100
5/28/10	EDZ0	12/15/10	3/16/11	102	99.005	99.265	66,300
5/28/10	EDH1	3/16/11	6/15/11	101	98.875	99.170	74,488
5/28/10	EDM1	6/15/11	9/21/11	109	98.705	99.040	91,288
5/28/10	EDU1	9/21/11	12/21/11	100	98.495	98.880	96,250
5/28/10	EDZ1	12/21/11	3/15/12	94	98.245	98.660	97,525
Total price appreciation, including accrued interest							201,558
Repo loan of \$101.588mm @.22% for 35 days							21,729
Grand Total							223,287

*Weighted average sales price as hedge is reduced over time

Table 15.16 describes the P&L components of the short $1\frac{3}{8}\%$ TED spread trade initiated on May 28, 2010, and unwound on July 2, 2010, corresponding to bond settlement dates of June 1 and July 6, respectively. Not shown in the table are the initial and final TED spreads, 50 and 38 basis points respectively.

The hedge implemented on May 28, 2010, is to buy 10 FFM0, 103 EDM0, 102 EDU0, etc. The ED contracts come directly from the hedge derived earlier in this chapter while the 10 FFM0 are the implementation of the 17 contracts required for the stub of that hedge. That many ED contracts, at \$25 per basis point, is equivalent to $\frac{17 \times 25}{41.67}$ or about 10 FF contracts at \$41.67 per basis point. Furthermore, it is most sensible to buy FFM0 to cover the stub risk from June 1, 2010, to June 16, 2010, even though that fed fund contract is sensitive to rates over the whole of June. This is the usual difficulty of hedging with standardized contracts.

As time passes, two adjustments have to be made to the hedge that was established at the initiation of the trade. First, from June 1 to June 16 the FFM0 hedge has to be reduced gradually, along the lines of the discussion earlier in this chapter. For this reason, the “Ending Price” of FFM0 in Table 15.16 is the weighted average price at which these 10 contracts are sold over time.

The second adjustment to the original hedging program occurs at the maturity of EDM0 on June 14, 2010. The expiration of this contract means that the interest rate exposure of the short bond position over the 92 days from June 16 to September 15²⁰ is no longer hedged. The total number of fed funds contracts replacing the 103 ED contracts is $\frac{103 \times 25}{41.67}$ or about 62. Prorating these across the fed fund contracts by the number of days of exposure in each month gives the following: $62 \times \frac{15}{91}$ or 10 FFM0; $62 \times \frac{31}{92}$ or 21 of each of FFN0 and FFQ0; and $62 \times \frac{14}{92}$ or 10 of FFU0. The initial price of each of these contracts reported in Table 15.16 is as of their purchase on June 14.

The “Initial Price” and “Ending Price” columns of Table 15.16 show that fixed income prices have increased, or rates fallen, over the course of the trade, meaning that the short position in the bond lost money, while the hedge made money. The sum of the price changes is \$201,558 to which is added the \$21,729 in interest earned from the repo transaction used to short that bond.²¹ But the best way to make sense of the P&L of the trade is to recall that the TED spread of the bond fell 12 basis points, from 50 to 38, and that the '01 of the bond position, derived earlier in this chapter, was \$18,190. Therefore, the P&L was expected to be approximately $12 \times \$18,190$ or \$218,280.

²⁰This exposure includes the day of June 16 but not of September 15.

²¹When shorting a bond, a trader pays its coupon interest but earns repo interest on its sale price. See Chapter 12.

Swaps

Parts One and Two of this book have already devoted substantial attention to the valuation and risk of the fixed sides of interest rate swaps. This chapter, therefore, focuses more on the valuation and risk of the floating sides of swaps and on other assorted issues relating to swap markets. These latter topics include counterparty risk, recent legislative and regulatory efforts to “clear” swaps, basis swaps, and constant-maturity swaps (CMS).

SWAP CASH FLOWS

Through an interest rate swap two parties agree to exchange interest payments calculated at a fixed rate for interest payments calculated at a short-term rate that changes over time. For discussion, consider the following swap. On May 28, 2010, party A agrees to pay party B 1.235% on \$100 million semiannually for two years while party B agrees to pay party A three-month LIBOR quarterly on that same amount over that same period. In the terminology of the swap market, 1.235% is the *fixed rate* and three-month LIBOR is the *floating rate*. Party A *pays fixed* and *receives floating* while party B *receives fixed* and *pays floating*. The \$100 million is the *notional amount* of the swap, rather than face or principal amount, because it is used solely to calculate the interest payment: the notional amount is never itself exchanged. Table 16.1 gives the cash flows of this swap using illustrative levels of three-month LIBOR realized on future dates.

Swaps settle $T + 2^1$, so for this swap, traded on May 28, 2010, interest begins to accrue on June 2, 2010. Interest payments on the *fixed leg* of the swap, given in Column (4) of Table 16.1, are paid semiannually and computed on a 30/360 basis. The fixed payment dates of this swap are December 2 and June 2 of each year, so long as these dates are business days. If not, like June 2, 2012, which falls on a Saturday, the payment is

¹Swaps denominated in GBP settle $T + 0$.

TABLE 16.1 Cash Flows on a 1.235% Two-Year, \$100 Million Interest Rate Swap Against Three-Month LIBOR for Settle on June 2, 2010, with Illustrative Future Levels of Three-Month LIBOR

(1)	(2)	(3)	(4)	(5)	(6)	(7)
Cash Flows						
3-Month LIBOR	Payment Date	Actual Days	Actual		With Fictional Notional	
			A Pays	B Pays	A Pays	B Pays
.50%	6/2/10					
.75%	9/2/10	92		127,778		127,778
.75%	12/2/10	91	617,500	189,583	617,500	189,583
.75%	3/2/11	90		187,500		187,500
.75%	6/2/11	92	617,500	191,667	617,500	191,667
1.00%	9/2/11	92		191,667		191,667
1.00%	12/2/11	91	617,500	252,778	617,500	252,778
1.00%	3/2/12	91		252,778		252,778
	6/4/12	94	624,361	261,111	100,624,361	100,261,111

made on the following business day, i.e., June 4, 2012.² When payment does fall on the second day of the month, the 30/360 day count between payments will be 180 and the payment will be $\frac{180}{360}$ or one half times the fixed rate times the notional amount. In the present example, the payment is $\frac{1}{2} \times 1.235\% \times \$100\text{mm}$ or \$617,500. When the payment date does not fall on the second of the month, the interest payment is adjusted for the correct number of days. For the payment falling on June 4, 2010, there are two extra days of interest³ worth $\frac{2}{360} \times 1.235\% \times \$100\text{mm}$ or \$6,861, which, when added to the 180-day payment of \$617,500, gives \$624,361.⁴ Note that swaps are different from bonds in this respect. When a bond payment date falls on a non-business day and payment is pushed to a subsequent business day, the amount of interest is not increased. This difference is taken into account, of course, by careful pricing methodologies.

²In the modified following convention, which is the most common for swaps, a payment is delayed until the subsequent business day unless that delay pushes the payment into a different calendar month. In that case, the payment is, instead, brought forward to the nearest business day.

³With June 2 falling on a Saturday, the next payment would normally be on Monday, June 4. However, there is a special bank holiday in London on June 4 and June 5 of 2012, so this payment would actually be made on June 6 and the interest adjusted accordingly.

⁴Equivalently, there are 182 30/360 days between December 2, 2011, and June 4, 2012, so the payment is $\frac{182}{360} \times 1.235\% \times \$100\text{mm}$ or \$624,361.

The interest payments of the *floating leg* of the swap are given in Column (5) of Table 16.1. The LIBOR rates, given in Column (1), are used to compute these quarterly floating rate payments, which are set in *advance* but paid in *arrears*. For the first quarterly payment, on September 2, 2010, the payment is determined as follows. On May 28, 2010, three-month LIBOR is observed to be .50%. (This setting or fixing date is not shown in the table.) Next, two business days later, on June 2, 2010, interest begins to accrue at this observed rate on an actual/360 basis. Then, 92 days later, on September 2, 2010, the payment of $\frac{92}{360} \times .5\% \times \$100\text{mm}$ or \$127,778 is made. Note that setting in advance and paying in arrears implies that the amount of each quarterly floating rate payment is known three months before it is made. To take one more example, the .75% LIBOR observed two business days before June 2, 2011, is used for calculating the payment accruing over the 92 days between June 2, 2011, and September 2, 2011, i.e., $\frac{92}{360} \times .75\% \times \$100\text{mm}$ or 191,667.

It is noted here that the floating legs of EUR-denominated swaps are typically set off Euribor, not EUR LIBOR. (See Chapter 15.)

Before moving on to valuation issues, note that swaps are complex legal agreements, usually executed under the industry standard *ISDA Master Agreement*. Details are available on the website of the *International Swaps and Derivatives Association*.

THE VALUATION OF SWAPS

As first mentioned in Chapter 2, when valuing interest rate swaps it is convenient to add a fictional payment of the notional amount to both the fixed and floating legs so that the payment schedule of the swap is as depicted in Columns (6) and (7) of Table 16.1. Under the assumption that swap payments are default-free, an assumption that will be discussed later in this chapter, adding this fictional payment from party A to party B and an equal payment from party B to party A has no effect on the value of the swap. But, with these fictional payments, the fixed leg of the swap resembles a fixed rate coupon bond and the floating leg resembles a floating rate note.

Until the onset of the 2007–2009 financial crisis, it had been traditional to assume that one could invest risklessly at the LIBOR index, or equivalently, for example, that arbitrage arguments justify discounting a three-month cash flow at three-month LIBOR. Under that assumption, as shown later in this section, the floating leg of a swap (with its notional principal payment) is worth par at reset dates. It then follows immediately that the fixed leg of a swap, which is exchanged for the floating leg without any other exchange of cash, must also be worth par. With the value of the fixed sides of par swaps pinned down in this way, their valuation proceeds along the lines of Part One. More specifically, discount factors, spot rates, or forward

rates are derived from the cash flows of the fixed legs of par swaps and their par prices, and then the fixed legs of non-par swaps are priced by arbitrage.

Since the onset of the financial crisis, however, this traditional valuation methodology has been overturned. Federal (fed) funds rates and their international equivalents, or Overnight Indexed Swap (OIS) rates, are now recognized as better proxies for riskless investable rates. This change voids the conclusion that the floating legs of LIBOR swaps are worth par and also voids the valuation methodology for the fixed legs of swaps presented in Part One. Furthermore, market participants now explicitly account for a related flaw in the methodology of Part One. As will be shown in Chapter 17, since the collateral posted to ensure the performance of swaps earns fed funds, the value of non-par swaps cannot be computed from the value of par swaps along the lines derived in Part One.

The widespread change in swap valuation methodology followed the crisis because, as discussed in Chapter 15, the spread between LIBOR and OIS, which had been reliably below 10 basis points for years before the crisis, exploded to hundreds of basis points during the crisis. Hence, a valuation methodology based exclusively on LIBOR quantities, which had been deemed accurate enough before the crisis, was deemed so no longer. A watershed moment for the shift away from the traditional methodology was in June 2010, when LCH.Clearnet, the most significant swap clearinghouse, moved from LIBOR to OIS discounting.

To allow the reader to digest the fundamentals of swap valuation before having to face all its nuances, this section presents the traditional, LIBOR-based valuation, which is reasonably accurate when LIBOR-OIS spreads are low. Chapter 17 explains and justifies the OIS-based approach.

The valuation of the floating side of a swap seems difficult at first because the LIBOR settings on future dates are not known as of the initial valuation date. With the fictional notional payment, however, the floating leg can, in fact, be valued. Start at maturity and work backward, three months at a time. Let the swap maturity be T years, the three-month LIBOR set three months before maturity $L_{T-.25}$, and the number of days in that three-month period $d_{T-.25}$. Then, taking LIBOR as the appropriate discount rate, the value of 100 notional of the floating leg in $T - .25$ years is its date- T cash flows discounted over those $d_{T-.25}$ days:⁵

$$\frac{100 + 100 \times \frac{L_{T-.25} \times d_{T-.25}}{360}}{1 + \frac{L_{T-.25} \times d_{T-.25}}{360}} = 100 \quad (16.1)$$

⁵If the number of days is not exactly three months, then the interest payment is still based on three-month LIBOR but the rate used for discounting is not exactly three-month LIBOR. In this case the value of the floating leg will be very slightly different from par.

In words, three months before maturity the floating leg is worth par. Intuitively, since the rate earned over the payment period is the same as the rate used for discounting, its present value is just the face amount.

Now value the swap six months before maturity, in $T - .5$ years. Since, by (16.1), the floating leg is worth 100 in $T - .25$ years, its value in $T - .5$ years is 100 plus its interest payment in $T - .25$ years all discounted at the LIBOR set in $T - .5$ years. But this value is again par:

$$\frac{100 + 100 \times \frac{L_{T-.5} \times d_{T-.5}}{360}}{1 + \frac{L_{T-.5} \times d_{T-.5}}{360}} = 100 \quad (16.2)$$

Continuing to roll backward in this way shows that the value of the floating leg at initiation of the swap is worth par as well.

The argument just made proves that the value of the floating leg at the beginning of each accrual period, including the initial settlement date of the swap, is par. But what is the value of the floating leg on some pricing date between initial settlement and the first payment date? Say that the first LIBOR setting was L_0 , that there are d days from the pricing date to the next payment, and that the appropriate rate for discounting over those d days is $\bar{L}$.⁶ Then, since the floating leg will be worth par immediately after the next payment, the value of the floating leg as of the pricing date is

$$\frac{100 + 100 \times \frac{L_0 \times d_0}{360}}{1 + \frac{\bar{L} \times d}{360}} \quad (16.3)$$

To summarize the pricing of the floating leg of a swap, the instant the LIBOR setting is observed, the value of the floating leg as of the beginning of the next interest accrual period is par. But a moment later, once the next floating payment is fixed, (16.3) says that the value of the floating leg equals the value of a short-term bond with coupon equal to the last observed LIBOR setting and with maturity equal to the days until the next floating payment date.

To illustrate the points made so far in this section, Table 16.2 prices, as of May 28, 2010, 100 face amount of a fixed *versus* three-month LIBOR swap that originally settled on December 15, 2009. The fixed rate is 1.386%, the swap matures on December 15, 2011, and the previous three-month

⁶In practice, the rates for discounting swap cash flows that occur before the first observable par swap are determined by a curve-fitting algorithm like that of Chapter 21. Note that LIBOR of other terms cannot be used for this purpose without adjustment. For example, payments of a fixed *versus* three-month LIBOR swap that will be made in one month cannot be discounted by one-month LIBOR. See the discussion of basis swaps and spreads later in this chapter, and Chapter 17.

TABLE 16.2 Pricing a Swap as of May 28, 2010, that Originally Settled on December 15, 2009

Previous LIBOR Setting		.257%	
Date	Discount Factor	Fixed-Leg “Cash Flows”	Floating-Leg “Cash Flows”
6/15/10	.999824	$\frac{1}{2} 1.386$	$100 + \frac{92}{360} .257$
12/15/10	.996185	$\frac{1}{2} 1.386$	
6/15/11	.990908	$\frac{1}{2} 1.386$	
12/15/11	.983968	$100 + \frac{1}{2} 1.386$	
Present Value		101.15	100.05
Net Present Value		1.10	

LIBOR setting, in the middle of March 2010, was .257%. The discount factors used in the table are those derived from a curve-fitting methodology, like that of Chapter 21, as of May 28, 2010.

The two “Cash Flows” columns in Table 16.2 give the fictional cash flows used for pricing the two legs of the swap. For the fixed leg this means including the fictional notional at maturity. For the floating side this means taking the value of the floating leg on the next floating payment date, i.e., on June 15, 2010, to be par, which means that its present value is found by discounting par plus the coupon accrued over the 92 days since March 15, 2010.

The last two rows of Table 16.2 give the present value (PV) of each leg of the swap and the *net present value*, or NPV of the swap as a whole. NPV is defined, for the receiver of fixed, as the PV of the fixed leg minus the PV of the floating leg. From the perspective of the receiver of floating, the NPV is the PV of the floating leg minus the PV of the fixed leg, which is just the negative of the NPV for the receiver of fixed. If not otherwise specified, however, NPV is conventionally calculated from the perspective of the receiver of fixed.

As mentioned in Part One, most swaps are initiated at par, that is, such that the value of the fixed leg is par. And, by the analysis of this section, the value of the floating leg at initiation is usually par as well. Hence, swaps are usually initiated without any initial exchange of cash. If, however, a swap is initiated with a rate different from the par rate (e.g., as part of an asset swap transaction; see Chapter 19) or if, as a result of day count idiosyncracies, the value of either leg is not exactly par (see Chapter 21), then the NPV of the swap is paid by one party to the other at the time of initiation. Note, however, that the receiver of any such up-front cash payment generally has to post that NPV as collateral to ensure performance

of obligations under the swap. Collateral posting is governed by the *credit support annex (CSA)* of an ISDA Master Agreement. The economics of posting collateral against swaps and the related pricing issues are a focus of Chapter 17.

Once traded, a swap may be kept in place until its maturity or it may be unwound early. Should both parties agree to unwind the swap, one party would pay the other the NPV of the swap at that time. Say, for example, that the parties to the swap described in Table 16.2 were to terminate the swap on the pricing date of the table, i.e., on May 28, 2010. For both parties to be willing to “tear up” this swap, the fixed payer would have to pay the NPV, or 1.10 per 100 face amount, to the fixed receiver. The fixed payer is content to pay 1.10 and to give up the floating payments worth 100.05 in exchange for no longer having to make fixed payments worth 101.15. Similarly, the fixed receiver is content to give up the fixed payments worth 101.15 in exchange for receiving 1.10 and no longer having to make floating payments worth 100.05.

To complete the discussion about unwinding swaps, it should be mentioned that if only party A wants to unwind a swap between parties A and B, party A can, with the consent of party B, *assign* the swap to a third party, C. In that case, after party A pays the NPV to or receives the NPV from party C, as appropriate, the original swap becomes an agreement between parties B and C. Market custom is for party B to agree to such an assignment request unless there is a reasonable objection, most often related to the credit of counterparty C. See the discussion of credit risk and interest rate swaps later in this chapter.

A NOTE ON THE INTEREST RATE RISK OF SWAPS

Part Two developed methodologies to measure and hedge interest rate risk and applied them to the fixed sides of swaps. With the discussion of the previous section, these same methodologies can be applied to the floating sides as well. In particular, the interest rate risk of the floating leg at any particular time is the same as that of a bond making a single payment of par plus interest on the next floating payment date. This leads to a repeating pattern. Just after the setting of the next floating rate payment, the *DV01* (or duration) of the floating leg is approximately equal to the time to that next payment. *DV01* then falls with time, reaching nearly zero just before the next payment date, after which payment it jumps back up to approximately the time between resets. The generic name for the interest rate risk of the floating leg is *reset risk*.

Strictly speaking, the *DV01* of receiving fixed in a swap is equal to the *DV01* of its fixed leg minus the *DV01* of its floating leg. This observation is not particularly useful, however, because it makes much more sense to

hedge the fixed and floating legs separately: the value of the fixed leg depends on relatively long-term rates while the value of the floating leg depends on short-term rates. Put another way, subtracting the *DV01s* and hedging the net exposure with a relatively long-term security implicitly hedges the floating leg with that same long-term security, which hedge results in unnecessary curve risk. In practice, traders and portfolio managers who use a relatively limited number of swaps do hedge the portfolio of fixed legs and the portfolio of floating legs separately. The portfolio of the fixed legs is hedged using any of the methods described in Part Two while the portfolio of floating legs is hedged with short-term rate derivatives along the lines of Chapter 15. Swap trading desks, on the other hand, with portfolios of very many swap contracts, hedge using partial or forward bucket '01s, as described in Chapter 5, which methodologies automatically separate exposures arising from different parts of the curve.

ON CREDIT RISK AND INTEREST RATE SWAPS

It is very important to distinguish between the credit risk that is inherent in the LIBOR index and the credit risk of a swap agreement. The former, as discussed in Chapter 15, arises because LIBOR is the rate on an unsecured, short-term loan between large financial institutions. Since there is some risk that a financial institution will default on its borrowings, LIBOR is above the rates of equivalent maturity, safer loans, like U.S. Treasury bills or the secured loans of repo transactions. Furthermore, since an interest rate swap exchanges a fixed rate for future LIBOR rates, expectations and risk premia with respect to future LIBOR rates will have an impact on observed swap rates.

Completely separate from the risk of the LIBOR index, however, is the counterparty risk of a swap agreement, i.e., the risk that one party to the swap will default on its obligations to make contracted fixed or floating payments. A swap of fixed *versus* the three-month Treasury bill rate with a small corporation would have a substantial amount of counterparty risk even though the floating rate index of the swap has almost no credit risk component. Conversely, a swap of fixed *versus* a three-month commercial paper rate with an agency of the U.S. government would have almost no counterparty risk even though the level of the index reflects corporate credit risk.

The next important point to make in this section is that the counterparty risk of a swap is relatively small, certainly in comparison with the credit risk of a corporate bond. First, as a swap agreement requires the exchange of only interest rate payments, there is no settlement risk with respect to a terminal exchange of the notional amount (i.e., the risk that a payment of notional is made but that no offsetting payment is received). Second, should one counterparty default on an interest payment, the other counterparty

need no longer make interest payments either. Hence, each counterparty is at risk only for the NPV of the swap, i.e., the difference between the value of receiving contracted cash flows and the value of paying contracted cash flows. In the example of Table 16.2, the fixed receiver is at risk for 1.10 per 100 notional amount. By contrast, the holder of a corporate bond with the same coupon and maturity, would be at risk for the full 101.15.⁷ The NPV of a swap can certainly be substantially higher than 1.10—imagine a 20-year swap that was initiated 10 years ago when rates were much higher—but the NPV is always substantially smaller than the notional amount. This is an important point to bear in mind when the swap exposure of a corporation or a financial institution is quoted in terms of notional amount: the true exposure is the NPV of the swap book, which is a much smaller quantity. This discussion also reveals why adding fictional payments of the notional amounts to both legs of the swap, so useful a device for valuation in the absence of counterparty risk, can be misleading in the presence of counterparty risk.

Some swaps, particularly those in which one counterparty is a non-financial corporation, are subject to the counterparty risk described in the previous paragraphs. Between dealers, other financial institutions, and many other counterparties, however, this counterparty risk is largely eliminated by means of collateral requirements. The details and mechanics of these requirements are discussed in Chapter 17, but the salient point is as follows. Say that, in a swap between parties A and B, the NPV with respect to party A is negative, i.e., A owes money in PV terms. In this situation, the swap agreement would require A to post collateral to party B equal to this amount owed. If A should then default, B can sell the collateral and be made whole with respect to the value of the swap.

Ignoring bid-ask spreads, swaps that have collateral agreements are usually initiated at prevailing or quoted market rates. Swaps without collateral agreements, on the other hand, often include a *credit-value adjustment* (CVA) in the fixed rate. This means that the weaker counterparty would receive less than the prevailing rate if receiving fixed or pay more than the prevailing rate if paying fixed. The stronger counterparty can collect these spreads *versus* prevailing rates and hold them against future credit losses. Or, if the credit of the counterparty can be traded in credit default swap (CDS) markets (see Chapter 19) or private insurance markets, the stronger counterparty might buy insurance against a counterparty default. This is not a straightforward exercise, by the way, because the market for trading individual credits is not particularly liquid and the exposure that has to be

⁷The credit-risk analogue to this swap in the bond context is actually a repo transaction in which the borrower of the bond or lender of cash defaults. The lender of the bond then loses the 101.15 value of the bond, but no longer has to repay borrowings plus accrued interest of 100.05, for a total loss of 1.10.

insured, namely the NPV of the swap, changes constantly. Finally, it is worth noting that most institutions set limits on the amount of swap exposure they have with each of their swap counterparties.

MAJOR USES OF INTEREST RATE SWAPS

The first major use of swaps to be discussed in this section is to hedge future issuance of corporate debt. Once a corporation has decided to sell bonds to fund its capital expenditures or operations, it may want to hedge against rising interest rates from the time of its decision to raise funds until the actual sale of the bonds. To hedge the future issuance of 10-year debt, for example, it can pay fixed on a 10-year swap at the time of the decision and unwind the swap at the time of the bond sale. Or, perhaps even more effectively, it can pay fixed on a 10-year swap with forward settlement at the time of the bond sale. In either case, if rates rise between the decision and the bond sale, the corporation will have to pay a higher coupon rate on its debt but will have gained from its swap position. Of course, if rates fall then the corporation will not benefit from selling debt at a lower rate because the PV of that gain will have been offset by losses from its swap position. While hedging the risk of future issuance with swaps is perfectly reasonable, substantial basis risk remains. First, swap rates reflect the short-term credit risk of the banking system rather than long-term, generic, corporate credit risk. Second, swap rates cannot possibly hedge the risk that the credit of a particular corporation worsens relative to short-term bank credit.

A second use of interest rate swaps, also by corporations, is to create synthetic floating rate debt. Consider a corporation that has decided to borrow at short-term, floating rates, perhaps because it believes that interest rates will not increase or perhaps because its assets are also short-term in nature. In any case, the corporation can consider two options. First, issue and plan to roll over short-term debt, like commercial paper. Second, issue long-term floating rate debt. Both of these options have the desired interest rate exposure, but each has serious drawbacks. With respect to issuing short-term debt, only the most creditworthy companies can tap into the commercial paper market. Even a company with the ability to sell commercial paper, however, might be leery of bearing too much liquidity risk, i.e., the risk of having to replace maturing short-term debt at a time of stress in its own financial situation or in the wider financial system. Issuing long-term floating rate debt solves for the problem of liquidity risk,⁸ but the market

⁸Naturally the maturities of the floating rate issues would have to be staggered so that rolling over large amounts of debt never has to happen over a short period of time.

for long-term floating rate debt is very small. A popular solution for such corporations, therefore, is to issue fixed rate debt and then receive fixed and pay floating in a swap. The net effect of this strategy is to achieve floating rate funding at some spread above three-month LIBOR or, if different, above the floating rate index of the swap.

A third use of interest rate swaps is to hedge mortgage-related interest rate risk. (Mortgages and mortgage-backed securities are the subject of Chapter 20.) As discussed in the Overview the mortgage market is the single, largest fixed income market in the United States. This means that there is a lot of interest rate risk to be hedged. Furthermore, because the interest rate risk of mortgages and mortgage-backed securities changes over time, hedges have to be adjusted over time. Significant hedging demand of this sort comes from the portfolio businesses of the government-sponsored entities (GSEs), in which they sell fixed rate debt and buy mortgages. Additional demand comes from mortgage servicers because the value of their future fees, collected in compensation for processing mortgage payments, is quite sensitive to the level of interest rates.

The demand to hedge long positions in mortgages with swaps is so great, in fact, that the market for swaps is sometimes distorted by mortgage-related hedging activity. For example, rising mortgage rates, which usually increase the interest rate risk of outstanding mortgages (see Chapter 20), generate a demand to hedge this increased risk by paying fixed in swaps. This demand has been substantial enough at times to push swap rates down significantly relative to other rates (e.g., governments). A natural question is why this effect is not offset by the demand to hedge short positions in mortgages by receiving fixed in swaps. The answer is that homeowners are the main shorts in the mortgage market, and they simply do not hedge their interest rate risk.

The final use of swaps to be mentioned here is to manage mismatches between the durations of assets and liabilities. As discussed in the Overview, the liabilities of pension funds and insurance companies tend to have longer durations than the assets they wish to purchase. As a result of internal and regulatory pressures, however, these institutions need to limit the mismatch between their asset and liability durations. Receiving fixed in swap is an effective way to lengthen asset duration and achieve this goal.

THE REGULATORY AND LEGISLATIVE MANDATES TO CLEAR OVER-THE-COUNTER DERIVATIVES

Historically, interest rate swaps had been an over-the-counter (OTC) product. The term OTC literally means that the product does not trade on an organized exchange, in contrast with, to cite examples covered in earlier chapters, bond and note futures, ED futures, and fed funds futures that trade under the rules and procedures of a futures exchange. More broadly

defined, OTC products are *bilateral agreements* in which two counterparties independently set the terms of a trade, manage cash flows and any collateral requirements, and bear all of the market and counterparty risks.

During and following the financial crisis of 2007–2009, a narrative emerged in which the OTC derivatives market as a whole played a significant role in the crisis. In particular, it was argued that the complex interconnectedness of derivative counterparties, in combination with the inability of regulators to map the resulting counterparty risks across the financial system, allowed, caused, or exacerbated the crisis. It is beyond the scope of this book to debate the merits of this narrative, but the bottom line is that regulators, and then the Dodd-Frank law, now mandate the *clearing* of various OTC derivatives.

Clearing has many meanings,⁹ but, in the current context, the regulatory and legislative intent is for a *clearinghouse*, among other more technical responsibilities, 1) to set and manage collateral requirements, and 2) to be the *central counterparty* (CCP) for the trades it clears. To explain the concept of a CCP, say that counterparties A and B agree that A will pay fixed and receive floating in an interest rate swap that is cleared by a CCP. In that case, A will enter into a contract with the CCP to pay fixed and receive floating while B will enter into a contract with the CCP to receive fixed and pay floating. Thus, A and B have the same market risk as in a bilateral agreement, but counterparty risk to the CCP rather than to each other. Note now that it makes perfect sense for a clearinghouse that acts as a CCP to protect itself from counterparty defaults by setting and managing collateral requirements.

A clearinghouse would want to set collateral requirements such that, in a stress scenario, collateral that has been posted by defaulting counterparties is sufficient to cover their losses, i.e., to make other counterparties whole. But what if collateral posted by defaulting counterparties is not sufficient to cover their losses? It is impractical to require enough collateral so that losses are covered in every conceivable scenario: posting collateral is expensive¹⁰ and market participants would simply not trade through the clearinghouse if collateral requirements were too onerous. Therefore, in a particularly simple structure, a clearinghouse sets prudent collateral requirements and the owners or members of the clearinghouse contribute capital to cover any losses not covered by posted collateral. To summarize the *waterfall* in this structure, should some number of trading counterparties default, the clearinghouse uses their collateral to settle their obligations. If this collateral

⁹ For a book-length treatment, see Tina Hasenpusch, *Clearing Services for Global Markets*, Cambridge University Press, 2009.

¹⁰ Even if cash collateral earns interest or if collateral can be posted in high-quality, interest-bearing securities, having to post collateral almost always involves a suboptimal allocation of scarce resources.

is not sufficient, the clearinghouse uses its own capital to cover the residual obligations. And if even this capital is not sufficient, the clearinghouse itself fails.¹¹

An important policy question is whether clearinghouses, operating as just described, reduce systemic risk relative to a system of bilateral agreements. Regulatory and legislative action seem to rest on the assumption that the answer is yes, but the question is hardly resolved.¹² First, while it is often claimed that a CCP “eliminates” counterparty risk, in actuality, it *mutualizes* counterparty risk; i.e., it spreads the risk across clearinghouse members. To see this, suppose that each potential clearinghouse member diversifies its business and, therefore, its counterparty risk, across all other potential members. In that case, the counterparty risk of each, from the collection of its bilateral trades, is the same as its counterparty risk from being a member of a CCP. Second, there is no reason to believe that a regulated CCP will manage counterparty risk better than individually regulated institutions doing bilateral trades, particularly for less liquid and harder-to-price derivatives. Third, if a CCP should ever default, the systemic damage could be substantial. Fourth, forcing an institution to transfer some of its exposures to a counterparty to a CCP while leaving other bilateral exposures to that same counterparty intact would increase the total counterparty risk of the institution so long as the portfolio of the two sets of exposures benefits from any diversification.¹³ Fifth, while CCPs do have outstanding records with respect to handling systemic disruptions, these records have been built in the context of particularly liquid products, whether liquid inherently or liquid because of a combination of inherent properties and CCP efforts. There is no evidence, however, that these successful records can be replicated in the case of less liquid products. Furthermore, even in the case of liquid products, CCPs have come close to failure, e.g., after the stock market crash of ’87.¹⁴

A policy issue distinct from but very much related to clearing is the *standardization* of contracts. Consider the highly successful U.S. note and bond futures contracts described in Chapter 14. The market seems content to trade U.S. Treasury interest rates through only a handful of contracts. More

¹¹ Clearinghouse structures can be more complex than this simple description. For example, clearinghouses might have the right to require owners or members to contribute additional capital or even the right to use other counterparty collateral to settle outstanding obligations.

¹² The discussion in this paragraph was first presented in Bruce Tuckman, “Amending Safe Harbors to Reduce Systemic Risk in OTC Derivatives Markets,” Center for Financial Stability policy paper, April 22, 2010.

¹³ See Darrell Duffie and Haoziang Zhu, “Does a Central Clearing Counterparty Reduce Counterparty Risk?” Stanford University, July 1, 2009.

¹⁴ See Leo Melamed, *For Crying out Loud*, John Wiley & Sons, Inc., 2009, pp. 149–151.

specifically, rather than design their own bilateral contracts with tailor-made expiration dates, delivery rules, and delivery prices, market participants choose to trade a few highly liquid contracts with a very limited number of expiration dates, very particular delivery rules, and a single delivery price at any time for each ticker. By contrast, trading in options on interest rate swaps, or swaptions (see Chapter 18), is highly *customized* in the sense that the buyer and seller choose from a very large menu of strikes, option expiration dates, and terms of the underlying swap.

There are many market conventions used when initiating interest rate swaps (e.g., day-count conventions, payment date schedules), but swaps, for the most part, are still regarded as more customized than standardized. First, some market participants do request particular fixed rates or payment dates that deviate from convention, e.g., a corporate issuer of fixed rate debt that wants to receive fixed in a swap at a rate and on payment dates that correspond exactly to the coupon rate and payment dates of its debt. Second, because swaps are initiated at par rates prevailing at the time of trade and with payment schedules in intervals from the settlement date (as discussed earlier in this chapter), there are an enormous number of distinct swap contracts trading at any time.

Many swap market customers (i.e., non-dealers) and regulators argue for increased standardization of swaps. As the most liquid trading happens at the par rate with conventional payment dates (i.e., in regular intervals from the settlement date), unwinding an existing swap proves expensive. The dealer with whom the swap was initiated can demand an exit premium or, equivalently, finding a counterparty willing to take on that particular swap and willing to arrange an assignment from the initiating dealer proves costly. As a result, rather than unwind an existing swap, most customers (and, in fact, trading desks) hedge market risk by adding new swap trades. This is far from ideal, however, for two reasons. First, if an original swap and its offsetting swap are with different counterparties, market risk may be hedged but counterparty risk remains. Second, the size and complexity of trading books grow to extents that can be particularly problematic in a crisis.

Some swap customers and dealers take the other side of this argument, claiming that the ability to customize trades is worth the inconveniences set out in the previous paragraph. It is often countered that dealers are simply acting as an oligopoly to maintain their pricing power, particularly in the matter of profits from unwinding trades. A natural compromise would be to trade both customized and standardized swaps, although wide-scale standardization is very difficult to achieve without the active participation of the dealer community.

As of the time of this writing, regulators have succeeded in pushing dealers to clear most interest rate swaps and the most liquid CDS with each other, although dealer-customer trades remain bilateral. Regulators have also

succeeded in standardizing liquid CDS contracts, which is discussed further in Chapter 19. Much less standardization characterizes the interest rate swap market, although customers and dealers have begun to trade swaps with payments on International Money Market (IMM) dates (see Chapter 15), which dramatically reduces the number of payment-date schedules across swaps and, therefore, may improve their liquidity.

BASIS SWAPS AND SPREADS

Because short-term borrowing and lending in fixed income markets are keyed off several distinct short-term interest rates, e.g.—fed funds, one-month LIBOR, three-month LIBOR—market participants often find themselves bearing basis risk across these rates. For example, a bank might borrow money in the fed funds market but lend to clients at spreads to three-month LIBOR; an investment bank might fund securities at repo, which is highly correlated with fed funds, but finance customers at spreads to one-month LIBOR; and a corporation might issue floating rate debt at a spread off six-month LIBOR but invest temporary cash balances at rates keyed off three-month LIBOR. Furthermore, apart from the need to hedge basis risks, there is also demand to trade these bases, e.g., to bet that LIBOR rates will widen relative to fed funds. Finally, both hedging and speculative demand to trade bases increased through the financial crisis of 2007–2009, as spreads that were small and not very volatile for the longest of times rose to unprecedented levels and volatilities—recall the LIBOR-OIS spread in Figure 15.5.

Bases are traded through *basis swaps*, in which market participants swap interest payments keyed off one short-term rate for interest payments keyed off another short-term rate. The most important permutations are trading three-month LIBOR against one-month LIBOR, six-month LIBOR, or fed funds, but many less common pairs trade as well, e.g., a LIBOR index *versus* a short-term municipal rate index. In any case, to take one important example for discussion, party A might agree to pay party B compounded one-month LIBOR plus 25 basis points over three months on some notional amount while party B agrees to pay three-month LIBOR to party A on that same notional amount.¹⁵

Why is the spread in this last example 25 basis points rather than zero? A particularly easy way to understand this is to consider a basis swap of some high-yield corporate bond index rate against fed funds compounded over

¹⁵ For a more detailed treatment of basis swaps, see Carl Lantz, Michael Chang, and Sonam Lekhi Dorji, “A Guide to the Front-End and Basis Swap Markets,” Credit Suisse Fixed Income Research, February 18, 2010.

three months. The high-yield rate has always been and will almost certainly be above realized compounded fed funds. Therefore, it cannot possibly be fair to receive the high-yield rate and pay only compounded fed funds flat in exchange: it can be fair, however, to receive the high-yield rate and pay compounded fed funds plus a basis swap spread in exchange. Viewing this example from another perspective, the receiver of the high-yield rate through a derivative contract is receiving that rate without bearing any of the credit risk associated with high-yield bonds. Hence, the receiver of that relatively high rate has to pay, in exchange, a positive spread over the relatively low fed funds rate.

Returning now to the basis swap of three-month LIBOR against one-month LIBOR, three-month LIBOR is very much expected to be greater than one-month LIBOR compounded over three months. First, three-month loans have more credit and liquidity risk than one-month loans. Second, over a three-month horizon, three-month LIBOR is the average borrowing rate over a fixed set of banks, while one-month LIBOR is a rate on a *refreshed* credit. To explain, should the creditworthiness of one of the banks in the LIBOR survey deteriorate so that its borrowing rate increases, that rate will probably be dropped from the calculation of LIBOR. (See Chapter 15.) Hence, over an horizon of three months, three-month LIBOR can be the average borrowing rate over a worse set of credits than one-month LIBOR, which is another reason to expect three-month LIBOR to exceed compounded one-month LIBOR.

The fact that basis swaps trade at a spread rather than *flat*, i.e., at a spread of zero, raises the following issue under the traditional approach of valuing the floating legs of swaps. If one assumes that one-month LIBOR is the riskless investable rate, then a floating leg paying one-month LIBOR is worth par. But, since the basis swap of one-month LIBOR against three-month LIBOR does not trade flat, it cannot also be the case that a floating leg paying three-month LIBOR is worth par. By implication, then, the fixed payments of three-month LIBOR swaps could not be valued using the approach of Part One. Similarly, if three-month LIBOR is taken as the riskless investable rate, then the fixed payments of one-month LIBOR swaps could not be valued as in Part One. Chapter 17 shows how basis swap spreads are used explicitly or implicitly to price swaps of one floating rate index when another floating rate index is taken as the riskless investable rate. Consistent with the main focus of that chapter and current industry practice, however, the OIS curve is taken as the collection of riskless rates.

CONSTANT MATURITY SWAPS

This section discusses the pricing of a CMS in which one counterparty agrees to pay fixed and receive a swap rate at some periodicity over the life of the

swap on some notional amount. For example, quarterly payments at a fixed rate might be exchanged for quarterly payments at the five-year swap rate for two years. To begin this discussion, consider an agreement to exchange a single cash flow in two years, in particular, to pay a fixed rate S on a unit notional amount and to receive the prevailing five-year rate at that time on the same notional amount. What is the fair value of this fixed rate today?

In the notation of earlier chapters, in particular Chapters 2 and 13, $C_t(T)$ denotes the T -year par swap rate at time t and $C(t, t+T)$ denotes today's T -year swap rate t years forward. Similarly, $A_t(T)$ denotes the T -year annuity factor at time t , at the appropriate periodicity, and $A(t, t+T)$ denotes today's T -year annuity factor t years forward. Note that, in this notation, today's quantities are denoted without time subscripts. In any case, the payoff in two years of the single-payment CMS described in the previous paragraph is

$$C_2(5) - S \quad (16.4)$$

The first point to make about the CMS rate S is that it has to exceed the five-year swap rate two years forward, $C(2, 7)$. This will be demonstrated by showing that if $S = C(2, 7)$, a riskless arbitrage profit is available. Hence, S has to be greater than $C(2, 7)$ so that the payoff from the CMS swap in (16.4) is not so large.

To construct the arbitrage trade, consider hedging the single-payment CMS by receiving in a five-year swap two years forward on a notional of $\frac{1}{A(2, 7)}$. The value of this forward swap in two years is

$$\frac{1}{A(2, 7)} \times [C(2, 7) - C_2(5)] A_2(5) \quad (16.5)$$

The total payoff of the trade, then, is the value of the CMS in (16.4), with the assumption that $S = C(2, 7)$, plus the value of the hedge in (16.5), for a total of

$$[C(2, 7) - C_2(5)] \left[\frac{A_2(5)}{A(2, 7)} - 1 \right] \quad (16.6)$$

But this quantity is always positive. If the five-year rate has fallen relative to its initial forward over the two years since the initiation of the trade so that $C(2, 7) > C_2(5)$ and $A(2, 7) < A_2(5)$, both terms of (16.6) are positive. On the other hand, if the five-year rate has risen relative to its initial forward so that $C(2, 7) < C_2(5)$ and $A(2, 7) > A_2(5)$, both terms of (16.6) are negative

so that the total payoff is positive. Intuitively, when rates have fallen, the CMS loses and the hedge gains, but because the annuity factor has risen as well, the hedge gains more than the CMS loses. Though, when rates have risen, the CMS gains and the hedge loses, but because the annuity factor has fallen as well, the hedge loses less than the CMS gains. Put another way, the convexity of the payoff from the hedge in (16.5), which is not present in the payoff from the CMS in (16.4), causes the profit and loss (P&L) from the hedged position to be positive no matter which way rates move. In any case, as reasoned earlier, it must be the case that $S > C(2, 7)$.

The fair value of S relative to $C(2, 7)$ is called a convexity correction, terminology that follows from the intuition given in the previous paragraph. Letting t be the maturity of the single-payment CMS and T the tenor of the underlying swap rate, the appendix in this chapter shows that the fair convexity correction of S , with semiannual compounding, is given approximately by

$$S - C(t, t + T) \approx \text{Var}[C_t(T)] \left(\frac{1}{C(t, t + T)} - T \frac{\left[1 + \frac{C(t, t + T)}{2}\right]^{-2T-1}}{1 - \left[1 + \frac{C(t, t + T)}{2}\right]^{-2T}} \right) \quad (16.7)$$

This convexity correction appears model dependent since it depends on the variance of the swap rate $C_t(T)$. However, this variance can be approximated by a normal swaption volatility (see Chapter 18) or calculated more precisely from the prices of swaptions across a range of strikes.¹⁶ The convexity correction in (16.7) can be expressed even more simply, however. Let $D(T, C(t, t + T))$ denote the yield-based duration of a T -year par bond at a yield of $C(t, t + T)$, the formula for which is given in equation (4.45). Let $Cnvx(T, C(t, t + T))$ denote the yield-based convexity of that same par bond at that same yield, the formula for which is given in equation (4.52). Then,

$$S - C(t, t + T) \approx \text{Var}[C_t(T)] \frac{\frac{1}{2} Cnvx(T, C(t, t + T))}{D(T, C(t, t + T))} \quad (16.8)$$

As an example of using the convexity correction in (16.7), consider pricing a CMS to pay the five-year swap rate annually for four years starting in two years. For simplicity in this example, assume that the curve is flat at

¹⁶ See, for example, Jim Gatheral, *The Volatility Surface*, John Wiley & Sons, 2006, Chapter 11.

TABLE 16.3 An Example of Pricing a CMS Swap

(1)	(2)	(3)	(4)	(5)
Payment Date (years)	Discount Factor	5-Year Swap Rate Volatility (bps)	Convexity Correction (bps)	Discounted Correction (bps)
2	.923845	84.85	1.88	1.74
3	.887971	103.92	2.83	2.51
4	.853490	120.00	3.77	3.22
5	.820348	134.16	4.71	3.86
Sum	3.485655			11.33
Running Correction			3.25	

a semiannually compounded rate of 4% and that the volatility of all rates is 60 basis points per year.

Because the curve in this example is flat, the only term on the right-hand sides of (16.7) or (16.8) that changes with the date of a payment is the variance term: the convexity and duration terms depend only on the tenor of the underlying swap, here five years, and on the forward swap rate to the various payment dates, which, because the curve is flat, are all 4%. Hence, in this example,

$$\begin{aligned} S - C(t, t + T) &\approx Var[C_t(T)] \left(\frac{1}{4\%} - 5 \frac{[1 + \frac{4\%}{2}]^{-11}}{1 - [1 + \frac{4\%}{2}]^{-10}} \right) \\ &= 2.616047 \times Var[C_t(T)] \end{aligned} \tag{16.9}$$

Table 16.3 reports the remaining calculations. Column (1) lists the payment dates of the sample CMS swap. Column (2) gives the discount factors for payments on that date calculated at the semiannually compounded rate of 4%. Column (3) gives the volatility of the five-year swap rate under the simplifying assumption that all rates have a volatility of 60 basis points. For example, the terminal distribution of the swap rate in four years has a volatility of $60\sqrt{4}$ or 120 basis points. Column (4) gives the convexity correction for each payment using the expression (16.9). For example, for the payment in four years, the correction is $2.616047 \times \left(\frac{120}{10,000}\right)^2 = .000377$ or 3.77 basis points.

Column (4) in Table 16.3 gives the convexity corrections for each payment, but the CMS swap, as a whole, makes four payments. To convert

the four individual corrections to a single running rate, divide the sum of their present values by the appropriate annuity factor. Column (5) gives the present value of each correction, simply by multiplying each correction by the appropriate discount factor, and the sum of those present values is 11.33 basis points. From Column (1), the value of an annuity of one unit of currency on each of the payment dates is 3.486. Hence, the value of receiving the corrections indicated in Column (4) at each of their respective payment dates is the same as receiving $\frac{11.33}{3.486}$ or 3.25 basis points on each of those same dates.

Finally, then, the fair CMS rate in this example has to be 4.0325%. First, since the curve is assumed flat at 4%, paying 4% plus the convexity corrections in Column (4) of Table 16.3 is fair against receiving the five-year swap rate on the dates indicated. Second, the value of paying those individual corrections is the same as the value of paying 3.25 basis points on each payment date. Therefore, 4% plus 3.25 basis points, or 4.0325%, is fair against the five-year swap rate.

To compute the fair CMS rate when the term structure is not flat, begin by finding the fair fixed rate on each payment date, which equals the appropriate forward swap rate plus convexity correction for that payment date. Then find the single fixed rate such that paying that fixed rate on each payment date has the same present value as paying the fair fixed rates for each payment date just computed.

APPENDIX: DERIVATION OF CONVEXITY CORRECTION FOR CMS SWAPS

Chapter 18 shows that there exists a probability distribution for $C_t(T)$ such that two conditions hold. First, the T -year swap rate t years forward equals the expectation today of the T -year swap rate in t years, i.e.,

$$C(t, t+T) = E_0[C_t(T)] \quad (16.10)$$

Second, the fair market value of the CMS swap today, which is zero, is given by

$$E_0 \left[\frac{C_t(T) - S}{A_t(T)} \right] = 0 \quad (16.11)$$

Solving for S in (16.11), using a property of covariance, and rearranging terms,

$$S = E_0 [C_t (T)] + \frac{\text{Cov} \left(C_t (T), \frac{1}{A_t(T)} \right)}{E_0 \left(\frac{1}{A_t(T)} \right)} \quad (16.12)$$

$$= C(t, t+T) + \frac{\text{Cov} \left(C_t (T), \frac{1}{A_t(T)} \right)}{E_0 \left(\frac{1}{A_t(T)} \right)} \quad (16.13)$$

where (16.13) follows from (16.10).

To approximate (16.13), assume that the term structure is flat so that

$$C_t (T) A_t (T) + Z[C_t (T)] = 1 \quad (16.14)$$

where $Z(y)$ denotes the price of a zero coupon bond with yield y and maturity T years. Then, when compounding n times per year,

$$Z(y) = \left(1 + \frac{y}{n} \right)^{-2T} \quad (16.15)$$

Solving (16.14) for $A_t (T)$ and substituting into (16.13),

$$S - C(t, t+T) = \frac{\text{Cov} \left(C_t (T), \frac{C_t(T)}{1 - Z[C_t(T)]} \right)}{E_0 \left(\frac{C_t(T)}{1 - Z[C_t(T)]} \right)} \quad (16.16)$$

Now using the approximation that $\text{Cov}[X, H(X)] = E \left[\frac{d}{dX} H(X) \right] \times \text{Var}[X]$,¹⁷

$$S - C(t, t+T) = \text{Var}[C_t (T)] \frac{E_0 \left(\frac{d}{dC_t(T)} \frac{C_t(T)}{1 - Z[C_t(T)]} \right)}{E_0 \left(\frac{C_t(T)}{1 - Z[C_t(T)]} \right)} \quad (16.17)$$

¹⁷ This approximation relies on an expansion of $H(X)$ as the relationship is exact by Stein's lemma if X is normally distributed.

Performing the differentiation in the numerator,

$$\frac{d}{dC_t(T)} \frac{C_t(T)}{1 - Z[C_t(T)]} = \frac{1}{1 - Z[C_t(T)]} - \frac{TC_t(T) \left[1 + \frac{C_t(T)}{n}\right]^{-nT-1}}{(1 - Z[C_t(T)])^2} \quad (16.18)$$

Substituting (16.18) into (16.17) and relying on another approximation¹⁸ together with the expectation in (16.10),

$$S - C(t, t+T) = Var[C_t(T)] \frac{\frac{1}{1 - Z[C(t, t+T)]} - \frac{TC(t, t+T) \left[1 + \frac{C(t, t+T)}{n}\right]^{-nT-1}}{(1 - Z[C(t, t+T)])^2}}{\frac{C(t, t+T)}{1 - Z[C(t, t+T)]}} \quad (16.19)$$

Finally, simplifying and using (16.15),

$$S - C(t, t+T) = Var[C_t(T)] \left(\frac{1}{C(t, t+T)} - T \frac{\left[1 + \frac{C(t, t+T)}{n}\right]^{-nT-1}}{1 - \left[1 + \frac{C(t, t+T)}{n}\right]^{-nT}} \right) \quad (16.20)$$

¹⁸ In particular, let x be a random variable with mean $\bar{x}$. Then $E[g'(x)]/E[g(x)] \approx g'(\bar{x})/g(\bar{x})$.

Arbitrage with Financing and Two-Curve Discounting

This chapter covers two topics, which turn out to be related. The first topic is arbitrage pricing in bond and swap markets under realistic financing arrangements. The second is pricing swaps when the riskless investable rate (e.g., fed funds) is not the same as the swaps' floating rate index (e.g., three-month *London Interbank Offered Rate (LIBOR)*).

So as not to overwhelm the reader with the nuances of this chapter at the start of the book, Chapter 1 showed that discounting is shorthand for arbitrage pricing of bonds under simplified financing arrangements. In particular, it was assumed that the proceeds from shorting one set of bonds could be used to purchase other bonds. However, as described in Chapter 12, bonds are shorted through repos that require sale proceeds to be posted as collateral. This means that the long side of an arbitrage has to be financed by borrowing money in the repo market and introduces the complication that the financing rates on the short and long side of an arbitrage need not be the same. This chapter¹ begins by deriving the connection between discounting and arbitrage pricing for bonds under these more realistic financing arrangements. It turns out that the results of Chapter 1, which comprise the basic tool box for fixed income practitioners, hold true under the following additional conditions. First, all bonds finance at the same rate. Second, interim mispricings can be financed at that same rate. When these conditions do not hold, the enforcement of the law of one price by arbitrage is not so straightforward as presented in Chapter 1.

To price a fixed rate *versus* LIBOR swap, Chapters 2 and 16, relying on the assumption that LIBOR is the riskless investable rate, conclude that the floating leg of the swap is worth par. Then, the fixed leg exchanged for that floating leg with no other exchange of cash must also be worth

¹The authors would like to recognize Jean-Baptiste Homé for his important contributions to this chapter.

par. As a result, the fixed legs of par swaps can be treated like bonds with prices equal to their face values, meaning that discount factors, spot rates, or forward rates can be derived and that, along the lines of Chapter 1, all other swaps, i.e., non-par swaps, can be priced by discounting their fixed cash flows at those derived discount factors or rates. But this framework ignores the financing arrangements of the swap market, in particular, that one party has to post the net present value (NPV) of the swap to the other and that this NPV earns some collateral rate. As it turns out, non-par swaps can be priced off of a par swap curve under the following conditions. First, all swaps have the same floating rate index. Second, the collateral rate equals this floating rate index. Third, interim mispricings can be financed at this floating rate index as well. These conditions can be seen as analogous to those in the bond case upon realizing that the notional amount of a swap always finances at its floating rate index while its NPV finances at the collateral rate.

The second topic of the chapter is how to price fixed rate *versus* LIBOR swaps if the riskless investable rate is taken to be not LIBOR but fed funds, or the equivalent in another currency, with a term structure given by Overnight Indexed Swap (OIS). Market participants have always recognized that a rate like fed funds is closer to the ideal of a riskless investable rate than is LIBOR, but the difference between LIBOR and OIS had been so small for so many years that few cared about this seemingly subtle pricing issue. Similarly, market participants have been aware that the collateral rate on fixed *versus* LIBOR swaps is almost always fed funds, not LIBOR, so that even if LIBOR were the riskless investable rate it is not theoretically correct to price non-par LIBOR swaps from par swaps. In any case, attitudes changed completely through the 2007–2009 crisis as the LIBOR-OIS spread rose to hundreds of basis points (see Chapter 15). The resulting consensus to take fed funds as the riskless investable rate means that the floating leg of LIBOR swaps is not worth par and that the pricing methodology in Chapters 2 and 16 unravels.

For pricing fixed rate *versus* LIBOR swaps, this chapter lays out the following framework, which is consistent with revised market practice. First, because the fed funds rate is taken as riskless and investable, the floating leg of an OIS swap is worth par, as is the fixed leg that is exchanged for this floating leg with no other exchange of cash. Second, under the conditions described above, i.e., collateral posted earns fed funds, etc., all non-par OIS swaps can be priced from OIS par swaps. Third, under the same conditions, a fixed *versus* LIBOR swap can be priced by a replication argument using a (likely non-par) OIS swap and a LIBOR-fed funds basis swap. Fourth, there is a shorthand for this arbitrage pricing of fixed *versus* LIBOR swaps that requires two curves, the OIS curve for discounting and another curve, specially constructed for the purpose, for projecting LIBOR cash flows.

BOND TRADING WITH FINANCING

Chapter 1 implicitly assumed that an arbitrageur could short a portfolio of bonds that trades rich simply by promising to make its scheduled payments and that the proceeds from that short could be used to buy a portfolio of bonds that trades cheap. It followed, then, that the arbitrageur could pocket the difference between the prices of the rich and cheap portfolios. In practice, however, shorting the portfolio that trades rich is done through repo, which generates no cash today. Therefore, the cash to buy the portfolio of bonds that trades cheap has to be borrowed through repo. Note, by the way, that for the purposes of this chapter, repo haircuts are ignored. Note too that the discussion here assumes that the reader has fully digested the material of Chapter 12.

Table 17.1 describes the results of shorting a bond for one period using realistic financing arrangements, i.e., repo. The bond’s initial price is P_0 , its terminal price is P_1 , and its repo rate over the period is r .

The net result of the one-period short position, shown in Table 17.1, is

$$P_0 (1 + r) - P_1 \tag{17.1}$$

which is positive if

$$\frac{P_1 - P_0}{P_0} < r \tag{17.2}$$

Hence, a trader profits from a one-period short if the return of the bond over that period is less than the repo rate.

In the case that the bond pays a coupon c over the term of the repo agreement, then the trader has to make a payment of that amount to the

TABLE 17.1 Cash Flows from Shorting a Bond for One Period

Date	Transaction	Cash Flow
0	Buy the repo at rate r	$-P_0$
	Sell bond for P_0	P_0
	Total	0
1	Buy bond for P_1	$-P_1$
	Unwind the repo	$P_0 (1 + r)$
	Total	$P_0 (1 + r) - P_1$

TABLE 17.2 Cash Flow from a Single-Period Financed Long Position

Date	Transaction	Cash Flow
0	Buy bond for P_0	$-P_0$
	Sell the repo at rate r	P_0
	Total	0
1	Sell bond for P_1	P_1
	Unwind the repo	$-P_0(1+r)$
	Total	$P_1 - P_0(1+r)$

lender of the bond. In the special case that this coupon payment is due at the end of the period considered in Table 17.1, the terminal cash flow changes to

$$P_0(1+r) - (P_1 + c) \tag{17.3}$$

Once again, the short makes money if the total return of the bond is less than the repo rate.

If an arbitrageur is shorting one bond and buying another, it is clear from Table 17.1 that shorting a bond or portfolio does not generate any cash. Hence, to purchase the long side of the trade, an arbitrageur has to borrow cash through a repo. For a financed long position of a single bond over a single period, Table 17.2 shows that the net result is

$$P_1 - P_0(1+r) \tag{17.4}$$

Expression (17.4) can easily be used to show that this position is profitable if the bond's return over the period exceeds the repo rate. If the bond pays a coupon c at the end of the period, then the terminal cash flow becomes

$$P_1 + c - P_0(1+r) \tag{17.5}$$

which is positive if the total return of the bond is greater than the repo rate.

As an aside, the arbitrage-free price of a single-period bond, i.e., with a terminal coupon payment of c and a terminal price of $P_1 = 1$, is immediately evident from either (17.3) or (17.5). The short and long positions described in this section require no cash at initiation. Therefore, to rule out arbitrage opportunities, terminal profits must be zero as well. Mathematically,

$$1 + c - P_0(1+r) = 0 \tag{17.6}$$

Or,

$$P_0 = \frac{1+c}{1+r} \tag{17.7}$$

In terms of the forward rates of Chapter 2, equation (17.7) says that the forward rate from date 0 to date 1, $f(1)$, must equal the repo rate, r .

The cash flows described here can easily be extended to a multi-period setting. If the unit price of a bond, its coupon rate, and its repo rate at time t are P_t and r_t , then the cash flow from a financed long position in the bond from time $t - 1$ to time t is

$$P_t + c - P_{t-1}(1 + r_{t-1}) \tag{17.8}$$

BOND ARBITRAGE WITH FINANCING

To illustrate the impact of financing arrangements on bond arbitrage, return to the example presented in the section “Arbitrage and the Law of One Price” in Chapter 1. As of May 28, 2010, the $\frac{3}{4}$ s of November 30, 2011, with a price of 100.190, were trading cheap relative to the set of bonds used to calculate the discount factors for that date. As a result, a replicating portfolio of those bonds could be constructed that made the same payments as the $\frac{3}{4}$ s but that could be sold for 100.255. The arbitrage trade proposed in Chapter 1 was to buy the $\frac{3}{4}$ s, sell the replicating portfolio, pocket the .065 difference, and owe nothing on any future date. In light of the realities of financing described in the previous section, however, this arbitrage trade would actually be initiated through the transactions in Table 17.3. Note that, unlike the narrative of Chapter 1, the initial cash flow of the trade is zero, not the mispricing of .065.

Financing realities also cause the performance of the arbitrage trade over time to differ from the account of Chapter 1. To explore these differences, the next three subsections track the performance of the trade over a single six-month period, i.e., through November 30, 2010, under three different scenarios. Table 17.4 summarizes these scenarios and the respective proceeds to the arbitrageur.

TABLE 17.3 Initial Transactions of the Arbitrage Trade of Table 1.5 on May 28, 2010, With Financing

Transaction	Cash Flow
Buy the $\frac{3}{4}$ s for 100.190	−100.190
Sell the repo of the $\frac{3}{4}$ s	100.190
Buy the repo of the replicating portfolio	−100.255
Sell the replicating portfolio	100.255
Total	0

TABLE 17.4 The Arbitrage Trade of Table 17.3 After Six Months in Three Illustrative Scenarios

	Scenarios		
	I	II	III
Repo 5/28: $\frac{3}{4}$ s	.149%	.149%	.149%
Replicating portfolio	.149%	0%	.149%
Prices 11/30: $\frac{3}{4}$ s	100.50	100.50	100.40
Replicating portfolio	100.50	100.50	100.50
Proceeds on 11/30	$100.255 \left(1 + \frac{.149\%}{2}\right)$ $-100.190 \left(1 + \frac{.149\%}{2}\right)$ $= .065 \left(1 + \frac{.149\%}{2}\right)$	$100.255 \left(1 + \frac{0\%}{2}\right)$ $-100.190 \left(1 + \frac{.149\%}{2}\right)$ $= -.00964$	$.065 \left(1 + \frac{.149\%}{2}\right)$ $+100.40 - 100.50$ $= -.03495$

Scenario I assumes that the $\frac{3}{4}$ s and all of the bonds in the replicating portfolio share the same repo rate and that the prices of the $\frac{3}{4}$ s and of the replicating portfolio converge over the period. The lesson of this scenario is that, because of financing, the arbitrage profit of the trade is not the initial mispricing of .065 at initiation of the trade but rather the future value of .065 when the trade is closed, where the future value is computed using the repo rate.

Scenario II assumes that the repo rate paid to finance the $\frac{3}{4}$ s exceeds the repo rate earned on shorting the replicating portfolio and that, like Scenario I, the prices of the two sides of the trade converge in six months. The lesson here is that when financing rates differ across bonds, arbitrage trades of this sort may not prove profitable at all. And this in turn implies that, with different financing rates, the law of one price and the arbitrage pricing results of Chapter 1 do not apply.

Scenario III returns to the assumption of equal repo rates across bonds, but assumes that the prices of the $\frac{3}{4}$ s and of the replicating portfolio diverge over the period. The lesson of this final scenario is that, given the institutional details of financing, if prices diverge an arbitrageur has to come up with additional cash. If the arbitrageur cannot do so, the trade has to be unwound at a loss which, if large enough, could threaten the viability of the arbitrageur or of the sponsoring financial entity. If the arbitrageur can finance interim divergences, then the ultimate profitability of the trade depends on the rates, over time, at which these divergences are financed. In the special case that these rates, over time, match the repo rates of the bonds, it turns out that the terminal profit of the trade, whatever the time series of divergences, is the future value of the initial mispricing at those realized repo rates.

The three subsections that follow elaborate on the arbitrage trade example in these scenarios. A proof in a more general setting is given in Appendix A in this chapter.

Scenario I

All repo rates the same and single-period convergence. This scenario assumes that all of the bonds trade at a six-month general collateral (GC) repo rate of .149%, which is consistent both with the price of the six-month bond in the replicating portfolio, as presented in Chapter 1, and with equation (17.7) that links the first forward rate and the single-period repo rate.

By construction of the arbitrage portfolio, the .375 coupon payment received from the $\frac{3}{4}$ s exactly matches the coupon owed on the replicating portfolio. Hence, after one period, the arbitrageur is left collecting the repo loan proceeds from having shorted the replicating portfolio and repaying the repo borrowing from having financed the purchase of the $\frac{3}{4}$ s. The net payoff, given in the last rows of the Scenario I column of Table 17.4 is $.065 (1 + \frac{.149\%}{2})$. This is just the future value of the initial mispricing computed at the repo rate.

Scenario II

Repo rate of the $\frac{3}{4}$ s exceeds that of the replicating portfolio and single-period convergence. This scenario assumes that the bonds in the replicating portfolio are all trading special so that their repo rate is relatively low. An obvious motivation for this assumption is that the short sellers are attracted by the richness of these bonds, thus increasing the demand to borrow them in the repo market and causing them to trade special. In particular, this scenario assumes that the $\frac{3}{4}$ s trade at the GC rate of .149%, as in Scenario I, while the bonds in the replicating portfolio trade at a special rate of 0%.

The one-period payoff in this scenario is calculated as that of Scenario I, except that the repo loan from shorting the replicating portfolio earns the special rate of 0%. The result of $-.00964$ is shown in the last rows of the Scenario II column of Table 17.4. Hence, the arbitrageur loses money even though the prices of the replicating portfolio and the $\frac{3}{4}$ s converge from 6.5 cents to zero. Put another way, the expense of borrowing the specials (i.e., lending at 0%) overwhelms the initial relative mispricing.

A broader interpretation of this scenario is that apparent mispricings may be explained by repo specialness. It might be that a bond trades rich because it can be lent out profitably in the repo market or, without imposing that causality, that arbitrageurs cannot take advantage of a bond's richness because it is expensive to borrow in the repo market.

Scenario III

All repo rates the same but convergence is preceded by a period of divergence. This scenario assumes that the price of the replicating portfolio rises to 100.50 but that the price of the $\frac{3}{4}$ s rises only to 100.40. In other words, the initial 6.5 cents of relative richness of the replicating portfolio diverges to 10 cents instead of converging to something less than 6.5 cents. The one-period payoff from the arbitrage trade is comprised of unwinding the two repos, at the same rates as in Scenario I, but, in addition, of selling the position in the $\frac{3}{4}$ s and purchasing (i.e., covering the short of) the replicating portfolio for a total of $100.40 - 100.50$ or $-.10$. The net payoff of $-.03495$ is given in the last rows of the Scenario III column of Table 17.4.

Of course, the arbitrageur need not unwind the trade at this point but could hold the position until maturity, when convergence in price is guaranteed. In fact, since the mispricing has increased from 6.5 cents to 10 cents in the first six months, the arbitrageur might even consider putting on more of the trade. In any case, to maintain the trade at its current size, the arbitrageur has to roll the two repo positions instead of exiting the bond positions. Selling the repo of the $\frac{3}{4}$ s generates 100.40 while buying the repo of the replicating portfolio requires 100.50. Therefore, rolling the arbitrage position, which means settling the expiring repo trades for $.065 (1 + \frac{149\%}{2})$ and initiating the new ones for $100.40 - 100.50$, results in the same cash flow as the one-period payoff computed for this scenario, i.e., $.03495$. Put another way, maintaining the trade requires a cash infusion of the same 3.5 cents.

The intuition behind the arbitrageur's need to put up more cash is best seen from the point of view of the arbitrageur's repo counterparties. Since the price of the replicating portfolio has risen to 100.50, the lender of this portfolio requires 100.50 in cash collateral to continue lending the bonds. However, since the price of the $\frac{3}{4}$ s has risen to only 100.40, the lender of cash is willing to lend only 100.40 against the $\frac{3}{4}$ s. Hence, the arbitrageur has to come up with $.10$ minus the net proceeds from the repo transactions over the first six months of the trade.

The conclusion from this discussion is that to enforce the pricing relationships derived and used in earlier chapters an arbitrageur needs to be able to finance cash requirements resulting from interim mispricings. The most important source of cash for this purpose is the arbitrageur's initial capital. An important example would be a hedge fund that has collected capital from principals and investors in order to conduct arbitrage and other investment activities. If a trade should require cash, the hedge fund can dip into its capital base to stay in the trade. Of course, if the trade loses too much before ultimate convergence, the hedge fund might have to unwind the trade at a loss to preserve its capital buffer. Worse, in an extreme case, a fund might fail because it has used so much of its capital in an unwind that its stability is no longer tenable.

It is a difficult problem to assign an interest rate to financings required by interim mispricings. Since, as just discussed, this cash often comes from an institution’s capital base, the most appropriate rate is often the cost of that capital. An analysis of cost of capital, however, would take the discussion of this chapter far afield. Instead, relying on the admittedly strong assumption that these mispricings can be financed at the repo rate of the bonds in the trade, a conceptually useful result emerges. Namely, the profit of an arbitrage trade at convergence is the future value of the initial mispricing, where this future value is computed at the realized, single-period repo rates over the life of the trade. This result, a generalization of the single-period result of Scenario I, is proved in Appendix A in this chapter.

SWAP TRADING WITH FINANCING

The cash flows of receiving fixed on a swap are qualitatively the same as those from a financed bond purchase, with one key difference. First, to see the similarities, assume that the bond price and the present value of the fixed leg of the swap are always par and denote the bond coupon rate by c , the repo rate at time t by r_t , and the floating rate of the swap at time t by L_t . Also, for simplicity, assume that the fixed and floating payment dates are the same. Then, for unit notionals, Table 17.5 shows that the cash flows of the long financed bond and swap positions are identical except for the fact that the bond finances at the repo rate while the fixed cash flows of the swap finance at the floating rate.

Second, to see the key difference between receiving fixed and financing a bond purchase, drop the assumption of constant par values and denote the price of the bond at time t per unit face amount by P_t and the net present value at time t of receiving fixed on a unit notional by NPV_t . This NPV is, as in Chapter 16, the difference between the value of the fixed and floating legs, including the fictional notional payments, where discounting is done with rates derived from par swaps of fixed against L_t . Furthermore, because this analysis looks only at payment dates, the value of the floating leg is par and

TABLE 17.5 Cash Flows from a Long Financed Bond Position and Receiving Fixed in a Swap if Prices and Present Values Are Always Par

Cash Flows	Buy the Bond	Sell the Repo	Bond Total	Receive Fixed	Pay Floating	Swap Total
Initial	−1	1	0	0	0	0
Time t	c	$−r_{t−1}$	$c − r_{t−1}$	c	$−L_{t−1}$	$c − L_{t−1}$
Final	1	−1	0	0	0	0

TABLE 17.6 Cash Flows from a Long Financed Bond Position and Receiving Fixed in a Swap

Cash Flows	Buy the Bond	Sell the Repo	Swap Flows	
			Cash	Collateral
Initial	$-P_0$	P_0	$-NPV_0$	NPV_0
Time t	c	$-P_{t-1}(1 + r_{t-1}) + P_t$	$c - L_{t-1}$	$-NPV_{t-1}(1 + r_{t-1}) + NPV_t$
Time T	$1 + c$	$-P_{T-1}(1 + r_{T-1})$	$c - L_{T-1}$	$-NPV_{T-1}(1 + r_{T-1})$

the NPV is the present value of the fixed leg minus par. Then, Table 17.6 contrasts the cash flows of the long financed bond with those of receiving fixed. (Note that the row with time T cash flows gives the sum of all flows at the maturity date of the bond and swap, not just the principal flows as in the last row of Table 17.5.)

The bond flows in Table 17.6 are just as given in expression (17.8). As for the swap flows, the fixed receiver pays NPV_0 to the fixed payer so that both parties are willing to enter into the swap. Of course, if the fixed rate is less than the par rate, this “payment” will be negative so the fixed receiver is, in fact, receiving cash at the initial date. In any case, to ensure the fixed payer’s performance on the swap contract, which has a net initial present value of NPV_0 , the fixed receiver immediately takes NPV_0 from the fixed payer as collateral. Hence, the total cash flow at initiation is zero no matter what the initial NPV.

At time t , the net cash flow to the fixed receiver is, of course, $c - L_{t-1}$. But there are collateral flows as well. As customary in financial markets, posted collateral earns interest. Let r_t be the rate earned on swap collateral from $t - 1$ to t , to emphasize that this rate and the repo rate are both rates earned on collateral. At every date t , then, the fixed receiver returns collateral taken the previous period with interest, for a flow of $-NPV_{t-1}(1 + r_{t-1})$, and takes the appropriate amount of collateral for date t , i.e., NPV_t . At time T , of course, the swap matures, the NPV is identically zero, and collateral posting ends, so the collateral cash flow is just the return of the previous period’s collateral with interest.

To interpret the results of Table 17.6, rewrite the total bond cash flow at time t as

$$c + P_t - P_{t-1} - P_{t-1}r_{t-1} \quad (17.9)$$

To rewrite the total swap cash flow at time t similarly, define the present value of the fixed side of the swap at time t (with the fictional notional payment) as PV_t so that $NPV_t = PV_t - 1$. Then, from Table 17.6, the total swap cash flow at time t can be written as

$$c + PV_t - PV_{t-1} - 1 \times L_{t-1} - NPV_{t-1} \times r_{t-1} \quad (17.10)$$

Both the bond and swap expressions, (17.9) and (17.10) respectively, have a fixed payment c and the price change of the fixed side, i.e., $P_t - P_{t-1}$ or $PV_t - PV_{t-1}$ respectively. But the expressions differ in financing. The total value of the bond finances at the repo rate. The total value of the fixed side of the swap, however, finances at a blended rate of the floating rate and the collateral rate. More specifically, the total value of the fixed leg is $PV_t = 1 + NPV_t$. The face amount, i.e., 1, finances at the floating rate while the NPV at time t , i.e., NPV_t , finances at the collateral rate.

SWAP ARBITRAGE WITH FINANCING

It was concluded earlier that, in the bond context, discounting and arbitrage pricing were equivalent only if all bonds financed at the same repo rate. What is the analogous requirement ensuring that the fixed leg of all swaps finance at the same rate? Based on the discussion of the previous section, the floating rate would have to be common for all swaps, and the collateral rate would have to equal that same floating rate. Only in that way would the blended financing rate, highlighted in expression (17.10), be the same for all swaps, regardless of their NPVs. Appendix A in this chapter proves that discounting and arbitrage pricing for swaps are equivalent if these conditions hold, along with the condition that interim mispricings can be financed at the floating rate as well. For concreteness, however, this section works through two examples of arbitrage relationships across swaps. Scenario I shows that, under the conditions just specified, a high-coupon, one-year swap can indeed be replicated by six-month and one-year par swaps. Scenario II shows that this replication fails if the collateral rate does not equal the floating rate index. For simplicity, as in the previous section, both fixed and floating legs are assumed to make cash flows in six-month intervals.²

Both scenarios start from the USD swap curve as of May 28, 2010, originally shown in Table 2.1 and reproduced here, in part, in Table 17.7. According to the discounting methodology of Chapters 1 and 2, and the Chapter 16 result that the floating leg is worth par, the NPV of 100 face amount of a 1-year swap to receive 5% against LIBOR is

$$\frac{2.5}{1 + \frac{.705\%}{2}} + \frac{102.50}{\left(1 + \frac{.705\%}{2}\right) \left(1 + \frac{1.046\%}{2}\right)} - 100 = 4.0998 \quad (17.11)$$

But is there an arbitrage trade, which takes financing into account, that supports the valuation given in (17.11)? To begin, Table 17.8 uses the

²As the numerical examples are taken from the fixed *versus* three-month LIBOR curve, the precise assumption here is that the floating leg pays compounded three-month LIBOR every six months.

TABLE 17.7 Selected Data from the USD Swap Curve
as of May 28, 2010, from Table 2.1

Term in Years	Par Rate	Forward Rate
0.5	.705%	.705%
1.0	.875%	1.046%

methodology of Chapter 1 to find the portfolio of .5-year and 1-year par swaps that replicates the cash flows of 100 face amount of a 5% swap. The candidate arbitrage trade, therefore, is to receive fixed on the 5% swap and to pay fixed on the replicating portfolio given in Table 17.8.

Scenario I

No relative mispricing and the rate earned on collateral equals the first LIBOR set of .705%. In six months, LIBOR is .9%. In this scenario Table 17.9 shows that, taking financing into account, the one-year, 5% swap is, in fact, replicated by the portfolio given in Table 17.8. Dates 1 and 2 are six months and one year from the trade date, respectively.

Row (i) of Table 17.9 applies expression (17.10) to the 5% swap to determine its flows on date 1. Column (4) of this row gives the date 1 present value of the swap. Since its date 2, fixed-leg cash flow is 102.5 and the six-month rate in six months is .9% in this scenario, this present value is $\frac{102.5}{1 + \frac{.9\%}{2}}$ or 102.0408. The date 0 present value of the fixed leg, used in column (5), is simply 100 plus the NPV calculated in (17.11), which NPV is also used in column (7).

Rows (ii) and (iii) apply expression (17.10) to find the date 1 cash flows of the swaps in the replicating portfolio, with the face amounts given in Table 17.8. Since the 6-month, .705% swap matures on date 1, its date 1 present value in column (4) is simply its face amount. The date 2 cash flow of the one-year .875% swap, by construction, matches the 102.5 date 2 cash flow of the 5% swap. Its date 1 present value in column (4), therefore, is also 102.0408. Finally, because the present value of the fixed legs of par swaps

TABLE 17.8 Portfolio of .5- and 1-Year Par Swaps that Replicates a 5%, 1-year Swap as of May 28, 2010

	Replicating Portfolio		Swap
Fixed Rate	.705%	.875%	5%
Term in Years	.5	1.0	1.0
Face Amount	2.0463	102.0535	100

TABLE 17.9 Replicating a One-Year Swap With Two Other Swaps, With Financing

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Date 1	Terms of Expression (17.10)						
	Rate	Face	c	PV_1	$-P V_0$	$-\text{face} \times L_0$	$-NPV_0 \times r_0$
(i)	5.000%	100.0000	2.5000	102.0408	-104.0998	$-100 \frac{.705\%}{2}$	$-4.0998 \frac{.705\%}{2}$
(ii)	.705%	2.0463	.0070	2.0463	-2.0463	$-2.0463 \frac{.705\%}{2}$	0
(iii)	.875%	102.0535	.4465	102.0408	-102.0535	$-102.0535 \frac{.705\%}{2}$	0
(iv)	Replicating Portfolio		2.5 + 102.0408		-104.0998	$-104.0998 \frac{.705\%}{2}$	
Date 2	Terms of Expression (17.10)						
	Rate	Face	c	PV_2	$-P V_1$	$-\text{face} \times L_1$	$-NPV_1 \times r_1$
(v)	5.000%	100.0000	2.5000	100.0000	-102.0408	$-100 \frac{.9\%}{2}$	$-2.0408 \frac{.9\%}{2}$
(vi)	.875%	102.0535	.4465	102.0535	-102.0408	$-102.0535 \frac{.9\%}{2}$	$-(102.0408$ $-102.0535) \frac{.9\%}{2}$
(vii)	Replicating Portfolio		102.5		-102.0408	$-102.0408 \frac{.9\%}{2}$	

are, by definition, equal to their notional amounts, the entries in column (5) are simply the negatives of these notional amounts.

Row (iv) of Table 17.9 adds rows (ii) and (iii) to obtain the total flows from the replicating portfolio. By inspection, rows (iv) and (i) are equal, demonstrating that the date 1 flows, accounting for financing, are the same for the 5% swap and its replicating portfolio.

Rows (v) and (vi) of Table 17.9 apply expression (17.10) to the date 2 flows of the 5% and .875% swaps, the .705% swap having already matured. Row (vii), the flows from the replicating portfolio, which simplifies row (vi), is, by inspection, equal to row (v). Hence, the date 2 flows, accounting for financing, are also the same for the 5% swap and its replicating portfolio.

Scenario II

No relative mispricing and the rate earned on collateral is .25%, less than the floating rate index of .705%. In six months the floating rate index is .9%.

In this scenario, the rate r_0 used in column (7) of Table 17.9 is .25% instead of .705%. Thus, the financing of the 5% swap, given by the sum of columns (6) and (7) of row (i), becomes

$$-100 \frac{.705\%}{2} - 4.0998 \frac{.25\%}{2} \quad (17.12)$$

while the financing of the replicating portfolio, given by columns (6)–(7) of row (iv) remains at

$$-104.0998 \frac{.705\%}{2} \quad (17.13)$$

Expressions (17.12) and (17.13) are not equal because the NPV part of the 5% swap finances at the now lower collateral rate while the replicating portfolio, which is composed of par swaps, finances fully at LIBOR. Hence, accounting for financing, the flows of the 5% swap and those of the so-called replicating portfolio are not the same.

PRICING A USD LIBOR SWAP WITH FED FUNDS AS THE INVESTABLE AND COLLATERAL RATE

As recounted in the introduction to this chapter, the crisis of 2007–2009 pushed market practice to taking fed funds, or international equivalents, as the riskless investable rate. In addition, collateral on swaps, including LIBOR swaps, continues to earn fed funds. Using arbitrage arguments, this section

derives the price of a fixed *versus* LIBOR swap under these conditions. The resulting shorthand of two-curve discounting has rapidly become the standard across the industry.

The three swaps used in the arbitrage argument are the following:

1. The fixed *versus* LIBOR swap, which exchanges a fixed rate c^L for LIBOR, L_t , paid at time t , and which earns fed funds on collateral. The initial NPV of this swap is denoted by NPV_0^L . But, most importantly, this NPV is not calculated using discount factors or rates from the par swap curve. As shown in the previous sections, that methodology is equivalent to arbitrage pricing only if collateral earns LIBOR, which is not the case here. Instead, NPV_0^L is derived in this section by arbitrage arguments.
2. An OIS or fixed *versus* fed funds swap, which exchanges a fixed rate c^{FF} for fed funds r_t , paid at time t .³ Collateral posted against this swap, consistent with market practice, earns fed funds. Then, since fed funds is taken as the riskless investable rate, the floating leg of the swap is worth par. Furthermore, since collateral earns fed funds, the arbitrage pricing of any OIS can be found by discounting its fixed-leg cash flows at rates derived from the par OIS curve. Putting all of this together, let NPV_0^{FF} denote the NPV of the fixed and floating legs of the fed funds swap calculated using the OIS curve.
3. A T -year basis swap (see Chapter 16) which exchanges fed funds r_t , paid at time $t + 1$, plus a fixed spread $X(T)$, for LIBOR L_t , also paid at time t . The basis swap spread $X(T)$ is determined such that the basis swap is fair, i.e., such that the parties do not need to exchange an up-front payment when initiating the swap. Collateral posted against the basis swap is also assumed to earn fed funds.

Proceeding to the arbitrage pricing of the LIBOR swap, Table 17.10 shows the cash flows of the various swaps. Row (iii) gives the flows from the portfolio of receiving fixed *versus* fed funds and of receiving fed funds *versus* LIBOR. Row (iv) gives the flows from receiving fixed *versus* LIBOR. By inspection, the interim flows of these two rows are identical for $c^{FF} = c^L - X(T)$. Hence, for that c^{FF} , arbitrage pricing requires that the NPV of the portfolio equal the NPV of the fixed *versus* LIBOR swap, or, mathematically, $NPV_0^L = NPV_0^{FF}$.

In conclusion, then, to price a T -year swap of a fixed rate c^L *versus* LIBOR that earns fed funds on collateral, compute the NPV of a fixed *versus* fed funds swap with fixed rate $c^L - X(T)$, where $X(T)$ is the T -year basis swap spread of fed funds *versus* LIBOR.

³More precisely, compounded fed funds over the period t to $t + 1$ is paid at time $t + 1$.

TABLE 17.10 Equivalence of a Combination of a Fixed Rate Swap vs. Fed Funds and a Swap of Fed Funds vs. LIBOR to a Swap of a Fixed Rate vs. LIBOR

	Transaction	Initial NPV	Initial Collateral	Interim Flows
(i)	Rec Fixed vs. Fed Funds	$-NPV_0^{FF}$	NPV_0^{FF}	$c^{FF} - r_t$
(ii)	Rec Fed Funds vs. LIBOR	0	0	$r_t + X(T) - L_t$
(iii)	Total	$-NPV_0^{FF}$	NPV_0^{FF}	$c^{FF} + X(T) - L_t$
(iv)	Rec Fixed vs. LIBOR	$-NPV_0^L$	NPV_0^L	$c^L - L_t$

For concreteness, consider the special case of unit notional, two-period swaps. Let $f(t)$ denote forward rates from the fed funds curve. Then, by the argument just made,

$$NPV_0^L = NPV_0^{FF} = \frac{c^L - X(2)}{(1 + f(1))} + \frac{1 + c^L - X(2)}{(1 + f(1))(1 + f(2))} - 1 \quad (17.14)$$

This pricing rule defines a relationship between par swap rates and basis spreads. By definition, the NPV of par swaps is 0. Hence, by equation (17.14), the two-year par fixed *versus* LIBOR swap rate, $C^L(2)$, is such that

$$\frac{C^L(2) - X(2)}{(1 + f(1))} + \frac{1 + C^L(2) - X(2)}{(1 + f(1))(1 + f(2))} - 1 = 0 \quad (17.15)$$

Or,

$$\begin{aligned} & \frac{C^L(2)}{(1 + f(1))} + \frac{1 + C^L(2)}{(1 + f(1))(1 + f(2))} \\ &= 1 + \frac{X(2)}{(1 + f(1))} + \frac{X(2)}{(1 + f(1))(1 + f(2))} \end{aligned} \quad (17.16)$$

It is worth pointing out a qualitative difference between equation (17.16) and the pricing of Chapter 16. In the latter, with LIBOR taken as the riskless investable rate, the values of the fixed and floating legs of par LIBOR swaps are both worth par. Contrast this with (17.16). The left-hand side of this equation gives the present value of the fixed side of a LIBOR par swap when the riskless investable rate and the collateral rate are fed funds. But since LIBOR-OIS basis swap spreads are positive, the right-hand side is greater than par. Equivalently, viewing the right-hand side of (17.16) as the value of the floating side of the LIBOR swap, this value also exceeds par. This latter view is quite intuitive: receiving fed funds is worth par when discounting at

fed funds but, because $X(T)$ is positive in the basis market, receiving LIBOR is worth more than receiving fed funds.

Two-Curve Pricing

In practice, when implementing the pricing methodology of earlier chapters, swaps are valued by discounting both the fixed and floating legs, where future floating rates are taken to be the forward rates of the curve. In this way the value of the floating side of the swap is normally par, but the methodology is flexible enough to incorporate cash flow details that can slightly increase or decrease the value of the floating leg. What is the analogous procedure for implementing the correct pricing of (17.16)? The answer is to use two curves, one for discounting and one for projecting future rates. Details are presented in Appendix B in this chapter but the basic idea is presented in this subsection. Define projected floating rates L'_t such that the present value of the floating side of any maturity equals that maturity's equivalent of the right-hand side of (17.16). For a two-period swap, for example, L'_1 and L'_2 have to be defined such that

$$\begin{aligned} \frac{L'_1}{(1+f(1))} + \frac{1+L'_2}{(1+f(1))(1+f(2))} \\ = 1 + \frac{X(2)}{(1+f(1))} + \frac{X(2)}{(1+f(1))(1+f(2))} \end{aligned} \quad (17.17)$$

The computation of the L'_t can be done via (17.17), but it is even simpler to find the L'_t such that the value of the fixed sides of all par swaps, like the left-hand side of (17.16), equals the values of the floating sides, like the left-hand side of (17.17). In this way, the basis swaps are implicitly, but not explicitly, used in the calculations. In any case, after building up these projected rates L'_t for every date, swaps can be valued without any reference to the basis swap rates. The fixed side of any swap is valued by discounting its cash flows at the rate corresponding to the riskless investable rate and collateral rate (e.g., the left-hand side of (17.16)). The floating side is valued by discounting the basis-adjusted projected floating rates by that same curve (e.g., the left-hand side of (17.17)).

The Effect of LIBOR-OIS Spreads on Swap Valuation

This chapter has established that discounting using a par swap curve does not accurately value non-par swaps when the two-curve approach is appropriate. But how large an error is made by proceeding in this theoretically incorrect way? To answer this question, assume that trading desk #1, which values all swaps off the par swap curve, incorrectly calculates NPV^* as the

TABLE 17.11 Swap Valuation Errors from Failing to Account for a Fed Funds Collateral as of May 28, 2010. Spreads and errors are in basis points.

(1)	(2)	(3)	(4)
Valuation Error in Terms of Swap Rates			
Term	LIBOR-OIS Basis	200bp over Par Rate	100bp over Par Rate
2	48	1.2	.6
5	38	2.3	1.2
10	30	3.6	1.8

value of receiving the fixed rate c^1 *versus* three-month LIBOR. Trading desk #2, which accurately values swaps using the methodology of this section, computes that the fixed rate that truly generates a value of NPV^* is c^2 . Then, define the valuation error of trading desk #1 as $c^1 - c^2$.

Table 17.11 shows these errors calculated as of May 28, 2010, meaning that the par swap curve and the LIBOR-OIS basis curve are taken as of that date. Column (1) gives selected swap terms. Column (2) gives the LIBOR-OIS money market basis swap spread of that term. Note that, according to Figure 15.5, the level of these spreads is above that of the long-term history but below that prevailing over most of the crisis. Columns (3) and (4) give the valuation errors as defined in the previous paragraph. The higher the swap rate, the larger the NPV and, therefore, the larger the error from incorrectly accounting for the interest rate earned on collateral. The longer the term, the larger the error from discounting obligations at LIBOR that should be discounted at fed funds. While the orders of magnitude of these errors might appear small at first, they should be understood in the context of a market in which bid-ask spreads are .25 or .5 basis points.

Table 17.12 shows the errors calculated using the par swap curve as of May 28, 2010, but with LIBOR-OIS basis at levels more consistent with the crisis. Not surprisingly, the valuation errors are commensurately larger.

This subsection concludes with a discussion of the sign of the valuation errors. Under a valuation that does not take into account the collateral rate, it is advantageous to pay the NPV, which essentially invests it through the swap at a relatively high rate, i.e., LIBOR, and to take that same NPV as collateral, which essentially borrows it at a relatively low rate, i.e., fed funds. Hence, the receiver of an above-par rate swap and the payer of a below-par rate swap, who both receive collateral, value those swaps more than would be indicated by discounting at par swap rates. In terms of the errors in Tables 17.11 and 17.12, c^2 , the correct fixed rate associated with NPV^* , which reflects the value of taking collateral, is lower than the incorrect rate

TABLE 17.12 Swap Valuation Errors from Failing to Account for a Fed Funds Collateral; Curve as of May 28, 2010; LIBOR-OIS Basis Representative of Crisis Levels. Spreads and errors are in basis points.

(1)	(2)	(3)	(4)
Valuation Error in Terms of Swap Rates			
Term	LIBOR-OIS Basis	200bp over Par Rate	100bp over Par Rate
2	100	2.1	1.1
5	60	4.2	2.1
10	45	5.8	2.9

associated with NPV^* , which does not reflect the value of taking collateral. Hence $c^1 > c^2$ and the errors in the tables are all positive.

Bases other than LIBOR-Fed Funds

This section has focused on swap valuation when the floating rate is three-month LIBOR and the collateral rate is fed funds. But the logic here applies to other bases as well. For example, to value a fixed *versus* one-month LIBOR swap, with collateral rate of fed funds, discounting has to be done at fed funds while projected floating cash flows are determined implicitly or explicitly by one-month LIBOR-fed funds basis swap spreads.

APPENDIX A: ARBITRAGE RELATIONSHIPS ACROSS BONDS AND SWAPS WITH FINANCING

Proposition 1: Consider two bond portfolios constructed so that their coupon and principal cash flows are identical. When longs and shorts have to be financed by collateralized borrowing and lending, respectively, the arbitrage arguments that the two portfolios must have the same price obtains if, in addition to the usual assumption of trivial transaction costs:

- the two portfolios always finance at the same rate;
- interim mispricings can always be financed at that same rate.

Furthermore, under these conditions, an arbitrage that is long one portfolio and short the other generates a cash flow at maturity equal to the realized value of rolling over the intial mispricing to the maturity date at the financing rate.

Proof: Consider portfolios A and B that have been constructed so that their coupon and principal cash flows are identical. Set the following notation:

- P_t^A, P_t^B : prices of the portfolios at time $t, t = 1, \dots, T$
- r_t^A, r_t^B : financing rates on the portfolios from time $t - 1$ to time t . If $r_t^A = r_t^B$ then denote the financing rate by r_t
- ϵ_t : the market richness ($\epsilon_t > 0$) or cheapness ($\epsilon_t < 0$) of portfolio A relative to portfolio B at time t in price terms, i.e.,

$$P_t^A = P_t^B + \epsilon_t \quad (17.18)$$

Assume without loss of generality that the trader decides to buy portfolio A and short portfolio B. At initiation, at time 0, the trader does the following:

- Buy portfolio A for P_t^A , borrow its price P_t^A , and deliver it as collateral against the loan.
- Sell portfolio B, lend its price P_t^B , taking it as collateral against the loan.

Note that the total cash flow from these trades is zero.

Subsequently, on any date t , the trader does the following:

- Collect the coupon from portfolio A and pay the coupon to portfolio B, which payments, by construction, are perfectly offsetting.
- Roll over the financing of portfolio A over the previous period $t - 1$ to t by paying off the loan of $P_{t-1}^A (1 + r_t^A)$ and taking out a new loan of P_t^A .
- Roll over the short of portfolio B over the previous period by collecting loan proceeds of $P_{t-1}^B (1 + r_t^B)$ and making a new loan of P_t^B .

The total cash flow on date $t \leq T$ is, therefore,

$$P_t^A - P_{t-1}^A (1 + r_t^A) - P_t^B + P_{t-1}^B (1 + r_t^B) \quad (17.19)$$

To focus first on the financing rates, assume for the moment that the prices of these portfolios with identical coupon and principal cash flows are the same, i.e., $\epsilon_t = 0$ for all t . Then, (17.19) becomes

$$P_{t-1}^B [r_t^B - r_t^A] \quad (17.20)$$

But the quantity (17.20) will not be zero, as required for replication, unless $r_t^A = r_t^B$ for all t . In other words, even if the prices of the two portfolios are identical over time, the total cash flows from carrying a long in one and a short in the other will not cancel out unless the financing rates are the same.

Having established that the financing rates over time have to be the same for the arbitrage strategy to work, let that single rate be simply r_t and

allow for the possibility of interim mispricing, i.e., that ϵ_t may be positive or negative. Substituting (17.18) into (17.19), the cash flow on date t becomes

$$\epsilon_t - \epsilon_{t-1} (1 + r_t) \quad (17.21)$$

Equation (17.21) demonstrates that the arbitrage trade of buying one portfolio and shorting another may require cash injections or allow for cash withdrawals if there are any market mispricings. Hence the condition in the proposition that interim mispricings have to be financed in some way.

Now consider the special case that these mispricings are financed at rates r_t . The date 1 cash flow of (17.21), carried forward to date 2, is

$$[\epsilon_1 - \epsilon_0 (1 + r_1)] (1 + r_2) \quad (17.22)$$

And the date 2 cash flow of (17.21) is

$$\epsilon_2 - \epsilon_1 (1 + r_2) \quad (17.23)$$

Adding (17.22) and (17.23) together, the terms with ϵ_1 cancel, so that the accumulated cash on date 2 is

$$\epsilon_2 - \epsilon_0 (1 + r_1) (1 + r_2) \quad (17.24)$$

Continuing forward in this way, the accumulated cash at the maturity date T is

$$- \epsilon_0 (1 + r_1) (1 + r_2) \cdots (1 + r_T) \quad (17.25)$$

In words, the accumulated cash is the future value of the initial mispricing. Since in this exposition the trader bought portfolio A and sold portfolio B, portfolio A was presumably cheap at date 0, that is, $\epsilon_0 < 0$. Thus, the quantity available at maturity given by (17.25) is positive and depends only on the initial mispricing upon which the trader has originally decided to do the arbitrage trade.

Proposition 2: Consider two swap portfolios against a single floating rate index constructed so that their fixed-side coupon and (fictional) principal cash flows are identical. The arbitrage arguments that the two portfolios must have the same price obtains if, in addition to the usual assumption of trivial transaction costs:

- the riskless investable rate and the rate earned on collateral posted against the NPV of the swap portfolios equal that same floating rate index;
- interim mispricings can always be financed at that floating rate index.

Furthermore, under these conditions, an arbitrage that is long one portfolio and short the other generates a cash flow at maturity equal to the realized value of rolling over the initial mispricing to the maturity date at the floating rate index.

Proof: Consider portfolios A and B that have been constructed so that their coupon and principal cash flows are identical. Set the following notation:

- P_t^A, P_t^B : prices of the portfolios at time $t, t = 1, \dots, T$, where discounting is done at the curve corresponding to the floating rate index.
- F_t^A, F_t^B : total face or notional amount of the swaps in portfolios A and B. Note that this amount can change over time as swaps in the portfolios reach maturity.
- NPV_t^A, NPV_t^B : NPV of the portfolios, defined as the present value of the fixed side minus the present value of the floating side, where all discounting is done at the curve corresponding to the floating rate index. For the fixed sides, the present values have already been defined as P_t^A and P_t^B . For the floating side, as shown in Chapter 16, the present values are simply the face amounts. Hence,

$$NPV_t^A = P_t^A - F_t^A \quad (17.26)$$

$$NPV_t^B = P_t^B - F_t^B \quad (17.27)$$

- L_t : the floating rate index over the period $t - 1$ to t . This need not necessarily be a LIBOR index.
- r_t : rate earned on collateral posted from time $t - 1$ to time t
- ϵ_t : the market richness ($\epsilon_t > 0$) or cheapness ($\epsilon_t < 0$) of portfolio A relative to portfolio B at time t in price terms, defined as in (17.18)

Assume without loss of generality that the trader receives fixed on portfolio A and pays fixed on portfolio B. At time 0, then, the trader does the following:

- Agree to receive fixed and pay floating with portfolio A, pay NPV_0^A , take NPV_0^A as collateral, paying interest at the rate r_1 .
- Agree to pay fixed and receive floating with portfolio B, receive NPV_0^B , give NPV_0^B as collateral, receiving interest at the rate r_1 .

Note that the total cash flow from these trades is zero.

Subsequently, on any date t , the trader does the following:

- Receive fixed (including fictional principal payments) from A and pay fixed (including fictional principal payments) through B, which payments, by construction, are perfectly offsetting.

- Pay floating through A and receive floating through B (including fictional principal payments) at L_t on face amounts F_{t-1}^A and F_{t-1}^B , respectively. Note that the fictional maturing principal of A paid is $F_{t-1}^A - F_t^A$ while the fictional maturing principal of B received is $F_{t-1}^B - F_t^B$. These amounts have to be included because the fictional principal amounts are included in the fixed payments of the two portfolios.
- Pay interest on NPV_{t-1}^A at the rate r_t , return the collateral NPV_{t-1}^A and take the revised collateral amount NPV_t^A .
- Receive interest on NPV_{t-1}^B at the rate r_t , return the collateral NPV_{t-1}^B and take the revised collateral amount NPV_t^B .

Therefore, the total cash flow on date t is, after offsetting the fixed payments,

$$\begin{aligned} & -F_{t-1}^A L_t - [F_{t-1}^A - F_t^A] - NPV_{t-1}^A (1 + r_t) + NPV_t^A \\ & + F_{t-1}^B L_t + [F_{t-1}^B - F_t^B] + NPV_{t-1}^B (1 + r_t) - NPV_t^B \end{aligned} \quad (17.28)$$

Substituting prices for NPVs using equations (17.26) and (17.27),

$$\begin{aligned} & -F_{t-1}^A L_t - [F_{t-1}^A - F_t^A] - [P_{t-1}^A - F_{t-1}^A] (1 + r_t) + P_t^A - F_t^A \\ & + F_{t-1}^B L_t + [F_{t-1}^B - F_t^B] + [P_{t-1}^B - F_{t-1}^B] (1 + r_t) - P_t^B + F_t^B \end{aligned} \quad (17.29)$$

And simplifying,

$$\begin{aligned} & F_{t-1}^A [r_t - L_t] + P_t^A - P_{t-1}^A (1 + r_t) \\ & - [F_{t-1}^B [r_t - L_t] + P_t^B - P_{t-1}^B (1 + r_t)] \end{aligned} \quad (17.30)$$

To focus first on the rate earned on collateral posted, assume for the moment that the prices of these portfolios with identical fixed-side payments are the same, i.e., $\epsilon_t = 0$ for all t . Then, (17.30) becomes

$$[F_{t-1}^A - F_{t-1}^B] [r_t - L_t] \quad (17.31)$$

Ignoring the trivial case of $F_{t-1}^A = F_{t-1}^B$, in which the portfolios contain exactly the same swaps so that a long and short position is identically no position, the quantity (17.31) will not be zero, as desired, unless $r_t = L_t$ for all t . In other words, even if the prices of the two fixed-side, cash-matched portfolios are identical over time, the total cash flows from the arbitrage position will not cancel out unless the rate earned on collateral posted equals the floating rate index.

Having established that the rate on collateral must equal the floating rate index for the arbitrage strategy to work, allow for the possibility of interim mispricing, i.e., that $\epsilon(t)$ may be positive or negative. With $r_t = L_t$, substituting (17.18) into (17.30) gives a date- t cash flow of

$$\epsilon_t - \epsilon_{t-1} (1 + L_t) \quad (17.32)$$

But this is completely analogous to (17.21), so continuing the argument along the lines of the proof of Proposition 1 shows that the cash accumulated by the arbitrage trade through the maturity date T is the future value of the initial pricing with the floating index rates of compounding:

$$-\epsilon_0 (1 + L_1) (1 + L_2) \cdots (1 + L_T) \quad (17.33)$$

APPENDIX B: PRICING SWAPS WITH THE TWO-CURVE APPROACH

Proposition 3: Suppose that the riskless investable and collateral rates on all swaps are r_t and that the basis swap spread of $X(T)$ makes payments of $r_t + X(T)$ fair against the LIBOR index, L_t . Let the forward rates of the curve corresponding to the short-term rate realizations r_t be $f(t)$. Then, the NPV of a T -year swap that exchanges the fixed rate, $c(T)$, for L_t is given by

$$NPV = [c(T) - X(T)] A(T) + d(T) - 1 \quad (17.34)$$

where $A(T)$ and $d(T)$ are the annuity factors and discount factors, respectively, from the curve described by the forward rates, $f(t)$.

Proof: To price the swap of $c(T)$ against LIBOR, note that this swap is equivalent to the portfolio of the following two swaps:

- Receive $c(T) - X(T)$ and pay r_t
- Receive $r_t + X(T)$ and pay L_t

By definition of the basis swap spread, $X(T)$, the NPV of the second swap is zero. Hence, the NPV of the swap to be priced must equal the NPV of the first swap. But, from Proposition 2, the NPV of this first swap is priced by arbitrage arguments as

$$[c(T) - X(T)] A(T) + d(T) - 1 \quad (17.35)$$

Proposition 4: Suppose that the riskless investable and collateral rates on all swaps are r_t and that the basis swap spread of $X(T)$ makes payments of $r_t + X(T)$ fair against the LIBOR index, L_t . Then, swaps of fixed rates against the LIBOR index can be priced by the following two-curve approach:

- Discount all cash flows, fixed and projected floating, at the forward rates $f(t)$.
- Project floating rate payments using adjusted LIBOR rates, $\widehat{L}(t)$, which are constructed iteratively for $t = 1, \dots, T$ from the following equation:

$$\sum_{s=1}^t \frac{\widehat{L}(s)}{(1 + f(1)) \cdots (1 + f(s))} = 1 + X(t)A(t) - d(t) \quad (17.36)$$

where $A(t)$ and $d(t)$ are the annuity factors and discount factors, respectively, from the curve formed by the forward rates, $f(t)$.

Proof: Under the proposed methodology, the NPV of a swap with maturity T , the value of the fixed leg at the rate $c(T)$ minus the value of the floating leg, would be computed as follows:

$$NPV(T) = c(T)A(T) + d(T) - \left[\sum_{t=1}^T \frac{\widehat{L}(t)}{(1 + f(1)) \cdots (1 + f(t))} + d(T) \right] \quad (17.37)$$

Substituting for the adjusted LIBOR rates using (17.36), this becomes

$$NPV(T) = c(T)A(T) + d(T) - [1 + X(T)A(T)] \quad (17.38)$$

But, rearranging terms, equation (17.38) is exactly the correct NPV, as given by (17.34) in Proposition 3. Hence, the two-curve methodology does price all swaps correctly.

Corollary 1: The projected floating rate payments of the two-curve method in Proposition 4 can be computed using the par swap rates against the LIBOR index, $C(t)$, instead of using the basis swap spreads $X(t)$, through the following equation:

$$\sum_{s=1}^t \frac{\widehat{L}(s)}{(1 + f(1)) \cdots (1 + f(s))} = C(t)A(t) \quad (17.39)$$

Proof: Par swaps, by definition, have an NPV of zero. Hence, from (17.34),

$$[C(T) - X(T)] A(T) + d(T) = 1 \quad (17.40)$$

Rearranging terms,

$$1 + X(T)A(T) - d(T) = C(T)A(T) \quad (17.41)$$

But substituting (17.41) for the right-hand side of (17.36) gives the desired result (17.39).

Fixed Income Options

This chapter discusses some of the most popular fixed income options, namely, caps and floors, swaptions, bond options, Eurodollar (ED) and Euribor futures options, and bond futures options. These options can be, and some often are, priced using the term structure models of Part Three. When the objective is not to compute relative cheapness or richness, however, but to interpolate between known market prices or to calculate hedge ratios, practitioners often take the much simpler approach of using some form of the Black-Scholes (BS) option pricing model.

First, the chapter describes the fixed income options listed in the previous paragraph. The practice of applying BS in each context is explained, and the less justifiable practices critiqued in favor of other methodologies. For easy reference, common practitioner applications of BS are collected and summarized in Tables 18.8 through 18.10.

The chapter focuses next on the *skew*, the fact that BS implied volatilities vary with strike. The existence of the skew means that the BS model, even when applied to European-style options of a single maturity, cannot capture the richness of option prices in the market. The section on skew focuses on swaptions and presents two commonly used approaches to model the skew.

The final section of the chapter presents the theory that justifies the application of BS to the fixed income options described in this chapter. Consistent with the rest of this book, every effort has been made to use a minimum amount of advanced mathematics.

CAPS AND FLOORS

The easiest way to explain *caps* is to start by explaining *caplets*, even though caps are the more traded derivative. At the end of a given accrual period, a caplet pays the greater of zero and *London Interbank Offered Rate (LIBOR)* minus a strike over the accrual period, where LIBOR is set at the beginning of the accrual period. Consider, for example, a caplet with a three-month LIBOR reset date of November 28, 2012, a payment date of February 28, 2013, and a strike of .97%. Note that there are 92 days over this accrual

period. Then, if LIBOR on November 28, 2012, turns out to be L , then a unit notional of the caplet will pay, on February 28, 2013,

$$\frac{92}{360} [L - .97\%]^+ \quad (18.1)$$

where $[L - .97\%]^+$ is another way of writing $\max [L - .97\%, 0]$. Note that the payoff of a caplet looks like that of an option, but the maximum of zero and the difference between the rate and the strike is part of the contract rather than a result of optimal exercise behavior.

Caplets are typically valued by practitioners under the assumption that forward LIBOR rates are normally distributed. Under this assumption, the final section of the chapter shows that the value of a caplet with a reset at time T and payment at time $T + \tau$ is given by

$$\tau d_0(T + \tau) \xi^N(S_0, T, K, \sigma) \quad (18.2)$$

where τ is the term of the reference rate, $d_0(T + \tau)$ is the discount factor to the payment date, S_0 is today's forward rate from T to $T + \tau$, K the strike, σ the basis-point volatility of the forward rate, and ξ^N the BS-style formula defined in Appendix A in this chapter. Table 18.1 applies (18.2) to 100 notional of the caplet introduced previously, as of February 28, 2011. Note that there are 639 days from the pricing date, February 28, 2011, to the reset date, November 28, 2012. (Note that calendar days are used here for easier readability, but business days are more commonly used in practice.) The appropriate discount factor to the payment date, derived from the swap curve, is .981801. Finally, a volatility of 77.22 basis points, the source of

TABLE 18.1 Pricing a Caplet with a LIBOR Reset on November 28, 2012, and a Payment Date on February 28, 2013, as of February 28, 2011

Quantity	Value
S_0	1.9077%
T	$\frac{639}{365} = 1.7507$
τ	$\frac{92}{360} = .2556$
K	0.97%
σ	.7722%
$d_0(T + \tau)$	.981801
$\xi^N(S_0, T, K, \sigma)$	.01037
$V_0^{Caplet} = 100 \times \tau d_0(T + \tau) \xi^N$	.2602

TABLE 18.2 The Structure and Pricing of a Two-Year Cap as of February 28, 2011

Cap Strike			.97%
Cap Volatility			.7722%
Dates			
Reset	Payment	Forward Rate	Caplet Premium
2/28/11	5/31/11	.310%	
5/31/11	8/31/11	.377%	.0027
8/30/11	11/30/11	.452%	.0128
11/28/11	2/28/12	.599%	.0300
2/28/12	5/29/12	.830%	.0611
5/29/12	8/29/12	1.176%	.1157
8/28/12	11/28/12	1.555%	.1864
11/28/12	2/28/13	1.908%	.2602
Sum			.6688

which will be made clear below, is used to derive the price of 26.02 cents per 100 notional amount.

Having explained a caplet, the discussion can turn to caps. A cap is a portfolio of caplets, with the value of the cap being the sum of the value of its component caplets. The implied volatility of a cap is the volatility that, when used to value every component caplet, results in the market price of the cap. This leads to some complexities, as will be discussed presently, because the term structure of caplet volatility is not flat. In other words, every caplet is appropriately valued at a particular volatility even though, when quoting the price of a cap, all of its component caplets are valued at the same volatility.

Table 18.2 illustrates the structure of a cap and how a cap price is quoted using a two-year USD at-the-money (ATM) cap as of February 28, 2011. This cap is ATM because its strike of .97% equals the rate of the corresponding swap which, in this case, is the two-year par swap rate. The cap strike of .97% means that every component caplet has a strike of .97%. The cap volatility of 77.22 basis points means that the price of the cap is the sum of the component caplet values when each caplet is valued at a volatility of 77.22 basis points. The reset and payment dates for each caplet are given in the table, though the exact date logic will not be covered here. The forward rates are derived from the swap curve and the caplet premiums are calculated from the normal BS formula, using each respective forward rate, a strike of .97%, and a volatility of 77.22 basis points, in addition to the appropriate date parameters and discount factors along the lines of

Table 18.1. In fact, the caplet in Table 18.2 that pays on February 28, 2013, is the same caplet that was valued in Table 18.1.

Note that what might have been the first caplet in Table 18.2, with a LIBOR reset at the start of the cap initiation and a payment on May 31, 2011, is omitted from the table. The payment from such a caplet would be known as of the start of the cap and, as such, would have no option-like premium—it would simply be worth its present value. In fact, in this example, where the initial LIBOR setting is below the strike, at .310%, the payment from this caplet would be zero. As a consequence of this line of reasoning, this first caplet payment is usually omitted from a cap.

The two-year cap in the example is a spot starting cap, i.e., putting aside the skipping of the first payment, the schedule of payments starts immediately. There is also an active market, however, in forward starting caps. In a 5×5 cap, the first reset would be in five years, the first payment in five years plus the length of the accrual period (e.g., five years and three months), and the last payment would be in 10 years. Finally, to close the discussion of the structure of caps, for USD caps the reference rate is always three-month LIBOR. For EUR caps, however, the reference rate is typically three-month LIBOR for shorter-term caps and six-month LIBOR for longer-term caps.

As mentioned previously, if caplets traded individually, they would be priced at individual volatilities, not at a single volatility as is the convention for quoting cap prices. In other words, there is a term structure of caplet volatilities. This term structure is interesting for use as another perspective on the market price of volatility and for comparison with, and perhaps trading opportunities against, similar volatility instruments, e.g., ED futures options.

In theory, the term structure of caplet volatility could be recovered from caps of various terms, with the volatility of a three-month cap giving the first caplet volatility, then the volatility of the six-month cap giving the second caplet volatility, etc. The problem is complicated, however, by the fact that the most traded and useful volatilities are ATM volatilities, which, in the case of caplets, correspond to caplets with strikes equal to their underlying forward rates. But the strikes of caplets that are part of caps all have a single strike that cannot, in general, equal the underlying forward for every component cap. And, as discussed in the next section, on swaptions, and later in this chapter on skew, volatilities of options that are not ATM can be significantly different from those ATM. Hence, the extraction of caplet volatilities from caps is often combined with some adjustment for the caplet strikes not being ATM.

Floorlets and floors are analogous to caplets and caps. Using the same notation as in the rest of this section, the payment of a floorlet at time $T + \tau$, determined by the LIBOR rate set at time T , is

$$\tau [K - L]^+ \quad (18.3)$$

Assuming normal forward rates, the value of a floorlet is given by

$$\tau d_0(T + \tau) \pi^N(S_0, T, K, \sigma) \quad (18.4)$$

where the function $\pi^N(\cdot)$ is given in Appendix A in this chapter. The value of a floor is the sum of the values of its component floorlets.

Applying put-call parity, the prices of an ATM caplet and an ATM floorlet with the same expiration are equal, as are the prices of matched-date ATM caps and floors. Say, for example, that the five-year swap rate five years forward is 5%. Then paying fixed on this forward starting swap and buying a 5×5 floor with a strike of 5% has exactly the same cash flows as a 5×5 cap with a strike of 5%. But since, by definition, the value of the forward swap is zero, the values of the cap and the floor must be the same.

SWAPTIONS

A *swaption* is an over-the-counter (OTC) contract that gives the buyer the right, at expiration, to enter into a fixed-for-floating interest rate swap at the maturity and strike rate agreed to in the contract. A *receiver swaption* gives the buyer the right to receive fixed and pay floating while a *payer swaption* gives the buyer the right to pay fixed and receive floating.

As an example, consider a \$100 million 5.28% “5-year-5-year” or “5y5y” receiver swaption traded on February 9, 2011. This option gives the buyer the right, in five years, on February 9, 2016, to receive 5.85% and pay LIBOR on \$100 million for five years, that is, until February 9, 2021. What is the value of this swaption at expiration? Following the notation of Chapter 13, let $C_5(5, 10)$ denote the par swap rate from year five to year 10, i.e., the five-year par swap rate, five years from today. Also, let $A_5(5, 10)$ denote the value, five years from today, of an annuity of \$1 per year, paid on the fixed rate payment dates of a swap from year five to year 10. Then, at the expiration of the swaption, in five years, the value of receiving 5.28% for five years is

$$\$100 \text{ mm} \times [5.28\% - C_5(5, 10)]^+ \times A_5(5, 10) \quad (18.5)$$

Inspection of the payoff (18.5) reveals that a 5y5y receiver swaption is a put on the five-year par swap rate, five years forward. More generally, a T -year- τ -year receiver swaption is a T -year put option on the τ -year par swap rate, T -years forward. Similarly, a T -year- τ -year payer swaption is a T -year call option on the τ -year par swap rate, T -years forward.

Table 18.3 applies BS to the example of this section. As just discussed, the rate underlying the 5.28% 5y5y receiver option traded on February 9, 2011, is the forward par rate on a swap from February 9, 2016, to February

TABLE 18.3 Calculating the Receiver Swaption
Price per 100 Notional Amount of Swaps

Quantity	Value
S_0	5.28%
T	5
τ	5
K	5.28%
σ	1.145%
$A_0(T, T + \tau)$	3.8285
$\pi^N(S_0, T, K, \sigma)$	.010213
$V_0^{Receiver} = \$100 \text{ mm} \times A \times \pi^N$	\$3.91 mm

9, 2021. As of the pricing date, this forward rate was 5.28%. Hence S_0 of BS is 5.28% and the swaption of the example, with its strike at 5.28%, was ATM. (A higher strike would have been *in-the-money* while a lower strike would have been *out-of-the-money*.) For the 5.28% 5y5y, the other parameters are clearly $T = 5$, $\tau = 5$, and $K = 5.28\%$. The value of the annuity on the swap from February 9, 2016, to February 9, 2021, as of February 9, 2011, was 3.8285. Finally, the cost of this receiver option on the pricing date was 391 cents per 100 notional amount or \$3.91 million on \$100 million. The final section of this chapter shows that the value of a receiver swaption, per unit notional, when the underlying forward swap rate is normally distributed, is

$$A_0(T, T + \tau) \times \pi^N(S_0, T, K, \sigma) \tag{18.6}$$

where $\pi^N(\cdot)$ is once again from Appendix A in this chapter. Setting the market price of \$3.91 million equal to \$100 million times (18.6), Table 18.3 shows that the implied volatility of this swaption is 1.145%.

The analogous formula for a payer swaption, per unit notional, is

$$A_0(T, T + \tau) \times \xi^N(S_0, T, K, \sigma) \tag{18.7}$$

ATM swaption prices, which are by far the most commonly traded swaptions, are quoted in a matrix of either premia or implied normal volatilities. Table 18.4 is an example of the latter for USD swaptions as of February 9, 2011. For example, the ATM option to receive or to pay in a 10-year swap in two years, a 2y10y, is priced with an implied volatility of 116.4 basis points.¹ The price of the 5y5y option introduced in the previous

¹ATM calls and puts have the same BS values. This is easy to verify from Equations (18.88) through (18.94).

TABLE 18.4 USD ATM Swaption Normal Volatilities as of February 9, 2011

Option Exp.	Swap Tenor											
	1y	2y	3y	4y	5y	6y	7y	8y	9y	10y	15y	30y
1m	56.1	80.1	99.1	107.5	116.0	115.5	114.6	113.9	112.7	112.3	105.6	97.0
3m	53.8	85.5	102.4	110.1	118.5	117.9	117.4	116.2	114.7	113.3	108.7	100.9
6m	69.2	96.2	106.9	113.4	119.9	119.0	118.1	117.2	116.2	115.1	110.3	104.0
1y	97.5	113.4	117.7	119.7	121.5	120.8	119.7	118.7	117.8	116.9	110.6	104.8
2y	121.6	123.3	123.9	122.6	122.0	121.0	119.8	118.7	117.6	116.4	109.2	102.8
3y	122.4	122.6	121.7	121.1	120.8	119.5	118.3	117.0	115.6	114.4	106.8	100.2
4y	120.5	119.9	119.2	118.3	117.6	116.4	115.3	114.2	113.0	112.0	104.0	97.1
5y	119.7	118.9	117.4	116.0	114.4	113.3	112.0	110.8	109.7	108.4	100.0	93.0
10y	99.4	99.1	97.7	96.5	95.4	94.2	92.9	91.7	90.4	89.1	81.4	75.2

paragraph, also recorded in the table, is quoted at a volatility of 114.4 basis points.

Swaption *skew*, discussed later in this chapter, refers to the fact that implied volatilities for ATM swaptions in Table 18.4 are valid only for ATM swaptions. The broader swaptions market, therefore, actually trades a volatility *cube*, where the third dimension represents strike, usually in 50 basis-point increments away from the forward par swap rate corresponding to each entry of the swaption matrix. For a 5y5y as of February 9, 2011, with the underlying par forward swap at 5.28%, a volatility cube would show volatilities for the 5y5y at higher strikes of 5.78%, 6.28%, etc., and for lower strikes of 4.78%, 4.28%, etc.

The skew applies not only for trading swaptions that are not ATM but also for valuing or marking swaptions in position that were initiated ATM. Consider the 5.28% 5y5y receiver swaption traded on February 9, 2011. The underlying swap of this swaption, from initiation to expiration, is an unchanging 5.28% swap from February 9, 2016, to February 9, 2021. As of February 9, 2011, this swap was a par swap, but over time this will no longer be the case. Say, for example, that one month later, on March 9, 2011, the rate of the forward par swap corresponding to those dates is 5.40%. In that case, the 5.28% receiver swaption can be characterized as a 4-year-11-month-5-year that is 12 basis points out-of-the-money. Valuing the option, therefore, requires an interpolation between the ATM and the 50 basis-point, out-of-the-money volatilities, as well as an interpolation between the 4y5y and 5y5y option expirations. In practice, these interpolations are carried out by means of a stochastic volatility model, discussed later in this chapter.

This section has described swaptions as if they are physically settled, meaning that, at expiration, the counterparties enter into a swap at the appropriate rate and maturity. In fact, however, swaptions in the United States are almost always cash settled and, in Europe, approximately half of all swaptions are cash settled. In the United States, the cash settlement feature has no material effect on valuation because the cash settlement value is found by multiplying the appropriate annuity factor, evaluated along the swap curve, by the difference between the par rate and the strike in the case of payers or the difference between the strike and the par rate in the case of receivers. In Europe, however, where the annuity factor is computed at a flat rate equal to the appropriate par swap rate, there can be very minor valuation differences between the two forms of settlement.

As a final note, as of September 2010, some swaptions in Europe are structured so that the premium is paid at the expiration date rather than the trade date. This practice makes swaption valuation less sensitive to the choice of the discounting curve, that is, to the issues raised in Chapter 17.

BOND OPTIONS

Options on bonds, particularly government bonds, do trade over-the-counter, but the market is not particularly active. On the other hand, *embedded options*, i.e., options that are included as part of the agreement between issuer and bondholders, are quite common. To take one example, on August 4, 2010, the Federal National Mortgage Association (FNMA) issued a 2% bond maturing on November 4, 2014. Coupon payments were set for the 4th of November and the 4th of May of each year from November 4, 2010, to November 4, 2014. Note that the first coupon, due only three months after issuance, was for only 1% of the face amount and is known as a *short coupon*.² As part of the bond indenture, however, FNMA had the right to *call* the bonds on February 4, 2011, at a price of 100. This means that FNMA had the right, on February 4, 2011, to buy back every bond it had issued for a price of par and, in so doing, have no further payment obligations.

The structure of the call option in the FNMA bond just described is a bit unusual in the context of the corporate bond market. More usual structures can be described as follows. First, after an initial period of noncallability or *call protection*, the issuer has the right to call a bond on every coupon payment date until maturity. A “10 non-call five,” for example, would be a 10-year bond first callable in five years.³ Second, the call price on the first call date typically includes a premium, frequently equal to half the coupon, so that a 5% coupon bond might be callable on the first call date at a price of 102.50. Third, the schedule of call prices would typically decline linearly from the first call price on the first call date to 100 at maturity.

The embedded call feature was initially designed for issuers to maintain flexibility with respect to managing their debt rather than for them to purchase call options when issuing debt. Issuers might want to retire debt so as to alter their capital structures or so as to eliminate restrictive covenants.⁴ But tracking down all holders of an issue and negotiating the repurchase of their holdings is simply not practical. The solution was the call feature, through which an issuer could instruct the trustee of an issue to pay the call price and extinguish the debt. As time passed, however, and as interest rates became more volatile, the value of the call feature as an interest rate option was recognized and incorporated into the pricing of corporate securities.

The discussion now turns to pricing bond options. The typical structure, which allows the issuer to exercise its call on many coupon payment dates

²The first coupon is a *long coupon* if a bond otherwise making semiannual payments sets the accrual period for the first coupon payment at more than six months.

³An option with several exercise dates, which is between a European- and American-style option, is called a *Bermuda*-style option.

⁴Chapter 19 lists some common bond covenants.

and at a declining schedule of call prices, is handled most precisely through the methods of Part Three. More specifically, in the context of a one-factor model, the valuation procedure for a callable bond can be described as follows:

1. Create an appropriate risk-neutral tree of short-term rates. In some cases the short-term rate process would be designed to price bonds of the issuer selling the callable bond, but in most cases the process would correspond to swap or government bond benchmarks. Either way, but almost certainly in the latter case, a bond-specific spread or option adjusted spread (OAS) would be added to the rates for valuation.
2. Using the methodology of Part Three, calculate the value along the tree of an otherwise identical, noncallable bond. In the case of the FNMA 2s of 2014, the otherwise identical noncallable bond would have exactly the same schedule of coupon and principal payments as the FNMA 2s, but would not be callable.
3. Calculate the value along the tree of the call option embedded in the callable bond. Consistent with the well-known methods for pricing an American-style option along a tree, start from the maturity date of the option and work backwards, using the rule that the value of the option at any node equals the maximum of the value of immediately exercising the option and the value of holding the option for another period. And, of course, the value of holding the option for another period at any node is its expected discounted value.
4. The value of the callable bond equals the value of the noncallable bond minus the value of the option. This reflects the fact that the issuer has effectively sold the noncallable bond and bought a call option on that bond or, equivalently, that bondholders have bought the noncallable bond and sold a call option on that bond.

With this valuation procedure in place, other computations are straightforward. Given a market price for the callable bond, an OAS can be computed along the lines of Chapter 7. Also, the interest rate sensitivity of the bond can be calculated by perturbing the short-term rate factor and repeating steps 1 to 4 or, for a metric more similar to yield-based *DV01*, by perturbing the initial term structure and repeating the valuation procedure.

While term structure models are best suited for pricing callable bonds, there are occasional uses for the BS model. First, there are special cases like the FNMA 2s of 2014, introduced earlier, that are simpler, European-style calls. Second, if the first call date is relatively distant, most of the value of the call feature is the value of calling on that first call date. For example, the embedded call in a 10-non-call-five structure is greater than, but approximately equal to, the option to call the bond in exactly five years, i.e., on the first call date only.

TABLE 18.5 Calculating the Value of the Embedded Option in the FNMA 2s of 2014, as of August 20, 2010

Quantity	Value
S_0	101.9193
T	$\frac{163}{365} = .446575$
K	100
σ	2.813%
$d_0(T)$	$(1 + \frac{.2888\% \times 163}{360})^{-1} = .998694$
$\xi^{LN}(S_0, T, K, \sigma)$	2.074061
$V_0^{BondCall} = d_0(T) \xi^{LN}$	2.0714
LIBOR OAS	−.0644%
Noncallable Price	102.6964
Callable Price	100.625

To illustrate the use of BS for callable bonds, consider the FNMA 2s of 2014, which traded at 100.625 as of August 20, 2014. The final section of this chapter shows that, under the assumption that the forward bond price is lognormally distributed, the value of a call option on a bond is $d_0(T) \xi^{LN}(S_0, T, K, \sigma)$, where S_0 is the forward price of the bond for delivery on date T , the other parameters are as usually defined, and ξ^{LN} is as given in Appendix A in this chapter. But the call option is not on the FNMA 2s of 2014, but on the otherwise identical noncallable bond whose price is not observed. Hence, both the forward price and the volatility are missing from the list of needed parameters. Since FNMA bonds are essentially backed by the U.S. government and have little credit risk, one strategy for pricing is to rely on the swap market for rates and volatility. In particular, model the forward price of the noncallable bond as an OAS to the swap curve and take the volatility of the forward bond as the volatility of a forward swap with the same cash flows. Then use BS to imply the OAS from that volatility and the price of the callable bond.

Table 18.5 shows the results. With the noncallable forward bond trading at an OAS of −6.44 basis points to the swap curve, its value is 101.9193. Then, assigning that forward bond price a volatility of 2.813%, which is taken from the swaptions market,⁵ and calculating a discount factor from a swap market rate plus that same OAS, the BS price of the option is 2.071. Subtracting this option value from the 102.6964 value of the noncallable bond gives a callable price of 100.625, as desired.

⁵The appropriate swaption volatility was 78.43 basis points, the forward DV01 of the noncallable bond .03656, and the forward noncallable price 101.9193. Hence, the volatility of the forward price is $78.43 \times .03656 \div 101.9193$ or 2.813%.




FIGURE 18.1 Price-Rate Curves of the Callable FNMA 2s of 2014, and of Two Reference Bonds

This section concludes by discussing the interest rate behavior of callable bonds. Figure 18.1 graphs the price-rate functions of the FNMA 2s of 2014, as of August 20, 2010, and of two fictional reference bonds in a BS framework with a flat term structure. The first reference bond, labeled “NC to Maturity,” pays exactly the same cash flows as the FNMA bonds but is not callable. These are the cash flows an investor would realize on the FNMA bonds were they not to be called. The second reference bond, labeled “To Be Called,” pays the same coupons as the FNMA to its call date, on February 4, 2011, and then returns principal plus accrued interest. These are the cash flows an investor would realize on the FNMA bonds were they to be called on February 4, 2011.

The NC to Maturity is a longer maturity bond than the To Be Called bond, so its price-rate graph is steeper. For relatively high rates, the option to call the FNMA bond is not likely to be exercised so the value of the FNMA bond approximately equals that of the NC to maturity bond. For relatively low rates, the option to call the FNMA bond is very likely to be exercised so the value of the FNMA bond approximately equals that of the To Be Called bond. For intermediate rates, the value of the FNMA bonds is significantly below that of the two reference bonds. The difference between the value of the NC to maturity bond and the value of the FNMA bond is just the value of the call option.

Figure 18.1 also illustrates the negative convexity of callable bonds. This can be seen directly from the shape of the FNMA price-rate curve. Equivalently, the *DV01* of the FNMA bond at high rates resembles the

relatively high $DV01$ of the NC to Maturity bond while at low rates it resembles the relatively low $DV01$ of the To Be Called bond. Hence the $DV01$ of the FNMA bond declines with rate, i.e., it is negatively convex.

EURODOLLAR AND EURIBOR FUTURES OPTIONS

Chicago Mercantile Exchange *VERSUS* NYSE Euronext Futures Options

The markets for ED futures options and Euribor futures options⁶ are extremely active. Like the underlying futures contracts, these American-style call and put options are exchange traded and are highly standardized with respect to expiration dates, strikes, maturities, and contract size.

Futures options are all, of course, options on futures contracts, but there are two distinct types of futures options.

The first type, like the ED futures options traded on the Chicago Mercantile Exchange (CME) in the United States, can be described as follows. The buyer of an option pays a premium and the seller of an option receives a premium. Upon exercise or expiration of an in-the-money call option, the buyer receives the intrinsic value of the option and a long position in the underlying futures contract at the prevailing market price. But since a futures contract at the prevailing market price has no value, this total payoff equals the intrinsic value. Similarly, upon exercise of an in-the-money put, the buyer receives the intrinsic value of the option and a short position in the underlying futures contract at the prevailing market price.⁷

The second type of futures options, like the Euribor futures options traded on NYSE Euronext, are themselves futures contracts. This means that there is no initial payment or receipt of a premium. Daily settlement payments are made and received as the futures option price changes and the final settlement price equals the intrinsic value of the option. Expressed another way, for this second type of futures option there is no initial exchange of a premium or final exchange of intrinsic value; all cash flows are in the form of daily settlement payments.

The difference between the two types of futures options results in different pricing formulae and different rules for early exercise. Appendix B in

⁶Euribor futures are the EUR equivalent of ED futures, with the final settlement rate depending on Euribor rather than USD LIBOR.

⁷Futures options of this type are “marked-to-market” in the sense that position values are calculated daily for the purposes of computing margin requirements. But there are no daily settlement payments flowing from option buyers to sellers or *vice versa*. This is a good example of the why the terms “daily settlement” and “marked-to-market” are confusing when used interchangeably.

this chapter shows, in fact, that it is optimal to exercise ED futures options early only when they are very much in-the-money and that it is not optimal to exercise Euribor futures options early. This section, therefore, treats both types of options as if they were European-style options.

CME ED Futures Options

Quarterly ED options expire at the same time as their underlying ED futures contract. An example, to be discussed presently, is an option to purchase a June 2011 ED futures contract at its expiration in June 2011. *Serial* options expire a month or two before the expiration of the underlying, e.g., an option expiring in May 2011 to purchase a June 2011 contract. Lastly, *mid-curve* options expire one, two, or five years before the underlying futures contract, e.g., an option expiring in June 2011 to buy a June 2012 contract.

As the quarterly options are the most popular, an example of that kind is discussed here. EDM1C 99.25 is the ticker for a call option on one EDM1 (i.e., the June 2011 contract) with a strike equal to 99.25 and an expiration date the same as that of EDM1, namely, June 13, 2011.

Recall from Chapter 15 that ED futures are scaled to change in value by \$25 per basis point, and that the futures price F is really just a means of quoting the futures rate f , that is,

$$F = 100 - 100f \quad (18.8)$$

Then, if the final settlement price is F , the payoff of EDM1C 99.25 at expiration is

$$\max \left[\frac{F - 99.25}{100}, 0 \right] \times 10,000 \times \$25 \quad (18.9)$$

The factor 10,000 turns the result of the maximum function into basis points. For example, if $F = 99.50$, then the max is .25%, which, multiplied by 10,000, gives 25 basis points. Finally, substituting (18.8) into (18.9), the payoff to EDM1C 99.25 becomes

$$\$250,000 \times \max [.75\% - f, 0] \quad (18.10)$$

Therefore, while EDM1C 99.25 is referred to as a call option on EDM1's price struck at 99.25, it is really a put option on EDM1's rate struck at .75%. Similarly, references to put options on ED futures prices are really call options on ED futures rates.

The justification for applying a BS-type formula in the case of ED futures options is not so satisfying as in the case of the other options discussed in this chapter. This is discussed further in the final section of this chapter.

Other methods of pricing ED futures options can be justified readily, of course, like the application of one of the term structure models of Part Three. Nevertheless, practitioners do seem to apply a BS-type formula to ED futures options as well.

When applying BS to ED options, practitioners assume that the futures rate is normally distributed and use the pricing formula

$$\$250,000 \times d_0(T) \times \pi^N(S_0, T, K, \sigma) \tag{18.11}$$

where $\pi^N(\cdot)$ is as defined in Appendix A in this chapter, S_0 is the underlying futures rate, T is the expiration time of the option, in years, K is the strike rate, σ is the annual basis-point volatility of the futures rate, and $d_0(T)$ is the current discount factor to the expiration of the option.

The analogous formula for an ED put option (which is a call on rates) is

$$\$250,000 \times d_0(T) \times \xi^N(S_0, T, K, \sigma) \tag{18.12}$$

Table 18.6 illustrates the application of BS to EDM1C 99.25 as of February 15, 2011. The price of EDM1 was 99.567, so $S_0 = .433\%$. There are 118 days from February 15 to option expiration on June 13, so $T = \frac{118}{365}$. An appropriate discount rate to expiration was .316%, so the discount factor is .998965, as computed in the table. Finally, the price of EDM1C 99.25 was \$881.25. With all of these quantities set, the volatility that sets (18.11) to the market price of the option, i.e., \$881.25, can be solved to be about 62 basis points.

TABLE 18.6 Calculating the Price of EDM1C 99.25 as of February 15, 2011

Quantity	Value
S_0	0.433%
T	$\frac{118}{365}$
K	0.75%
σ	.6227%
$d_0(T)$	$(1 + \frac{.316\% \times 118}{360})^{-1} = .998965$
$\pi^N(S_0, T, K, \sigma)$	.00352865
$V_0^{ED\ Call} = \$250,000 \times d_0(T) \pi^N$	\$881.25

NYSE Euronext Euribor Futures Options

Unlike the case of CME ED futures options, pricing NYSE Euronext Euribor futures options in a BS framework is easily justified, as demonstrated in the final section of this chapter. The resulting formula for calls and puts on the contracts (which are puts and calls on rates) are

$$€250,000 \times \pi^N(S_0, T, K, \sigma) \quad (18.13)$$

$$€250,000 \times \xi^N(S_0, T, K, \sigma) \quad (18.14)$$

Note that unlike other option pricing formulae in this chapter, expressions (18.13) and (18.14) do not include a discount or annuity factor. The formal derivation is given below, but the intuition is as follows. Since Euribor futures options are futures, there is no up-front payment of a premium and the option value does not have to earn a return from initiation to expiration in the pricing measure. The same reasoning was used to explain why a futures price equals its expected future value rather than its expected discounted future value; see Chapter 13.

BOND FUTURES OPTIONS

Options on bond futures are exchange-traded and highly standardized, like the underlying bond futures (see Chapter 14). To illustrate bond futures options, consider an American-style put option on a Japanese bond futures. The seven-year Japanese government bond (JGB) futures has a size of ¥100,000,000 and, as of February 11, 2011, had a price of 137.88. A 135 put on this contract matures on May 31, 2011, and had a price of ¥200,000.

Options on bond futures can be valued in the frameworks of Part Three and Chapter 14. These methods describe how to create a risk-neutral tree for the futures price. And from this tree, the value of a futures option can be computed: start from the maturity date of the option and work backwards, using the rule that the value of the option at each node is the maximum of the value of holding the option and of exercising it immediately. Of course, because the delivery option depends on the slope of the term structure as well, many futures pricing models use a two-factor rather than a one-factor model.

While simple conceptually, it should be clear from Chapter 14 that building a futures model takes a good deal of effort. Therefore, practitioners who do not otherwise need such a model tend to use BS for bond futures options. Assume for the moment that the bond futures option is a European-style option and that the bond futures contract has no delivery option. Under these assumptions, it is shown below that applying BS

TABLE 18.7 Calculating the Price of JBM1P as of February 11, 2011

Quantity	Value
S_0	137.88
T	$\frac{109}{365} = .29863$
K	135
σ	3.635%
$d_0(T)$	$(1 + \frac{.216\% \times 109}{360})^{-1} = .999346$
$\pi^{LN}(S_0, T, K, \sigma)$	.200131
$V_0^{FutPut} = 100,000,000 \times d_0(T) \times \frac{1}{100} \pi^{LN}$	200,000

requires that the discount factor to the option expiration date be uncorrelated with the futures price of the bond. While not literally true, this is not a bad assumption when applied to the most commonly occurring situation, namely, a short-term option on a relatively long-term bond futures. In this situation the volatility of the short-term discount factor is small relative to the volatility of the bond futures price and the correlation of the short-term discount factor with the long-term bond futures price is indeed relatively small.

Making all of these necessary assumptions, in addition to assuming that the futures price is lognormal, the final section of this chapter shows that a put on a bond futures is approximated by $d_0(T) \times \pi^{LN}(S_0, T, K, \sigma)$, where S_0 is the futures price. From the description of the JGB futures put option previously mentioned, along with an appropriate discount rate of .216%, the appropriate parameters in this application are given in Table 18.7. Note that there are 109 days from February 11, 2011, to the contract's expiration on May 31, 2011. Then using the formula just cited, the implied price volatility is 3.635%.

Return now to the two assumptions made a moment ago to justify the application of BS to bond futures options. First, assuming that these options are European in style is not a serious problem. As shown in Appendix B in this chapter, options on futures are exercised early only when they are very much in-the-money. Hence, ignoring their American feature is usually reasonable in practice. The second assumption, however, to ignore the delivery option, is more serious. The BS framework assumes that the volatility of the futures price is constant. But as the bond most likely to be delivered changes, the DV01 of the futures contract changes and so does its volatility. While less of a problem when the delivery option is significantly out-of-the-money, as it happens to be in the low-rate environment at the time of this writing (see Chapter 14), the objection remains serious. After all, one of the main motivations for using BS in the first place is to obtain accurate deltas.

SUMMARY OF APPLYING BS TO FIXED INCOME OPTIONS

Tables 18.8 through 18.10 summarize the applications of BS in the fixed income context, as used in the previous sections and as derived in the final section of this chapter. As in those sections, ξ^N and π^N assume normality of the underlying, ξ^{LN} and π^{LN} assume lognormality, and all four functions are given in Appendix A in this chapter. The function $d_0(T)$ gives today's discount factor to time T while $A_0(T, T + \tau)$ gives today's annuity factor for an annual payment of one unit of currency, with the appropriate frequency, from time T to $T + \tau$. The distributions listed give the most common practitioner choice. The ED and Euribor futures options prices are scaled to the standardized contract size. Caplets, floorlets, and swaptions prices are per unit notional amount. Bond and bond futures options prices are in the units of the price of the underlying. Finally, note that call options on ED futures are puts on futures rates, so that π^N is appropriate, while puts on ED futures are calls on rates and ξ^N is appropriate.

SWAPTION SKEW

Skew and Smile

If the BS normal model were true for all swaptions of a given expiration and tenor, a single basis-point volatility would price swaptions of all strikes. Equivalently, the implied volatility of swaptions of all strikes would be the same. As Figure 18.2 illustrates, however, this is not the case. As of November 30, 2010, the 2y2y par forward rate was 2.01% and the ATM

TABLE 18.8 Parameters for Applying the BS Formula to Fixed Income Options, I

	CME ED Futures Option	NYSE Euro Next Euribor Futures Option	Caplet/Floorlet
S_0	Futures rate	Futures rate	Fwd rate from T to $T + \tau$
T	Option expiration	Option expiration	Rate observation date
τ	Term of futures rate	Term of futures rate	Term of forward rate
σ		Vol of rate	
Distribution		Rate normal	
Call/Caplet	$\$250,000d_0(T)\pi^N$	$\text{€}250,000\pi^N$	$\tau d_0(T + \tau)\xi^N$
Put/Floorlet	$\$250,000d_0(T)\xi^N$	$\text{€}250,000\xi^N$	$\tau d_0(T + \tau)\pi^N$

TABLE 18.9 Parameters for Applying the BS Formula to Fixed Income Options, II

	Bond Option	Bond Futures Option
S_0	Forward bond price, delivery T	Futures price, delivery T
T	Bond option expiration	Futures option expiration
σ	Vol of price	
Distribution	Forward bond price lognormal	Futures price lognormal
Call Price	$d_0(T) \xi^{LN}$	$d_0(T) \xi^{LN}$
Put Price	$d_0(T) \pi^{LN}$	$d_0(T) \pi^{LN}$

basis-point volatility was 99 basis points. But the solid line in Figure 18.2 shows that implied basis-point volatility was less than 99 basis points for swaptions with strikes below 2.01% and greater than 99 basis points for swaptions with strikes above 2.01%. The phenomenon that basis-point volatility is not constant with strike is known as the *skew*. (The lognormal model curve in the figure will be discussed shortly.) As an additional example, Figure 18.3 presents the analogous picture for the 5y10y as of the same date, with a par forward rate of 4.57% and an ATM basis-point volatility of 106 basis points. In this figure, the extent of the skew varies with strike, with implied volatilities significantly higher than the ATM volatility for high strikes but only somewhat lower than the ATM volatility for low strikes. A pattern like this is sometimes referred to as a *smile*.

Without going into more detail here, the implied volatility of an option can be thought of as the expected gamma-weighted average of instantaneous volatilities over possible paths of the underlying forward rate from its current level to the strike. From this perspective, the market-implied volatilities in Figures 18.2 and 18.3 show how instantaneous volatilities are expected to vary directly with the forward rate.

TABLE 18.10 Parameters for Applying the BS Formula to Fixed Income Options, III

	Swaption	
	Receiver	Payer
S_0	Fwd swap rate, T to $T + \tau$	
T	Swaption expiration	
τ	Tenor of fwd swap	
σ	Vol of rate	
Distribution	Rate normal	
Price	$A_0(T, T + \tau) \pi^N$	$A_0(T, T + \tau) \xi^N$




FIGURE 18.2 2y2y Skew as of November 30, 2010

The lognormal version of BS exhibits the property that instantaneous basis-point volatility is proportional to the forward rate. (See the discussion of lognormal term structure models in Chapter 10.) Therefore, lognormal BS should match market implied volatilities more closely than normal BS. And this is indeed the case. The lognormal model curves in Figures 18.2 and 18.3 are computed in three steps. First, choose the volatility parameter




FIGURE 18.3 5y10y Skew as of November 30, 2010

of the lognormal model so that the ATM volatility matches the market. Second, with that parameter, use the lognormal model to price swaptions with different strikes. Third, use those prices and the normal model to back out a basis-point volatility for each strike. The results in this rate environment show that the lognormal model matches the market well for high strikes, but not so well for low strikes.

The Shifted-Lognormal Model

The interpretation of implied volatility as reflecting the dependence of the instantaneous volatility on the forward rate suggests formulating models to generalize this dependence. The dynamics of the underlying rate in the normal and lognormal BS models respectively can be written as

$$dS_t = \phi(S_t) \sigma \epsilon_t \sqrt{dt} \quad (18.15)$$

$$\epsilon_t \sim N(0, 1) \quad (18.16)$$

where $\phi(S_t) = 1$ for the normal version and $\phi(S_t) = S_t$ for the lognormal version. A popular alternative specification is the *shifted-lognormal* model, which sets

$$\phi(S_t) = a + S_t \quad (18.17)$$

with $a \geq 0$. The BS-style call option pricing formula resulting from (18.17) is

$$C^{SLN}(S_0, T, K, \sigma) = (a + S_0) N(d_1) - (a + K) N(d_2) \quad (18.18)$$

$$d_1^{SLN} = \frac{\ln\left(\frac{a+S_0}{a+K}\right) + \frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}} \quad (18.19)$$

$$d_2^{SLN} = d_1^{SLN} - \sigma\sqrt{T} \quad (18.20)$$

The shifted lognormal model is between the normal and lognormal models in the following sense. When $a = 0$, (18.17) and (18.15) result in an instantaneous basis-point volatility that is proportional to S_t , as in the lognormal model. At the other extreme, as a approaches infinity, the instantaneous volatility approaches a constant, as in the normal model. To see this, define $\bar{\sigma} = a\sigma$ so that $\phi(S_t)\sigma = (1 + \frac{S_t}{a})\bar{\sigma}$. Then, letting a approach infinity while $\bar{\sigma}$ stays constant, which requires that σ approach 0, $\phi(S_t)\sigma$ approaches the constant $\bar{\sigma}$.

It turns out that the shifted-lognormal model does match market prices more accurately than the lognormal model, but not accurately enough for




FIGURE 18.4 5y10y Skew on November 30, 2010, with the Shifted Lognormal Distribution

use across all strikes. Figure 18.4 adds to Figure 18.3 the 5y10y basis-point volatilities from a shifted lognormal model with $a = 1\%$. A value of a that fits the market better, on average, could have been selected. But taking $a = 1\%$, which generates volatilities approximately equal to implied volatilities for high strikes but much too low for low strikes, highlights the need for a different function $\phi(S_t)$ to match the market across strikes. Of course, matching market prices does not ensure the accuracy of the resulting deltas. To that end, as discussed at the end of this section, the specified relationship between forward rates and volatilities has to be adequate.

The SABR Model

The shifted lognormal model is an example of trying to fit the skew through the dependence of instantaneous basis-point volatility on the level of the underlying. Another approach is taken by *stochastic volatility* models, which specify volatility as a factor in its own right. And if volatility is itself volatile, then the underlying can experience relatively large random fluctuations as large realized shocks to the underlying, e.g., ϵ_t in (18.16), can occur when volatility is particularly high. Put another way, volatility of volatility leads to probability distributions with relatively *fat tails*.

The SABR model⁸ is a particularly popular stochastic volatility model in the context of interest rate options. The dynamics of the model are given by the following equations:

$$dS_t = S_t^\beta \sigma_t \epsilon_t \sqrt{dt} \quad (18.21)$$

$$\epsilon_t \sim N(0, 1) \quad (18.22)$$

$$d\sigma_t = \alpha \sigma_t v_t \sqrt{dt} \quad (18.23)$$

$$v_t \sim N(0, 1) \quad (18.24)$$

$$E[\epsilon_t v_t] = \rho \quad (18.25)$$

with $0 < \beta < 1$, $\alpha \geq 0$, and $0 \leq \rho \leq 1$. There are a few things to note about this formulation. First, the SABR model approaches the normal model as α and β approach zero, and is the same as the lognormal model as α approaches zero and β approaches one. Second, the initial value of volatility, the parameter σ_0 , is most naturally used to match ATM swaption volatility. Third, since Figures 18.2 and 18.3 show that implied volatility does not increase as quickly for high strikes as indicated by the lognormal model, it makes sense to have, in terms of (18.15), $\phi(S_t) = S_t^\beta$ with $\beta < 1$. Fourth, (18.25) says that the correlation between changes in the underlying and changes in volatility is ρ .

An approximate solution of option prices resulting from (18.21) through (18.25) is to use the lognormal BS formula with an implied volatility $\hat{\sigma}$ equal to

$$\hat{\sigma} = \frac{\sigma_0}{S_0^{1-\beta}} \left\{ 1 - \frac{(1-\beta-\rho\lambda) \ln\left(\frac{K}{S_0}\right)}{2} + \frac{[(1-\beta)^2 + (2-3\rho^2)\lambda^2] \ln^2\left(\frac{K}{S_0}\right)}{12} \right\}$$

$$\lambda = \frac{\alpha}{\sigma_0} S_0^{1-\beta} \quad (18.26)$$

From the specification of the SABR model in (18.25), the parameters β and ρ control the skew, or the relationship between the underlying and volatility, so that higher values of these parameters lead to more steeply upward-sloping basis-point volatilities as a function of strike. The parameter α controls the smile, or fat tails, so that higher values of α increase the basis-point volatility of low and high strikes relative to intermediate strikes. In

⁸P.S. Hagan, D. Kumar, A.S. Lesniewski, and D.E. Woodward, "Managing Smile Risk," *Wilmott*, September 2002, pp. 84–108. SABR is an acronym for Stochastic Alpha, Beta, Rho.

TABLE 18.11 Fitted Beta Parameter of the SABR Model for USD Swaptions as of November 30, 2010

Option		Underlying Swap Tenor						
Exp.	1y	2y	3y	4y	5y	7y	10y	30y
1y	.925	.945	.967	.972	.956	.945	.847	.129
2y	.798	.841	.867	.883	.882	.849	.763	.180
3y	.836	.842	.840	.827	.806	.760	.671	.202
4y	.823	.801	.783	.760	.737	.690	.608	.230
5y	.786	.744	.713	.685	.655	.619	.560	.234
7y	.639	.604	.572	.539	.507	.469	.413	.205
10y	.512	.466	.425	.386	.349	.320	.277	.173

practice, the SABR model is flexible enough to capture the shape of market-implied volatilities. In fact, given the similar effects of β and ρ , it is often the case that only one of these parameters is necessary for a good fit. In the 5y10y example of this section, fixing $\rho = 0$ there are values of α and β that result in basis-point volatilities almost identical to market levels: $\alpha = 0.35$ and $\beta = 0.56$. This value of α says that a 35% volatility of volatility is required to fit the market, while the value of β indicates that the market relationship between the underlying and volatility is about halfway between that of the normal and lognormal models.

While this discussion has emphasized the flexibility of the SABR model in fitting the skew, all of the discussion has been in the context of the 5y10y. When describing the relationship between basis-point volatility and strike for other expirations and tenors, different parameters of the SABR model are required. In other words, there is no BS-style model that describes the skew across the entire volatility cube. To emphasize this point, Tables 18.11 and 18.12 give the α and β parameters of the SABR model, with ρ fixed at zero, that describe the USD swaption skew as of November 30,

TABLE 18.12 Fitted Alpha Parameter of the SABR Model for USD Swaptions as of November 30, 2010

Option		Underlying Swap Tenor						
Exp.	1y	2y	3y	4y	5y	7y	10y	30y
1y	4%	16%	23%	27%	29%	33%	37%	38%
2y	21%	28%	31%	32%	34%	35%	37%	35%
3y	30%	33%	34%	35%	35%	36%	36%	35%
4y	34%	35%	35%	35%	35%	36%	36%	35%
5y	34%	35%	35%	35%	35%	35%	35%	34%
7y	34%	34%	33%	33%	33%	33%	32%	32%
10y	31%	31%	31%	30%	30%	29%	29%	30%

2010. Shorter expiries and tenors trade at higher values of β , i.e., closer to the lognormal model, while longer expiries and tenors trade at lower values of β , i.e., closer to the normal model. The fitted values of α , by contrast, are relatively stable across expiries and strikes, although there are some notable exceptions at short expiries and tenors.

As a final note on the SABR model, the fact that the model is flexible enough to fit the skew for a given expiry and tenor with $\rho = 0$ does not mean that this correlation parameter should be set to zero. After all, the point of the endeavor is to obtain accurate deltas, not to fit market prices. Hence, in addition to fitting market prices, it might very well be worthwhile to optimize simultaneously the choice of all three parameters, α , β , and ρ . One popular choice is to constrain the parameters to match the empirical *backbone*, i.e., the empirical relationship between ATM basis-point volatility and forward rates.

Deltas from Different Skew Models

As mentioned at the start of this chapter, BS-style models are often used by taking market prices as given and computing deltas. To understand the significance of this use in the swaption context, consider calculating the delta of a swaption using normal BS. The model assumes that as the underlying forward rate changes, the basis-point volatility stays the same. But, as evident from Figures 18.2 and 18.3, as the underlying rate changes, the swaption moves further away from or closer to ATM and volatility rises or falls. Furthermore, ATM volatility itself rises or falls with rates. The resulting change in the option price due to these changes in basis-point volatility, however, is not picked up by the normal BS delta. Hence, accurate deltas require a model that accurately captures the change in basis-point volatility as the underlying changes.

Figure 18.5 shows deltas of 5y10y payer swaptions as of November 30, 2011, as a function of strike for the normal BS, lognormal BS, and SABR models, with the latter calibrated as in the previous section. (A delta of .5 here means that \$100 million notional of a swaption would be hedged with \$50 million notional of the 5y10y ATM forward swap.) While all the models are calibrated to have the same ATM basis-point volatility, the deltas from the three models can differ significantly, even for the ATM payers.

THEORETICAL FOUNDATIONS FOR APPLYING BLACK-SCHOLES TO SELECTED FIXED INCOME OPTIONS

The purpose of this section is to justify the application of BS to the European-style options earlier in this chapter. The outline of the argument is as follows:




FIGURE 18.5 Deltas of 5y10y Payer Swaptions as of November 30, 2011, Across Three Models

1. The next subsection shows that, given the functional form of a probability distribution (e.g., normal, lognormal),⁹ there exist parameters (e.g., the mean) of that distribution such that V_0 , the arbitrage-free price of any asset today, is given by

$$\frac{V_0}{N_0} = E_0 \left[\frac{V_t}{N_t} \right] \quad (18.27)$$

where N_t is the price at time t of an asset chosen as the *numeraire*, V_t is value at time t of an asset being priced today, including reinvested cash flows, and $E_t[\cdot]$ gives expectations as of time t under the appropriately parameterized probability distribution. Equation 18.27 is known as the *martingale property* of asset prices.

2. Say that the rate or security price underlying an option at time t is S_t . It follows from 1. that the value of a call option with strike K and time to expiry T is

$$V_0^{Call} = N_0 E_0 \left[\frac{(S_T - K)^+}{N_T} \right] \quad (18.28)$$

⁹Probability distributions actually have to satisfy certain technical conditions for this statement to apply. As the normal and lognormal distributions satisfy these conditions, however, this point receives no further attention here.

while the value of a put is

$$V_0^{Put} = N_0 E_0 \left[\frac{(K - S_T)^+}{N_T} \right] \quad (18.29)$$

3. The subsection on choosing the numeraire and BS pricing shows that, in most of the contexts covered in this chapter, it is possible to choose the numeraire such that

$$S_0 = E_0 [S_T] \quad (18.30)$$

and such that the valuation equations (18.28) and (18.29) can be written, respectively, as

$$V_0^{Call} = h_0 E_0 [(S_T - K)^+] \quad (18.31)$$

$$V_0^{Put} = h_0 E_0 [(K - S_T)^+] \quad (18.32)$$

for some h_0 that is known as of time 0.

4. If it is further assumed that S_t has a normal distribution with volatility parameter σ , then, as described in Appendix A in this chapter, equations (18.30) through (18.32) become the normal BS-style formulae

$$V_0^{Call} = h_0 \xi^N(S_0, T, K, \sigma) \quad (18.33)$$

$$V_0^{Put} = h_0 \pi^N(S_0, T, K, \sigma) \quad (18.34)$$

for the functions $\xi^N(\cdot)$ and $\pi^N(\cdot)$ defined in Appendix A in this chapter. If, on the other hand, it is assumed that S_t has a lognormal distribution with volatility parameter σ , equations (18.30) through (18.32) become the lognormal BS-style formulae

$$V_0^{Call} = h_0 \xi^{LN}(S_0, T, K, \sigma) \quad (18.35)$$

and

$$V_0^{Put} = h_0 \pi^{LN}(S_0, T, K, \sigma) \quad (18.36)$$

where, again, the functions $\xi^{LN}(\cdot)$ and $\pi^{LN}(\cdot)$ are as defined in Appendix A.

Numeraires, Pricing Measures, and Martingales

Definitions and Statement of Result To begin, define the following two concepts:

- *Gains process.* The gains process of an asset at any time equals the value of that asset at that time plus the value of all its cash flows reinvested to that time. For this purpose, all cash flows are reinvested and then rolled at prevailing short-term interest rates.
- *Numeraire asset.* Given a particular numeraire asset, the gains process of any other asset can be expressed in terms of the numeraire asset by dividing that gains process by the gains process of the numeraire asset. The gains process of a security in terms of the numeraire asset is called the *normalized gains process* of that asset.

For concreteness, Table 18.13 gives an example of these concepts. The asset under consideration is a long-term, 4% coupon bond with a face amount of 100. The numeraire is a two-year zero coupon bond with a unit face amount. The gains process is observed today, after one year, and after two years.

The realization of the short-term rate, which in this example is the one-year rate, is given in row (i). The realization of the bond price over time is given in row (ii). A 4% coupon on 100 is paid on dates 1 and 2 and shown in rows (iii) and (iv). The payment on date 1 is reinvested for one year at the short-term rate on date 1, i.e., 2%. The gains process of the bond given in row (v) is the sum of its price and reinvested cash flows, i.e., the sums of rows (ii) through (iv). The price realization of the two-year zero coupon bond, which is the chosen numeraire, is given in row (vi). Finally, the normalized bond gains process, given in row (vii), is the bond gains process divided by the price of the numeraire, i.e., row (v) divided by row (vi).

TABLE 18.13 Example of the Calculation of a Normalized Gains Process

		End-of-Year Realizations		
		0	1	2
(i)	Short-Term/1-Year Rate	1%	2%	1.5%
(ii)	Bond Price	100	95	97.50
(iii)	Date 1 Reinvested Coupon		4	4 (1.02) = 4.08
(iv)	Date 2 Reinvested Coupon			4
(v)	Gains Process	100	99	105.58
(vi)	Price of 2-Year Zero/Numeraire	0.9612	0.9804	1.0
(vii)	Normalized Bond Gains Process	104.04	100.98	105.58

These definitions allow for the statement of the main result of this subsection: in the absence of arbitrage opportunities, there exists a parameterization of a given probability distribution, or a *pricing measure*, such that the normalized gains of any asset today equals the expected value of that asset's normalized gains in the future. Technically, there exist probabilities such that the normalized gains process is a *martingale*. As the goal here is intuition rather than mathematical generality, this result will be proven in the context of a single-period, binomial process.

Setup and Arbitrage Pricing The starting point is state 0 of date 0, after which the economy moves to either state 0 or state 1 of date 1. Three assets will be considered, A, B, and C, with current prices A_0 , B_0 , and C_0 , and date 1, state i prices of A_1^i , B_1^i , and C_1^i .¹⁰ Without loss of generality here, the date 1 prices include any cash flows of the securities on date 1.

In this framework, any asset can be priced by arbitrage relative to the other two assets. The method is just as in Chapter 7. To price asset C by arbitrage, construct its replicating portfolio, in particular, a portfolio with α of asset A and β of asset B such that

$$C_1^0 = \alpha A_1^0 + \beta B_1^0 \quad (18.37)$$

$$C_1^1 = \alpha A_1^1 + \beta B_1^1 \quad (18.38)$$

Then, to rule out risk-free arbitrage opportunities, it must be the case that

$$C_0 = \alpha A_0 + \beta B_0 \quad (18.39)$$

Now let asset A be the numeraire and rewrite equations (18.37) through (18.39) in terms of the normalized gains processes of assets B and C. To do this, simply divide each of the equations by the corresponding value of the numeraire asset A, i.e., divide (18.37) by A_1^0 , (18.38) by A_1^1 , and (18.39) by A_0 . Furthermore, denote the normalized gains processes of the assets by $\bar{B}$ and $\bar{C}$. Then, equations (18.37) through (18.39) become

$$\bar{C}_1^0 = \alpha + \beta \bar{B}_1^0 \quad (18.40)$$

$$\bar{C}_1^1 = \alpha + \beta \bar{B}_1^1 \quad (18.41)$$

$$\bar{C}_0 = \alpha + \beta \bar{B}_0 \quad (18.42)$$

¹⁰ The date 1 prices must be such that at least one of the assets is risky. For asset B, for example, this would mean that $B_1^0 \neq B_1^1$.

Furthermore, solving (18.40) and (18.41) for α and β ,

$$\alpha = \frac{\bar{B}_1^1 \bar{C}_1^0 - \bar{B}_1^0 \bar{C}_1^1}{\bar{B}_1^1 - \bar{B}_1^0} \quad (18.43)$$

$$\beta = \frac{\bar{C}_1^1 - \bar{C}_1^0}{\bar{B}_1^1 - \bar{B}_1^0} \quad (18.44)$$

Pricing Measure In the framework just described, it will now be shown that there exists a pricing measure such that the expected normalized gains process of each security is a martingale. More specifically, there is a probability p of moving to state 1 of date 1 (and $1 - p$ of moving to state 0 of date 1) such that the expected value of the normalized gain of each security on date 1 equals its normalized gain on date 0. Mathematically, it has to be shown that there is a p such that

$$\bar{C}_0 = p \bar{C}_1^1 + (1 - p) \bar{C}_1^0 \quad (18.45)$$

$$\bar{B}_0 = p \bar{B}_1^1 + (1 - p) \bar{B}_1^0 \quad (18.46)$$

Solving (18.46) for p gives

$$p = \frac{\bar{B}_0 - \bar{B}_1^0}{\bar{B}_1^1 - \bar{B}_1^0} \quad (18.47)$$

But this value of p also solves (18.45). To see this, start by substituting p from (18.47) into the right-hand side of (18.45):

$$p \bar{C}_1^1 + (1 - p) \bar{C}_1^0 = \frac{\bar{B}_0 - \bar{B}_1^0}{\bar{B}_1^1 - \bar{B}_1^0} \bar{C}_1^1 - \frac{\bar{B}_0 - \bar{B}_1^1}{\bar{B}_1^1 - \bar{B}_1^0} \bar{C}_1^0 \quad (18.48)$$

$$= \bar{B}_0 \frac{\bar{C}_1^1 - \bar{C}_1^0}{\bar{B}_1^1 - \bar{B}_1^0} + \frac{\bar{B}_1^1 \bar{C}_1^0 - \bar{B}_1^0 \bar{C}_1^1}{\bar{B}_1^1 - \bar{B}_1^0} \quad (18.49)$$

$$= \bar{B}_0 \beta + \alpha \quad (18.50)$$

$$= \bar{C}_0 \quad (18.51)$$

Equation (18.49) just rearranges the terms of (18.48); combining (18.49) with (18.43) and (18.44) gives (18.50); and (18.50) with (18.42) gives (18.51). Hence, as was to be shown, there is a pricing measure, in this

case the probability p , such that the normalized gains processes of B and C are martingales. And, of course, since nothing distinguishes A from the other assets, a probability with the same properties could have been found had B or C been chosen as the numeraire instead.

Summary While for pedagogical purposes this subsection proved its result only in the binomial context, this section will use the result in a more general form, namely, that arbitrage-free pricing implies that a probability distribution can be parameterized such that equation (18.27) holds. In words, arbitrage-free pricing implies that a given probability distribution can be parameterized such that the normalized gains process of any asset is a martingale.

Choosing the Numeraire and BS Pricing

As mentioned in the introduction to this section, in most of the contexts of this chapter it is possible to choose a numeraire such that the underlying is a martingale and such that the value of a call is given by $h_0 \xi^N(S_0, T, K, \sigma)$ or $h_0 \xi^{LN}(S_0, T, K, \sigma)$, in the normal or lognormal cases, respectively, and the value of a put by $h_0 \pi^N(S_0, T, K, \sigma)$ or $h_0 \pi^{LN}(S_0, T, K, \sigma)$ in the normal or lognormal cases. In the cases for which this is possible, this subsection gives the appropriate definition of the underlying, the appropriate numeraire, and the resulting quantity h_0 . The functions ξ^N , ξ^{LN} , π^N , and π^{LN} are all given in Appendix A in this chapter.

Caplets Caplets that mature at time T are written on a forward rate from time T to $T + \tau$, whose value, at time t , is denoted by $f_t(T, T + \tau)$. It will first be shown that taking a $T + \tau$ -year zero coupon bond as the numeraire makes this forward rate a martingale. Let $d_t(T)$ be the time- t price of a zero coupon bond maturing at time T . By the definition of a forward rate of term τ ,

$$\begin{aligned} f_t(T, T + \tau) &= \frac{1}{\tau} \left(\frac{d_t(T)}{d_t(T + \tau)} - 1 \right) \\ &= \frac{1}{\tau} \left(\frac{d_t(T) - d_t(T + \tau)}{d_t(T + \tau)} \right) \end{aligned} \quad (18.52)$$

Next, consider a portfolio that is long a T -year zero and short a $T + \tau$ -year zero. Taking the $T + \tau$ -year zero as the numeraire, the normalized

gains process of this portfolio is a martingale by the results of the previous subsection. Mathematically,

$$\begin{aligned} \frac{1}{\tau} \left(\frac{d_t(T) - d_t(T + \tau)}{d_t(T + \tau)} \right) &= \frac{1}{\tau} E_t \left[\frac{d_T(T) - d_T(T + \tau)}{d_T(T + \tau)} \right] \\ &= \frac{1}{\tau} E_t \left[\frac{d_T(T)}{d_T(T + \tau)} - 1 \right] \\ &= E_t [f_T(T, T + \tau)] \end{aligned} \quad (18.53)$$

where the last line of (18.53) just uses the definition of the forward rate. Combining (18.52) and (18.53) shows that the forward rate is a martingale under the chosen numeraire:

$$f_t(T, T + \tau) = E_t [f_T(T, T + \tau)] \quad (18.54)$$

Turning to the valuation of the caplet, its normalized gains process is a martingale as well. Hence, taking expectations of its normalized gain as of $T + \tau$,

$$\frac{V_0^{Caplet}}{d_0(T + \tau)} = E_0 \left[\frac{\tau (f_T(T, T + \tau) - K)^+}{d_{T+\tau}(T + \tau)} \right] \quad (18.55)$$

$$= E_0 [\tau (f_T(T, T + \tau) - K)^+] \quad (18.56)$$

Finally, assuming that the forward rate $f_T(T, T + \tau)$ is normal with variance $\sigma^2 T$, and knowing from (18.54) with $t = 0$ that its mean is $f_0(T, T + \tau)$, the results of Appendix A apply and

$$V_0^{Caplet} = d_0(T + \tau) \tau \xi^N(f_0(T, T + \tau), T, K, \sigma) \quad (18.57)$$

Swaptions The underlying of a T -year into τ -year swaption is the forward par swap rate from T to $T + \tau$, which, at time t , is denoted by $C_t(T, T + \tau)$. It will first be shown that taking an annuity from T to $T + \tau$ as the numeraire makes this forward par swap rate a martingale. Denote the price of this annuity by $A_t(T, T + \tau)$.

Consider receiving the fixed rate K on a swap from T to $T + \tau$. Its value at time t is

$$[K - C_t(T, T + \tau)] A_t(T, T + \tau) \quad (18.58)$$

Applying the martingale property with this annuity as numeraire,

$$\frac{[K - C_t(T, T + \tau)] A_t(T, T + \tau)}{A_t(T, T + \tau)} = E_t \left[\frac{[K - C_T(T, T + \tau)] A_T(T, T + \tau)}{A_T(T, T + \tau)} \right]$$

$$C_t(T, T + \tau) = E_t [C_T(T, T + \tau)] \quad (18.59)$$

Hence, as claimed, the forward par swap rate is a martingale under this numeraire.

To price a receiver swaption, note that the payoff is $[K - C_T(T, T + \tau)]^+$ times $A_T(T, T + \tau)$. Therefore, its value can be calculated as the expectation of its normalized payoff using the same numeraire:

$$\begin{aligned} \frac{V_0^{Receiver}}{A_0(T, T + \tau)} &= E_0 \left[\frac{(K - C_T(T, T + \tau))^+ A_T(T, T + \tau)}{A_T(T, T + \tau)} \right] \\ &= E_0 [(K - C_T(T, T + \tau))^+] \\ V_0^{Receiver} &= A_0(T, T + \tau) \pi^N(C_0(T, T + \tau), T, K, \sigma) \end{aligned} \quad (18.60)$$

The last line of (18.60) follows from (18.59), the assumption that the forward par swap rate is normal with variance $\sigma^2 T$, and from the appropriate result from Appendix A in this chapter.

Similarly, a payer option under the assumption of normality has the value

$$V_0^{Payer} = A_0(T, T + \tau) \xi^N(C_0(T, T + \tau), T, K, \sigma) \quad (18.61)$$

Bond Options Start with a European-style option, expiring on date T , written on a longer-term bond. The underlying of this option is a forward position in the bond for delivery on date T . It will first be shown that taking the zero coupon bond maturing at time T to be the numeraire makes this forward bond price a martingale. Proving this is a bit more complex than the analogous results derived previously because the gains process of a bond includes reinvested coupons. Therefore, to keep the presentation simple, the martingale result will be derived in a three-date, two-period setting. The current date is date 0 and the expiration or forward delivery date is date 2. The bond is assumed to pay a coupon c on each of dates 1 and 2 and its price at time t is denoted B_t . The numeraire is the zero coupon bond maturing on date 2 with a price, on date t , of $d_t(2)$. Lastly, let r_1 denote the current one-period rate; r_2 the one-period rate, realized one period from now; and f the current one-period rate, one-period forward.

Under these assumptions, and the expression of the zero coupon bond price on various dates in terms of the prevailing one-period rates, the gains process of the bond on the three dates is given by the following expressions:

- Date 0: $\frac{B_0}{d_0(2)} = B_0 (1 + r_1) (1 + f)$;
- Date 1: $\frac{B_1 + c}{d_1(2)} = (B_1 + c) (1 + r_2)$;
- Date 2: $\frac{B_2 + c(1 + r_2) + c}{d_2(2)} = B_2 + c (1 + r_2) + c$

Therefore, the martingale property for the bond says that

$$\begin{aligned} \frac{B_0}{d_0(2)} &= E_0 \left[\frac{B_2 + c (1 + r_2) + c}{d_2(2)} \right] \\ B_0 (1 + r_1) (1 + f) &= E_0 [B_2 + c (1 + r_2) + c] \end{aligned} \quad (18.62)$$

The term $c (1 + r_2)$ in the expectation on the right-hand side of (18.62) requires some attention since r_2 is not known as of date 0. The date-0 value of a payment of $c (1 + r_2)$ on date 2 is, however, by the definition of forward rates,

$$\frac{c (1 + f)}{(1 + r_1) (1 + f)} = \frac{c}{1 + r_1} \quad (18.63)$$

So, applying the martingale property under the numeraire to a payment of $c (1 + r_2)$ on date 2 requires that

$$\begin{aligned} \frac{\frac{c}{1 + r_1}}{d_0(2)} &= E_0 \left[\frac{c (1 + r_2)}{d_2(2)} \right] \\ c (1 + f) &= E_0 [c (1 + r_2)] \end{aligned} \quad (18.64)$$

With this result, the discussion can return to the martingale property of the bond in (18.62). Substituting (18.64) into (18.62),

$$\begin{aligned} B_0 (1 + r_1) (1 + f) - c (1 + f) - c &= E_0 [B_2] \\ \left[B_0 - \frac{c}{1 + r_1} - \frac{c}{(1 + r_1) (1 + f)} \right] (1 + r_1) (1 + f) &= E_0 [B_2] \\ B_0(2) &= E_0 [B_2] \end{aligned} \quad (18.65)$$

The left-hand side of the second line of (18.65) is the date 0 forward price of the bond for delivery on date 2. (See Chapter 13.) The third line, then, simply denotes this forward price by $B_0(2)$. Hence, taking the zero coupon

bond of maturity T as a numeraire, the forward price of a bond for delivery on date T is a martingale.

Turning now to the price of an option on the bond, consider a call with payoff $(B_T - K)^+$. Applying the martingale property to the option price and assuming that the forward bond price is lognormal with volatility parameter σ , the call option is priced as follows:

$$\frac{V_0^{BondCall}}{d_0(T)} = E_0 \left[\frac{(B_T - K)^+}{d_T(T)} \right] \quad (18.66)$$

$$= E_0 [(B_T - K)^+] \quad (18.67)$$

$$V_0^{BondCall} = d_0(T) \xi^{LN}(B_0(T), T, K, \sigma) \quad (18.68)$$

An analogous argument for a put shows that

$$V_0^{BondPut} = d_0(T) \pi^{LN}(B_0(T), T, K, \sigma) \quad (18.69)$$

Eurodollar Futures Options As mentioned in the text, the justification for using a BS-type model for ED futures options is not straightforward. To see the problem, focus on quarterly ED options. For these, the futures and option expire on the same date so that, at expiration, the futures rate equals the forward rate. Then proceed along the lines of a caplet to reach equation (18.55), which is reproduced here with V_0^{EDPut} substituted for V_0^{Caplet} .

$$\frac{V_0^{EDPut}}{d_0(T + \tau)} = E_0 \left[\frac{\tau (f_T(T, T + \tau) - K)^+}{d_T(T + \tau)} \right] \quad (18.70)$$

In the case of the caplet, the cash flow occurs at time $T + \tau$ so the denominator inside the expectation equals one and the equation simplifies to (18.56). In the case of the ED futures option, however, the cash flow occurs at time T and no such simplification is possible. One approach around this is to rewrite (18.70) as follows:

$$\begin{aligned} \frac{V_0^{EDPut}}{d_0(T + \tau)} &= E_0 \left[\frac{\tau (f_T(T, T + \tau) - K)^+}{d_T(T + \tau)} \right] \\ &= E_0 \left[\frac{\tau (f_T(T, T + \tau) - K)^+}{d_T(T)} \frac{d_T(T)}{d_T(T + \tau)} \right] \\ &= E_0 [\tau (f_T(T, T + \tau) - K)^+ \{1 + f_T(T, T + \tau)\}] \end{aligned} \quad (18.71)$$

The third line of (18.71) follows from the fact that, as of time T , $d_T(T) = 1$, and from the definition of the forward rate $f_T(T, T + \tau)$.

Now, making the heroic assumption that the two terms inside the expectation are uncorrelated—even though they both include the same forward rate—equation (18.71) can be further rewritten as follows:¹¹

$$\begin{aligned}
 \frac{V_0^{EDPut}}{d_0(T+\tau)} &= E_0 [\tau (f_T(T, T+\tau) - K)^+] E_0 [1 + f_T(T, T+\tau)] \\
 &= E_0 [\tau (f_T(T, T+\tau) - K)^+] [1 + f_0(T, T+\tau)] \\
 V_0^{EDPut} &= d_0(T) E_0 [\tau (f_T(T, T+\tau) - K)^+] \\
 &= \tau d_0(T) \xi^N(f_0(T, T+\tau), T, K, \sigma)
 \end{aligned} \tag{18.72}$$

The second line of (18.72) follows from the fact that, with the $T + \tau$ -year zero coupon bond as the numeraire, the forward rate is a martingale. The third line follows from the relationship between forward rates and discount factors and the fourth line from the definition of ξ^N in Appendix A in this chapter together with the assumption made in the caplet case that the forward rate is normally distributed. As mentioned earlier in this chapter, the bottom line of (18.72) is used in practice despite the lack of a straightforward theoretical justification.

Euribor Futures Options Following the earlier discussion in this chapter leading to equation (18.10), the terminal payoff of a Euribor futures call option with strike K and expiring at time T , per unit notional, is

$$[K - f_T(T, T+\tau)]^+$$

Given the daily settlement feature of Euribor futures options, the numeraire of choice is the money market account, the value of one unit of currency invested and then rolled every period, at the prevailing short-term rate. Denoting the money market account by $M(t)$ and the short-term rate from time $t - 1$ to t by r_t ,

$$M(0) = 1 \tag{18.73}$$

$$M(T) = (1 + r_1)(1 + r_2) \cdots (1 + r_T) \tag{18.74}$$

¹¹ Another commonly made assumption to obtain the same result is that the discount factor in the denominator of the first line of (18.71) is nonstochastic, i.e., a constant. This assumption, that the forward rate determining the ED option value is a random variable while the relatively short-term discount factor to the option expiration date is constant, can also be classified as heroic.

The first point to make about the money market account is that it is the numeraire of the risk-neutral short-term rate process described and used in Part Three. To see this, apply the martingale property with the numeraire to an arbitrary gains process V_t at time t :

$$\begin{aligned}\frac{V_0}{M(0)} &= E_0 \left[\frac{V_T}{M(T)} \right] \\ V_0 &= E_0 \left[\frac{V_T}{(1+r_1)(1+r_2)\cdots(1+r_T)} \right]\end{aligned}\quad (18.75)$$

But the second line of (18.75) is just the condition derived in Part Three that the value of a claim today equals its expected discounted value. (Part Three valued interim cash flows by expected discounted value, but they could just as easily have been invested at the short-term rate to a terminal date and then discounted back to the initial date, which would more readily resemble using (18.75) and the gains process.)

The second point to make about the money market as numeraire is that futures prices are martingales under this numeraire. This was demonstrated in Chapter 13 in terms of the language of Part Three, but is shown again in Appendix C in this chapter in a manner consistent with the language of this chapter, namely, gains processes and numeraires.

Turning now to Euribor futures options, because they are subject to daily settlement and are futures contracts, their prices are also martingales with the money market account as numeraire. Furthermore, if F_t is the underlying futures price at time t then, at the expiration of a put option (call on rates) at time T , the option is worth $(F_T - K)^+$. Putting together the martingale property of the futures, (18.76), the martingale property of futures options, (18.77), and the final settlement price of the futures options, (18.78), results in the price of the Euribor futures put option at time t , denoted V_t^{EBPut} :

$$F_0 = E[F_T] \quad (18.76)$$

$$V_0^{EBPut} = E_0[V_T^{EBPut}] \quad (18.77)$$

$$= E_0[(F_T - K)^+] \quad (18.78)$$

Assuming now that F_T is normally distributed, applying Appendix A in this chapter to equations (18.76) and (18.78) shows that

$$V_0^{EBPut} = \xi^N(F_0, T, K, \sigma) \quad (18.79)$$

Similarly, for the Euribor futures call option (put on rates),

$$V_t^{EBCall} = \pi^N(F_0, T, K, \sigma) \quad (18.80)$$

Bond Futures Options As discussed in Chapter 13 and previously in this chapter, in the context of Euribor futures options, futures prices are a martingale in the risk-neutral measure, i.e., when the numeraire is the money market account, $M(t)$. Hence, with F_t the underlying bond futures price at time t ,

$$F_0 = E[F_T] \quad (18.81)$$

By the martingale property, the price of a put option on the futures is

$$\frac{V_0^{FutPut}}{M(0)} = E_0 \left[\frac{(K - F_T)^+}{M(T)} \right] \quad (18.82)$$

Then, by the definition of the money market account,

$$V_0^{FutPut} = E_0 \left[\frac{(K - F_T)^+}{(1 + r_1)(1 + r_2) \cdots (1 + r_T)} \right] \quad (18.83)$$

To continue, it has to be assumed that the discount factor is uncorrelated with the futures price. This assumption was discussed and defended earlier in this chapter. With this assumption equation (18.83) becomes

$$V_0^{FutPut} = E_0 \left[\frac{1}{(1 + r_1)(1 + r_2) \cdots (1 + r_T)} \right] E_0 [(K - F_T)^+] \quad (18.84)$$

$$= d_0(T) E_0 [(K - F_T)^+] \quad (18.85)$$

where (18.85) follows from the risk-neutral pricing of a zero coupon bond. (See Chapter 13.)

Finally, applying Appendix A to (18.81), (18.85) with the assumption that the bond futures price has a lognormal distribution,

$$V_0^{FutPut} = d_0(T) \pi^{LN}(F_0, T, K, \sigma) \quad (18.86)$$

For calls, the analogous result is

$$V_0^{FutCall} = d_0(T) \xi^{LN}(F_0, T, K, \sigma) \quad (18.87)$$

APPENDIX A: EXPECTATIONS FOR BLACK-SCHOLES-STYLE OPTION PRICING

As the results in this appendix are part of the option pricing literature, they are presented here for easy reference but without proof. Let $E^N[\cdot]$

and $E^{LN}[\cdot]$ denote the expectations operators under the normal and log-normal distributions, respectively, and let $N(\cdot)$ denote the standard normal cumulative distribution.

If S_T is normally distributed with mean S_0 and variance $\sigma^2 T$, then

$$\begin{aligned}\xi^N(S_0, T, K, \sigma) &\equiv E_0^N[(S_T - K)^+] \\ &= (S_0 - K) N(d) + \frac{\sigma\sqrt{T}}{\sqrt{2\pi}} e^{-\frac{1}{2}d^2}\end{aligned}\quad (18.88)$$

$$\begin{aligned}\pi^N(S_0, T, K, \sigma) &\equiv E_0^N[(K - S_T)^+] \\ &= (K - S_0) N(-d) + \frac{\sigma\sqrt{T}}{\sqrt{2\pi}} e^{-\frac{1}{2}d^2}\end{aligned}\quad (18.89)$$

$$d = \frac{S_0 - K}{\sigma\sqrt{T}} \quad (18.90)$$

If S_T is lognormally distributed with mean S_0 and variance $S_0^2(e^{\sigma^2 T} - 1)$, then

$$\begin{aligned}\xi^{LN}(S_0, T, K, \sigma) &\equiv E_0^{LN}[(S_T - K)^+] \\ &= S_0 N(d_1) - K N(d_2)\end{aligned}\quad (18.91)$$

$$\begin{aligned}\pi^{LN}(S_0, T, K, \sigma) &\equiv E_0^{LN}[(K - S_T)^+] \\ &= K N(-d_2) - S_0 N(-d_1)\end{aligned}\quad (18.92)$$

$$d_1 = \frac{\ln\left(\frac{S_0}{K}\right) + \frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}} \quad (18.93)$$

$$d_2 = d_1 - \sigma\sqrt{T} \quad (18.94)$$

APPENDIX B: EARLY EXERCISE OF AMERICAN-STYLE FUTURES OPTIONS

Early Exercise of CME-type Futures Options

It is optimal to exercise CME-type American-style futures options early only when they are very much in-the-money or, equivalently, when their time value is very low. The rule for exercising American-style options on cash assets is different and not equivalently symmetric across calls and puts. For non-dividend-paying assets, call options are never exercised early while put options are exercised early if interest earned on the strike, from the time of exercise to expiration, is large relative to time value.

To demonstrate these results and isolate the differences between options on futures and options on cash assets, consider two dates, dates 0 and 1, with options expiring on date 1. As of date 0, date 0 prices are known but date 1 prices are stochastic. Denote futures and cash prices on date i by F_i and P_i , respectively, and denote the short-term rate over the period by r . In the risk-neutral measure (see Chapters 7 and 13), date 0 futures and cash prices are

$$F_0 = E[F_1] \quad (18.95)$$

$$P_0 = E\left[\frac{P_1}{1+r}\right] = \frac{E[P_1]}{1+r} \quad (18.96)$$

The intuition behind the difference between (18.95) and (18.96) is that futures position are entered into without an initial investment, so they do not earn an expected return in the risk-neutral measure. Cash assets, by contrast, do require an initial investment and, therefore, earn an expected return equal to the short-term rate in the risk-neutral measure. The second equality in (18.96) follows because the short-term rate is known as of date 0.

On date 0 the holder of an American-style option needs to decide whether to hold the option to date 1 or to exercise early, on date 0. By definition, it is optimal to hold the option if the value of holding exceeds the value of immediate exercise.

Beginning with call options, let the strike price be K , and assume that both F_0 and P_0 are greater than K so that the option is in-the-money on date 0 and might potentially be exercised early. In the risk-neutral measure, the value of holding an option is the expected discounted value of its terminal payoff. Hence, the condition for holding a call option on a futures is that

$$\frac{1}{1+r} E[(F_1 - K)^+] > F_0 - K \quad (18.97)$$

By the properties of the maximum function, the expectation in (18.97) can be written as

$$E[(F_1 - K)^+] = E[F_1] - K + \Omega^{FC} \quad (18.98)$$

$$\Omega^{FC} \geq 0 \quad (18.99)$$

The quantity Ω^{FC} can be thought of as measuring the time value of the option: it is the expected value of being able to choose whether to exercise or not on date 1 compared with having to exercise on that date. The quantity Ω^{FC} is near zero for options that are way in-the-money and increases as options moves out-of-the-money.

Substituting (18.95) and (18.98) into (18.97) and rearranging terms shows that a futures call option should be held if

$$\Omega^{FC} > r(F_0 - K) \quad (18.100)$$

In words, a futures call should be held if the time value of the option exceeds the interest earned on the intrinsic value of the option. On the other hand, if the call option is sufficiently in-the-money today, Ω^{FC} will be close enough to zero so that (18.100) will not obtain and it will be optimal to exercise early.

The analogous argument for a call option on a cash asset proceeds as follows. The condition for holding the call is

$$\frac{1}{1+r} E[(P_1 - K)^+] > P_0 - K \quad (18.101)$$

and the time value Ω^{CC} is defined such that

$$E[(P_1 - K)^+] = E[P_1] - K + \Omega^{CC} \quad (18.102)$$

$$\Omega^{CC} \geq 0 \quad (18.103)$$

Then, substituting (18.96) and (18.102) into (18.101) gives the condition for holding a call option on a cash asset:

$$\begin{aligned} \frac{E[P_1] - K + \Omega^{CC}}{1+r} &> \frac{E[P_1]}{1+r} - K \\ \Omega^{CC} + Kr &> 0 \end{aligned} \quad (18.104)$$

But (18.104) is always true, so a call option on a cash asset (with no intermediate dividend payments) is never exercised early.

The contrast between the futures result in (18.100) and the cash result in (18.104) can be explained as follows. Exercising a call on a cash asset requires funds equal to the strike. Therefore, delaying exercise effectively earns interest on the strike from exercise to expiration. By contrast, exercising a call on a futures generates the intrinsic value: upon exercise the holder of the option receives a settlement payment equal to the intrinsic value, $F_0 - K$, and is then put into a futures contract at the prevailing market price, F_0 , which position has no market value at that time. Therefore, exercising a call on a futures early earns interest on the intrinsic value from exercise to expiration. If the time value of reserving the right to change one's mind about exercise is small, this interest earning dominates and early exercise is optimal.

Conditions for holding a put option on futures and cash assets can be derived analogously, the details of which are left to the reader. For puts, assume that $K > F_0$ and $K > P_0$ so that early exercise is potentially optimal. Then, the results are that a futures put option should be held if

$$\Omega^{FP} > r (K - F_0) \tag{18.105}$$

$$\Omega^{FP} = E [(K - F_1)^+] - (K - E [F_1]) \geq 0 \tag{18.106}$$

and that a put option on a cash asset should be held if

$$\Omega^{CP} > Kr \tag{18.107}$$

$$\Omega^{CP} = E [(K - P_1)^+] - (K - E [P_1]) \geq 0 \tag{18.108}$$

Contrasting the two conditions for put exercise, futures puts are held if the time value exceeds the value of earning interest on the intrinsic value of the put from exercise to expiration. Puts on cash assets are held if the time value exceeds the interest earned on the strike from exercise to expiration. Since the strike is much greater than the intrinsic value, it is much more likely to hold a futures put than a put on a cash asset. Or, equivalently, it is much more likely to exercise a put option on a cash asset early than to exercise a futures put.

Note too that the rules for holding a futures option are symmetric across calls and puts. Early exercise of either futures option results in the realization of the intrinsic value that can be invested from exercise to expiration. By contrast, the rules for holding options on cash assets are not symmetric across calls and puts. The exercise of a call requires cash equal to the strike and discourages early exercise. The exercise of a put generates cash equal to the strike and encourages early exercise.

Table 18.14 summarizes the results of this subsection in terms of early exercise conditions, which are found by reversing the inequalities of the conditions for optimally holding options over the subsequent period.

TABLE 18.14 Summary of Conditions for Optimal Early Exercise of CME-Type American-Style Options on Non-Dividend Paying Cash Assets and on Futures

	Cash Asset	Futures
Call	Never	$r (F_0 - K) > \Omega^{FC}$
Put	$r K > \Omega^{CP}$	$r (K - F_0) > \Omega^{FP}$

Early Exercise of NYSE Euronext-type Futures Options

After the daily settlement of a futures option of this type, the option has no market value or intrinsic value. Therefore, unlike the futures options described in the previous section, there is no intrinsic premium to be realized from these options and no incentive to exercise early.

APPENDIX C: FUTURES PRICES ARE MARTINGALES WITH THE MONEY MARKET ACCOUNT AS A NUMERAIRE

Recall from Chapter 13 that the initial value of a futures contract is zero; that the subsequent cash flows of a futures contract are its daily settlement flows; and that, at maturity, the futures price is determined by some final settlement rule. Consider a two-period, three-date framework for simplicity, and let the futures price on date t be F_t . Then, the normalized gains process is as follows:

- Date 0: $\frac{V_0}{M(0)} = 0$;
- Date 1: $\frac{V_1}{M(1)} = \frac{F_1 - F_0}{1 + r_1}$;
- Date 2: $\frac{V_2}{M(2)} = \frac{(F_1 - F_0)(1 + r_2) + F_2 - F_1}{(1 + r_1)(1 + r_2)}$

Since the value of a futures contract on date 0 is zero, the martingale property implies that the expectation of the normalized gains at any future date is zero. In particular, for date 1,

$$0 = E_0 \left[\frac{F_1 - F_0}{1 + r_1} \right] \quad (18.109)$$

But since r_1 is known as of date 0, it follows from (18.109) that

$$F_0 = E_0 [F_1] \quad (18.110)$$

As for date 2, the martingale property says that

$$0 = E_0 \left[\frac{(F_1 - F_0)(1 + r_2) + F_2 - F_1}{(1 + r_1)(1 + r_2)} \right] \quad (18.111)$$

Using the law of iterated expectations, and the fact that r_1 is known as of date 0,

$$0 = E_0 \left[\frac{1}{1 + r_2} E_1 [F_2 - F_1] \right] \quad (18.112)$$

But, since F_1 is known as of date 1, (18.112) implies that

$$F_1 = E_1 [F_2] \quad (18.113)$$

Finally then, combine (18.110) and (18.113) to see that

$$F_0 = E_0 [E_1 [F_2]] = E_0 [F_2] \quad (18.114)$$

Together with (18.110), (18.114) shows that the futures price is a martingale under the money-market account or risk-neutral measure, as desired.

Corporate Bonds and Credit Default Swaps

Credit risk, the risk that the promised cash flows from an asset will not be paid as promised, is a primary risk when investing in corporate bonds. Say that a corporation sells \$100 million principal amount of a bond issue, contracting with bondholders to make coupon payments of 7.50% for 10 years and then to return the \$100 million principal amount. There is a risk that the corporation will experience financial difficulties before the bonds mature and default on its contractual agreements. The result might be a reorganization or liquidation in which bondholders not only fail to receive the promised 7.50% interest payments but also fail to recover the full \$100 million principal amount.

Because corporate bonds are characterized by credit risk, investors demand a higher promised return on corporate bonds than on safer forms of investments, like U.S. Treasury bonds. While the corporation in the previous paragraph was selling its 7.50% 10-year bonds, the U.S. Treasury might have been selling 10-year bonds at 3.50%. Part of the higher return paid by corporations compensates investors for the expected losses due to default and part is a risk premium for bearing default risk.

Since the late 1990s, corporate credit risk has traded not only through corporate debt, but also through derivative contracts known as *Credit Default Swaps* or CDS. The exposure to corporate default through CDS is in many ways similar to a *cash* or direct exposure to corporate debt. There are, however, two particularly important differences. First, since a CDS contract is between two counterparties, each is exposed to the counterparty risk of the other in addition to the corporate credit risk that is the purpose of trading the contract. This counterparty risk can be mitigated, however, by the taking and posting of collateral.¹ Second, the financing risks of a CDS position and a cash position are quite different. A counterparty in a CDS contract can maintain exposure to credit through the CDS

¹See Chapter 16 for details in the context of interest rate swaps.

maturity date by meeting any collateral calls. By contrast, maintaining a position in corporate bonds requires financing, either with capital, which can be expensive, or through repo markets, which is subject to significant liquidity risk.

This chapter begins with a description and discussion of the corporate bond market, including ratings, the ratings agencies, and empirical data on defaults and recovery rates. Spread measures of credit risk are next, from the application of spread metrics introduced earlier in this book to *asset swap spreads*. The rest of the chapter is dedicated to credit default swaps (CDS), which have become enormously important in credit markets: how they work, how they are quoted, how CDS spreads are compared with and traded against bond spreads, and how implied hazard rates can be used to compute the *DV01* or duration of a bond subject to credit risk.

CORPORATE SECURITIES

The Overview described the corporate life-cycle of debt financing, which culminates in a corporation's being large enough and creditworthy enough to borrow funds in public markets. Almost all of the secondary trading in corporate debt is over-the-counter.²

For short-term public borrowing, corporations sell commercial paper (CP). CP is typically a discount (i.e., zero-coupon) security that can be unsecured, backed by a letter of credit from a bank, or backed by assets. From the perspective of the issuer, CP is attractive because it is relatively inexpensive and generally liquid. In addition, CP in the United States is exempt from registration with the Securities and Exchange Commission (SEC), with its attendant costs, prospectus disclosures, and other requirements, so long as the CP issue matures in less than 270 days, and the corporation can argue that the proceeds of the CP issue are being used for short-term purposes (rather than, for example, building a factory). The disadvantage of selling CP, of course, is the liquidity risk of having to roll short-term borrowing as it matures.

For public borrowing with customized payments terms, corporations sell medium-term notes (MTNs). Historically these were of intermediate maturity and were so named to distinguish them from shorter-term CP and longer-term corporate bonds. Currently, however, MTNs are just as well characterized by their customization to suit the needs of issuers and investors. MTNs first became popular in the United States in the early

²Secondary trading means trading after the initial sale of the issue. In over-the-counter trading two parties set their own terms and conditions, in contrast with trading that occurs under the auspices of an organized exchange.

1980s when the SEC introduced shelf registration. This allowed issuers to register once to sell bonds gradually over a two-year period at payment terms that can be set, based on market conditions, at the time of each sale.

For public borrowing with relatively standard payment terms and with maturities longer than those of CP, firms sell corporate bonds. In the United States, these issues have to be registered with the SEC. Corporate bonds are typically coupon-bearing securities with relatively standard payment terms that often include a call option for an issuer to repurchase the securities at some schedule of prices over time. (See Chapter 18.) A much smaller market exists for corporate *floating rate notes (FRNs)*. Each interest payment on these bonds is usually *London Interbank Offered Rate (LIBOR)*, observed at the start of that payment's accrual period, plus a fixed spread, although sometimes LIBOR might be multiplied by a factor or *leverage* and the spread might depend on the credit rating of the bond at the start of the accrual period.

When selling a debt issue, a corporation enters into a contract with debtholders, called an *indenture*, which is enforced by a trustee. Aside from payment terms, the indenture specifies the priority of the issue in the event of default. For example, one bond issue of a company might be secured by a particular set of assets, a second might be unsecured but "senior," and yet another might be "subordinated." In this example, should a corporation be reorganized or liquidated according to "strict priority," proceeds from selling the ring-fenced assets would first be applied to satisfy the claims of the secured bondholders. Any excess proceeds, together with other assets and cash of the corporation, would next be applied to satisfy the claims of the senior debtholders. Finally, whatever of value remains after that would be used to satisfy the claims of the subordinated debtholders. Of course, reorganizations and even liquidations often involve negotiation and the courts, so that strict priority is not always applied to the settlement of claims.

Indentures also include *covenants* to protect the claims of debtholders. Examples include these: maintaining various financial ratios; restricting the amount of cash that can be paid to stockholders; requiring a corporation to repurchase a debt issue after a change of control; limiting the total amount of new debt incurred by the corporation; and preventing the sale of debt with higher seniority than a particular debt issue.

RATINGS, DEFAULT, AND RECOVERY

While investors very much need to understand the credit quality of individual corporate bonds, this analysis requires substantial information and expertise. Furthermore, it would seem inefficient for all investors to start

TABLE 19.1 Cumulative Default Rates by Original Rating, 1970–2009

Age in Years	Aaa	Aa	A	Baa	Ba	High Yield
1	0.0%	0.0%	0.1%	0.2%	1.2%	4.5%
2	0.0%	0.1%	0.2%	0.5%	3.2%	9.3%
3	0.0%	0.1%	0.3%	0.9%	5.6%	13.9%
4	0.0%	0.2%	0.5%	1.4%	8.1%	17.9%
5	0.1%	0.2%	0.7%	1.9%	11.9%	21.4%
10	0.5%	0.5%	2.0%	4.9%	20.0%	34.0%
15	0.9%	1.2%	3.6%	8.8%	29.7%	43.3%
20	1.1%	2.5%	5.9%	12.3%	37.2%	49.6%

Source: Moody's.

from scratch when investigating the credit quality of a particular corporate issuer or issue. Not surprisingly, therefore, *rating agencies* have developed to help investors assess the credit quality of debt issues. The three major rating agencies are Moody's, Standard and Poor's (S&P), and Fitch. They are normally paid by corporations to rate particular debt issues and by investors to access the resulting ratings and analyses. Corporate bond ratings³ range from Aaa by Moody's and AAA by S&P and Fitch for the most creditworthy issues to C by Moody's and D by S&P and Fitch for issues already in default. Issues with ratings of Baa and above by Moody's and BBB- and above by S&P and Fitch are considered investment grade; issues with lower rating are considered speculative grade or, euphemistically, high yield.

Historical Averages of Default and Recovery Rates

For orders of magnitude with respect to default rates and the variation of default rates across ratings, Table 19.1, from Moody's, shows cumulative default rates historically, by rating, as a function of age. For example, 29.7% of issues rated Ba defaulted over the subsequent 15 years. As expected, cumulative default rates increase as ratings decline.

As indicated in the introduction to this chapter, realized returns from investing in corporate bonds depend not only on defaults but also on losses given default. The *recovery rate* of an issue after default is defined as the fraction of the principal amount ultimately returned to debtholders. Table 19.2, also from Moody's, shows average historical recovery rates for senior unsecured debt as a function of rating. For the most part, recovery rates do decline with rating, but, apart from Aaa-rated securities, the decline

³The scales for CP are different, e.g., for Moody's, P-1, P-2, P-3, and NP, where "P" is for Prime.

TABLE 19.2 Average Recovery Rates for Senior Unsecured Debt, 1982–2009

Rating	Aaa	Aa	A	Baa	Ba	B	Caa-C
Recovery	62%	44.4%	41.4%	43.8%	42.4%	37.5%	34.9%
Grade		Investment				Speculative	
Recovery		43.5%				37.5%	

Source: Moody’s.

is not particularly dramatic. This table also supports the industry practice of using a recovery rate of 40% for assorted credit-market calculations.

Variability of Default and Recovery Rates

While it is tempting to focus on average default rates, like those presented in Table 19.1, Figure 19.1 demonstrates that annual default rates can vary significantly over time. Using S&P data, the heavy line in the figure shows a time series of overall default rates, i.e., across rating classes and sectors, while the three lighter, dashed lines show time series of high-yield default rates for the United States, Europe, and emerging markets. The overall default rates vary from less than .5% to about 4% and the high-yield default rates vary from less than 1% to over 15%. Apart from their volatility over time, two other remarks might be made about the default rates in the figure. One, defaults are highly correlated across the three sectors, which is indicative




FIGURE 19.1 Historical Default Rates, Overall and for High Yield in the United States, Europe, and Emerging Markets, 1990–2009

Source: Standard & Poor’s.

of a systemic component. Two, the magnitude of the default rates over the 2007–2009 crisis were last matched at the time of the Enron and WorldCom defaults in December 2001 and July 2002 respectively.

The double-line in Figure 19.1 gives the defaulting notional amount each year, in billions of USD. From this perspective, the crisis of 2007–2009 was many times more severe than any other episode over the last two decades.

Default rates and the mapping from ratings to default rates are very important to investors, but are not a complete description of creditworthiness over time. The most common path to default is not a single jump from a top rating to default but a sequence of smaller downgrades. Investors and rating agencies constantly assess the creditworthiness of issues, with the latter adjusting ratings up or down as appropriate. Figure 19.2, using S&P data, presents a time series of the percentage of issuers experiencing upgrades and downgrades, as well as the commonly-cited ratio of upgrades to downgrades. The percentage of upgrades or downgrades in a given year is economically significant and volatile, ranging from about 5% to about 19% over the sample period. Also, systemic effects are reflected in the ratio of upgrades and downgrades, just as they are in the default rates of Figure 19.1. In fact, the ratio of upgrades to downgrades hits a trough or reaches a peak at the same times that default rates reach a peak or hit a trough.

Just as default rates in a particular year can be very different from average historical default rates, so can recovery rates in a particular year




FIGURE 19.2 Percent of Issuers Upgraded or Downgraded and the Ratio of These Percentages, 1990–2009

Source: Standard & Poor's.

be very different from the average historical recovery rates, like those in Table 19.2. In fact, not surprisingly, default and recovery rates are highly correlated over time. There is an additional problem, however, in using average recovery rates, even when conditioned on the priority class of a debt issue. To take a simple example, subordinated debt can recover a lot more from an issuer with a relatively small amount of senior debt than from an issuer with a relatively large amount of senior debt. More generally, the value of the priority order of a debt issue can be very idiosyncratic to the issuer's capital structure, which can often be complex. As an illustration of a particularly complicated priority structure, consider the following excerpt with respect to a revised Chapter 11 plan for Lehman Brothers Holdings Inc., filed at the end of January 2011:⁴

... [Senior unsecured creditors in Class 3 with claims against the holding company should have a 21.4 percent recovery. ... Senior intercompany claims against the holding company in Class 4a are in line for 16.6 percent. Intercompany claims against the holding company in Class 8a are to have 15 percent. ... Senior third-party guarantee claims against the holding company in Class 5a are estimated to see 12.9 percent. ... For Class 7 general unsecured claims against the holding company, the recovery is an estimated 19.8 percent. ...

The recovery on derivative claims and unsecured claims against Lehman Commercial Paper Inc. is an estimated 51.9 percent. ... For derivative and general unsecured claims against Lehman Brothers Special Financing, the recovery is 22.3 percent.

Policy Issues with Respect to the Rating Agencies

For a long time there has been controversy surrounding the role of rating agencies in the financial system. Some of the major facets of this controversy are as follows. First, regulatory bodies outsource some of their responsibilities to privately-run rating agencies by making rules that depend on ratings. The most prominent examples include international bank capital rules under the Basel Accords; U.S. broker-dealer capital requirements; and quality standards for the security holdings of U.S. money market funds. Second, regulatory bodies choose which rating agencies can be used for regulatory

⁴Bill Rochelle, "Lehman, Summit, OTB, Townsends, Vitro: Bankruptcy," Bloomberg, January 26, 2011.

purposes. These selections might very well confer special status and competitive advantage on the chosen rating agencies and create undue concentration in the ratings industry. Third, since issuers pay to have their issues rated, rating agencies are consistently open to charges of conflicts of interest. Controversy around rating agencies waxes and wanes over time, but becomes more vociferous after rating agencies “miss” a significant default, e.g., keeping investment-grade ratings on Enron and WorldCom until shortly before they defaulted.

Not surprisingly, controversy flared dramatically through the 2007–2009 financial crisis. The rating agencies had been earning a larger and larger percentage of their revenues from rating mortgage-related structured products, which performed very poorly through the crisis, including securities that were rated Aaa/AAA. In response, the Dodd-Frank Act of 2010 in the United States made several changes to the *status quo*, two of which will be mentioned here. First, regulatory bodies were given two years to remove all references to ratings from their rules and replace these ratings with their own credit standards. This is a substantial undertaking, which, at the time of this writing, is ongoing. Furthermore, this does not seem consistent with the latest international banking accord, Basel III, which does not excise references to ratings in the determination of bank capital requirements. Second, Dodd-Frank changed the law so that rating agencies can be held liable for ratings that are used as part of a security’s registration statement. And, because of this potential liability, rating agencies have to agree to have their ratings so used. This provision misfired. The SEC requires that the registration of asset-backed securities (ABS) include a rating. But soon after Dodd-Frank became law in July 2010, the rating agencies, eager to avoid potential liabilities, refused to attach their ratings to ABS registrations. As a result, the ABS market ground to a halt for several days. The SEC resolved the impasse by temporarily suspending the requirement that ratings be attached to ABS registration statements. At the time of this writing, this temporary measure is still in place.⁵

CREDIT SPREADS

Credit spreads are the differences between the relatively high rates earned on bonds that are subject to credit risk and the relatively low rates on securities subject to little or no credit risk. The simplest measure of credit spreads is the *yield spread*, which is the difference between the yield on the bond and the yield or rate on a similar maturity (highly creditworthy) government bond

⁵The corporate bond market experienced a similar although less dramatic sequence of events, resulting in ratings requirements being relaxed there as well.

or par swap. While yields spreads are computed in practice, they suffer from two drawbacks. First, they confound differences in the structure of cash flows with credit risk. (See “Yield Curves and the Coupon Effect” in Chapter 3.) Second, to the extent that the issuer of a bond has embedded options (see Chapter 18), yield will be higher than it would be otherwise and the yield spread will indicate a misleadingly high level of credit risk.

A better measure of credit spread will be referred to in this chapter as the *bond spread*. This term includes the spread defined in Chapter 3, the Treasury Euro Dollar (TED) spread in Chapter 15, and the more general option-adjusted spread (OAS) defined in Chapter 7. In the credit context, a bond spread is computed by assuming no default and finding the spread (or term structure of spreads) over a benchmark curve that prices the bond as it is priced in the market. As the market price incorporates the risk of default while the pricing methodology just described does not, the resulting spread is an indicator of credit risk. Bond spreads, unlike yields spreads, properly account for any differences between the structure of cash flows of a bond and those of the benchmark securities. Hence, the spread in Chapter 3 is suitable for bonds without embedded options. For bonds with embedded options, the OAS in Chapter 7 is appropriate since its computation is designed to incorporate the value of embedded options and to attribute to the OAS only the remaining price difference, which, in the present application, is chiefly due to credit spreads.

It is a sign of the times, as of this writing, that the market is following the bond spreads of European government bonds. Table 19.3 gives bond spreads for 5- and 25-year issues of several European governments, both with respect to the LIBOR and Overnight Indexed Swap (OIS) benchmark curves. Because LIBOR is higher than OIS (see Chapter 15), bond spreads

TABLE 19.3 Bond Spreads of European Government Bonds, in Basis Points, as of December 10, 2010

Issuer	5-Year		25-Year	
	LIBOR	OIS	LIBOR	OIS
Germany	-28.2	2.0	-5.0	16.1
Finland	-21.7	8.7	-10.4	12.6
Netherlands	-15.3	14.9	0.6	19.3
France	-8.0	22.4	25.5	46.3
Belgium	82.0	112.4	79.7	100.7
Italy	114.1	144.3	168.6	190.0
Spain	213.6	244.0	243.9	263.3
Portugal	295.1	325.1	281.9	301.1
Ireland	455.7	486.2	438.6	462.3
Greece	931.5	961.7	525.9	545.8




FIGURE 19.3 A Par-Par Asset Swap with Financing

against LIBOR are lower than bond spreads against OIS. In fact, as the best sovereign credits are better than short-term bank credits, several of the bond spreads against LIBOR are negative. In any case, the spreads in the table dramatically indicate the perceived pecking order of creditworthiness in European government issuers, from the solid credits of Germany, Finland, the Netherlands, and France, to the intermediate credits of Belgium and Italy, to the relatively impaired credits of Portugal, Ireland, Greece, and Spain.

The plan for the rest of this section is the following. The next subsection introduces the popular measures of credit spreads known as *asset swap spreads*. The following two subsections then illustrate differences across the various measures of credit spread, first in the context of a simple example and then in the context of a particular credit-impaired bond. Finally, the last subsection describes credit spread measures for floating rate notes.

Asset Swaps and Asset Swap Spreads

The point of an *asset swap* is to use interest rate swaps to transform a fixed-coupon bond into an asset that earns a spread over LIBOR. One flavor of asset swaps, called the *par-par asset swap*, or, more simply, the *par asset swap*, is illustrated in Figure 19.3. The light lines indicate cash flows at initiation of the asset swap, the heavy lines indicate cash flows during its life, and the dashed lines indicate cash flows at its termination. At initiation, the purchaser of the bond or asset swapper buys the bond for P per 100 face amount, earning a periodic coupon payment of c per 100 face amount. The purchase of the bond is financed with 100 from the repo desk (or, not shown, with some combination of repo and capital financing) and with $P - 100$ that comes from an up-front payment from an interest rate swap.⁶

⁶To keep the focus on assets swap spreads, this discussion ignores collateral requirements and how they change over time, for both repo and swap agreements.

Finally, through that interest rate swap, the asset swapper agrees to pay the bond coupon c in exchange for receiving LIBOR plus the spread s^{PAR} on 100 face amount and for receiving the up-front payment just mentioned of $P - 100$. Note that this trade requires no cash at initiation, and, so long as the bond does not default, earns LIBOR plus s^{PAR} minus the repo rate on 100 over the life of the trade. The trade also neither generates nor requires cash at initiation since the principal payment from the (non-defaulting) bond is used to pay off the repo borrowing. Hence, so long as the bond does not default, the asset swapper has converted the fixed cash flows of the bond into floating payments of LIBOR plus s^{PAR} on 100 notional amount. (Repo or capital financing costs would be incurred with or without the swap.)

What is the fair value of the asset swap spread, s^{PAR} ? Let d be the discount factor corresponding to the maturity date of the swap and let A^{Fixed} and A^{Float} be annuity factors from the swap curve, adjusted for payment schedules, so that A^{Fixed} times the coupon payment gives the present value of those coupon payments and $100s^{PAR}$ times A^{Float} gives the present value of the floating payments. Then the swap depicted in Figure 19.3 is fair if the present value of all the payments received by the asset swapper equals the present value of all the payments made by the asset swapper. Using the device of the fictional notional payment at maturity and the result that payments of LIBOR together with that fictional notional amount is worth par (see Chapter 16), the fair-pricing condition is

$$(P - 100) + 100 + 100s^{PAR}A^{Float} = cA^{Fixed} + 100d$$

$$s^{PAR} = \frac{cA^{Fixed} + 100d - P}{100A^{Float}} \quad (19.1)$$

Note that s^{PAR} depends on the credit risk of the bond through the bond's price, P : the lower this price relative to the present value of the same fixed payments from an essentially default-free swap (see "On Credit Risk and Interest Rate Swaps" in Chapter 16), the higher the spread. This is the sense in which the par asset swap is a measure of credit risk.

The par asset swap package has minimal interest rate risk conditional on no bond default. The floating side of the swap does have interest rate risk to the next reset date, but even this small risk might very well be hedged by the financing of the bond⁷ or with the addition of short-term rate derivatives

⁷Borrowing cash through repo to the term of the next payment would hedge the interest rate risk of the next—and already set—LIBOR-based floating payment. If repo borrowing is overnight, however, then ED futures or another rate derivative might be used to hedge that risk.

(see Chapter 15). The credit risk of the bond, however, is most certainly retained. If the bond price falls due to a credit event, the value of the package will fall. And, of course, if the bond defaults, the coupon and principal payments from the bond might not be made in full while the asset swapper still owes coupon payments through the swap in addition to the repayment of repo borrowings (or return of capital).

To be more precise about the credit risk of the par asset swap position, consider the profit and loss (P&L) impact at some time, in the future, of two events: 1) the par asset swap spread jumps to $\tilde{s}^{PAR}$; and 2) the bond defaults. Note that for the purpose of quantifying the credit risk of the position going forward, income earned from the initiation of the asset swap through the event date is ignored.

If the new asset swap spread jumps to $\tilde{s}^{PAR}$, an investor can sell the par asset swap, i.e., sell the bond, receive fixed in the swap, etc. Using the analysis of this subsection, this set of trades has no net initial or final payments, but generates intermediate payments of $\tilde{s}^{PAR}$ plus LIBOR minus repo on the face amount of the bond. But the existing long asset swap position earns s^{PAR} plus LIBOR minus repo on the face amount. Hence, the P&L from the jump in the asset swap spread is the present value of the difference between these interim cash flows, i.e., $100(s^{PAR} - \tilde{s}^{PAR})\tilde{A}^{Float}$, where $\tilde{A}^{Float}$ is the relevant annuity factor at the time of the jump in the asset swap spread. If the creditworthiness of the bond improves so that the asset swap spread falls, the original, long asset swap spread position makes money; if the creditworthiness of the bond deteriorates so that the spread increases, the long asset swap position loses money.

Moving to the default scenario, let R be the recovery rate. If the bond defaults and the trade is unwound, the P&L is as follows. First, the bond is worth its notional amount times the recovery rate, i.e., $100R$, while it was worth some price $\tilde{P}$ just before default. Second, 100 is owed on the repo loan. Third, the swap has to be unwound at some net present value (NPV) with respect to the asset swapper, to be denoted as $\widetilde{NPV}$. Note that, since the swap is fair at initiation, as soon as the asset swapper takes the swap's initial payment of $P - 100$, the NPV of the swap to the asset swapper is $100 - P$. This quantity will trend to zero over time as any premium or discount that arises from fixed cash flows. In any case, just before default the total asset swap position is worth $\tilde{P} - 100 + \widetilde{NPV}$ while just after default it is worth $100R - 100 + \widetilde{NPV}$ for a net change of $100R - \tilde{P}$. This P&L expression certainly shows that the asset swap loses money in the event of a default when the recovery is low relative to the prevailing price. But the expression also shows that the asset swap trade does have interest rate risk in the event of default: the prevailing price, $\tilde{P}$, can differ from the original price because of changes in interest rates.




FIGURE 19.4 A Market Value Asset Swap with Financing

Delaying further discussion of the par asset swap until the numerical example, Figure 19.4 depicts another flavor of asset swaps, called the *market value asset swap*. In this case the purchaser of the bond or asset swapper finances the purchase price P from the repo desk (or, again not shown, with some combination of repo and capital financing). Through the swap, the asset swapper pays the coupon c as in the par asset swap, but receives LIBOR plus the spread s^{Mkt} on the notional amount P rather than on 100. Then, at termination, the repo loan of P is repaid with the 100 principal payment of the bond and a terminal payment from the swap of $P - 100$. Netting all the pieces, if the bond does not default, the market value asset swap converts the fixed payments of the bond into floating payments of LIBOR plus s^{Mkt} on the notional amount P . Taking care to recall that the notional amount of the swap is P , so that LIBOR plus the fictional notional is worth P ; that the spread is earned on P ; and that the fictional notional on the fixed side is also P , the fair pricing condition of the swap in this trade requires that

$$\begin{aligned}
 P + s^{Mkt} P A^{Float} + (P - 100) d &= c A^{Fixed} + P d \\
 s^{Mkt} &= \frac{c A^{Fixed} + 100 d - P}{P A^{Float}} \\
 &= \frac{100 s^{PAR}}{P}
 \end{aligned} \tag{19.2}$$

Equations (19.1) and (19.2) show that the par and market value swaps are very closely related. The asset swapper can transform the cash flows of the bond so as to earn s^{PAR} on 100 or so as to earn $s^{Mkt} = \frac{100 s^{PAR}}{P}$ on P , which amount to the same size payments. The choice between the two

TABLE 19.4 Various Measures of Credit Spreads for a Two-Year, 4.25% Bond with Forward Rates of 1% Over the First Year and 2% Over the Second Year

Bond Yield	10.013%
Par Swap Rate	1.495%
Yield Spread	8.518%
Bond Spread/OAS	8.525%
Par-Par Asset Swap Spread	7.855%
Market Value Asset Swap Spread	8.728%

kinds of asset swaps, therefore, is not one of economics but of collateral or counterparty risk considerations.⁸

Numerical Example of Credit Spreads

Consider a simple, two-period, two-year example. Let the one-year forward swap rates be $f(1) = 1\%$ and $f(2) = 2\%$. A corporate bond has a coupon of 4.25%, matures in two years, and sells at a price of 90. Table 19.4 reports various rates and spreads for this example. Note that the bond in this example is perceived to be subject to a lot of credit risk. Despite having a very much above-market coupon, its price is only 90. This perceived credit risk is reflected in the large magnitudes of all of the credit spread measures reported in the table.

The yield spread is simply the bond yield minus the par swap rate. As mentioned in the introduction to this section, the coupon effect lowers the bond yield and reduces the yield spread relative to the bond spread or OAS, although the effect is small in this example.

⁸Begin with the case of a premium bond, i.e., $P > 100$. In the par asset swap the swap desk advances money to the asset swapper in exchange for the promise of future payments. If the asset swapper does not post collateral, this exposes the swap desk to counterparty risk at the time of initiation. If the asset swapper does post collateral, this poses an opportunity cost of posting collateral on the asset swapper. Over time, however, as the asset swapper makes coupon payments, the counterparty risk or collateral requirements decline. In the market asset swap, by contrast, there is no initial swap payment and, therefore, no initial counterparty risk or collateral requirement. Over time, however, as the asset swapper makes payments, the obligation of the swap desk to pay $P - 100$ at termination becomes counterparty risk for the asset swapper or a collateral requirement for the swap desk. In the case of a discount bond, i.e., $P < 100$, the obligations flip: the par asset swap initially exposes the asset swapper to counterparty risk or the swap desk to collateral requirements while the market value swap still has no initial counterparty risk or collateral requirements but, over time, exposes the swap desk to counterparty risk or the asset swapper to collateral requirements.

Since this bond has no embedded options, its spread is computed along the lines of Chapter 3 and is the same as the bond's OAS. In particular, the spread s is such that

$$90 = \frac{4.25}{1.01 + s} + \frac{104.25}{(1.01 + s)(1.02 + s)} \quad (19.3)$$

Solving, $s = 8.525\%$.

Rearranging the terms of (19.1), s^{PAR} is such that

$$100s^{PAR} \left[\frac{1}{1.01} + \frac{1}{1.01 \times 1.02} \right] = \frac{4.25}{1.01} + \frac{104.25}{1.01 \times 1.02} - 90 \quad (19.4)$$

And solving, $s^{PAR} = 7.855\%$. Hence, a bondholder can transform the fixed coupon payments of 4.25% into floating rate payments of the short-term rate plus 7.855%. The market value asset swap spread is, according to (19.2), simply the par spread normalized to price, i.e., $7.855\% \frac{100}{90}$ or 8.728%.

Comparing (19.3) and (19.4) shows that the bond spread adds a spread to the denominator in order to capture the difference between the bond price, i.e., 90, and what the bond price would be without credit risk, i.e., the present value of the cash flows at the market forward rates. The asset swap spread, on the other hand, explains this difference with a spread in the numerator.

Some additional intuition about the difference between the bond and asset swap spreads are gleaned from considering how these quantities are related to return given that a bond does not default. As discussed in Chapter 3, the bond spread is a component of return over the life of the bond in the following sense: investing the bond price by compounding returns at the forward rates plus the bond spread is equivalent to investing in the bond so long as all coupon payments can be reinvested at those same rates and spreads. In terms of this example,

$$\begin{aligned} 90(1.01 + 8.525\%)(1.02 + 8.525\%) &= 4.25(1.02 + 8.525\%) + 104.25 \\ &= 108.947 \end{aligned} \quad (19.5)$$

For the market value asset swap, the relationship to return is somewhat different: investing the bond price at the forward rates plus the bond spread while compounding these returns at the forward rates (without spread) is equivalent to investing in the bond so long as all coupon payments are invested at the forward rates (without spread). In terms of this example,

$$\begin{aligned} 90[1 + (1\% + 8.728\%)1.02 + (2\% + 8.728\%)] &= 4.25(1.02) + 104.25 \\ &= 108.586 \end{aligned} \quad (19.6)$$

TABLE 19.5 Various Measures of Credit Spreads for the KB Home 5 $\frac{3}{4}$ s of February 1, 2014, on November 10, 2008

Bond Yield	15.089%
Par Swap Rate	3.714%
Yield Spread	11.375%
Bond Spread/OAS	11.492%
Par-Par Asset Swap Spread	9.230%
Market Value Asset Swap Spread	13.443%

Application: KB Home 5 $\frac{3}{4}$ s of February 1, 2014

This subsection presents credit spread measures for a particular distressed bond, the 5 $\frac{3}{4}$ s of February 1, 2014, issued by KB Home, a U.S. residential home-construction company.

Table 19.5 gives various rates and spread for the KB Home bond on November 10, 2008. The bond is clearly distressed because, while the matched-date par swap rate was 3.714%, this relatively high-coupon bond sold for a full price of only 68.66. As a result, the credit risk measures in the table are quite significant. Also, as in the numerical example of the previous subsection, the yield spread is not very far from the bond spread and the market value asset swap spread is greater than the bond spread.

This KB Home bond will be revisited in the context of CDS basis trades in the next section of this chapter.

Credit Spreads for Floating Rate Notes

The spread on a floating rate note that is priced at par is a particularly pure form of a credit spread—it is the spread over the short-term benchmark received for bearing credit risk with no other consideration. Over time, however, as the credit quality of the issuer changes, the price of a floating rate note with a fixed spread will change. As a result, the spread is not so pure a measure of credit risk because an investor is paying a premium or getting a discount to face amount in addition to that spread. Hence, there is a need to measure credit spreads even for floating rate notes.

One way to do this is to quote an *effective spread* that converts the premium or discount into a run rate and adds it to the actual spread. Say that the actual spread of the floater is s^{Float} , that the price of the floater is P , and that the annuity corresponding to payment dates is, as before, A^{Float} . Furthermore, let s^{Eff} be the effective spread. Then,

$$100s^{Eff} A^{Float} - 100 = 100s^{Float} A^{Float} - P \quad (19.7)$$

In words, investors are indifferent between receiving the spread s^{Eff} for a price of 100 and receiving the spread s^{Float} for a price of P . Rearranging terms,

$$s^{Eff} = s^{Float} + \frac{100 - P}{100 A^{Float}} \quad (19.8)$$

Another method of quoting an effective spread on a previously issued floating rate note is to fix its cash flows at the benchmark forward rates plus the actual spread. Then, find the bond spread that equates the present value of these cash flows to the price of the floater.

CREDIT SPREADS AND DEFAULT RATES

The measures of credit spread developed in the previous section represent, in various ways, the additional return an investor gets from holding a bond that does not default. But, of course, bonds with credit risk sometimes do default. It is natural to ask, therefore, whether, on average, the magnitude of credit spreads offered in the market compensates for events of default.

To analyze this question, consider the bond spread. As explained in Chapters 3 and 7, in the context of bonds that cannot default, the short-term return of an interest-rate hedged position in a bond with a constant bond spread equals the short-term rate plus that bond spread. But what if the bond might default and, in particular, that the probability of default over the next small time interval, dt , is λdt ? Let r be the short-term rate, s the bond spread, and R the recovery rate in the event of default. Then, if the bond does not default over the next instant, which happens with probability $(1 - \lambda dt)$, the return on an interest-rate hedged investment is $(r + s) dt$. However, if the bond does default, principal is lost but for the recovered fraction R , implying a return of $-(1 - R)$. Thus, the expected return of an interest-hedged position in the bond over the next instant is

$$(1 - \lambda dt) \times (r + s) dt - \lambda dt \times (1 - R) \quad (19.9)$$

As investors are risk averse, they prefer a riskless return of $r dt$ to a risky return with an expectation of $r dt$.⁹ But, for present purposes, ignore risk aversion and assume that investors are content to earn a spread on corporate bonds that compensates for the expected losses due to default. In

⁹More precisely, from asset pricing theory, this reasoning applies when the risk is positively correlated with the wealth of the economy. This is certainly the case here since corporate bond returns are higher when the economy is doing well.

other words, the required expected return in (19.9) is equal to rdt . Equating these two quantities, while ignoring the very small, higher-order terms, i.e., those with a factor of dt^2 , results in the following spread requirement:

$$\begin{aligned} rdt &= (r + s) dt - \lambda dt \times (1 - R) \\ s &= \lambda (1 - R) \end{aligned} \tag{19.10}$$

Since data are available on cumulative default rates, as in Table 19.1, rather than on instantaneous default rates, the spread in equation (19.10) has to be expressed in terms of the cumulative default probability to some time T , denoted here as $CD(T)$, instead of λ . It is shown in Appendix A in this chapter that if the instantaneous default rate is constant at λ , then

$$CD(T) = 1 - e^{-\lambda T} \tag{19.11}$$

Finally, then, substitute λ from (19.10) into (19.11) and solve for the spread:

$$s = -\frac{1 - R}{T} \ln [1 - CD(T)] \tag{19.12}$$

The top panel of Table 19.6 uses equation (19.12), data on five-year cumulative default rates, and an assumed recovery rate of 40% to imply the

TABLE 19.6 Top: Spreads, in Basis Points, Required to Compensate for Realized Default Risk Over Five-Year Periods; Bottom: the Widest and Tightest Market Spreads Over the Sample, by Rating Category

Period	Investment Grade	Rating		High Yield
		B	Caa	
1970–74	10	89		60
1975–79	13	227		94
1980–84	20	384		293
1985–89	15	641		448
1990–94	2	219	292	172
1995–99	16	557	1,136	450
2000–04	8	108	473	132
2005–09	16	265	631	268
Market Spreads				
Widest	619	1,825	2,992	1,857
Tightest	82	239	395	256

Source: Moody’s and Deutsche Bank.

spreads that would have been necessary to compensate investors for realized defaults over sequential five-year periods since 1970. The bottom panel indicates the widest and tightest spreads over the sample. For the lower ratings, the orders of magnitude of the *ex-post* required spreads are for the most part bracketed by the widest and tightest spreads, although there are some periods for which spreads implied by realized defaults are lower than the tightest market spreads. Nevertheless, the widest market spreads are much larger than anything justified by realized defaults. This might be due to the market's overestimation of default risk or to the existence of a significant credit risk premium. For investment grade debt, however, market spreads seem particularly wide: even the tightest market spread significantly overcompensates investors for realized defaults over any of the five-year periods. Once again, this might be due to overestimation of default risk or a risk premium. Another possible explanation is that, because the default rate of investment grade debt is so low, very long observation periods are required to witness periods of significant defaults.

CREDIT DEFAULT SWAPS

Definitions and Mechanics

Through a single-name CDS, a *protection buyer* or *CDS buyer* pays a *protection seller* or *CDS seller* in exchange for a *compensation payment* in the event that an issuer of bonds defaults. These derivatives are useful for hedging credit exposures to particular issuers and for betting on the credit-worthiness of an issuer relative to implied market prices. The last subsection discusses CDS on indexes.

A CDS contract is defined by a *reference entity*, a list of *credit events*, a *term* or maturity, a *reference obligation*, and a notional amount. To take one example, consider a five-year CDS on \$1 million of the Senior Unsecured 7½s of May 15, 2016, issued by Hovnanian Enterprises (HOV), a U.S. home-building company. A payment by the seller of this CDS would be triggered if, before the maturity of the CDS in five years, the reference entity, HOV, experiences a credit event, which typically includes bankruptcy, failure to pay, obligation acceleration, repudiation, moratorium, and restructuring. Furthermore, should such a credit event occur, the payment by the CDS seller, or the CDS settlement, would be determined, in a manner to be described presently, with reference to \$1 million face amount of the 7½s of May 15, 2016.

The purpose of the compensation payment of a CDS is to make a holder of the reference obligation whole in the event of a default. Continuing with the example, should HOV default and the value of the 7½s plummet to 20 per 100 face amount, the ideal compensation payment would be 80 per 100 face




FIGURE 19.5 Physical and Cash Settlement of CDS Contracts

amount or \$800,000 based on the notional amount of this particular CDS. The actual payment is more complicated because of the difficulty in precisely determining the market price of a corporate bond at any time, let alone after a credit event. One solution to this difficulty is *physical settlement*, described in the lower half of Figure 19.5 for a notional amount of 100. In this process the protection buyer delivers 100 notional amount of a bond to the protection seller and receives 100 in exchange. In terms of the example, the buyer can deliver \$1 million face amount of the 7½s to the seller and receive \$1 million in exchange. Hence, no matter what the market price of the bond, the protection buyer has recovered the full face amount of the bonds despite the credit event.

Requiring physical settlement of a particular issue might subject the buyer to a squeeze. This same fear was described in Chapter 14 in the context of note and bond futures, and the solution for CDS is the same as the solution for futures, i.e., to permit delivery of any security in a list of eligible securities. In the example, the buyer would typically be able to deliver any HOV bond with the seniority of senior unsecured or better. These seniority criteria for eligibility usually work well because, in a default, seniority is the most important determinant of value.¹⁰ Of course, as in

¹⁰ In the case of a restructuring, the seniority criteria do not necessarily work well. Short-term debt is often restructured advantageously relative to long-term debt, resulting in cheaper, long-term debt being delivered through CDS contracts. Many market participants wanted to avoid this side-effect of the delivery option and, as a result, some contracts now restrict eligibility to relatively short-term debt in the case of restructurings while other contracts have eliminated restructuring as a credit event.

the case of futures, the flexibility to deliver any of a list of eligible securities creates a delivery option and a cheapest-to-deliver that impacts the valuation of CDS contracts.

A second solution to the difficulty of pricing a corporate bond is to hold an auction to determine bond prices and then use the resulting prices for *cash settlement*. In the example, an industry organization would hold an auction in which market participants could bid to buy and offer to sell senior unsecured HOV bonds. If the clearing price turned out to be 20 per 100 face amount, then all sellers of CDS contracts on senior unsecured debt of HOV would owe compensation payments of $100 - 20$ or 80 per 100 face amount. More generally, as shown in Figure 19.5, if the auction clearing price is the recovery rate, R , per unit face amount, then the compensation payment is $100(1 - R)$. The auction process worked relatively well through the 2007–2009 financial crisis. In October 2008, auctions set the price of Lehman Brothers secured debt at about 9 per 100 face amount and of Wahington Mutual senior debt at 57.¹¹

In exchange for the compensation payment in the event of default, the CDS buyer pays an *up-front* amount at the initiation of the trade and then an actual/360 *fee*, *premium*, or *coupon* paid quarterly, until the earlier of the maturity of the CDS or the issuer's default. If the up-front payment is zero, the premium is also known as the *default swap spread* or the *CDS spread*. In the HOV example, as of November 10, 2008, the 7.5% reference obligation was quite distressed: the price of five-year CDS protection was an up-front payment of 55.5 per 100 face amount and a coupon of 500 basis points per year. Note that in the event of default, the buyer owes accrued interest on the coupon from the last payment date to the credit event, as indicated in the center of Figure 19.5.

The structure of the coupon and up-front payments in CDS markets changed dramatically since the crisis of 2007–2009. Previously, the coupon of newly initiated CDS trades was set such that the up-front payment was zero. This convention made it difficult to unwind trades. To take a simple example, if counterparty A bought five-year protection from B at a spread of 200 basis points and the market moved immediately to 180 basis points, A could not unwind by selling protection to B at the new market level of 180 basis points—in that case A would still owe net cash flows of 20 basis points. Hence, to “tear up” the original trade, A would have to make a lump-sum payment to B that both accepted as representing the value of the

¹¹ The auctions of the Federal National Mortgage Association (FNMA) and Federal Home Loan Mortgage Corporation (FHLMC) debt in October 2008 also worked in the sense of settling CDS contracts in an orderly way. However, the price results were somewhat perverse. The auction prices of the senior debt issues were about 92 and 94 for FNMA and FHLMC, respectively, but the prices of the subordinated debt issues were 99 and 100!

20 basis-point change in market levels. The unwind is even more difficult to value, of course, if six months passed and the now 4.5 years of protection at 200 basis points has to be compared with 180 basis points of five-year protection.

Since the crisis, in response to the prompting of regulators, CDS contracts have become more standardized, particularly between dealers. First, most trading occurs in CDS with terms of approximately five and 10 years where exact maturity dates are limited to pseudo-International Money Market (IMM) dates,¹² i.e., the 20th days of March, June, September, or December. So, for example, all five-year contracts traded between June 21 and September 20, 2010, will mature on September 20, 2015; all five-year contracts traded between September 21 and December 20, 2010, will mature on December 20, 2015; etc. Second, contracts have been standardized by setting the coupon at either 100 or at 500 basis points annually, with an up-front amount adjusting to credit conditions as appropriate. This convention makes it particularly easy to unwind trades. If A bought protection from B for an up-front payment of 5 and a spread of 100 basis points and then the market moved to an up-front payment of 10 and a spread of 100 basis points, A would simply sell protection to the same maturity to B for 10 and no net position would remain.

The push for standardization of CDS and, more broadly, of derivatives, is part of a larger discussion about the clearing of derivatives. See the section “Regulatory and Legislative Mandates to Clear OTC Derivatives” in Chapter 16.

Quoting CDS Spreads and Calculating Up-Front Payments

For low-quality credits, CDS are quoted in terms of up-front payments. For high-quality credits, however, CDS are quoted in terms of spread from which up-front payments are calculated and subsequently paid. This spread represents what the market CDS coupon would be if there were no up-front payment. In this section, therefore, the terms spread and quoted coupon will be used interchangeably. Note that, for both low- and high-quality credits, a CDS quoted coupon is more intuitive than the combination of a standardized coupon and an up-front payment. For example, in the illustration to follow, it is more useful to say that €10 million of credit protection on Deutsche Bank can be bought at an annual cost of \$93,550 than to say that this protection costs \$100,000 per year after generating an initial up-front payment of \$22,272. Similarly, it is easier to compare two credits if protection for both are expressed as pure running costs. In any case, this section

¹² See the subsection “Eurodollar Futures” in Chapter 15 for a description of IMM dates.

describes the market convention for converting quoted coupon to up-front payments or *vice versa*.

The value of a CDS can be thought of as having two legs. The fee leg is the payment of the coupon until the earlier of the event of default or the maturity of the CDS. Note that this includes accrued coupon from a previous payment date to the time of default. The contingent leg is the payment of one minus the recovery rate per unit face amount, $1 - R$, in the event of default. To express the value of each of these legs mathematically, some notation needs to be set. Let $\bar{C}$ be the standardized coupon rate on the CDS, let $C(T)$ be the quoted coupon rate, and let $UF(T)$ be the up-front payment. Let $d(T)$ be the discount factor to time T . As in the previous section, let the hazard rate be λ and the cumulative default probability to time T , $CD(T)$. Then, denote the cumulative survival probability by $CS(T)$, which equals one minus the cumulative default probability, or, from (19.11),

$$CS(T) = e^{-\lambda T} \quad (19.13)$$

With this notation, the value of the fee leg, V^{Fee} , can be defined in terms of a quoted quarterly coupon.¹³ The CDS actually pays the standardized coupon $\bar{C}$, but for the purposes of quoting a coupon rate, the value of the fee leg is expressed in terms of this quoted rate:

$$V^{Fee} = \frac{C(T)}{4} \sum_{i=1}^{4T} CS(t_i) d(t_i) + \frac{1}{2} \frac{C(T)}{4} \sum_{i=1}^{4T} [CS(t_{i-1}) - CS(t_i)] d(t_i) \quad (19.14)$$

The first term of (19.14) is the expected value of the coupon payments. The quarterly coupon payment at time t_i is $\frac{C(T)}{4}$ and is made with probability $CS(t_i)$, i.e., the probability that the reference entity has not defaulted by time t_i . Hence, $\frac{C(T)}{4} CS(t_i)$ is the expected coupon payment at time t_i and $\frac{C(T)}{4} CS(t_i) d(t_i)$ is its discounted value. Summing across payment dates gives the discounted expected value of all coupon payments conditional on no default.

The second term of (19.14) is the expected value of the accrued coupon payments made at the time of a default. First, $CS(t_{i-1}) - CS(t_i)$, the difference between the probability of surviving to time t_{i-1} and the probability of surviving to time t_i , is equal to the probability of a default happening

¹³ For simplicity, equation (19.14) assumes that the length of accrual periods is exactly .25. More precise calculations are illustrated in the numerical example to follow.

between time t_{i-1} and time t_i . The accrued coupon payment in the event of a default during a period depends on when the default happens within that period. For simplicity, the convention is to assume that defaults happen in the middle of coupon periods. Hence, the accrued coupon payment in the event of a default between t_{i-1} and t_i is $\frac{1}{2} \frac{\mathbb{C}(T)}{4}$, the expected value of that payment is $\frac{1}{2} \frac{\mathbb{C}(T)}{4} \times [CS(t_{i-1}) - CS(t_i)]$, and the discounted value of that expectation is $\frac{1}{2} \frac{\mathbb{C}(T)}{4} \times [CS(t_{i-1}) - CS(t_i)] d(t_i)$. Summing across payment dates gives the total discounted expected value of accrued coupon payments conditional on defaults in each period.

The value of the contingent leg, V^{Cont} , is defined as follows:

$$V^{Cont} = (1 - R) \sum_{i=1}^{4T} [CS(t_{i-1}) - CS(t_i)] d(t_i) \quad (19.15)$$

The CDS pays $1 - R$ in the event of default. The probability of a default between t_{i-1} and t_i is $CS(t_{i-1}) - CS(t_i)$. Hence, the discounted expected payment of the contingent leg of the CDS between t_{i-1} and t_i is $(1 - R)$ times $[CS(t_{i-1}) - CS(t_i)]$ times $d(t_i)$. Summing across periods gives the total discounted expected value of contingent leg payments.

With these definitions of the value of the two legs, the CDS is fair if the value of the fee leg, received by the seller of protection and paid by the buyer of protection, equals the value of the contingent leg, received by the buyer of protection and paid by the seller of protection. To ensure that the CDS is fair in this framework, therefore, find the hazard rate λ^* such that (19.14) equals (19.15). As part of this convention, by the way, the recovery rate is usually set at 40%.

The link between the quoted coupon and the up-front fee—defined as paid to the seller of protection—can now be made through the following equation:

$$UF(T) = \frac{\mathbb{C}(T) - \bar{\mathbb{C}}}{4} \left[\sum_{i=1}^{4T} CS(t_i) d(t_i) + \frac{1}{2} \sum_{i=1}^{4T} [CS(t_{i-1}) - CS(t_i)] d(t_i) \right] \quad (19.16)$$

The right-hand side of (19.14) is the value of receiving the fee $\mathbb{C}(T)$. By inspection then, the right-hand side of (19.16) is the value of receiving a fee of $\mathbb{C}(T) - \bar{\mathbb{C}}$. But this value is exactly what the up-front payment represents, namely, the quantity that makes the seller of protection willing to accept the standardized coupon $\bar{\mathbb{C}}$ instead of the market quoted coupon $\mathbb{C}(T)$ (with no up-front payment).

Before proceeding to a numerical example, it is important to emphasize that all of the expressions in this subsection are conventions to quote up-front payments from quoted spreads or *vice versa*. In fact, the CDSW function in Bloomberg performs these calculations. The manner in which individual investors or traders determine the value of a CDS, however, need not resemble these expressions at all.

The following CDS illustrates the calculation of an up-front payment from a quoted spread. On March 21, 2011, an EUR denominated five-year CDS on Deutsche Bank, AG, with a coupon of 100 basis points, was quoted at a spread of 95.33 basis points. Table 19.7 details the calculations used in determining the up-front spread. Column (1) gives the schedule of payment dates. By the convention described in the previous section, this five-year

TABLE 19.7 Calculating the Up-Front Payment for the Five-Year 100 Basis Point Deutsche Bank CDS as of March 21, 2011

Hazard Rate 1.61091%					
(1)	(2)	(3)	(4)	(5)	(6)
Payment Date	Term (years)	Accrual Days	Discount Factor	Cumulative Survival Probability %	Period Default Probability %
6/20/11	.249315	92	.996859	99.5992	.4008
9/20/11	.501370	92	.992676	99.1956	.4036
12/20/11	.750685	91	.987929	98.7980	.3976
3/20/12	1.000000	91	.982593	98.4020	.3960
6/20/12	1.252055	92	.976739	98.0033	.3987
9/20/12	1.504110	92	.970392	97.6061	.3971
12/20/12	1.753425	91	.963769	97.2149	.3912
3/20/13	2.000000	90	.956923	96.8295	.3854
6/20/13	2.252055	92	.949771	96.4372	.3924
9/20/13	2.504110	92	.942489	96.0464	.3908
12/20/13	2.753425	91	.935213	95.6614	.3850
3/20/14	3.000000	90	.927875	95.2822	.3792
6/20/14	3.252055	92	.920277	94.8961	.3861
9/20/14	3.504110	92	.912542	94.5116	.3845
12/20/14	3.753425	91	.904811	94.1328	.3788
3/20/15	4.000000	90	.897056	93.7596	.3732
6/20/15	4.252055	92	.889084	93.3797	.3799
9/20/15	4.504110	92	.881015	93.0013	.3784
12/20/15	4.753425	91	.873001	92.6285	.3728
3/20/16	5.002740	91	.864934	92.2572	.3713
6/20/16	5.254795	92	.856794	91.8834	.3738

CDS matures on June 20, 2016, with scheduled quarterly payments starting from June 20, 2011. Column (2) gives the term of each payment, in years, from the settlement date of March 21, 2011. This column will be used to calculate the survival and default probabilities. Column (3) gives the actual number of days in each quarterly period. Since payments follow the actual/360 convention, each payment will equal the coupon of 1% times the appropriate accrual days divided by 360. Column (4) gives discount factors derived from the EUR swap curve for March 21, 2011. Column (5) gives the cumulative survival probabilities, to each payment date based on the hazard rate given in the first row, i.e., 1.61%. The derivation of this hazard rate will be described presently. Given this hazard rate, however, the cumulative survival probability is given by (19.13). Hence, for the payment on June 20, 2012, $CS(1.252055) = e^{-1.61091\% \times 1.252055}$, which is 98.0033%. Finally, column (6) gives the probability of a default over each of the payment periods. The probability of a default from settlement to June 20, 2011, is .4008%; from June 20, 2011, to September 20, 2011, is .4036%; etc. These are derived simply by subtracting sequential cumulative survival probabilities: the probability of a default from June 20, 2011, to September 20, 2011, is the probability of surviving to June 20, 2011, minus the probability of surviving to September 20, 2011. Mathematically, 99.5992% – 99.1956% is .4036%.

The hazard rate given in Table 19.7 is calculated so that the value of the fee leg of the CDS equals the value of the contingent payment leg. More specifically, with $T = 5.25$, $C(T) = .9533\%$, $R = 40\%$, and discount factors as given in the table, the hazard rate is found such that the resulting set of $CS(t_i)$ set V^{Fee} in (19.14) equal to V^{Cont} in (19.15), except that, in this applied example, actual accruals are used instead of the constant $\frac{1}{4}$ accrual term in those equations. At the resulting hazard rate of 1.61091%, the value of each leg of the CDS is worth 4.5464% of face amount. Note that almost all of the value of the fee leg in (19.14) comes from the first term of the right-hand side of that equation. The value of the coupon payments given survival is 4.5372% while the value of accrued coupon payments in the event of default is only .0092%.

Once the CDS-implied hazard rate has been found, the up-front payment can be solved using equation (19.16). The set of $CS(t_i)$ are taken from Table 19.7 and $\bar{C} = 1\%$. For €10 million notional of CDS, the up-front payment is €10 million times the result of (19.16), which turns out to be –€22,272. Since the market quoted spread of .9533% is below the standardized coupon of 1%, the buyer receives €22,272 for paying this 1% annually in exchange for contingent default benefits. Finally, upon entering into this CDS the buyer has to pay one day of accrued interest since the first accrual period of this standardized contract began on March 20, 2011. The amount of accrued interest is $\frac{1}{360} \times 1\% \times €10$ million or €278.

CDS-Bond Basis

A trader or investor wanting to buy exposure to a particular credit can do so by purchasing bonds issued by that credit or by selling protection on the corresponding reference entity. Similarly, a trader or investor wanting to sell exposure to a particular credit can short bonds or buy protection. The natural question to ask, therefore, is whether exposure to this particular credit is cheaper in the CDS market than in the bond market or *vice versa*. The generic name for the spread between the cost of protection in the CDS market and some equivalent measure implied from bond prices is the *CDS-bond basis*.

The framework of the previous section provides one methodology for computing a CDS-bond basis. Given the price of a bond, compute the constant hazard rate such that the discounted expected value of its cash flows equals its price. Then, with this bond-implied hazard rate, compute a CDS-equivalent bond spread, denoted $\mathbb{C}^{Bond}$, such that a CDS with that hazard rate would be fair. Mathematically, from equations (19.14) and (19.15), find $\mathbb{C}^{Bond}$ such that

$$\frac{\mathbb{C}^{Bond}}{4} = \frac{(1 - R) \sum_{i=1}^{4T} [CS(t_{i-1}) - CS(t_i)] d(t_i)}{\sum_{i=1}^{4T} CS(t_i) d(t_i) + \frac{1}{2} \sum_{i=1}^{4T} [CS(t_{i-1}) - CS(t_i)] d(t_i)} \quad (19.17)$$

As an illustration of this measure of basis, consider the Deutsche Bank floater that pays Euribor plus 50 basis points quarterly, matures on April 11, 2018, and, as of March 21, 2011, sold at a full price of 100.8335. Taking the floater's promised cash flows as EUR forward swap rates plus 50 basis points and proceeding along the lines of Table 19.7 implies a hazard rate of .62%, which is significantly below the 1.61% hazard rate implied by the CDS market at that time. To continue, however, use the bond-implied hazard rate of .62% and a recovery rate of 40% to solve that $\mathbb{C}^{Bond}$ in (19.17) is 36.5 basis points. Comparing this CDS-equivalent bond spread with the quoted CDS spread of 95.33 basis points indicates that, in this case, the bond is rich to the CDS: the implied cost of protection from the price of the bond is below the cost of protection in the CDS market. Hence, relying on this metric alone, an investor wanting to buy exposure to Deutsche Bank credit would sell protection in CDS while an investor wanting to sell exposure to Deutsche Bank would short the bond.

Another measure of CDS-bond basis that has been very popular, although more so before the standardization of CDS contracts, is the difference between the CDS quoted coupon (or the CDS coupon itself if the up-front payment is zero) and the par asset swap spread. For the purposes of explaining the appeal of this measure, assume that counterparties can still enter into a CDS at a market premium such that the up-front payment is zero. In that case, so long as a bond does not default, writing CDS

protection earns the CDS premium while a par asset swap, inclusive of financing, earns the par asset swap spread plus LIBOR minus the cost of financing. Neglecting the LIBOR-financing spread as small, though this is not always the case (see Chapter 15), one might conclude that the CDS premium should approximately equal the par asset swap spread and, therefore, that the difference between the CDS premium and the par asset swap spread is a good measure of relative value. In fact, some practitioners believe that there is an arbitrage relationship linking the two quantities. As it turns out, this is not usefully true. (See Appendix B in this chapter) And in any case, since the standardization of CDS contracts, it is no longer even practical to trade a CDS with no up-front payment.

The main difficulty with any of the commonly used measures of the CDS-bond basis, including the two introduced in this subsection, is that they do not account for a fundamental difference between a CDS position and a bond position. A CDS position does not require financing while a bond position does. For discussion, fix a horizon of five years, a typical CDS maturity. Maintaining a long bond position over that horizon requires committing the purchase price of the bond for five years or borrowing the purchase price through repo for five years. The former would result in a very high implicit or explicit cost of capital while the latter, even if it were possible to find a willing counterparty, would result in a very high borrowing rate. Basically, lenders are generally unwilling to commit funds for a long term when they might need those funds back in the interim, e.g., in times of financial stress. Similarly, maintaining a short bond position over a five-year horizon would require finding a counterparty willing to lend that bond long term and face the risk that the bond would be needed, perhaps to raise funds, at some interim time. This all implies that bonds can trade cheap to CDS (negative basis) when funding is expensive and rich to CDS (positive basis) when financing shorts is expensive. In fact, bonds did trade very cheap to CDS during the crisis of 2007–2009, when funding was particularly difficult, with the investment grade CDS-bond basis falling from near zero to negative 250 basis points.

Another difficulty in comparing CDS quoted coupons with bond spreads is the CDS delivery option. Because buyers of protection have this option, they are willing to pay higher quoted coupons than would otherwise be the case. A naive comparison of CDS and bond spreads, therefore, would erroneously conclude that exposure is cheaper through CDS than through bonds.

Example of a Negative Basis Trade

This subsection gives an example of a class of trades that have enjoyed popularity (when they are making money) and notoriety (when they are losing money), namely negative basis trades, which simultaneously buy a

bond and CDS protection. The flavor of the trade considered here is to profit from cash flows that are immunized to the event of default. Another flavor, not illustrated here, is to buy protection on more or less than the face amount of the bonds purchased so as to bet on the default outcome or on the recovery rate.

The KB Home $5\frac{3}{4}$ s of February 1, 2014, was introduced earlier in this chapter and reported to trade at a par asset swap spread of 9.23% as of November 10, 2008. On the same date, the full price of the bond was 68.66 and an investor could purchase CDS protection for 664 basis points and no up-front payment. The 259 basis-point difference between the par asset swap spread of 9.23% and the CDS spread of 664 basis points attracted interest in buying the bond at its relatively high asset swap spread and buying CDS protection at its relatively low cost, i.e., this spread attracted interest in the negative basis trade.

Table 19.8 shows the cash flows from 100 notional of this trade assuming a repo haircut of 50%, a repo rate of 2%, a capital cost rate of K , and a recovery rate R . The bond costs 68.66 and then makes coupon payments. If it matures it makes a final payment of 100 while, if it defaults, it is worth $100R$. Half of the purchase price is funded in repo and half with capital, each half at its own running cost. Finally, buying protection through the CDS costs a running 664 basis points, which, should the bond default, results in a payment of $100(1 - R)$.

According to Table 19.8, the trade pays 31.34 either at maturity or when the bond defaults. The annual interim cash flow totals $I(K) \equiv -1.5766 - 34.33K$. To simplify the illustration, assume that, conditional on no default, one quarter of the annual payment is made on February 1, 2009, and is then made annually until maturity on February 1, 2014. Hence, the worst outcome of the trade is for the bond not to default: the terminal payment is 31.34 whether or not there is a default but the negative, interim cash flows have to be made only until default.

To determine whether the trade is worthwhile, the investor might assume the worst case of no default and that interim payments can be financed

TABLE 19.8 KBH $5\frac{3}{4}$ s of February 1, 2014, Negative Basis Trade as of November 10, 2008

Position	Initiation	Interim	Maturity: No Default	Default
Bond	-68.66	5.75	100	100R
Repo	34.33	$-34.33 \times 2\% = -.6866$	-34.33	-34.33
Capital	34.33	$-34.33K$	-34.33	-34.33
CDS	0	-6.64	0	$100(1 - R)$
Total	0	$-1.5766 - 34.33K$	31.34	31.34

to maturity at the same cost of capital K . In that case, the breakeven cost of capital such that the trade is profitable is such that

$$31.34 + \frac{1}{4}I(K)(1+K)^5 + I(K)\sum_{j=0}^4(1+K)^j = 0 \quad (19.18)$$

Solving (19.18), the breakeven cost of capital is about 9.6%. Investors with a cost of capital below that can be sure of profiting from the negative basis, i.e., the cheapness of the bond relative to the cost of protection, provided that the trade can indeed be held until default or maturity. In particular, this caveat requires that both the repo and capital financing be maintained at the rates indicated no matter what happens to the mark-to-market or intermediate value of the trade and no matter what happens to the general level of interest rates. In other words, the financing of the bond matters to these trades and to assessing the difference between bond and CDS spreads.

The DV01 or Duration of a Bond with Credit Risk

The framework of this section can be used to account for credit risk when computing the interest rate risk of a bond. Yield-based DV01 and duration assume that all of a bond's cash flows will be paid on schedule. For a bond with a significant probability of default, however, there may very well be an early payment of whatever can be recovered from promised principal. Whatever this might mean for the bond's price, it would seemingly shorten the life of the bond and, therefore, decrease its interest rate sensitivity in a way not captured by yield-based metrics of interest rate risk.

A more appropriate way to calculate the bond's sensitivity to interest rates would be to use a hazard rate, perhaps implied by CDS markets, to compute the bond's price both before and after a shift of the benchmark rate curve. Along the lines of earlier subsections, the value of the bond would be the sum of its discounted expected cash flows given no default and its discounted expected recovery given default. More specifically, using the notation of the rest of this section and a bond coupon rate of c paid semiannually for T years, the bond value would be

$$\begin{aligned} & \frac{100c}{2} \sum_{i=1}^{2T} CS(t_i) d(t_i) + 100CS(t_{2T}) d(t_{2T}) \\ & + 100R \sum_{i=1}^{2T} [CS(t_{i-1}) - CS(t_i)] d(t_i) \end{aligned} \quad (19.19)$$

Shifting the benchmark rate curve to obtain a new discount function, calculating a new price, and then computing a *DV01* or duration would provide a credit-risk adjusted measure of interest rate risk.

As a simple example, assume that the benchmark rate curve is flat at 4%, and consider a 10-year bond with a coupon of 6% and a price of 57.62. This is clearly a very credit-impaired security. In any case, its yield can be calculated to be 14% and its yield-based *DV01* and duration to be .0365 and 6.34, respectively.

While a hazard rate might be available from CDS markets, it is also reasonable to calculate a hazard rate such that, for a recovery rate of 40%, expression (19.19) gives the bond's market price. In the present example, this hazard rate turns out to be 23%. Shifting the benchmark rate by one basis point, applying (19.19) to get a shifted bond price, and then computing sensitivities gives a *DV01* of .0228 and a duration of 3.95. Hence, for this bond, accounting for the probability that a default would result in an early, partial payment of principal reduces duration dramatically from 6.34 to 3.95.

Index CDS

A significant part of trading in credit markets is through CDS indexes, which are simply portfolios of CDS written on individual names. Through these products investors can take a broad and diversified exposure to credit risk. The two main indexes are iTraxx and CDX, which cover 125 investment grade names domiciled in Europe and North America, respectively. The buyer of an index receives a fee in exchange for offering protection while the seller of an index pays a fee to purchase protection. After a credit event with respect to one of the names in the index, the buyer of the index makes a compensation payment and the seller receives a compensation payment. The affected name is dropped from the index without replacement, which reduces the notional amount of the contract.

The most popular CDS indexes are very much standardized. New CDS indexes or index *series* are issued semiannually on pseudo-IMM dates with maturities of three, five, seven, and 10 years, although the five-year maturity is the most liquid. As an example of the issuance cycle, the protection from the five-year iTraxx series 14 started on September 20, 2010, and ends on December 20, 2015: it is the on-the-run series, i.e., the most recently issued, from September 20, 2010, to March 20, 2011, when a new series is issued. Therefore, the actual maturity of a five-year index is 63 months at issuance and 57 months when a newer series becomes the on-the-run. The indexes are also standardized with respect to coupon, like single-name CDS, along the lines of the first subsection in this section.




FIGURE 19.6 Cost of Protection of Five-Year CDS Indexes

Figure 19.6 shows the cost of protection for the on-the-run, five-year iTraxx Main and CDX IG indexes, two investment grade indexes, and for the iTraxx Crossover (XO) and CDX Crossover (XO) indexes, two indexes of lower quality. The unfolding of the crisis of 2007–2009 is clearly evident from the quoted cost of protection, with the spikes in the below-investment grade indexes particularly dramatic.

The 2007–2009 crisis also saw the inversion of the term structure of credit spreads. Figure 19.7 shows the difference between the iTraxx Main




FIGURE 19.7 Term Structure of Credit Spreads from iTraxx Main Series and Cost of Protection from Five-Year iTraxx Main

costs of 10-year and five-year protection, and of five-year and three-year protection, along with the absolute five-year cost of protection. The term structure of credits became downwardly sloping or inverted shortly after the turmoil of September 2008 and stayed inverted well into 2009.

APPENDIX A: CUMULATIVE DEFAULT RATES

Proposition: If the default rate is constant at λ , then the cumulative default probability to time t , $CD(t)$, is $1 - e^{-\lambda t}$.

Proof: Let $V(t)$ be the survival probability to time t , i.e., the probability of no default to time t . Then, the probability of no default to time $t + \Delta t$, i.e., $V(t + \Delta t)$, is the probability that there is no default to time t and that there is no default from then to time $t + \Delta t$. Mathematically,

$$V(t + \Delta t) = V(t) \times (1 - \lambda \Delta t) \quad (19.20)$$

Rearranging terms,

$$\lambda V(t) = -\frac{V(t + \Delta t) - V(t)}{\Delta t} \quad (19.21)$$

The limit of the right-hand side as Δt approaches zero is the derivative of $V(t)$, denoted $V'(t)$. Hence,

$$\lambda V(t) = -V'(t) \quad (19.22)$$

The solution to (19.22) is $V(t) = e^{-\lambda t}$. The cumulative default probability to time t is $1 - V(t)$ or $1 - e^{-\lambda t}$, as was to be proved.

APPENDIX B: CDS-BOND BASIS AS THE DIFFERENCE BETWEEN THE CDS SPREAD AND THE PAR ASSET SWAP SPREAD

The theoretical justification for concluding that the CDS spread should approximately equal the par asset swap spread is not strong. This appendix reviews the arbitrage arguments linking the two quantities. By showing how many strong assumptions have to be used to demonstrate the equivalence of the two quantities, this appendix effectively shows that the quantities need not be equivalent. In addition, as mentioned in the text, since the standardization of CDS coupons, trading the CDS spread, i.e., a coupon with no up-front payment, is no longer practical.

Consider the following trades:

- Buy a bond and pay fixed on a swap, as in the par asset swap trade of the asset swap section earlier in this chapter. So long as the bond does not default, this package earns LIBOR, L , plus s^{PAR} quarterly and 100 at maturity. If the bond does default, the asset swap subsection showed that the position is worth $100R + \widetilde{NPV}$, where R is the recovery rate and $\widetilde{NPV}$ is the NPV of the swap at the time of default. Recall too that $\widetilde{NPV}$ is $100 - P$ just after the initiation of the asset swap trade and zero at its termination.
- Raise 100 as required in the par asset swap trade of the asset swap subsection. Instead of raising all 100 in short-term repo, however, do the following:
 - Sell repo to borrow $100(1 - h)$ for a term equal to the remaining maturity of the bond at a fixed spread of ρ^* over LIBOR. The quantity h represents the haircut applied to repo borrowing. (See Chapter 12.)
 - Raise or use $100h$ of capital at a fixed spread k over LIBOR.
- Buy protection on 100 face amount of this particular bond though CDS at a spread of s^{CDS} (with no up-front payment).

Table 19.9 shows the net results of these trades. The total cash flows are zero at initiation and at maturity in the case of no default. The payoff in the case of default, however, depends on the NPV of the interest rate swap at that time. The discussion will now make the sequential assumptions necessary to draw the conclusion that the CDS spread equals the asset swap spread.

1. To make the arbitrage argument, the payoff in case of default has to equal zero. But this cannot be true since $\widetilde{NPV}$ trends from $100 - P$ to zero over time, fluctuating with interest rates. Neglecting this payoff somehow, perhaps by considering only par bonds and positing small changes in rates, the total value of the trades in Table 19.9 is zero at

TABLE 19.9 Arbitrage Pricing of CDS Spread

Position	Initiation	Interim	Maturity: No Default	Default
Floater	-100	$100(L + s^{PAR})$	100	$100R + \widetilde{NPV}$
Repo	$100(1 - h)$	$-100(1 - h)(L + \rho^*)$	$-100(1 - h)$	$-100(1 - h)$
Capital	$100h$	$-100h(L + k)$	$-100h$	$-100h$
CDS	0	$-100s^{CDS}$	0	$100(1 - R)$
Total	0	$100(s^{PAR} - [hk + (1 - h)\rho^*] - s^{CDS})$	0	$\widetilde{NPV}$

initiation, at maturity, and in the case of default. Therefore, by arbitrage, the sum of the interim cash flows must be zero as well. This leads to the following expression for the CDS-bond basis:

$$s^{CDS} - s^{PAR} = -[hk + (1 - h)\rho^*] \quad (19.23)$$

In words, equation (19.23) says that the CDS-bond basis is the negative of the weighted average cost of financing over LIBOR.

2. Add the unlikely assumptions that the asset swap position can be financed fully in repo, i.e., $h = 0$, and that a proxy for the term repo spread, ρ^* , is simply the current spread between the short-term repo rate r and LIBOR, i.e., $\rho^* = r - L$. Substituting these values of h and ρ^* into (19.23) gives another expression for the basis:

$$s^{CDS} - s^{PAR} = L - r \quad (19.24)$$

In words, equation (19.24) says that the basis is the LIBOR-repo spread. In fact, this special case motivates some practitioners to define the basis not as the difference between the CDS and par asset swap spreads, but as that difference minus the LIBOR-repo spread.

3. Finally, add the assumption that the bond finances at LIBOR, i.e., $r = L$. Substituting that condition into (19.24),

$$s^{CDS} - s^{PAR} = 0 \quad (19.25)$$

Mortgages and Mortgage-Backed Securities

The Overview introduced and highlighted the importance and size of the mortgage market in the United States. This chapter describes mortgage loans and mortgage-backed securities (MBS), presents the most popular methods used for valuation and hedging, and illustrates how prices behave as a function of the relevant variables.

MORTGAGE LOANS

Mortgage loans come in many different varieties. They can carry fixed or variable rates of interest and they can be extended for residential or commercial purposes. This chapter will focus almost exclusively on fixed rate residential mortgages. Residential mortgages typically mature in 15 or 30 years and constitute 80% of the total principal of securitized mortgages in the United States.

Given the importance of the securitization process, which will be discussed ahead, residential loans are typically classified by how they might be subsequently securitized. *Agency* or *conforming* loans are eligible to be securitized by such entities as Federal National Mortgage Association (FNMA), Federal Home Loan Mortgage Corporation (FHLMC), or Government National Mortgage Association (GNMA). The exact criteria vary by program, but these loans are relatively creditworthy¹ and limited in principal amount.

Non-agency or non-conforming loans have to be part of *private-label* securitizations. The relevant loan types include *jumbos*, which are larger in notional than conforming loans but otherwise similar; *Alt-A*, which deviate from conforming loans in one requirement; and *subprime*, which deviate

¹Typical criteria would be a Fair Isaac Corporation (FICO) score greater than 660, a loan-to-value ratio of less than 80%, and full documentation of three years of income. FICO scores and loan-to-value ratios are described in subsequent footnotes.

from conforming loans in several dimensions. About 80% of subprime loans are *adjustable-rate mortgages* (ARMs).

Given the role of subprime mortgages at the start of the 2007–2009 financial crisis, some further comment is in order. Borrowing and lending in the subprime market revolved around the following strategy. A relatively low-credit borrower would take out an ARM that carried a particularly low initial rate, called a *teaser*, which would reset higher after two or three years. In that time, however, should the credit of the borrower improve or should housing prices increase, the borrower would be able to pay off that first mortgage and borrow through a subsequent mortgage at a fixed rate that would have been unattainable at the start. This strategy worked well until the peak of housing prices in 2006. In fact, most subprime mortgage originations occurred between 2004 and 2006. In any case, the subsequent decline in housing prices and the resetting of ARMs to higher rates led to a significant number of defaults: by May 2008 the delinquency rate for ARMs reached 25%. The resulting foreclosures put further downward pressure on housing prices. By September 2008, the average home price had declined 20% from its 2006 peak. By September 2009, about 14.4% of all U.S. mortgages were either delinquent or in foreclosure, and, in 2009–2010, between 4% and 5% of the total number of mortgages ended in repossessions. Finally, by September 2010, principal balance exceeded home price for 23% of mortgages outstanding, with the percentages in the worst-performing real estate markets even worse (e.g., California at 32.8% and Florida at 46.4%).²

Fixed Rate Mortgage Payments

The most typical mortgage loan is a fixed rate, *level payment* mortgage. A homeowner might borrow \$100,000 from a bank at 4% and agree to make payments of \$477.42 every month for 30 years. The mortgage rate and the monthly payment are related by the following equation:

$$\$477.42 \sum_{n=1}^{360} \frac{1}{\left(1 + \frac{.04}{12}\right)^n} = \$100,000 \quad (20.1)$$

In words, the mortgage loan is fair in the sense that the present value of the monthly mortgage payments, discounted at the monthly compounded mortgage rate, equals the original amount borrowed. In general, for a monthly

²Source: Wells Fargo.

payment X on a T -year mortgage with a mortgage rate y and an original principal amount or loan balance of $B(0)$,

$$X \sum_{n=1}^{12T} \frac{1}{\left(1 + \frac{y}{12}\right)^n} = B(0)$$

$$X \frac{12}{y} \left[1 - \frac{1}{\left(1 + \frac{y}{12}\right)^{12T}} \right] = B(0) \quad (20.2)$$

which can be solved for X given y directly or y given X numerically as needed. Note that the second line of (20.2) uses the summation formula in Appendix D in Chapter 2.

The fixed monthly payment is often divided into its interest and principal components, a division interesting in its own right as well as for tax purposes; mortgage interest payments are deductible from income tax while principal payments are not. Letting $B(n)$ be the principal amount outstanding after the mortgage payment due on date n , the interest component on the payment on date $n + 1$ is

$$B(n) \times \frac{y}{12} \quad (20.3)$$

In words, the monthly interest payment over a particular period equals the mortgage rate times the principal outstanding at the beginning of that period. The principal component of the monthly payment is the remainder, that is,

$$X - B(n) \times \frac{y}{12} \quad (20.4)$$

In the example, the original balance is \$100,000. At the end of the first month, interest at 4% is due on this balance, which comes to $\$100,000 \times \frac{.04}{12}$ or \$333.33. The rest of the monthly payment, $\$477.42 - \$333.33$ or \$144.08, is payment of principal. This \$144.08 principal payment reduces the outstanding balance from the original \$100,000 to $\$100,000 - \$144.08$ or \$99,855.92 at the end of the first month. Then, the interest payment due at the end of the second month is based on the principal amount outstanding at the end of the first month, etc. Continuing in this way produces an *amortization table*, the first few rows of which are given in Table 20.1.

Figure 20.1 graphs the interest and principal components from the full amortization table of this mortgage. The height of each bar is the full monthly payment of \$477.42, the darkly shaded height is the interest component, and the lightly shaded height is the principal component. Early payments are composed mostly of interest while later payments are composed mostly of principal. This is explained by the phrase “interest lives off

TABLE 20.1 First Rows of an Amortization Table, in Dollars, of a 100,000 Dollar 4% 30-Year Mortgage

Payment Month	Interest Payment	Principal Payment	Ending Balance
			100,000.00
1	333.33	144.08	99,855.92
2	332.85	144.56	99,711.36
3	332.37	145.04	99,566.31
4	331.89	145.53	99,420.78
5	331.40	146.01	99,274.77

principal.” Interest at any time is due only on the then outstanding principal amount. As principal is paid off, the amount of interest necessarily declines.

While the outstanding balance of a mortgage on any date can be computed through an amortization table, there is an instructive shortcut. Discounting using the mortgage rate at origination, the present value of the remaining payments equals the principal outstanding. This is a fair pricing condition under the assumptions that the term structure is flat and that interest rates have not changed since the origination of the mortgage.

To illustrate this shortcut in this example, after 5 years or 60 monthly payments there remain 300 payments. The present value of these payments at the mortgage rate of 4% is

$$\begin{aligned} \$477.42 \sum_{n=1}^{300} \frac{1}{\left(1 + \frac{.04}{12}\right)^n} &= \$477.42 \frac{12}{.04} \left[1 - \frac{1}{\left(1 + \frac{.04}{12}\right)^{300}} \right] \\ &= \$90,448 \end{aligned} \tag{20.5}$$




FIGURE 20.1 Amortization of a \$100,000 4% 30-Year Mortgage

Hence, the scheduled principal amount outstanding after five years is also \$90,448.

This section describes the market convention of calculating the mortgage payment from a single mortgage rate or *vice versa*. This in no way contradicts the fact that the market values mortgages using an appropriate term structure of rates and spreads.

If rates or spreads rise after origination, the present value of the remaining mortgage payments will be worth less than the outstanding principal amount while, if rates fall, this present value will exceed the outstanding principal amount. The value of a mortgage, however, is not simply the present value of its payments because of the borrower's prepayment option, which is introduced in the next subsection.

The Prepayment Option

Mortgage borrowers have a *prepayment option*, that is, the option to pay the lender the outstanding principal at any time and be freed of the obligation to make further payments. In the example of the previous subsection, the mortgage balance at the end of five years is \$90,448. At that time, therefore, the borrower can pay the lender this balance and no longer have to make monthly payments.

The prepayment option is valuable when mortgage rates have fallen. In that case, as mentioned previously, the present value of the remaining monthly payments exceeds the principal outstanding. Therefore, the borrower gains in present value from paying the principal outstanding in exchange for not having to make further payments. When rates have risen, however, the present value of the remaining payments is less than the principal outstanding and prepayment would result in a loss of present value. By this logic, the prepayment option is an American call option on an otherwise identical, (fictional) nonprepayable mortgage. The strike of the option is the principal amount outstanding and, therefore, changes after every payment.

When pricing the embedded options in bonds issued by government agencies or corporations (see Chapter 18), it is reasonable to assume that a relatively efficient call policy will prevail. In terms of a term structure model, an efficient call policy means that an issuer will exercise a call option if and only if the value of immediately exercising the option exceeds the value of holding the option. If the mortgage borrowers faced as simple an optimization problem, so that their prepayments were as easily predictable, mortgages could be valued along the lines of Part Three of this book. However, prepayments of mortgages turn out to be much more difficult to model, which is discussed later in this chapter.

While the prepayment option refers to the choice borrowers can make to return outstanding principal, the term prepayment refers to any return of principal above the amount scheduled to be returned by the amortization table. When a mortgage borrower sells a property, for example, the principal

becomes due no matter what the level of interest rates. Hence, to value mortgages, prepayment models have to consider all forms of prepayments.

MORTGAGE-BACKED SECURITIES

Until the 1970s banks made mortgage loans and held them until maturity, collecting principal and interest payments until the mortgages were repaid. The primary market was the only mortgage market. During the 1970s, the *securitization* of mortgages began. The growth of this secondary market substantially changed the mortgage business. Banks that might have had to restrict mortgage lending, either because of limited capital or risk appetite, could now continue to make mortgage loans since these loans could be quickly and efficiently sold. At the same time, investors gained a new security type through which to lend their surplus funds. Of course, one of the policy questions raised by the 2007–2009 financial crisis was whether the mortgage securitization process, for any of several reasons, had created too much systemic risk.

Issuers of MBS gather mortgage loans into *pools* and then sell claims on those pools to investors. In the simplest structure, a *mortgage pass-through*, the cash flows from the underlying mortgages, that is, interest, scheduled principal, and prepayments, are passed from the borrowers to the investors with some short processing delay. Mortgage *servicers* manage the flow of cash from borrowers to investors in exchange for a fee taken from those cash flows. *Mortgage guarantors* guarantee investors the payment of interest and principal against borrower defaults, also in exchange for a fee. When a borrower does default, the guarantor compensates the pool with a lump-sum payment and then, through the servicer, pursues the borrower and the underlying property to recover as much of the amount paid as possible. By the way, in comparison with U.S. lenders, European lenders have easier recourse to borrower assets that are not part of the mortgaged property.

The Overview reported that U.S. mortgage debt was a little over \$14 trillion in 2010. Of this total, \$7.5 trillion had been securitized. This securitized amount is further subdivided into \$5.4 trillion of *agency securities*, i.e., securities guaranteed or issued by such entities as GNMA, FNMA, and FHLMC, and the remainder *private-label* securities issued by private financial institutions. These amounts outstanding are misleading, however, with respect to new issuance. Since the 2007–2009 crisis to the time of this writing, agency securities comprised almost all of new MBS issuance.

Mortgage Pools

Loans that are collected into a pool are usually similar with respect to loan type, mortgage rate, and date of origination. Table 20.2 gives some summary statistics, both at origination and as of December 2010, of a pool

TABLE 20.2 Summary Statistics for FNMA Pool FG A47828, 3.5% 2004 Vintage at Origination and as of December 2010

	Original	Dec 2010
Number of Loans	91	69
Principal Amount	\$13,635,953	\$9,326,596
WAC	3.940%	3.928%
WAM (months)	335	271

Source: Bloomberg.

of 30-year loans issued by FNMA in January 2005 of loans originated in 2004, i.e., of the 2004 “vintage.” The coupon of the pool, that is, the rate paid to investors, is 3.5%. According to the table, the pool was issued with 91 loans and a total principal amount of about \$13.6 million. The table next reports two weighted averages, where the weighting is based on loan size. The *weighted-average coupon* or WAC is the weighted average of the mortgage rates of the loans and was 3.94% at issuance. Note that, as a weighted average of loan rates, the term WAC somewhat confusingly uses the word “coupon.” It is best to think of there being only one “coupon” rate, namely the interest rate on the pool as a whole that is passed on to investors. In any case, returning to the pool of Table 20.2, note that the 3.5% coupon is less than the 3.94% original WAC: the difference between what the borrowers pay and what the investors receive is paid to the servicer and to the guarantor. Finally, the *weighted-average maturity* (WAM) of the loans was 335 months. This original WAM on a pool of “30-year” loans means that some of the loans were slightly seasoned (i.e., had been outstanding for some amount of time) when the pool was issued.

The summary statistics of the FNMA 3.5% 2004 pool as of December 2010 show that a significant fraction of the pool has paid down. The pool’s *factor* is the ratio of the current to the original principal amount outstanding, which in this case is about 68%. A good deal of this is due to prepayments rather than scheduled amortization. First, although the principal amount of each loan is not provided here, only 69 of the original 91 loans are still in the pool. Second, for an order of magnitude calculation, equation (20.5) calculated that the scheduled principal outstanding of a 4% 30-year loan after five years is a little over 90% of the original principal amount. The WAC here is slightly less than 4% and the pool is not exactly five years old, but the factor of 68% is significantly below 90%. Note that the WAC of the pool has fallen very slightly since origination, indicating that prepaying loans had slightly higher rates than loans remaining in the pool. Finally, the WAM has fallen by 64 months or a bit over five years, indicative mostly of the loans aging five years from issuance at the end of 2004 to December 2010.

TABLE 20.3 Agency Pool Issuance, in Billions of Dollars

	2010						Full Year		
	Dec*	Nov	Oct	Sep	Aug	Jul	2010*	2009	2008
Total	72	146	143	141	111	107	1,312	1,725	1,153
Issuer									
FHLMC	21.6	38.6	37.4	36.4	28.4	26.6	351	462	341
FNMA	42.0	73.5	69.8	70.6	48.0	42.9	586	806	541
GNMA1	8.0	14.9	16.0	13.4	12.2	13.4	156	288	146
GNMA2	.6	19.3	19.5	20.3	22.0	24.5	219	169	125
Loan Type									
30-Year	56.1	103	100	103	79.2	79.6	973	1,449	951
15-Year	11.9	25.3	24.7	21.7	15.3	13.4	187	181	93
ARM	1.6	7.1	6.8	4.9	6.8	8.0	67	33	78
Other	2.5	10.9	11.2	10.9	9.4	6.4	85	62	32
Coupon									
<4%	8.4	12.4	11.0	4.6	1.2	.5	40	3	0
4%–	35.6	63.1	59.6	51.0	16.0	7.8	250	211	0
4.5%–	9.2	20.9	24.0	39.7	45.7	46.0	428	715	18
5%–	2.2	5.1	4.5	6.5	14.6	23.5	233	375	201
>5%	.7	1.3	1.0	1.4	1.8	1.7	23	145	731

*To Dec 10.

Source: Bloomberg.

While coupon and age are the most important characteristics of loans and pools with respect to pricing, other characteristics are important as well, as will be discussed further in the section on modeling prepayments. As a result, issuers of MBS provide pool summary statistics on characteristics other than those listed in Table 20.2. Examples include FICO scores,³ loan-to-value (LTV) ratios,⁴ and the geographical distribution of the loans. For the FNMA 3.5% 2004 pool, it happens that 100% of the loans are in New Jersey.

Table 20.3 shows the issuance volumes of agency pools for the full years 2008, 2009, and 2010, along with monthly issuance for the second half of

³FICO scores, a product of Fair Isaac Corporation, measure a borrower's ability to pay based on credit history. The scores range from 300 to 850, with a score above 650 considered creditworthy by many lenders.

⁴The LTV ratio is the principal amount of the loan divided by the value of the mortgaged property.

2010. These volumes are also broken down by issuer, loan type, and coupon. Total issuance fell dramatically in 2010 relative to 2009, reflecting lower volumes of real estate transactions. Furthermore, the increase from 2008 to 2009 is in part due to the shift from private label to agency issuance mentioned earlier. The issuer breakdown reveals that FNMA is the largest issuer, and the breakdown by loan type reveals the dominance of the 30-year mortgage. Mortgage loans, and therefore pools, are issued at prevailing market rates, that is the rates that make them sell for approximately par. Thus, the shift of dominant volume from the >5% bucket in 2008, to the 4.5%–5% bucket in 2009, to the 4% bucket in September 2010, simply reflects the fall in mortgage rates, and interest rates generally, over this time period.

Calculating Prepayment Rates for Pools

In any given month, some loans in a pool will prepay completely, some will not prepay at all, and some—usually a small number—may *curtail*, i.e., partially prepay. For the purposes of valuation it is conventional to measure the principal amount prepaying as a percentage of the total principal outstanding. The *single monthly mortality rate* at month n , denoted SMM_n , is the percentage of principal outstanding at the beginning of month n that is prepaid during month n , where prepayments do not include scheduled, i.e., amortizing, principal amounts. The SMM is often annualized to a *constant prepayment rate* or *conditional prepayment rate (CPR)*. A pool that prepays at a constant rate equal to SMM_n has $1 - SMM_n$ of the principal remaining at the end of one month, $(1 - SMM_n)^{12}$ remaining at the end of 12 months, and, therefore, $1 - (1 - SMM_n)^{12}$ principal prepaying over those 12 months. Hence, the annualized CPR is related to SMM as follows:

$$CPR_n = 1 - (1 - SMM_n)^{12} \quad (20.6)$$

For example, if a pool prepaid .5% of its principal above its amortizing principal in a given month, it would be prepaying that month at a CPR of about 5.8%. Note that a pool has a CPR every month even though CPR is an annualized rate.

Specific Pools and TBAs

Agency mortgage pools trade in two forms: *specified pools* and TBAs. The latter is an acronym for *To Be Announced* and only the acronym is used by practitioners.

In the specified pools market, buyers and sellers agree to trade a particular pool of loans. Consequently, the price of a trade reflects the characteristics of the particular pool. For example, the next section of this chapter will argue that pools with relatively high loan balances are worth less to investors

TABLE 20.4 Bid Prices for Selected FNMA 30-Year TBAs as of December 10, 2010. Fractional prices are in 32nds; a “+” is half a 32nd or a 64th

	4%	4.5%	5%
Jan	98 – 30+	101 – 31+	104 – 15+
Feb	98 – 21	101 – 22	104 – 09
Mar	98 – 10+	101 – 12	104 – 01

Source: Bloomberg.

because these pools make relatively better use of their prepayment options. Therefore, in the specified pools market, relatively high loan-balance pools will trade for relatively low prices.

Much more liquid, however, is the TBA market, which is a forward market with a delivery option. Table 20.4 gives bid prices for selected FNMA 30-year TBAs as of December 10, 2010. Consider a trade on that date of \$100 million face amount of the FNMA 5% 30-year TBA for February delivery at a price of 104-09. Come February the seller chooses a 30-year 5% FNMA pool and delivers \$100 million face amount of that pool to the buyer for 104-09. Just as in the case of the delivery option in note and bond futures (see Chapter 14), the TBA seller will pick the cheapest-to-deliver (CTD) pool, that is, the pool that is worth the least subject to the issuer, maturity, and coupon requirements. For example, following up on the remark in the previous paragraph that pools with high loan balances are less valuable than other pools, the TBA seller might wind up delivering a pool with particularly high loan balances. In any case, *ex-ante*, TBA prices will reflect the fact that the CTD pools will be delivered. In fact, specified pools trade at a reference TBA price plus a *pay-up* that depends on the specified pools' characteristics *versus* those of the pools likely to be delivered.

As the TBA market is so liquid, especially the front contracts that trade near par, there is particular focus in the broader mortgage market on the contract that trades closest to, but below par. This contract is called the *current contract* and its coupon the *current coupon*. In Table 20.4, since the prices of the 4% and 4.5% January TBAs bracket par, 4% would be the current coupon. Furthermore, the term *current mortgage rate* is sometimes used to refer to the interpolated coupon at which a front TBA would sell for par.⁵ Using the prices in Table 20.4 for this purpose, the current mortgage rate would be about 4.17%.

While the TBA market is much more liquid than the specified pools market, the latter has grown rapidly in recent years. First, episodes in which

⁵The term “current mortgage rate” is also used to refer to the rate borrowers pay on newly originated mortgages.

the delivery option was particularly valuable have made traders and investors increasingly aware of the risks posed by the delivery option. Second, agencies have been supplying increasing amounts of granular data about the characteristics of loans in pools, which allows for more effective specified pools trading.

Dollar Rolls

Consider an investor who has just purchased a mortgage pool but wants to finance that purchase over the next month. One alternative is an MBS repo. Along the lines of Chapter 12, the investor could sell the repo, i.e., sell the pool today while simultaneously agreeing to repurchase it after a month. This trade has the same economics as a secured loan: the investor effectively borrows cash today by posting the pool as collateral, and, upon paying back the loan with interest after a month, retrieves the collateral.

An alternative for financing mortgages is the *dollar roll*. The *buyer of the roll* sells a TBA for one settlement month and buys the same TBA for the following settlement month. For example, the investor who just purchased a 30-year 4% FNMA pool might sell the FNMA 30-year 4% January TBA and buy the FNMA 30-year 4% February TBA. Delivering the pool just purchased through the sale of the January TBA, which raises cash, and purchasing a pool through the February TBA, which returns cash, is very close to the economics of a secured loan. There are, however, two important differences between dollar roll and repo financing.

First, the buyer of the roll may not get back in the later month the same pool delivered in the earlier month. In the example, the buyer of the Jan/Feb roll delivers a particular pool in January but will have to accept whatever eligible pool is delivered in February. By contrast, an MBS repo seller is always returned the same pool that was originally posted as collateral.

Second, the buyer of the roll does not receive any interest or principal payments from the pool over the roll. In the example, the buyer of the Jan/Feb roll, who delivers the pool in January, does not receive the January payments of interest and principal.⁶ By contrast, as described in Chapter 12, a repo seller receives any payments of interest and principal over the life of the repo. While the prices of TBA contracts reflect the timing of payments, so that the buyer of a roll does not, in any sense, lose a month of payments relative to a repo seller, the risks of the two transactions are different. The buyer of a roll does not have any exposure to prepayments over the month being higher or lower than what had been implied by TBA prices while the repo seller does.

⁶The record date for MBS is usually the last day of the month while pools delivered through TBA settle on the 15th or 25th of the month depending on the underlying issuer.

Chapter 13 presented the forward drop, the difference between a spot and forward price. The forward price is usually below the spot price because buying a security forward sacrifices the relatively high rate of interest earned on the security in exchange for the relatively low, short-term rate of interest earned by investing the funds that would have gone into the spot purchase. Put another way, the forward price is determined such that investors are indifferent between buying a security forward and buying it spot. In an important sense, the same reasoning applies to TBA prices and the roll: prices of pools for later delivery tend to be lower because pools earn a higher rate of interest than the short-term rate. Note how this rule characterizes the prices in Table 20.4. Once again, however, the TBA delivery option complicates the analysis. Consider the Jan/Feb roll as of January. If the delivery option had no value, the forward price for February would be determined along the lines of Chapter 13 and investors would be indifferent between: 1) buying the pool and the roll, which is essentially buying a pool forward for February delivery; and 2) buying a pool and holding it from January to February. But if the delivery option has value, the February TBA price would be lower and the forward drop would be larger than it would be otherwise.

In market jargon, the *value of the roll* is the difference in proceeds between 1) starting with a given pool and buying the roll and 2) holding that pool over the month. If the value of the roll is zero, as it would be if the forward pricing methodology of Chapter 13 applied, the roll is said to trade at *breakeven*. If the forward drop is larger so that the value of the roll is positive, the roll is said to trade *above carry*. Given the delivery option of TBAs, the roll would be expected to trade somewhat above carry without necessarily implying a value opportunity.

To make the roll more concrete, consider the following example. Suppose that the TBA prices of the Fannie Mae 5% for July 12 and August 12 settlements are \$102.50 and \$102.15, respectively. The accrued interest to be added to each of these prices is 12 actual/360 days of a month's worth of a 5% coupon, i.e., $100 \times (12/30) \times 5\%/12$ or .167. Let the expected total principal paydown, that is, scheduled principal plus prepayments, be 2% of outstanding balance and let the appropriate short-term rate be 1%.

If an investor rolls a balance of \$10 million, proceeds from selling the July TBA are $\$10\text{mm} \times (102.50 + .167)/100$ or \$10,266,700. Investing these proceeds to August 12 at 1% earns interest of $\$10,266,700 \times (31/360) \times 1\%$ or \$8,841. Then, purchasing the August TBA, which has experienced a 2% principal paydown, costs $\$10\text{mm} \times (1 - 2\%) \times (102.15 + .167)/100$ or \$10,027,066. The net proceeds from the role, therefore, are $\$10,266,700 + \$8,841 - \$10,027,066$ or \$248,475.

If the investor does not roll, the net proceeds are the coupon plus principal paydown, i.e., $\$10\text{mm} \times (5\%/12 + 2\%)$ or \$241,667.

In conclusion, then, the roll is trading above carry in this example, with the value of the roll at $\$248,475 - \$241,667$ or \$6,808.

Other Products

This chapter focuses on pass-through MBS, but a few other products will also be mentioned.

The properties of pass-through securities do not suit the needs of all investors. In an effort to broaden the appeal of MBS, practitioners have carved up pools of mortgages into different derivatives. One example is *planned amortization class (PAC)* bonds, which are a type of *collateralized mortgage obligation (CMO)*. A PAC bond is created by setting some fixed prepayment schedule and promising that the PAC bond will receive interest and principal according to that schedule so long as the actual prepayments from the underlying mortgage pools are not exceptionally large or small. In order to fulfill this promise, other derivative securities, called *companion* or *support* bonds, absorb the prepayment uncertainty. If prepayments are relatively high and PAC bonds receive their promised principal payments, then the companion bonds must receive relatively large prepayments. Alternatively, if prepayments are relatively low and PAC bonds receive the promised principal payments, then the companion bonds must receive relatively few prepayments. The point of this structure is that investors who do not like prepayment uncertainty can participate in the mortgage market through PACs. Dealers and investors who are comfortable with modeling prepayments and with controlling the accompanying interest rate risk can buy the companion or support bonds.

Other popular mortgage derivatives are *interest-only (IO)* and *principal-only (PO)* strips. The cash flows from a pool of mortgages are divided such that the IO gets all the interest payments while the PO gets all the principal payments. The unusual price rate behavior of these mortgage derivatives is illustrated later in this chapter.

Constant maturity mortgage (CMM) products allow investors to trade mortgage rates directly as a convexity-free alternative to trading prices of MBS that depend on mortgage rates. A CMM index is constructed from 30-year TBA prices to be the hypothetical coupon on a TBA for settlement in 30 days that trades at par. Market participants trade CMM mostly through Forward Rate Agreements (FRAs) (see Chapter 15).

Mortgage options are calls and puts on TBAs. The most liquid options are written on TBAs with delivery dates in the next three months.

PREPAYMENT MODELING

Earlier in this chapter it was noted that prepayment option is not as simply modeled as are the contingent claims priced by the methods of Part Three. Part of the reason for this is that some sources of prepayments are not determined exclusively or even predominantly by interest rates, e.g., selling

a home to buy a bigger or smaller one, divorce, default, and natural disasters that destroy a property. Another reason is that the cost of focusing on the prepayment problem, of figuring out the best action to take, and of navigating the process through financial institutions can be quite large. In any case, just because prepayments cannot be predicted by a simple optimization model does not mean that they are suboptimal from the point of view of mortgage borrowers. In any case, with the optimization problem across borrowers so difficult to specify, prepayment modeling relies heavily on empirical estimation of observed behavior.

A prepayment model uses loan characteristics and the economic environment (i.e., interest rates and sometimes housing prices) to predict prepayments. The most common practice identifies four components of prepayments, namely, in order of importance, *refinancing*, *turnover*, *defaults*, and *curtailments*. These components are typically modeled separately and their parameters estimated or calibrated so as to approximate available historical data.

Refinancing

In a refinancing a borrower pays off the principal of an existing mortgage with the proceeds of a new one. One major motive of refinancing is to reduce cost. A refinancing saves the borrower money if the rate on an available new mortgage has declined sufficiently relative to the rate on the existing mortgage and the transaction costs of refinancing. The most likely reason for a decline in the mortgage rate is that the general level of interest rates has declined. But there are other reasons as well: the spread of mortgage rates over benchmark rates has declined; the borrower's credit rating has improved; or the value of the mortgaged property has increased. Another important motive of refinancing is to extract home equity. If a property value has increased, a borrower might take out a new mortgage with a higher balance than that on the existing mortgage so as to pay off that existing mortgage and have cash remaining for other purposes. This is known as a *cash-out refinancing* and was used extensively in the run-up to the 2007–2009 crisis.

Modeling the refinancing component of prepayments often starts with an *incentive function* for a pool or group of loans in a pool and then defines prepayments due to refinancing as a nondecreasing function of that incentive. A simple example of an incentive might be

$$I = (WAC - R) \times WALs \times A - K \quad (20.7)$$

where WAC is the weighted average coupon of the pool, R is the current mortgage rate available to borrowers,⁷ $WALS$ is the weighted-average loan

⁷The Primary Mortgage Market Survey Rate, published weekly by FHLMC, is often used to represent the mortgage rate available to borrowers for conforming loans.

size of the pool, A is an annuity factor that gives the present value of an annual dollar payment from the average loan (i.e., from a loan with a remaining maturity equal to the average maturity of the loans being modeled), and K is an estimate of the fixed cost of refinancing. The current mortgage rate is actually lagged by a month or two in an incentive function to reflect lags in initiating and processing a refinancing application.

The logic of the incentive function (20.7) is that it estimates the present value of the dollar gains to the borrower from refinancing. Refinancing reduces the mortgage rate by $WAC - R$ on a principal amount of $WALS$. Then, to get the present value of this reduction, multiply by the appropriate annuity factor. Lastly, subtract the fixed cost of refinancing to get the net present value of refinancing. This theoretical argument in support of having incentive increase with loan size is quite persuasive, but the proposition is supported by empirical evidence as well. Average loan balances decline as pools age, indicating that loans with higher balances prepay more quickly. For orders of magnitude, average loan balances in newly issued agency pools are typically larger than \$175,000 but can be smaller than \$80,000 for older pools.

Having specified the incentive, prepayments, measured in terms of CPR , are typically modeled as an S -curve of that incentive. One example of such a function is

$$CPR(I) = T + \frac{1}{a + e^{-bI}} \quad (20.8)$$

where T is turnover, discussed in the next subsection, and a and b are parameters that are calibrated to fit the empirical prepayment behavior of pools, or groups of loans within pools, that are similar to the mortgages being modeled. Figure 20.2 graphs the function (20.8) with an incentive measured simply as the difference between the WAC and the current mortgage rate available to borrowers. The generic shape of the S -curve is popular since it reflects the empirical behavior that prepayments eventually flatten for very low (negative) and very high incentives.

To capture the complex behavior of actual prepayments the parameters a and b have to vary across loan types and also have to be functions of loan characteristics and the economic environment within loan types. There are very many examples. Since borrowers with relatively high creditworthiness prepay relatively quickly for a given incentive, parameters are made to depend on some proxy for credit, e.g.: *spread at origination* ($SATO$), which is WAC or mortgage coupon relative to current coupon at origination; original FICO; or original LTV. Since higher home prices make it easier for homeowners to refinance, parameters can depend on general or local measures of home price appreciation since origination, to the extent these data are available. Another example is having parameters vary by state or locality




FIGURE 20.2 An Example of S-Curve Prepayments as a Function of Incentive

to reflect observed differences in prepayment behavior across geographic regions.

An additional and extremely important reason that the parameters a and b cannot be constant is so that the prepayment function (20.8) can model *burnout*. Figure 20.3 shows a time series for the monthly CPR for the FNMA 30-year 7% 1995 along with the current coupon as a proxy for the mortgage rate faced by borrowers. The very broad story of the figure is consistent with prepayments increasing with incentive. For example, as the mortgage rate fell from 8% in the beginning of 2000 to less than 4.5% in




FIGURE 20.3 CPR of the FNMA 30-Year 7.0% 1995 and the Current Coupon

spring 2003, CPR increased and peaked at over 60%. But there is another story at work in the figure. When this 1995 vintage pool first experienced mortgage rates of between 6% and 6.50%, in fall 1998, CPR peaked at over 40%. But when the mortgage rate was between 6% and 6.50% in 2006 and 2007, average CPR was much lower. Similarly, CPR peaked at 60% when the mortgage rate was around 4.5% to 5.5%, but with rates below 4.5% after early 2009, CPR was mostly in the range of 10% to 15%. Finally, with mortgage rates eventually falling to historic lows of less than 3.5%, CPR essentially remained in that 10% to 15% range.

To explain the prepayment behavior just described, think about each borrower in the pool as having some set of characteristics that determines a propensity to prepay for a given incentive. For example, a financially sophisticated borrower with a relatively high credit rating, a large loan balance, and a home that has appreciated in price will be the most likely to refinance as mortgage rates decline. In terms of Figure 20.3, this borrower most probably refinanced when rates fell to between 6% and 6.50% in fall 1998. From then on, however, this and other borrowers who are most likely to prepay are no longer in the pool. Therefore, with rates in that same 6% to 6.50% range at a later date, like the period in 2006 and 2007 in the figure, prepayments will be determined by borrowers with a lower propensity to refinance and, therefore, CPR will be lower. The phenomenon of CPR being less responsive to incentive as a pool prepays is known as *burnout*. In terms of the prepayment model (20.8), capturing burnout requires that the parameters be a function of past levels of prepayment rates or mortgage rates.

To mention one more example of how complex models of refinancing can be, researchers have posited a *media effect*, in which a precipitous decline in mortgage rates or mortgage rates reaching a new low creates media reports and cocktail-party conversation that encourage even those borrowers with relatively low propensities to refinance to do so. Capturing this phenomenon in a model would require its parameters to depend on carefully chosen summary statistics that describe the historical path of mortgage rates, e.g., the current mortgage rate relative to the lowest mortgage rate over the last five years.

Turnover

Prepayments due to turnover occur when borrowers sell houses to relocate, to change to a bigger or smaller house, as a result of a divorce, or in response to other personal circumstances. This driver of prepayments typically accounts for less than 10% of overall prepayment rates.

A turnover model for a particular group of loans begins with a base rate that is adjusted to account for the seasonality of relocations, e.g., higher in summer, lower in winter. The model would then add a *seasoning ramp*.

Households are very unlikely to move just after taking out a mortgage. A typical average assumption would be that turnover starts at zero at the time of initiation and increases to the base rate after 30 months. The steepness of the seasoning ramp is often made to depend on several factors. For example, less creditworthy borrowers are more likely to prepay sooner after taking out a mortgage as some will experience improvements to their creditworthiness.

While prepayments classified as due to turnover are for the most part independent of interest rates, there is an interaction that cannot be ignored. Borrowers are less likely to move if they enjoy a below-market mortgage rate, or, put another way, if they would have to pay a higher rate on a new mortgage after selling their homes and moving. This behavior is known as the *lock-in* effect.

Defaults and Modifications

Defaults are a source of prepayments in the sense that mortgage guarantors pay interest and principal outstanding when a borrower defaults. Over the most recent cycle of increasing real estate values, modeling defaults had been less important and had received less attention. This changed dramatically, of course, in reaction to falling housing prices in the run-up to and progression of the 2007–2009 crisis. In addition, mortgage modifications, which did not exist previously, have become an important part of the landscape. From the modeling perspective, more effort is being dedicated to using pertinent variables, e.g., initial LTV ratios, FICO scores, and SATO (which are not usually updated after mortgage issuance), and to incorporating the dynamics of housing prices into the analysis.

Curtailments

Curtailments are partial prepayments by a particular borrower. These tend to be most important when loans are older and balances are low. This driver of prepayments is modeled as a function of loan age and can, with only a couple of years remaining to maturity, rise to a CPR of about 5%.

MBS VALUATION AND TRADING

This section describes how to combine models of the benchmark interest rate with mortgage-specific model components to value MBS. As will be explained presently, while the term structure models of Part Three are relevant for MBS valuation, the tree implementations of these models are not. Therefore, the section begins with an alternate implementation, namely, *Monte Carlo simulation*, to be followed by other valuation issues.

Monte Carlo Simulation

Suppose for a moment that a one-factor tree implementation of a term structure model was used to value MBS. The cash flows at any node of the tree would be determined by scheduled cash flows and the prepayment model. Then, the value of the MBS at any node would be the cash flow on that date plus the expected discounted value of the MBS on the subsequent date. The problem with this approach, however, is that it assumes that the cash flows at any node depend only on the short-term rate at that node, or, equivalently, on the term structure of interest rates at that node. But what if prepayments at particular nodes depend on the history of interest rates on the way to that node, as models of burnout require. In that case the tree implementation fails because it does not naturally recall, for example, whether a node five periods from the start was reached by two down moves followed by three up moves, by three up moves followed by two down moves, or by the sequence up-down-up-down-up. But the burnout effect says that prepayments at a particular node will be less if that node was reached by passing through a node with a relatively low interest rate. In the jargon of valuation models, the tree implementation assumes that cash flows are *path independent* while the cash flows from a burnout model are *path dependent*.

The most popular solution to pricing path-dependent claims is Monte Carlo simulation. To price a security in this framework, proceed as follows. First, generate a large number of paths of interest rates at the frequency and to the horizon desired. For this purpose paths are generated using a particular risk-neutral process for the short-term rate. Second, calculate the cash flows of the security along each path. In the mortgage context this would include the security's scheduled payments along with its prepayments. Note that burnout and media effects can be implemented because each path is available in its entirety as cash flows are calculated. Third, starting at the end of each path, calculate the discounted value of the security's cash flows along each path. Fourth, compute the value of the security as the average of the discounted values across paths.

Table 20.5 presents an extremely simple example of a 5% five-year, annually-paying mortgage pool to illustrate the process along a single path. The arrows at the top indicate that the process is moving forward in time, from date 0 to 5. The interest rate, used as the mortgage rate and as the discounting rate in this simple example, starts at 5%, is 5% at the end of the first year, 4% at the end of the second year, etc. The next rows give, per 100 of original notional, the pool's scheduled interest and principal payments based on the amount outstanding at the beginning of each period, the pool's prepayments from some model, and the total cash flows on each date. Note that the prepayment model can refer to the entire history of rates along the path when computing prepayments.

TABLE 20.5 Example of a Single Path in the Monte Carlo Framework in the Mortgage Context

Date	0	1	2	3	4	5
→ → → → → → → → → → → → → → → → → → → →						
Interest Rate	5%	5%	4%	3%	4%	
Starting Principal		100.00	80.00	54.00	21.70	10.79
Interest Due		5.00	4.00	2.70	1.09	0.54
Principal Due		18.10	18.56	17.13	10.59	10.79
Prepayments		1.90	7.44	15.17	0.33	0.00
Total Cash Flow		25.00	30.00	35.00	12.00	11.32
← ← ← ← ← ← ← ← ← ← ← ← ← ← ← ← ← ← ← ←						
Value	100.93	80.97	55.02	22.22	10.89	11.32

At this point the process starts from the last date and moves backwards in time. The value of the pool on date 5 is simply the cash flow paid on that date, which is 11.32. The value on date 4 is the present value of the date 5 cash flow, i.e., $11.32(1.04)^{-1}$ or 10.89. The value on date 3 is the present value of the date 4 value plus the date 4 cash flow, that is,

$$\frac{10.89 + 12}{1.03} = 22.22 \tag{20.9}$$

Continuing in this manner, the value of the MBS on date 0 along this path is 100.93. Having gone through this process for all of the paths, the value of the MBS is the average date 0 value across paths.

To reconcile Monte Carlo pricing with pricing using an interest rate tree, recall equation (13.17), which, derived in the context of interest rate trees, gives the price of a claim that is worth P_n in n periods. This equation is reproduced here for convenience:

$$P_0 = E \left[\frac{P_n}{\prod_{i=0}^{n-1} (1 + r_i)} \right] \tag{20.10}$$

In light of the discussion of this subsection, the term inside the brackets is analogous to the price of a security along one path. The expectation is analogous to the averaging across paths.

Two more comments will be made about the Monte Carlo framework. First, measures of interest rate sensitivity can be computed by shifting the initial term structure in some manner, repeating the valuation process, and calculating the difference between the prices after and before the interest rate shift. Second, while the Monte Carlo approach does accommodate path-dependent cash flows, it has two major drawbacks. One, it is more computationally and numerically challenging than pricing along a tree. Two,

it is difficult in the Monte Carlo framework to value American- or Bermuda-style options. (Examples in the mortgage context include mortgage options, mentioned earlier, and callable CMOs.) For these options, which allow early exercise, the value of the option at each node is the maximum of the value of exercising the option immediately and the value of the option not exercised. In a tree methodology, which starts at maturity and works backwards, both of these values are available at each node. Along a Monte Carlo path, however, the value of immediate exercise is always known, but the value of the unexercised option is very difficult to compute. Starting a new Monte Carlo pricing simulation at a particular date on a particular path so as to compute the value of the unexercised option for that date and path is possible, but doing so for every exercise date on every path is not computationally feasible.

Valuation Modules

Computing values for MBS require several modules. In no particular order, since they interact with another, these include a model of benchmark interest rates, the scheduled cash flows of the MBS, a model of the mortgage rate, a housing price model, and a prepayment model.

The model of benchmark interest rates can be along the lines of those in Part Three, but, as described in the previous subsection, Monte Carlo implementations usually replace tree implementations. The scheduled cash flows of the MBS are straightforward, as described in the first section of this chapter.

While glossed over in the example of the previous subsection, valuing an MBS along a path requires both the benchmark or discounting rate as well as the mortgage rate; discounting might be done at swap rates plus a spread, but the incentive of a prepayment model depends on the current mortgage rate. But determining the fair mortgage rate at a single date and on a single path of a Monte Carlo valuation is a problem of the same order of magnitude as the original problem of pricing a particular MBS! Common practice, therefore, is to build a simple model of the mortgage rate as a function of the benchmark rates, e.g., as a function of the 10-year swap rate. A particularly simple approach—some say simplistic—is to use a regression of the 30-year mortgage rate on the 10-year swap rate. Note, in any case, that it may not be trivial to compute any longer-term swap rate at points along a path of short-term rates for the same reason as highlighted in the context of pricing options with early exercise. But the problem of computing swap rates can often be handled by using a closed-form solution or a numerical approximation consistent with the process generating the path of short-term rates.

A model of the evolution of housing prices can be particularly useful in modeling the default component of prepayments or prepayments more

generally. The major difficulties, of course, are determining an appropriate probability distribution for housing prices and appropriate correlations for housing prices and interest rates.

Putting the modules together, cash flows are determined by the scheduled cash flows and the prepayment model. The prepayment model depends on the interest rate model, the mortgage rate model, and the housing price model. The mortgage rate model and the housing price model depend on the interest rate model. And finally, the interest rate model is used to value the cash flows.

MBS Hedge Ratios

As mentioned earlier, interest rate sensitivities and hedge ratios can be computed from MBS valuation models. Given the considerable investment required to build an MBS valuation model, however, some market participants, particularly those trading only the simplest products, e.g., TBAs, use empirical hedge ratios or deltas. These can be computed from market data using the tools of Chapter 6. Table 20.6 shows a major dealer's empirical hedge ratios as of December 2010 for various 30-year FNMA TBAs against 5- and 10-year U.S. Treasuries. For example, to hedge a long position in 100 face amount of the 4.0% TBAs, the current coupon, requires the sale of 66 face amount of on-the-run 10-year Treasuries or 115 face amount of 5-year Treasuries.

As expected, the hedge ratios in Table 20.6 fall with coupon. Since higher coupons prepay faster, they are effectively shorter-term securities and, as such, have lower interest rate sensitivities. Of course, this table says nothing about the curve exposure of TBAs. It may be better to hedge with a

TABLE 20.6 Empirical Hedges of TBAs with U.S. Treasuries as of December 9, 2010

FNMA 30-Year TBA Coupon	Treasury Hedge Ratios	
	10-Year	5-Year
3%	0.93	1.64
3.5%	0.80	1.40
4%	0.66	1.15
4.5%	0.54	0.94
5%	0.43	0.75
5.5%	0.35	0.61
6%	0.28	0.50
6.5%	0.23	0.40

Source: JPMorgan Chase.

combination of 5- and 10-year Treasuries, or even with a 7-year Treasury, than to hedge with either a 5- or 10-year Treasury.

Option Adjusted Spread

Option Adjust Spread (OAS) is the most popular measure of relative value for MBS.⁸ Chapter 7 described how to compute OAS in the context of interest rate trees. The method in a Monte Carlo framework is analogous: find the single spread such that shifting the paths of short-term rates by that spread results in a model value equal to the market price. To the extent that the model accounts correctly for scheduled cash flows and prepayments, the OAS represents the deviation of a security's market price from its fair value. Furthermore, as explained in Chapter 7, when OAS is constant the return on a security hedged by a correct model is the short-term rate plus the OAS. Of course, to the extent that a model does not correctly account for prepayments, the OAS will be a blend of relative value and left-out factors.

The practical challenge of using models and OAS to measure relative value is in determining when OAS really does indicate relative value and when it indicates that the model is misspecified. A particular security is most likely mispriced when its OAS is significantly positive or negative while, at the same time, all substantially similar securities trade at an OAS near zero. In practice, however, this is rarely the case. Much more common is the situation in which a model finds relative value across a segment of the market, e.g., finding that premium or high-coupon mortgages are relatively cheap. Deciding whether that segment is really mispriced or whether the model is miscalibrated is the art of relative value trading.

One useful approach in determining whether the OAS of a sector indicates trading opportunities is to graph OAS over time and look for mean reversion. It may prove profitable to buy high-coupon mortgages at high OAS if the model finds that the sector used to trade at zero or negative OAS or, even better, if the sector's OAS oscillates with relatively high frequency around zero. But if the OAS of high-coupon mortgages has been fixed at a particular level over a long period of time, it is likely that it is a feature of the market rather than a mispricing to be exploited. Another useful approach is to determine whether there are any institutional or technical reasons to explain why a particular segment of the market would trade rich or cheap.

⁸Another sometimes-used measure is the *zero-volatility spread*. This is computed by assuming that forward rates are realized, computing prepayments, discounting using those forward rates, and finding the spread above forward rates that results in a model price equal to the market price. While easy to compute, this measure has serious theoretical drawbacks. First, forward rates are not expected rates, so valuation is not taking place along the expected path. Second, even if it were, price equals the expected discounted value not the discounted expected value.

The combination of an empirical finding of relative value combined with a supporting story can be quite convincing.

Turning the discussion to hedging, it can be argued that OAS should be uncorrelated with interest rate movements: the valuation model is supposed to account completely for the effects of interest rates on cash flows and discounting. Furthermore, it is most convenient that OAS be uncorrelated with interest rates because, in that case, interest rate risk can be hedged with the exposures calculated by the model. On the other hand, if OAS is correlated with rates, then that correlation has to be hedged as well to construct a truly rate-neutral position. All in all, this line of reasoning suggests that relative value trading and hedging be restricted to models that produce OAS that are essentially uncorrelated with rates. The only counterargument would be that market mispricings or, alternatively, risk preferences, may, in fact, be correlated with the level of rates.

Figure 20.4 shows the OAS of the FNMA 30-year TBA, as computed by a major broker-dealer, along with the current coupon rate. The OAS of this benchmark mortgage security displays relative value fluctuations from cheap to rich and back, i.e., the series appears to be mean reverting. The OAS also seems to be relatively uncorrelated with the level of mortgage rates. In short, the model does seem like a good candidate for relative value trading. Turn then to the 2007–2009 crisis. With credit concerns rife, the TBA OAS broke out of its band, peaking at an unprecedented 100 basis point of cheapness. *Ex ante*, should a trader have bought TBAs as the OAS of the TBA broke out of its band during the crisis, reaching 60 or 70 basis points? Could the trade be sustained through the OAS peak of 100 basis points so as to reap




FIGURE 20.4 OAS of the FNMA TBA 30-Year, as Calculated by a Major Broker-Dealer, with the FNMA 30-Year Current Coupon

the profits of its eventually falling to zero? Or, *ex ante*, should the OAS have been considered a reasonably accurate reflection of deteriorating credit conditions and not an indicator of a relative value opportunity?

PRICE-RATE BEHAVIOR OF MBS

Figure 20.5 shows the rate behavior of a 5% 30-year MBS along with two other price curves for reference. The dotted curve is the price-rate curve of a (fictional) mortgage with scheduled interest and principal payments only, that is, with no prepayments. Not surprisingly, the curve looks like that of any security with fixed cash flows: it is decreasing in rates and positively convex. The dashed curve gives the price of mortgage with a constant CPR of 6%, which is the CPR of the S-Curve in Figure 20.2 for sufficiently negative incentives. (The portion of this curve in the right half of the graph coincides with the solid curve, which will be discussed presently.) Since a fixed CPR leads to just another set of fixed cash flows, the price behavior of the dashed line is, like the dotted line, qualitatively similar to any security with fixed cash flows. A mortgage with a CPR of 6%, however, is effectively a shorter-term security than an otherwise identical mortgage with a CPR of 0%. Hence, the *DV01* of the dashed curve is less than the *DV01* of the dotted curve at any given level of rates.

The solid curve in Figure 20.5 is the price-rate curve of a 5% 30-year MBS with prepayments governed by the S-curve in Figure 20.2. For very




FIGURE 20.5 Price-Rate Curve of a 5% 30-Year MBS with Prepayments from the S-Curve of Figure 20.2 along with Two Curves of 5% 30-Year Mortgages at Fixed CPRs

high rates, i.e., negative incentives, the CPR of the MBS is 6% and the solid line corresponds to the dashed line discussed in the previous paragraph. As rates fall, and the value of the scheduled cash flows rise, CPR increases. This means that principal is repaid at par and that the value of the MBS cannot continue increasing as rates fall. This qualitative price-rate behavior is very much like that of callable bonds with an important difference. Since the exercise of callable bonds is close to efficient, a corporation that can call its bonds at par does so: the bond's value cannot, therefore, rise much above par. In the case of mortgages, however, borrowers do not prepay when they "ought" to, in a strict present value sense, enabling the value of a mortgage at low rates to rise above par, as it does in the figure. Finally, note that, because of the prepayment option, the price-rate curve of the mortgage is negatively convex at lower rates. This is very much analogous to the negative convexity of the price-rate curve of a callable bond.

Figure 20.6 graphs the price of the same 5% 30-year MBS, labeled here as a pass-through, along with the prices of its associated IO and PO. When rates are very high and prepayments low, the PO is like a zero coupon bond, paying nothing until maturity. As rates fall and prepayments accelerate, the value of the PO rises dramatically. First, there is the usual effect that lower rates increase present values. Second, since the PO is like a zero coupon bond, it will be particularly sensitive to this effect. Third, as prepayments increase, some of the PO, which sells at a discount, is redeemed at par.




FIGURE 20.6 Price-Rate Curve of a 5% 30-Year MBS with Prepayments from the S-Curve of Figure 20.2 along with the Price-Rate Curves of Its Associated IO and PO

Together, these three effects make PO prices particularly sensitive to interest rate changes.

The price-rate curve of the IO is, of course, the pass-through curve minus the PO curve, but it is instructive to describe the IO curve independently. When rates are very high and prepayments low, the IO is like a security with a fixed set of cash flows. As rates fall and mortgages begin to prepay, the cash flows of an IO vanish. Interest lives off principal. Whenever some principal is paid off there is less available from which to collect interest. But, unlike callable bonds or pass-throughs that receive such prepaid principal, when prepayments cause interest payments to stop or slow the IO gets nothing. Once again, its cash flows simply vanish. This effect swamps the discounting effect so that, when rates fall, IO values decrease dramatically. The negative DV01 or duration of IOs, an unusual feature among fixed income products, may be valued by traders and portfolio managers in combination with more regularly behaved fixed income securities.

HEDGING REQUIREMENTS OF SELECTED MORTGAGE MARKET PARTICIPANTS

As mentioned earlier in the chapter, mortgage servicers are responsible for managing mortgage loans and passing cash flows from the borrowers to the lenders. Servicers are paid a fee for this service, typically between 20 and 50 basis points of the notional amount. If a loan is prepaid, the fee stream from that loan ends. Hence, while the valuation of *mortgage servicing rights* (MSR) is quite complex, some qualitative features of that business resemble the characteristics of IOs. From this perspective, mortgage servicers stand to lose revenue and value as rates fall. There is an offsetting effect, however: to the extent that borrowers refinance and servicers collect fees on the newly issued mortgages, and to the extent that lower rates actually increase the notional of mortgages outstanding, servicers might not lose very much from declining rates. But a servicer that has decided to hedge some of its revenue stream from falling rates faces a challenge. Hedging an IO-like security with a TBA would entail a severe convexity mismatch, conceptually similar to the discussion in the context of futures and options in the hedging application of Chapter 4, but quantitatively much worse a problem. Hedging with swaps also entails a convexity mismatch and suffers, in addition, from mortgage-swap basis risk, i.e., the risk that mortgage rates and swap rates move by different amounts or, worse, in opposite directions. The risk profile of the securities mentioned in the subsection “Other Products” in this chapter might be better suited to this hedging problem, but their relative lack of liquidity limits their usefulness to hedgers of the size of servicers.

Lenders in the primary market, meaning financial institutions that lend money directly to mortgage borrowers, also have interest rate risk to hedge. From the time that the lender and borrower agree on the terms of a loan until the time the lender sells the loan to be securitized, the lender is exposed to the risk that rates will rise and result in the loan's losing value. Selling TBAs is a fine solution to this hedging problem. Secondary market originators that buy mortgages from lenders in the primary market and sell these mortgages through securitizations face the same risk as primary lenders. Rates may rise between the time the mortgages are bought and the time they are sold.

Curve Construction

This chapter begins with an introduction to the goals of constructing curves of discount factors or rates and then recommends and presents in detail two popular methodologies, namely, flat forwards and a smoothing of those forwards based on piecewise quadratic interpolation. To present ideas and techniques, the focus here is on building a single *London Interbank Offered Rate (LIBOR)*-based curve. The techniques for implementing the two-curve methodology of Chapter 17 are essentially the same.

INTRODUCTION

In some very special cases a security can be priced by arbitrage relative to a set of other securities, e.g., a 5% 2-year swap can be priced relative to par swaps with maturities of six months, 1 year, 1.5 years, and 2 years. More frequently, however, arbitrage pricing is not possible because a security to be priced makes cash flows on one set of dates while benchmark securities make cash flows on another set of dates. Continuing with another swap example, it might be necessary to value a 5% 13.4-year swap relative to a set of more frequently traded swaps, none of which makes payments on exactly the same set of dates as the 5% swap. Or, to take another example, it might be necessary to calculate the spread to swaps of a 9% 7.3-year corporate bond with annual coupon payments. Problems of this sort could, in theory, be solved with the models of Part Three: a term structure model could be calibrated based on both historical data and selected current benchmark prices and then used to price other securities, possibly with a spread. In practice, however, securities with fixed cash flows are priced essentially by interpolating benchmark prices. More specifically, it is assumed that discount factors, spot rates, or forward rates are described by some mathematical function of term, i.e., some curve. This curve is then calibrated to price benchmark securities correctly and used to discount cash flows occurring at arbitrary dates in the future. Hence, a curve prices the 13.4-year swap not by

arbitrage arguments, but essentially by interpolating nearby benchmark swap rates, e.g., the 12- and 15-year rates.

In most applications there are three major objectives in building a curve. One, the chosen benchmark securities, which typically include many liquid securities along the term structure, should all be priced correctly with the resulting curve.¹ Two, the resulting curve is economically reasonable, e.g., the forward rates, which are being set to price the benchmarks, do not wind up exhibiting wild oscillations. Three, bucket exposures are relatively local, i.e., changing the rate of a benchmark security in one part of the curve does not ripple through prices of securities elsewhere on the curve. As will be discussed below, forcing a curve to be smooth can cause significant violations of this locality property.

Given the benchmark status of LIBOR-based securities, these markets are the most widely used to calibrate pricing curves. Deposit rates or forward-rate agreements (FRAs) might be used for the very short end of a curve, Eurodollar (ED) futures rates are typically used for the short end, and swap rates are the standard for the intermediate and long end. As discussed in Chapter 17, however, federal (fed) funds-based curves, constructed from OIS, are also used by practitioners.

Curve construction requires two steps: choosing a functional form for the curve and then fitting that functional form to benchmark prices. Part of the choice of functional form is whether the curve is expressed in terms of discount factors, spot rates, or forward rates. The most popular choice is forward rates. First, as has been made clear throughout this book, forward rates are interesting quantities that are monitored and often traded. Second, unlike discount factors and spot rates, forward rates are non-overlapping and, therefore, can be shifted relatively independently. Third, as pointed out in Part One, forward rates are essentially changes or derivatives of spot rates. But the calculation of derivatives can be numerically unstable so that small errors in determining spot rates can translate into large errors in implied forward rates. Hence, since forward rates are of interest, it is numerically safer to impose desirable curve properties directly on forward rates.

The chosen functional form of a curve is often fitted to benchmark security prices or rates by a process called *bootstrapping*. Basically, the first segment of the curve is fit to the first benchmark security. Then, given that

¹If benchmark prices are subject to significant observation errors, a curve might be designed to minimize errors when pricing these benchmarks as opposed to forcing all such errors to zero. For example, the most liquid government bond prices can be affected by idiosyncratic factors, e.g., on-the-run effects, that should not be incorporated into a curve that is used to price other securities. In any case, as the benchmarks have moved to LIBOR-based securities, many of which are quoted with very reliable rates or prices, the exact fitting of benchmark securities has become the dominant practice.

segment, the second segment is fitted to the second benchmark security, etc. A particularly easy example of bootstrapping was introduced in Chapters 1 and 2, where discount factors were extracted from government bond prices and spot and forward rates from par swap rates. The simple functional form of the curve in those chapters can be described as a set of sequential six-month forward rates. The bootstrapping process began by finding the six-month forward rate that priced a six-month security. Then, fixing that six-month rate, the six-month rate six months forward was found that priced a one-year security. Then, fixing those two rates, the six-month rate one year forward was found to price a 1.5-year security, etc.

In the simple process in Chapters 1 and 2, the granularity of the functional form was, for convenience, set at six months. In practice, however, this granularity has to match the availability of benchmark securities. If, for example, 10- and 12-year swaps are liquid enough to be taken as benchmarks, with no other sufficiently liquid swap between them, it would not be feasible to extract the six-month rate 10 years forward from traded securities. The 2-year rate 10 years forward, however, can be extracted from the 12-year swap rate given the prior extraction of rates up to 10 years.

FLAT FORWARDS

A very simple and widely-used functional form for pricing curves is that of *flat forwards*. The assumption is that, between terms t_{i-1} and t_i , the instantaneous or continuously compounded forward rate is constant at some value f_i . With $d(t)$ denoting the discount factor to term t , the flat forward assumption implies that

$$d(t) = d(t_{i-1}) e^{-f_i(t-t_{i-1})} \quad (21.1)$$

Hence, given the discount factor to t_{i-1} and the discount factor to t_i , the forward rate from t_{i-1} and t_i can be recovered. Put another way, with a curve built to term t_{i-1} and a security maturing at term t_i , bootstrapping can continue with the extraction of the forward rate over that incremental term.

Another property of flat forwards is that any forward rate can be written as a weighted average of the collection of flat forward rates. To illustrate, say that the flat forward between 10 and 12 years is 4% while the flat forward between 12 and 15 years is 4.25%. Then what is the 3-year forward rate, f , between 11 and 14 years? First, by the definition of the rate f ,

$$d(14) = d(11) e^{-3f} \quad (21.2)$$

Second, by the assumption of flat forwards and the rates supplied, the forward rate over the year 11 to year 12 is 4% while the forward rate over the two years 12 to 14 is 4.25%. Hence,

$$d(14) = d(11)e^{-1 \times 4\%}e^{-2 \times 4.25\%} \quad (21.3)$$

Putting (21.2) and (21.3) together,

$$-3f = -1 \times 4\% - 2 \times 4.25\% \quad (21.4)$$

$$f = \frac{1 \times 4\% + 2 \times 4.25\%}{3} \quad (21.5)$$

$$= 4.167\% \quad (21.6)$$

At first glance flat forwards might seem an odd assumption since the forward rates jump from segment to segment. With respect to the objectives of curve fitting, however, this functional form has desirable properties. It can be easily adapted to fit benchmark securities, as will be seen in the next section. The forward rates do jump across segments, but, in practice, the jumps do not behave wildly, that is, flat forwards do not jump dramatically higher and then dramatically lower, or *vice versa*. Finally, because flat forwards consist of independent segments, one part of the forward curve can move without affecting other parts of the curve. This feature has several benefits. In particular, as fitting securities change price, the flat forwards move in a stable and well-behaved manner. Also, as will be shown in a later section, the independence of flat forward segments generates intuitively appealing hedge ratios.

The benefits and weaknesses of the flat forward functional form will become clearer through the discussion of piecewise quadratics below. For now, however, the text turns to an example of constructing a flat forward curve.

FLAT FORWARDS FOR A USD LIBOR CURVE

This section presents a detailed example of building a flat forward curve using USD LIBOR derivatives as of May 28, 2010, for settlement on June 2, 2010. In particular, this curve matches the prices of the first 10 ED futures contracts, swap rates in annual maturity increments from 3 to 10 years, and then the 12-, 15-, 20-, 25-, 30-, and 40-year swap rates. Not discussed here, but very common in practice, is to use a short-term rate, e.g., three-month LIBOR, to fit a *stub* rate for the period from spot settlement to the beginning of the period covered by the first ED contract.

TABLE 21.1 Selected Fitted USD LIBOR Flat Forwards as of May 28, 2010

From (not inclusive)	To (inclusive)	Fwd Rate Cont. Comp. (%)
6/2/10	9/16/10	.5945
9/16/10	12/15/10	.8480
12/15/10	3/15/11	.9876
3/15/11	6/16/11	1.1138
6/16/11	9/15/11	1.2856
9/15/11	12/21/11	1.4821
12/21/11	3/21/12	1.7293
3/21/12	6/21/12	1.9553
6/21/12	9/20/12	2.1977
9/20/12	12/19/12	2.4363
12/19/12	6/3/13	2.7647

Table 21.1 lists the dates and continuously compounded forward rates of the segments comprising the short end of the fitted curve, from the settlement date to June 3, 2013, the maturity date of the 3-year swap.² By demonstrating how the forward rates in Table 21.1 correctly price the 10 ED futures contracts and the 3-year swap rate, it will be clear how to derive these rates in the first place and how to derive the longer-term flat forward rates.

Given the flat forward rates in Table 21.1, Table 21.2 describes the pricing of ED contracts. Column (1) gives the tickers of the contracts while columns (2) and (3) give the dates of the underlying LIBOR deposits, as described in Chapter 15. For example, since EDM1 expires on June 13, 2011, its final settlement price is determined by three-month LIBOR set on that expiration date, which, by definition, is the rate on a deposit settling two business days later, on June 15, 2011, and maturing three months after that, on September 15, 2011.

Column (4) of Table 21.2 gives the futures rate on the pricing date.³ Column (5) gives the forward rates corresponding to the same periods, which are obtained by applying a model-specific convexity correction to the futures rates in column (4). (See Chapter 13.) Column (6) simply transforms the actual/360 rate in Column (5) to a continuously compounded rate.

²Since the scheduled maturity date, June 2, 2013, falls on a weekend, the final cash flows are paid on the next business day, June 3, 2013.

³These rates do not correspond to settlement prices and, therefore, differ slightly from the rates given in Table 15.3. In building curves, practice is to take a snapshot of swap and ED futures rates when both markets are suitably liquid. This time need not correspond to the futures close.

TABLE 21.2 Pricing ED Futures Options with the USD LIBOR Flat-Forward Curve as of May 28, 2010

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Contract	LIBOR		Fut Rate	Fwd Rate	Fwd Rate	Days 1	Rate 1	Days 2	Rate 2
	Start	End							
	Act/360				Continuously Compounded				
	(%)				(%)		(%)		(%)
EDM0	6/16/10	9/16/10	.595	.5949	.5945	92	.5945		
EDU0	9/15/10	12/15/10	.8745	.8462	.8453	1	.5945	90	.8480
EDZ0	12/15/10	3/15/11	.9925	.9888	.9876	90	.9876		
EDH1	3/16/11	6/16/11	1.1225	1.1154	1.1138	92	1.1138		
EDM1	6/15/11	9/15/11	1.2975	1.2859	1.2838	1	1.1138	91	1.2856
EDU1	9/21/11	12/21/11	1.5025	1.4848	1.4821	91	1.4821		
EDZ1	12/21/11	3/21/12	1.7575	1.7331	1.7293	91	1.7293		
EDH2	3/21/12	6/21/12	1.9925	1.9602	1.9553	92	1.9553		
EDM2	6/20/12	9/20/12	2.2425	2.2012	2.1950	1	1.9553	91	2.1977
EDU2	9/19/12	12/19/12	2.4925	2.4412	2.4337	1	2.1977	90	2.4363

Pricing EDM0 from the flat forward rates in Table 21.1 is particularly easy. The forward rate relevant for pricing EDM0 is that from June 16 to September 16, 2010. But according to Table 21.1, the single continuously-compounded forward rate over that entire 92-day period is .5945%. Hence, the continuously compounded forward rate that prices EDM0 is also .5945%, as reported in Column (6) of Table 21.2. Columns (7) and (8) of that table simply indicate that the rate .5945% is applicable to all 92 days.

Pricing EDU0 is a bit more complicated. The underlying deposit covers the period September 15 to December 15. The applicable forward rate for the one day from September 15 to September 16 is, according to Table 21.1, .5945%. But the rate applicable over the rest of the period, the 90 days from September 16 to December 15 is, according to that table, .8480%. Columns (7) through (10) of Table 21.2 report this finding. Hence, by the discussion in the previous section, the forward rate for EDU0 must be the weighted average of the two relevant forward rates,

$$\frac{1 \times .5945\% + 90 \times .8480\%}{91} = .8452\% \quad (21.7)$$

which, to rounding error, is the continuously-compounded market forward rate reported in Column (6) of Table 21.2. Hence, the flat forward rates of Table 21.1 do price EDU0.

To take one more example, the forward rate for EDM1, covering the period from June 15, 2011, to September 15, 2011, is composed of one day at 1.1138%, the forward rate relevant from June 15 to June 16, and 91 days at 1.2856%, the forward rate relevant over the period June 16 to September 15. Hence, the forward rate for EDM1 is

$$\frac{1 \times 1.1138\% + 91 \times 1.2856\%}{92} = 1.2837\% \quad (21.8)$$

which, again to rounding error, is the given market forward rate.

In this manner it can be verified that the flat forward rates of Table 21.1 do correctly price the ED futures rate given in Table 21.2.

Before continuing to pricing swaps, note the logic of choosing the flat forward segments in Table 21.1. First, the segments are arranged so that each ED futures rate determines the forward rate of one segment. Six-month segments, for example, would not be a granular enough division of term to fit all ED futures rates. On the other hand, the forward rates corresponding to one-month segments could not be determined by ED futures rates alone. Second, because the dates of the deposits underlying ED futures overlap, there is no way to choose nonoverlapping flat forward segments that correspond exactly to underlying deposits.

An interesting alternate approach to choosing the segments is to set their endpoints to correspond to central bank meeting dates. This allows for an economic interpretation of the flat forwards as the short-term rate targeted by the central bank, although this interpretation is not as clean in the presence of a risk premium. Also, setting the segments in this way creates some implementation complications analogous to those found in the process of extracting implied Board of Governors of the Federal Reserve System (Fed) policy changes from market rates in Chapter 15; there might be no benchmark security or more than one benchmark security that determines the forward rate from one meeting date to the next.

The demonstration that the flat forward rates of Table 21.1 correctly price the fitting securities continues now with the 3-year par swap, which, as of the pricing date, carried a rate of 1.6695%. Table 21.3 presents the relevant calculations. Columns (1) and (2) give the accrual periods for the floating rate payments. Column (3) to (7) compute the continuously compounded forward rates over these accrual periods from the flat forward rates in Table 21.1, along the same lines as the analogous computations in Table 21.2. Column (8) converts the continuously compounded rates in Column (7) to actual/360 rates. Columns (9) and (10) give the accrual factors for the floating and fixed rate payments, with the former using the actual/360 day count and the latter using 30/360 day count. (See Chapter 16.) Finally, Column (11) gives the discount factor to each payment date from the calculated forward rates. All of this data allows for the calculation of the present value of the payments from the two legs of the swap.

The fixed side is straightforward. The payments, each equal to the swap rate of 1.6695% time the fixed-rate accrual factor, are multiplied by the respective discount factors and then added together. Note that the accrual factors are not equal to .5 when the start and end of the accrual periods do not fall on the same day of the month. In any case, the present value of the fixed cash flows in Table 21.3, without the fictional notional amount, is .0490 per unit notional amount, shown in the penultimate row of the table. With the present value of the fictional notional amount, which is just the discount factor to June 3, 2013, i.e., .9510, the value of the fixed side is 1.0.

The value of the floating leg is a bit tedious when calculated very precisely. Chapter 16 showed that the value of the floating side of a swap is par on reset dates because, essentially, all accrual periods correspond to the term of the LIBOR index. To be very precise, however, accrual periods and LIBOR terms do not always line up exactly. In the 3-year swap of Table 21.3, the payment date of June 2, 2012 is pushed from that Saturday to the next business day on Wednesday, June 6.⁴ Therefore, the subsequent accrual period begins on June 6, but ends on September 4, the first business

⁴Monday and Tuesday, June 4 and 5, are special London bank holidays in 2012 in honor of the Queen's Diamond Jubilee.

TABLE 21.3 Pricing a 1.6695% 3-Year Swap with the USD LIBOR Flat-Forward Curve as of May 28, 2010

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Accrual Start	Period End	Days 1	Rate 1	Days 2	Rate 2	Fwd	Fwd	Acc. Float	Factors Fixed	Disc. Factor
Continuously Compounded						Act/360				
			(%)			(%)	(%)			
6/2/10	9/2/10	92	.5945			.5945	.5949	.2556		.998482
9/2/10	12/2/10	14	.5945	77	.8480	.8090	.8099	.2528	.5000	.996442
12/2/10	3/2/11	13	.8480	77	.9876	.9675	.9686	.2500		.994035
3/2/11	6/2/11	13	.9876	79	1.1138	1.0960	1.0975	.2556	.5000	.991255
6/2/11	9/2/11	14	1.1138	78	1.2856	1.2595	1.2615	.2556		.988069
9/2/11	12/2/11	13	1.2856	78	1.4821	1.4540	1.4567	.2528	.5000	.984445
12/2/11	3/2/12	19	1.4821	72	1.7293	1.6777	1.6812	.2528		.980279
3/2/12	6/6/12	19	1.7293	77	1.9553	1.9105	1.9154	.2667	.5111	.975297
6/6/12	9/4/12	15	1.9553	75	2.1977	2.1573	2.1631	.2500		.970051
9/4/12	12/3/12	16	2.1977	74	2.4363	2.3939	2.4011	.2500	.4917	.964263
12/3/12	3/4/13	16	2.4363	75	2.7647	2.7070	2.7162	.2528		.957688
3/4/13	6/3/13	91	2.7647			2.7647	2.7744	.2528	.5000	.951018
Value Floating Side:			.0490							
Value Fixed Side:			.0490							

day after the regularly scheduled payment date of September 2. But the LIBOR deposit from June 6 matures on the valid business day of September 6, not September 4. Hence, the accrual period is different than the term of the LIBOR index and the discounted value as of the reset date will not be exactly par. Because of this (admittedly small) effect, instead of pricing the floating leg as in Chapter 16, practitioners use forward rates to project future LIBOR rates, multiply by the accrual factors, and discount to the present. (See Chapter 17.) Applying this methodology, Table 21.3 calculates the value of the floating leg by multiplying each actual/360 forward rate in Column (8) by its accrual factor in Column (9) and its discount factor in Column (11) and then summing the results. The total comes to .0490, shown in the last row of the table. The fact that the values of the fixed and floating sides of the swap are equal—with a fixed rate equal to the 3-year par rate of 1.6695%—means that the flat forward rates of Table 21.1 have indeed correctly priced the 3-year swap. Note, by the way, that the pricing of the floating leg here assumes that the day's LIBOR reset has not yet happened. If it had, the first projected forward rate would be replaced by that realized setting.

The illustration of pricing with flat forwards ends here, but it is clear how to continue. The next flat forward segment will cover the period June 3, 2013, to June 2, 2014, the maturity of the 4-year swap. Furthermore, the forward rate of this segment is set to ensure that the 4-year swap is priced by the flat forward curve, along the lines of Table 21.3. Flat forward segments continue to be added one at a time, with each corresponding forward rate set so as to match another swap rate, until all of the fitting securities have been used.

Figure 21.1 shows the resulting flat forward curve, along with a smoother curve to be discussed in the next section. Despite its jumps, the curve is well-behaved and reasonable. For many purposes the jumps hardly matter in the short end, where the forward segments are about three months in length. For intermediate and longer terms the jumps are noticeable but their amplitudes change in a gradual and consistent fashion.

SMOOTH FORWARDS BY PIECEWISE QUADRATICS

For some applications, when jumps across flat forward rates are undesirable, practitioners prefer a smoother forward rate function. One cost of moving to smooth forwards is that such curves are more difficult to construct and maintain. Illustrations will not be given here, but for a given functional form it may be very difficult to fit a smooth curve to a particular set of market data without introducing undesirable features. The cubic spline, for example, which is very smooth in that neither the function nor its first derivative nor its second derivative jumps, is notorious for oscillating wildly unless carefully suited to the data at hand. A very much related cost of moving to smooth forwards is that parts of the fitted term structure become




FIGURE 21.1 Fitted USD LIBOR Curves as of May 28, 2010

dependent, in artificial and unwanted ways, on other parts of the term structure. This locality property has been mentioned previously and will be better understood through this section and the next.

One curve-fitting methodology that is popular as a compromise between the advantages of flat forwards and the desirability of smoothness is to build piecewise quadratic functions from the flat forwards. More specifically, given a flat forward curve like the one built in the previous section, find a set of piecewise quadratic functions that essentially smooth out the flat forwards while maintaining the property that all benchmark securities are priced correctly.

Mathematically, let the midpoint of the flat forward segment between t_{i-1} and t_i be

$$t_{i-1}^m = \frac{t_{i-1} + t_i}{2} \quad (21.9)$$

Note that the flat forward rates are defined such that f_i is the rate from t_{i-1} to t_i . Hence, for example, the flat forward rate at t_1^m is f_2 .

Define smooth continuously compounded forward rates with the piecewise functions, $\phi_i(t)$, $i = 2, 3, \dots, N$. (The first function, $\phi_1(t)$, will be handled separately in a moment.) Each function gives the smooth forward rate from one midpoint to the next, so that $\phi_i(t)$ is valid over the range t_{i-1}^m to t_i^m . Also, each function is defined to be quadratic, so that, for some constants a_i , b_i , and c_i ,

$$\phi_i(t) = a_i + b_i(t - t_{i-1}^m) + c_i(t - t_{i-1}^m)^2 \quad (21.10)$$

These piecewise quadratics are now constrained in the following ways. First, the value of each quadratic function starts at a flat forward rate and ends at the subsequent flat forward rate. This means, for example, that the function ϕ_2 , which is valid from t_1^m to t_2^m , is constrained such that $\phi_2(t_1^m) = f_2$ and $\phi_2(t_2^m) = f_3$. More generally,

$$\phi_i(t_{i-1}^m) = f_i \quad (21.11)$$

$$\phi_i(t_i^m) = f_{i+1} \quad (21.12)$$

Second, the smooth forward rate over every flat forward segment equals the flat forward rate over that segment. Consider, for example, the flat forward segment from t_2 to t_3 with forward rate f_3 . The smooth forward rate from t_2 to t_2^m is determined by the function ϕ_2 while the smooth forward from t_2^m to t_3 is determined by the function ϕ_3 . Furthermore, analogous to the discussion in the previous section, the forward rate over the interval t_2 to t_3 using the smooth forward functions is simply the average of the instantaneous forward rates over that interval. Hence, for the smooth forward rate functions to give the forward rate f_3 over the interval t_2 to t_3 , it has to be the case that

$$\int_{t_2}^{t_2^m} \phi_2(t) dt + \int_{t_2^m}^{t_3} \phi_3(t) dt = (t_3 - t_2) f_3 \quad (21.13)$$

Or, more generally,

$$\int_{t_{i-1}}^{t_{i-1}^m} \phi_{i-1}(t) dt + \int_{t_{i-1}^m}^{t_i} \phi_i(t) dt = (t_i - t_{i-1}) f_i \quad (21.14)$$

Note that since the $\phi_i(t)$ are quadratic, these integrals are quite easy to evaluate.

The first of the piecewise quadratics, $\phi_1(t)$, is defined over the interval t_0 to t_1^m (instead of from t_0^m to t_1^m). Therefore, conditions (21.11), (21.12), and (21.14) become

$$\phi_1(t_0^m) = f_1 \quad (21.15)$$

$$\phi_1(t_1^m) = f_2 \quad (21.16)$$

$$\int_{t_0}^{t_1} \phi_1(t) dt = (t_1 - t_0) f_1 \quad (21.17)$$

Using the functional form (21.10) for $\phi_1(t)$, together with the restrictions (21.15) to (21.17), it is easy to show that⁵

$$\phi_1(t) = f_1 + \frac{f_2 - f_1}{(t_1^m - t_0^m)} (t - t_0^m) \quad (21.18)$$

Taken together across quadratic segments, the conditions (21.14) ensure that the forward rates over the flat forward segments are the same whether computed with the smooth forward functions or with the flat forward rates. But the flat forward rates were set so all of the benchmark securities are priced correctly. Hence, all benchmark securities are priced correctly using the smoothed forwards as well.

Given the flat forwards, the first quadratic function in (21.18), and the rules (21.11), (21.12), and (21.14) for $i = 2$, the function $\phi_2(t)$ can be found. Then, using that result, the flat forwards, and the rules with $i = 3$, $\phi_3(t)$ can be found, etc. In this manner, then, all of the $\phi_i(t)$ can be computed.

The results of this process for the USD LIBOR curve as of May 28, 2010, are shown in Figure 21.1. By construction the smooth forward rates pass through the midpoints of the flat forward segments at the flat forward rates.

It should now be clear how imposing smoothness links one part of a fitted curve with another. Say that the flat forward f_3 changes but that no other flat forward rate changes. Then, from (21.13), $\phi_3(t)$ has to change. (The function $\phi_2(t)$ cannot change without affecting earlier forwards.) But this means from (21.14), with $i = 4$, that the function $\phi_4(t)$ has to change so as to keep f_4 the same, which in turns means that $\phi_5(t)$ has to change to keep f_5 the same, etc.

While the piecewise quadratic forward rates created here are continuous, i.e., they do not jump like the flat forwards, they are not very smooth in a mathematical sense. In particular, the first derivative of the smooth forward curve jumps at the midpoints. Imposing more smoothness on the curve might be appealing from an economic view of forward rates, but, having to add conditions that equate the first derivatives of the segments would create even more dependency across the term structure of the kind described in the previous paragraph. It is in this sense that piecewise quadratics based on flat forwards is viewed as a compromise between smoothness and desirable locality properties. The next section illustrates the hedging implications of this tradeoff.

⁵To show that (21.17) is satisfied, integrate both sides of (21.18) and then use the definition of the midpoints in (21.9).

LOCALITY PROPERTIES AND HEDGING

This chapter has pointed out that a desirable property of flat forwards is that different parts of the term structure of forward rates can move independently. It was also pointed out that methodologies that produce smoother forwards, like the piecewise quadratics of the previous section or, as a more extreme example, like cubic splines, tend to sacrifice this locality property. This section illustrates an important manifestation of this sacrifice of locality by comparing hedge ratios across curve building methodologies.

Consider a portfolio that pays 2.61% on a \$100 million 5.5-year swap and 3.853% on a \$100 million 17-year swap. Table 21.4 gives two analyses of the partial '01s of the portfolio, one using the flat forward curve of the chapter and one using the smooth forward curve. Partial '01s were discussed in Chapter 5, but the basic idea is the following. To assess the risk of a portfolio to a particular benchmark security that is used to fit the curve, shift the rate on that benchmark security by one basis point, keeping the rates on all other benchmark securities unchanged. Then refit the curve and calculate the change in portfolio value from its value before the shift. Columns (2) and (5) report this change in value for the portfolio for an increase of one basis point in the respective benchmark. Using flat forwards, for example, increasing the 5-year par swap rate by one basis point increases

TABLE 21.4 Partial '01 Analyses of a Portfolio Paying 2.61% on USD 100mm 5.5-Year Swaps and 3.853% on USD 100mm 17-Year Swaps Under Different Curve Methodologies as of May 28, 2010

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Benchmark	Flat Forwards			Smooth Forwards			
Swap	P&L	DV01	Hedge	P&L	DV01	Hedge	
Term in Years	(\$)	\$/100	\$ Notional	(\$)	\$/100	\$ Notional	
3	0	.02934	0	5	.02935	15,569	
4	0	.03858	0	-2,281	.03859	-5,911,373	
5	23,517	.04748	49,533,497	26,140	.04749	55,040,920	
6	28,339	.05603	50,577,405	32,283	.05605	57,597,224	
7	0	.06422	0	-4,338	.06423	-6,753,941	
8	0	.07204	0	-4	.07206	-5,068	
9	0	.07955	0	-8	.07957	-9,945	
10	0	.08665	0	-7	.08667	-8,496	
12	0	.09993	0	-11,972	.09995	-11,978,509	
15	66,978	.11763	56,938,610	85,743	.11765	72,882,187	
20	61,205	.14214	43,060,038	65,029	.14213	45,753,641	
25	0	.16168	0	-10,697	.16166	-6,616,993	
30	0	.17735	0	0	.17734	0	

the value of the portfolio by \$23,517. Columns (3) and (6) give the *DV01* of each benchmark security, i.e., the change in the value of receiving \$100 notional amount in that benchmark should all rates fall by one basis point. Finally, Columns (4) and (7) give the notional amount of each benchmark security that, taken all together, hedges the partial '01 exposures of the portfolio. These numbers are just the profit and loss (P&L) sensitivity with respect to each benchmark divided by the *DV01* per unit notional amount of that benchmark. For example, hedging the \$23,517 P&L exposure of the portfolio to the 5-year swap rate under flat forwards requires a notional amount of 5-year swaps of

$$\frac{\$23,517}{.04748/100} \quad (21.19)$$

or about \$49.5 million.

The hedges calculated from the flat forward curve are local in that the portfolio of 5.5-year and 17-year swaps can be hedged with 5-year, 6-year, 15-year, and 20-year swaps alone. Basically, each of the 5.5-year and 17-year swaps can be hedged by the benchmarks surrounding it. By contrast, using the smooth forward curve to calculate partial '01s shows that the portfolio is significantly sensitive to the 4-, 7-, 12-, and 25-year benchmark rates in addition to the benchmark rates immediately surrounding the swaps in the portfolio. Because forward rates are interconnected in the construction of the smooth forward curve, changing the 25-year par rate and refitting the curve, for example, generates a change in forwards that has a significant effect on the value of a 17-year swap. It may well be that the best hedge of a 17-year swap does require holdings in each of the 12-year, 15-year, 20-year, and 25-year swaps. But it is unlikely that this best hedge has been generated by the quadratic and other assumptions used to construct the smooth forward curve.

References

- Aegon Global Pensions, "Pension Provision in Germany," 2010.
- Allianz Global Investors, "The German Pension System," Mayer-Brown, 2009.
- Andersen, L., and V. Piterbarg, *Interest Rate Modeling*, Volume II, Atlantick Financial Press, 2010.
- Baygun, B., J. Showers, and G. Cherpelis, "Principles of Principal Components," Salomon Smith Barney, January 31, 2000.
- Burghardt, G. and T. Belton, *The Treasury Bond Basis*, Third Edition, McGraw-Hill, 2005.
- Black, F., "Interest Rates as Options," *Journal of Finance*, Vol. 50, 1995, pp. 1371–1376.
- Bloomberg *BusinessWeek*, "Greece Pays Bond Investors 5 Times Spain Yield Spread (Update1)," Thursday, May 27, 2010.
- Board of Governors of the Federal Reserve System, "Flow of Funds Accounts of the United States," June 10, 2010.
- Board of Governors of the Federal Reserve System, "Guide to the Flow of Funds Accounts."
- Brace, A., D. Gatarek, and M. Musiela, "The Market Model of Interest Rate Dynamics," *Mathematical Finance*, 7 (2), 1997, pp. 127–154.
- Brigo, D., and Mercurio, F., *Interest Rate Models: Theory and Practice*, Springer, 2001.
- Cochrane, J. H., and M. Piazzesi, "Bond Risk Premia," *American Economic Review*, 95, 2005, pp. 138–160.
- Cochrane, J. H., and M. Piazzesi, "Decomposing the Yield Curve," Working Paper, March 13, 2008.
- Congressional Budget Office, "Updated Estimates of the Subsidies to the Housing GSEs," April 8, 2004.
- Duffie, D., and H. Zhu, "Does a Central Clearing Counterparty Reduce Counterparty Risk?" Stanford University, July 1, 2009.
- Ejsing, J.W., and J. Sihvonen, "Liquidity Premia in German Government Bonds," European Central Bank Working Paper Series, no. 1081, August 2009.
- Eurostat, "Financial Assets and Liabilities of Households in the European Union," 2009.
- Federal Housing Finance Agency, "U.S. Treasury Support for Fannie Mae and Freddie Mac," Mortgage Market Note 10-1, January 20, 2010.
- Fleming, M., and K. Garbade, "Repurchase Agreements with Negative Interest Rates," *Current Issues in Economics and Finance*, Volume 10, Number 5, April 2004.
- Friedman, P., testimony before the Financial Crisis Inquiry Commission, May 5, 2010.

- Gatheral, Jim, *The Volatility Surface*, John Wiley & Sons, 2006.
- Glasserman, Paul, *Monte Carlo Methods in Financial Engineering*, Springer, 2003.
- Hagan, P.S., D. Kumar, A.S. Lesniewski, and D.E. Woodward, "Managing Smile Risk," *Wilmott Magazine*, September 2002, pp. 84–108.
- Hasenpusch, T., *Clearing Services for Global Markets*, Cambridge University Press, 2009.
- Heath, D., R. Jarrow, and A. Morton, "Bond Pricing and the Term Structure of Interest Rates: A New Methodology for Contingent Claims Valuation," *Econometrica*, Vol 60, Issue 1, 1992, pp. 77–105.
- Ho, T., "Key Rate Duration: A Measure of Interest Rate Risk," *Journal of Fixed Income*, September, 1992.
- Homer, S., and R. Sylla, *A History of Interest Rates*, 3rd edition Revised, Rutgers University Press, 1996.
- Ingersoll, J., *Theory of Financial Decision Making*, Rowman & Littlefield, 1987.
- Lantz, Carl, Michael Chang, and Sonam Leki Dorji, "A Guide to the Front-End and Basis Swap Markets," Credit Suisse Fixed Income Research, February 18, 2010.
- Ludvigson, S. C., and S. Ng, "Macro Factors in Bond Risk Premia," The Society for Financial Studies, Oxford University Press, 2009.
- Martin, P.P., "Comparing Replacement Rates Under Private and Federal Retirement Systems," *Social Security Bulletin*, Vol. 65, No. 1, 2003–2004.
- McCormick, L.C., "Treasury Traders Paid to Borrow as Fed Examines Repos," Bloomberg, November 24, 2008.
- Melamed, L., *For Crying Out Loud*, John Wiley & Sons, Inc., 2009.
- Rebonato, R., *Modern Pricing of Interest Rate Derivatives: The LMM and Beyond*, Princeton University Press, 2002.
- Reuters, "GSE Debt Not Sovereign—U.S. Treasury Secy Geithner," Reuters, April 6, 2010.
- Rochelle, B., "Lehman, Summit, OTB, Townsends, Vitro: Bankruptcy," Bloomberg, January 26, 2011.
- Stigum, M., *The Money Market*, 3rd Edition, Dow Jones-Irwin, 1990.
- Tajika, E., "The Public Pension System in Japan: The Consequences of Rapid Expansion," The International Bank for Reconstruction and Development/The World Bank, 2002.
- Tuckman, B., "Systemic Risk and the Tri-Party Repo Clearing Banks," *Center for Financial Stability*, February 2010.
- Tuckman, B., "Amending Safe Harbors to Reduce Systemic Risk in OTC Derivatives Markets," *Center for Financial Stability*, April 2010.
- Vasicek, O., "An Equilibrium Characterization of the Term Structure," *Journal of Financial Economics*, 5, 1977, pp. 177–178.
- Watson Wyatt, "Global Pension Assets Study," 2008.
- Zubrow, B., testimony before the Financial Crisis Inquiry Commission, September 1, 2010.

Exercises

CHAPTER 1

- 1.1 What are the cash flow dates and the cash flows of \$1,000 face value of the U.S. Treasury $2\frac{3}{4}$ s of May 31, 2017, issued on May 31, 2010?
- 1.2 Use this table of U.S. Treasury bond prices for settle on May 15, 2010, to derive the discount factors for cash flows to be received in 6 months, 1 year, and 1.5 years.

Bond	Price
$4\frac{1}{2}$ s of 11/15/2010	102.15806
0s of 5/15/2011	99.60120
$1\frac{3}{4}$ s of 11/15/2011	101.64355

- 1.3 Suppose there existed a Treasury issue with a coupon of 2% maturing on November 15, 2011. Using the discount factors derived from Question 1.2, what would be the price of the 2s of November 15, 2011?
- 1.4 Say that the 2s of November 15, 2011, existed and traded at a price of 101 instead of the price derived from Question 1.3. How could an arbitrageur profit from this price difference using the bonds in the earlier table? What would that profit be?
- 1.5 Given the prices of the two bonds in the table as of May 15, 2010, find the price of the third by an arbitrage argument. Since the $3\frac{1}{2}$ s of 5/15/2020 is the on-the-run 10-year, why might this arbitrage price not obtain in the market?

Bond	Price
0s of 5/15/2020	69.21
$3\frac{1}{2}$ s of 5/15/2020	?
$8\frac{3}{4}$ s of 5/15/2020	145.67

CHAPTER 2

- 2.1 You invest \$100 for two years at 2% compounded semiannually. How much do you have at the end of the two years?
- 2.2 You invested \$100 for three years and, at the end of those three years, your investment was worth \$107. What was your semiannually compounded rate of return?
- 2.3 Using the discount factors in the table, derive the corresponding spot and forward rates.

Term	Discount Factor
.5	.998752
1	.996758
1.5	.993529

- 2.4 Are the forward rates above or below the spot rates in the answers to Question 2.3? Why is this the case?
- 2.5 Using the discount factors from question 2.3, price a 1.5-year bond with a coupon of .5%. If over the subsequent 6 months the term structure remains unchanged, will the price of the .5% bond increase, decrease, or stay the same? Try to answer the question before calculating and then calculate to verify.

CHAPTER 3

- 3.1 The price of the $\frac{3}{4}$ s of May 31, 2012 was 99.961 as of May 31, 2010. Calculate its price using the discount factors in Table 2.3. Is the bond trading cheap or rich to those discount factors? Then, using trial-and-error, express the price difference as a spread to the spot rate curve implied by those discount factors.
- 3.2 The yield of the $\frac{3}{4}$ s of May 31, 2012, was .7697% as of May 31, 2010. Verify that this is consistent with the price in Question 3.1.
- 3.3 The price of the $4\frac{3}{4}$ s of May 31, 2012, was 107.9531 as of May 31, 2010. What was the yield of the bond? Please solve by trial-and-error.
- 3.4 Did you get a higher yield for the $4\frac{3}{4}$ s from Question 3.3 than the yield of the $\frac{3}{4}$ s given in Question 3.2? Is that what you expected? Why or why not?
- 3.5 An investor purchases the $4\frac{3}{4}$ s of May 31, 2012 on May 31, 2010, at the yield given in Question 3.3. Exactly six months later the investor sells the bond at that same yield. What is the price of the bond on the

- sale date and what is the investor’s total return from the bond over those six months?
- 3.6 Interpret your answer to Question 3.5. a) In what way is the return significant or interesting? b) Explain why an investor would buy a premium bond when that bond is worth only par at maturity? How does this relate to your work in Question 3.5?
- 3.7 Re-compute the sample return decomposition of Tables 3.2 and 3.3 of the text, replacing the assumption of realized forwards with the assumption of an unchanged term structure.
- 3.8 Start with any upward-sloping term structure, e.g., from C-STRIPS prices or even some made-up rates. Then replicate the 0-coupon, par, and 9% coupon curves in Figure 3.2. Add a curve for a security that makes equal fixed payments to various maturities, i.e., a mortgage.
- 3.9 In the subsection “News Excerpt: Sale of Greek Government Bonds in March, 2010,” approximately what is the yield on seven-year Spanish debt?
- 3.10 Return to Table 1.7 in the text, which shows that the $3\frac{1}{2}$ s of May 15, 2020, are 2.076 per 100 face amount away from being correctly priced by C-STRIPS while the $8\frac{3}{4}$ s maturing on the same date are .338 per 100 face amount away. According to the discussion of the text, this difference is due to the on-the-run premium of the $3\frac{1}{2}$ s that is reflected in the price of its P-STRIPS. As of the same pricing date, however, the yields of the $3\frac{1}{2}$ s and $8\frac{3}{4}$ s were only a few basis points apart, i.e., nothing like the difference justified by the more than 2% premium on the price of the final principal payment. How is this possible?

CHAPTER 4

- 4.1 The following tables give the prices of TYU0 and of TYU0C 120 as of May 2010 for a narrow range of the 7-year par rate. Please fill in the other columns, ignoring cells marked with an “X.” Over the given range, which security’s price-rate function is concave and which convex? How can you tell?

	TYU0					
Rate	Price	DV01	Duration	Convexity	1st Deriv	2nd Deriv
3.320%	115.5712	X	X	X	X	X
3.412%	114.8731			X		X
3.504%	114.1715					
3.596%	113.4668			X		X
3.688%	112.7591	X	X	X	X	X

TYU0C 120						
Rate	Price	DV01	Duration	Convexity	1st Deriv	2nd Deriv
3.320%	.4564	X	X	X	X	X
3.412%	.3483			X		X
3.504%	.2619					
3.596%	.1940			X		X
3.688%	.1415	X	X	X	X	X

- 4.2 Using the data in Question 4.1, how would a market maker hedge the purchase of \$50 million face amount of TYU0C with TYU0 when the 7-year par rate is 3.596%? Check how well this hedge works by computing the change in the value of the position should the rate move instantaneously from 3.596% to 3.668%. What if the rate falls to 3.320%? Is the P&L of the hedged position positive or negative? Why is this the case?
- 4.3 Using the data from the answer to Question 4.1, how much would an investment manager make from \$100mm of TYU0C if the rate instantaneously fell from 3.504% to 3.404%? Use a duration estimate.
- 4.4 Using the data in Question 4.1, provide a 2nd order estimate of the price of TYU0C should the 7-year par rate be 3.75%.
- 4.5 The table below gives the prices, durations, and convexities of three bonds. a) What is the duration and convexity of a portfolio that is long \$50mm face amount of each of the 5- and 10-year bonds? b) What portfolio of the 5- and 30-year bonds has the same price and duration as the portfolio of part a)? c) Which of the two portfolios has the greater convexity and why?

Coupon	Maturity	Price	Duration	Convexity
2.50%	5 years	102.248	4.687	25.052
2.75%	10 years	100.000	8.691	86.130
3%	30 years	95.232	19.393	495.423

- 4.6 The following table gives yields, DV01s, and durations for three 15-year bonds. The three coupon rates are 0%, 3.5%, and 7%. Which coupon rate belongs to which bond? What is the shape of the term structure of spot rates underlying the valuation of these bonds?

Bond	Yield	DV01	Duration
#1	3.50%	.1159	11.59
#2	3.50%	.0876	14.75
#3	3.50%	.1443	10.26

CHAPTER 5

- 5.1 Using the following instructions, complete a spreadsheet to compute the two-year and five-year key-rate duration profiles of four-year bonds. For the purposes of this question, key-rate shifts are in terms of spot rates.
- (a) In Column A put the coupon payment dates in years, from .5 to 5 in increments of .5. Put a spot rate curve, flat at 3%, in Column B. Put the discount factors corresponding to this spot rate curve in Column C. Now price a 3% and an 8% four-year coupon bond under this initial spot rate curve.
 - (b) Create a new spot rate curve in Column D by adding a two-year key rate shift of 10 basis points. Compute the new discount factors in Column E. What are the new bond prices?
 - (c) Create a new spot rate curve in Column F by adding a five-year key rate shift of 10 basis points. Compute the new discount factors in Column F. What are the new bond prices?
 - (d) Use the results in parts (a) through (c) to calculate the key-rate duration profiles of each of the bonds.
 - (e) Sum the key-rate durations for each bond to obtain the total durations. Calculate the percentage of the total duration attributed to each key rate for each bond. Comment on the results.
 - (f) What would the key-rate duration profile of a four-year zero coupon bond look like relative to those of these coupon bonds? How about a five-year zero coupon bond?
- 5.2 Continue with the setting and results of Question 5.1. Verify that a 3% two-year bond has a duration of 1.925 that is completely concentrated as a two-year key-rate duration. How would one hedge the key-rate risk profile of the 8% four-year bond with the 3% two-year bond and the 3% four-year bond? Note that the total value of the 8% bond and of the hedge need not be the same. Comment on the result.
- 5.3 Use Table 5.6 for this question. A trader constructs a butterfly portfolio that is short €100mm of the 10-year swap and long 50% of the 10-year swap's total '01 in 5-year swaps and 50% of the 10-year swap's total '01 in 15-year swaps. What are the forward-bucket exposures of the resulting portfolio?

CHAPTER 6

The following introduction applies to Questions 6.1 through 6.5.

You are a market maker in long-term EUR interest rate swaps. You typically have to hedge the interest rate risk of having received from or paid to a customer on a 20-year interest rate swap. Given the transaction

costs of hedging with both 10s and 30s and the relatively short time you wind up having to hold any such hedge, you consider hedging these 20-year swaps with either 10s or 30s but not both. To that end you run two single-variable regressions, both with changes in the 20-year EUR swap rates as the dependent variable, but one regression with changes in the 10-year swap rate as the independent variable and the other with changes in the 30-year swap rate as the independent variable. The results over the period July 1, 2009, to July 3, 2010, are given in the following table.

Number of Observations	259			
Independent variable	Change in 10-year		Change in 30-year	
R-squared	89.9%		96.3%	
Standard Error	1.105		.666	
Regression Coefficients	Value	Std. Error	Value	Std. Error
Constant	−.017	.069	−.008	.042
Independent variable	1.001	.021	.917	.011

- 6.1 What are the 95% confidence intervals around the constant and slope coefficients of each regression?
- 6.2 Use the confidence intervals just derived. Can you reject a) the hypothesis that the constant in the 10-year regression equals 0? b) That the slope coefficient in the 30-year regression equals 1?
- 6.3 As the swap market maker, you just paid fixed in €100 million notional of 20-year swaps. The DV01s of the 10-, 20-, and 30-year swaps are .0864, .1447, and .1911, respectively. Were you to hedge with 10-year swaps, what would you trade to hedge? And with 30-year swaps?
- 6.4 Approximately what would be the standard deviation of the P&L of a hedged position of 20-year swaps with 10-year swaps? And if hedged with 30-year swaps?
- 6.5 If you were to hedge with one of either the 10- or 30-year swaps, which would it be and why?
- 6.6 Use the principal components in Table 6.5 and the par swap data in Table 6.6 to hedge 100 face amount of 10-year swaps with 5- and 30-year swaps with respect to the first 2 principal components.

CHAPTER 7

- 7.1 A fixed income analyst needs to estimate the price of an interest rate caplet that pays \$1,000,000 next year if the one-year Treasury rate exceeds 3% and pays nothing otherwise. Using a macroeconomic model developed in another area of the firm, the analyst estimates that the one-year Treasury rate will exceed 3% with a probability of 25%.

Since the current 1-year rate is 1%, the analyst prices the caplet as follows:

$$\frac{25\% \times \$1,000,000}{1.01} = \$247,525$$

Comment on this pricing procedure.

- 7.2 Assume that the true 6-month rate process starts at 5% and then increases or decreases by 100 basis points every 6 months. The probability of each increase or decrease is 50%. The prices of 6-month, 1-year, and 1.5-year zeros are 97.5610, 95.0908, and 92.5069. Find the risk-neutral probabilities for the six-month rate process over the next year (i.e., two steps for a total of three dates, including today). Assume, as in the text, that the risk-neutral probability of an up move from date 1 to date 2 is the same from both date 1 states. As a check to your work, write down the price trees for the 6-month, 1-year, and 1.5-year zeros.
- 7.3 Using the risk-neutral tree derive for Question 7.2, price \$100 face amount of the following 1.5-year *collared floater*. Payments are made every six months according to this rule. If the short rate on date i is r_i then the interest payment of the collared floater on date $i + 1$ is $\frac{1}{2}3.50\%$ if $r_i < 3.50\%$; $\frac{1}{2}r_i$ if $6.50\% \geq r_i \geq 3.50\%$; $\frac{1}{2}6.50\%$ if $r_i > 6.50\%$. In addition, at maturity, the collared floater returns the \$100 principal amount.
- 7.4 Using your answers to Questions 7.2 and 7.3, find the portfolio of the originally 1-year and 1.5-year zeros that replicates the collared floater from date 1, state 1, to date 2. Verify that the price of this replicating portfolio gives the same price for the collared floater at that node as derived for Question 7.3
- 7.5 Using the risk-neutral tree from Question 7.2, price \$100 notional amount of a 1.5-year *participating cap* with a strike of 5% and a *participation rate* of 40%. Payments are made every six months according to the following rule. If the short rate on date i is r_i then the cash flow from the participating cap on date $i + 1$ is, as a percent of par, $\frac{1}{2}(r_i - 5\%)$ if $r_i \geq 5\%$ and $\frac{1}{2}40\%(r_i - 5\%)$ if $r_i < 5\%$. There is no principal payment at maturity.
- 7.6 Question 7.3 required the calculation of the price tree for a collared floater. Repeat this exercise under the same assumptions, but this time assume that the OAS of the collared floater is 10 basis points.
- 7.7 Using the price trees from Questions 7.3 and 7.6, calculate the return to a hedged and financed position in the collared floater from dates 0 to 1 assuming no convergence (i.e., the OAS on date 1 is also 10 basis points.) Hint #1: Use all of the proceeds from selling the replicating portfolio to buy collared floaters. Hint #2: You do not need to know

the composition of the replicating portfolio to answer this question. Is your answer as you expected? Explain.

- 7.8 What is the return if the collared floater converges on date 1, i.e., its OAS equals 0 on that date?

CHAPTER 8

- 8.1 Describe as fully as possible the qualitative effect of each of these changes on the instantaneous rates 10 and 30 years forward.
- The market risk premium increases.
 - Volatility across the curve increases.
 - Rates are not expected to increase as much as previously.
 - The market risk premium falls and volatility falls in such a way as to keep the 10-year forward rate unchanged.

CHAPTER 9

- 9.1 Assume an initial interest rate of 5%. Using a binomial model to approximate normally distributed rates with weekly time steps, no drift, and an annualized volatility of 100 basis points, what are the two possible rates on date 1?
- 9.2 Add a drift of 20 basis points per year to the model described in Question 9.1. What are the two rates now?
- 9.3 Consider the following segment of a binomial tree with 6-month time steps. All transition probabilities equal .5.

		5.360%
	4.697%	
3.99%		4.648%
	3.964%	
		4.031%

Does this tree display mean reversion?

- 9.4 What mean reversion parameter is required to achieve a half-life of 15 years?

CHAPTER 10

- 10.1 The yield volatility of a short-term interest rate is 20% at a level of 5%. Quote the basis point volatility and the CIR volatility parameter.

10.2 You are told that the following tree was built with a constant volatility. All probabilities equal .5. Which volatility measure is, in fact, a constant?

		5.49723%
	4.69740%	
4.014%		4.63919%
	3.96420%	
		3.91507%

10.3 Use the closed form solution for the Vasicek model in Appendix A in Chapter 10 to compute the spot rate of various terms with the parameters $\theta = 10\%$, $k = .035$, $\sigma = .02$, and $r_0 = 4\%$. Comment on the shape of the term structure.

CHAPTER 13

- 13.1 As of a spot settlement date of June 1, 2010, find the forward price of the U.S. Treasury 3⁵/₈s of February 15, 2020, for delivery on September 30, 2010. The spot price is 102 – 21 and the repo rate is .3%.
- 13.2 Using your answer to Question 13.1, compute the forward yield of the 3⁵/₈s to September 30. Use equation (3.32) and a spreadsheet or other application.
- 13.3 Use the risk-neutral tree, with annual steps, developed in the section “Arbitrage Pricing in a Multi-Period Setting” in Chapter 9. Consider a 5% 10-year bond that, 2 years from the starting date, takes on the values 104.701, 98.126, and 92.061, corresponding to the nodes 4%, 5%, and 6%, respectively. What is the forward price of the bond for delivery in two years? What is the futures price to that same delivery date?

CHAPTER 14

- 14.1 The conversion factor of the 4s of August 15, 2018, into TYU0 is .8774. If the price of TYU0 is 121.2039, what is the (flat) delivery price of the 4s at that time? If the price of the 4s is 107.1652, what is their cost of delivery at that time?
- 14.2 The following table gives the prices of TYU0 and of its deliverable bonds in a particular term structure scenario on the delivery date that corresponds to a 7-year par yield of 3.32%. Conversion factors are also provided. Which bond is CTD in this scenario?

Futures Price:		117.2606	
Rate	Maturity	Price	Conv. Factor
$3\frac{1}{4}$	3/31/17	100.1567	.8538
$4\frac{1}{2}$	5/15/17	107.9783	.9202
$3\frac{1}{8}$	4/30/17	99.3447	.8471
$2\frac{3}{4}$	5/31/17	96.9980	.8272
$4\frac{3}{4}$	8/15/17	109.3542	.9314
$4\frac{1}{4}$	11/15/17	106.0361	.9012
$3\frac{7}{8}$	5/15/18	102.7277	.8732
4	8/15/18	103.2276	.8774
$3\frac{1}{2}$	2/15/18	100.6805	.8547
$3\frac{3}{4}$	11/15/18	101.0422	.8587
$3\frac{5}{8}$	8/15/19	98.9248	.8401
$3\frac{1}{8}$	5/15/19	95.5539	.8107
$2\frac{3}{4}$	2/15/19	93.3962	.7909
$3\frac{3}{8}$	11/15/19	96.8798	.8195
$3\frac{5}{8}$	2/15/20	98.6522	.8332
$3\frac{1}{2}$	5/15/20	97.6531	.8210

- 14.3 The $\frac{3}{4}$ s of May 31, 2012, are deliverable into TUU0, the September 2010 2-year note contract. Assume that the delivery date of the contract is September 30, 2010. The notional coupon of the contract is 6%. Approximately what is the conversion factor of the $\frac{3}{4}$ s for delivery into that contract?
- 14.4 The forward price of the $3\frac{1}{2}$ s of February 15, 2018, to September 30, 2010, is 103.1303. Its conversion factor for delivery into TYU0 is .8547. If the price of TYU0 is 120, what is the net basis of the $3\frac{1}{2}$ s in ticks?
- 14.5 A trader sells \$50 million $3\frac{1}{2}$ s net basis at the level you calculated in question 14.4. What is the trader's position in the bond and in TYU0? If the net basis of $3\frac{1}{2}$ s is 10 ticks as of the delivery date, what is the trader's profit or loss?
- 14.6 Figure 14.4 of the text graphs various net bases as option-like payoffs. Describe how each of the following deliverable bonds would look if added to this graph: the $3\frac{1}{4}$ s of 3/31/2017; the 4s of 8/15/2018; and the $3\frac{1}{2}$ s of 5/15/2022.
- 14.7 How would the graphs in Figure 14.4 change if the curve steepened as the 7-year par rate increased?

CHAPTER 15

- 15.1 As of May 28, 2010, you are financing \$100 million worth of inventory of bonds in the repo market on an overnight basis. You plan to hold these bonds until mid-September 2010. Using both the Eurodollar and Fed funds futures listed in Tables 15.3 and 15.11, what trades can you do to hedge against the risk that rates rise and increase your borrowing cost? Will you have to adjust the hedge at all between May 28 and mid-September?
- 15.2 Approximately what borrowing rate is locked in by the hedge in Question 15.1?
- 15.3 Instead of the hedge constructed in Question 15.1, you decide to use only Eurodollar contracts. How does the hedge change? How does the locked-in rate compare with the previous hedge? Is this new hedge riskier in any way than the previous hedge?
- 15.4 As of May 28, 2010, a 5% U.S. Treasury bond maturing on September 15, 2010, had a full price of 102.4055. Using the dates and rates of Table 15.8, calculate the TED spread of the bond.
- 15.5 As of the end of July 2004, the fed funds target rate stood at 1.25%. Say that the August fed fund futures rate at that time was 1.3516%. What is the market implied probability of a 25 basis-point increase at the August 10 meeting? If you're willing to assume a 50% chance of no change in policy, what are the implied probabilities of 25 and 50 basis-point increases?

CHAPTER 16

- 16.1 Recalculate swap cash flows as in Table 16.1 for a 1.5% swap rate and a LIBOR rate that starts at 50 basis points on June 2, 2010, but then increases by 50 basis points every three months, reaching 4% by March 2, 2012.
- 16.2 Under the same simplifying assumptions used to price the CMS swap in Table 16.3, what is the fair fixed rate against the 10-year swap rate paid annually for four years starting in year 2?

CHAPTER 17

- 17.1 Consider 1- and 2-year swaps of annually-paying fixed vs. annually-paying LIBOR with par rates of 2% and 2.75%, respectively. The investable and collateral rates, given by the OIS curve, are 1% for 1 year and 2% for 1 year, 1 year forward. What is the NPV of receiving 3% for two years against LIBOR?

- 17.2 Using the data from Question 17.1, what is the implied term structure of basis swap spreads of OIS vs. LIBOR?

CHAPTER 18

- 18.1 Using the formulae in the text, recalculate the value of the caplet in Table 18.1 by changing only the volatility from 77.22 basis points to 120 basis points.
- 18.2 Using the formulae in the text, calculate the value of an at-the-money 2y5y receiver swaption on \$100 million notional when the term structure is flat at 4% and the appropriate volatility is 100 basis points.

CHAPTER 19

- 19.1 Consider a 10-year corporate bond with a coupon of 6%. The semi-annual compounded swap curve is flat at 4% and the corporate bond is trading at a LIBOR OAS of 3%. Calculate the par-par asset swap spread and the market value asset swap spread.
- 19.2 Say that the cumulative default rate over a 10-year horizon for some category of corporate bonds is 5%. If the recovery rate is 40%, what is the spread that just compensates investors for expected losses?
- 19.3 The quoted spread on a one-year quarterly paying CDS is 110 basis points while the standardized coupon is 100 basis points. Let the assumed recovery rate be 40% and let the quarterly compounded term structure of swap rates be flat at 3%. What is the up-front payment for \$10 million notional of the CDS? You will need to construct a spreadsheet to perform these calculations.
- 19.4 Create a spreadsheet to recreate the duration calculations in the subsection “The DV01 or Duration of a Bond with Credit Risk.” Use this spreadsheet to compute the duration in the example of that subsection with a coupon of 8% instead of 6%, keeping the yield at 14%.

CHAPTER 20

- 20.1 Assume that the term structure of monthly compounded rates is flat at 6%. Find the monthly payment of a \$100,000 15-year level-pay mortgage.
- 20.2 For the mortgage in Question 20.1, what is the interest component of the monthly payment after five years?
- 20.3 An adjustable-rate mortgage (ARM) resets the interest rate periodically. How does the refinancing option of an ARM compare with the option to prepay a fixed rate mortgage?

- 20.4 Explain the intuition for each of the following results:
- (a) When interest rates fall, holding all else equal, POs outperform 30-year fixed rate securities.
 - (b) When interest rates rise by 100 basis points, mortgage pass-throughs fall by about 7%. When interest rates fall by 100 basis points, pass-throughs rise by 4%.
 - (c) When interest rates decline, IOs and inverse IOs decline in price, but IOs suffer more severely. (Like an IO, an inverse IO receives no principal payments but receives interest payments that float inversely with the level of rates.)
- 20.5 Recompute the value of the roll in the example of the text for a coupon of 6%, a paydown percentage of 3%, and an August TBA price of 102.1. Keep all other quantities the same.

A

Accrued interest, 51, 53, 62–65, 115, 116–117
 Agency securities, 8
 Federal reserve bank balance sheet, 17
 Annuity factor, 76
 CV01, 164
 Yield-based formula, 101–102
 Arbitrage-free models, 229–230, 259, 260–262, 282
 Arbitrage pricing, 48, 51, 67–68
 Derivatives, 209–211, 214–221, 226
 Financing arrangements, 325–326, 457–470, 475–480
 Forward agreements, 352–355
 Law of one price, 56–58
 LIBOR swaps with fed funds as the investable and collateral rate, 470–473
 Replicating portfolio, 56–57, 65–67, 202, 209–211, 219, 226, 461–462
 Asset-backed securities
 Issuers, 17–18
 Japan, 45
 On- and off-balance sheet securitizations, 17–18
 Asset swap spreads, 528

B

Bank of Japan (BOJ), 35, 42–43
 Banks
 Europe, 28–31
 Japan, 37–38
 United States, 13–16
 Barbell vs. bullet, 123, 150–152
 Basel II, 30

Basis point, 84

Basis swaps, *see Swaps, Basis*

Benchmark rates and securities, 48, 49

Curve construction, 591–592

Key rate analysis, 154

Partial '01s and PV01, 163

Bermudan option, 201, 491 fn. 3

Bilateral agreements, 446

Binomial tree, 207

Black-Karasinski model, 282–284

Black-Scholes models, 204, 226–227, 483

Bond futures options, 498–499, 501, 520

Bond options, 492–493, 501, 515–517

Caplets and caps, 484–486, 500, 513–514

Expectation formulae, 520–521

Euribor futures options, 498, 500, 518–519

Eurodollar futures options, 496–497, 500, 517–518

Floorlet, 486–487, 500

Skew, 483, 490, 500–503, 507–508

Swaptions, 487–490, 501, 514–515

Theoretical foundations for applications, 507–513

Bond covenants, 491

Bond options, 491–495, 501

Call protection, 491

Embedded, 491

Interest rate behavior, 494–495

Valuation with Black-Scholes models, 492–493, 515–517

Valuation with term structure models, 491–492

Broker-Dealers, 19

- Bullet portfolio, *see* *Barbell vs. bullet*
- Business days and Treasury coupon payments, 52
- C
- Capital asset pricing model (CAPM), 240
- Caplets and caps, 483–486, 500, 513–514
- Carry, *see* *Carry-Roll-Down*
- Carry-Roll-Down, 95, 106–114, 116, 223
 - Forward swap rates, May 28, 2010, 85–87
 - Scenarios, 110–114
- Cash carry, 95, 105, 106, 107, 116–117, 356–357, 384
- Central counterparty, 446–447
- CDS, *see* *Credit default swaps*
- Cheap securities, 55
- Chinese ownership of U.S. Treasuries, 2–3
- Clean price, *see* *Quoted price*
- Clearinghouse, 446
- Clearing mandates for over-the-counter derivatives, 445–449
- Commercial paper
 - Europe, 33
 - Financing risk implications, 14–15
 - Special purpose vehicles, 18
 - United States, 10
- Compounding, 69–71, 87–89
 - Annual 87–88
 - Continuous, 71, 88–89
 - Daily, 87–88
 - Monthly, 87–88
 - Semiannual, 70–71, 87–88, 115
- Concave curve, 125
- Consistent estimators, 184
- Constant-maturity Treasury swap, 219–221, 221–222
- Contingent claim, 201
- Convex curve, 124–125
- Convexity, 123, 132–135
 - Asset-liability management, 140
 - Correction for CMS, 452–456
 - Estimating price changes, 137–139
 - Futures prices and rates, 369–371, 389
 - Model 1, 256
 - Portfolio, 142
 - Short position, 136
 - Term structure shape, 231–236, 239–240, 241, 243–248
 - Vasicek model, 271–272
 - Yield-based, 120, 123, 149–150
 - Barbell vs. bullet, 151–152
 - Par bond, 150
 - Zero-coupon bond, 150
- Corporate bonds, 5, 49, 527–529
 - Asset swap spreads, 536–540
 - Bond spread, 535
 - Callable, 47
 - CDS-bond basis, 553–554, 559–561
 - Credit spreads, 534–543
 - Day-count convention, 65
 - DV01 or duration with credit risk, 556–557
 - Financing risk implications, 14
 - Hedging issuance, 444
 - Life insurance company assets, 18, 28
 - Negative basis trade, 554–556
 - Private placements, 13
 - Public issues, 13
 - Ratings, default, and recovery, 529–534
 - Refinancing, 5
 - Yield spread, 534
- Cox-Ingersoll-Ross model, 277–280, 285
- Coupon effect, 102–104
- Coupon value of an '01 (CV01), 164
- Courtadon model, 278
- Credit default swaps (CDS), 527–528
 - CDS-Bond basis, 553–554, 559–561
 - Default swap spread, 547
 - Definitions and mechanics, 545–548
 - Index, 557–559
 - Negative basis trade, 554–556
 - Physical vs. cash settlement, 546–547
 - Spread, 547
 - Standardization, 548
 - Up-front payment, 547–553

- Credit risk, 14, 527–528
 - Corporate bonds, 528–543
 - Credit spreads, 534–545
 - Default rates, 530–533, 543–545, 559
 - Downgrades and upgrades, 532
 - Recovery rates, 530–533
- Credit support annex (CSA), 441
- Credit value adjustment, (CVA), 443
- Curve construction, 591–605
 - Bootstrapping, 592–593
 - Flat forwards, 593–600
 - Locality properties and hedging, 604–605
 - Smooth forwards, 600–603
 - Stub rate, 594
- Curve fitting, 261
- CV01, *see* *Coupon value of an '01*
- D**
- Day-count conventions, 51, 65
 - Actual/actual, 63
 - Actual/360, 65
 - Money markets, 65, 164
 - Swaps, fixed side, 164
 - 30/360, 65
- Debt markets, *see* *Fixed income markets*
- Defined-benefit pension plans, *see* *Pension funds*
- Defined-contribution pension plans, *see* *Pension funds*
- Dependent variable, 174
- Deposits
 - Fed reserve system, 16
 - Financing risk implications, 14–15, 28
 - Household financial assets, 21, 36
 - Japanese banks, 37
 - Japanese postal savings system, 40
 - Monetary policy, 16–17
 - U.S. commercial banks, 13, 15
- Derivatives of the price-rate function, 126, 130, 133–135, 138
- Derivatives of the price-yield function, 142–143, 149
- Derivatives
 - Fixed income vs. equity, 225–227
 - Exchange-traded, notional amounts outstanding, 2–3
 - Over-the counter
 - Clearing mandates, 445–449
 - Notional amounts outstanding, 1–2
 - Dirty price, *see* *Full price*
 - Discount to face value, 101
 - Discount factors, 48, 51, 53–55, 58, 67–68, 73
 - Treasury STRIPS, 59–60
 - Relationship to rates, 69, 74–78
 - Discount securities, 65
 - Dollar value of an '01 (DV01), 123–130, 153, 184
 - Asset-liability management, 140
 - Estimating price changes, 137–139
 - Forward rate agreements, 359–361
 - Negative, 129, 589
 - Portfolio, 141
 - Relationship to convexity, 132–137
 - Vs. duration, 130–132
 - Yield-based, 120, 123, 126, 142–149
 - Hedge of TIPS vs. nominal bonds, 172–174, 177–178
 - Par bonds, 146, 148–149
 - Perpetuities, 146, 148
 - Premium bond, 148–149
 - Vs. key-rate '01 totals, 159
 - Zero-coupon bonds, 145–146, 149
 - Drift, 212, 217, 251, 262
 - Ho-Lee Model, 259–260
 - LMM, 300–307
 - Model 1, 251–253
 - Model 2, 257–259
 - Original Salmon Brothers model, 280
 - Vasicek model, 262–265
 - Time-dependent, 259
 - Duration, 123, 130–132, 154–155
 - Adjusted, 143
 - Asset-liability management, 140
 - Effective, 132
 - Estimating price changes, 137–139
 - Macauley, 143 fn. 8
 - Modified, 143

- Duration (*Continued*)
 Portfolio, 141–142
 Risk premium, 237–239, 242–244
 Yield-based, 120, 123, 142–149
 Par bonds, 146–147
 Perpetuities, 146–147
 Vs. key-rate '01 totals, 159
 Zero-coupon bonds, 145–147
 DV01, *see Dollar value of an '01*
- E
- Efficient estimators, 184
- Embedded options, *see Bond options*
- EONIA, *see Euro Overnight Index Average*
- Equilibrium models, 229, 259, 260–262
- Expectations, 230–231
- Euribor, 407, 437
- Euribor futures, *see Futures, Euribor*
- Eurodollar futures, *see Futures, Eurodollar*
- Euro Overnight Index Average (EONIA), 33, 431
- European Central Bank (ECB), 31–32
 Bank funding, 28, 30, 32
 Deposit facility, 31–32
- Eurozone
 Definition, 1
 Debt markets, 1, 2, 19–34
- Exotic derivative, 201, 203, 298
- F
- Factor models, 119–121, 123, 202–205, 218, 226–227
- Federal Deposit Insurance Corporation (FDIC) insurance, 14
- Federal Funds, 417–419
 Commercial banking source of funds, 13
 Effective rate, 418
 Financing risk implications, 14–15
 Target rate, 417
 Vs. general collateral rates, 340
- Fed funds futures, *see Futures, Fed funds*
- Federal Home Loan Banks (FHLB), *see Government sponsored entities*
- Federal Home Loan Mortgage Corporation (FHLMC), *see Government sponsored entities*
- Federal National Mortgage Association (FNMA), *see Government sponsored entities*
- Federal Open Market Committee (FOMC), 417–418
 Market expectations at the start of the 2004–2006 tightening cycle, 424–429
- Federal Reserve System, 16–17
- Financed bond positions, 97–98, 105, 106, 107, 110, 116–117, 459–461
- Financial crisis of 2007–2009
 Bear Stearns, 334–336
 Cash-out mortgage refinancing, 576
 European bank real estate investments, 30
 European Central Bank response, 32
 Federal reserve bank balance sheet, 17
 Fed funds rate, 419
 Impact on balance sheets, 12, 13, 15
 Japanese bank holdings of government debt, 38
 JPMorgan Chase's exposure to Lehman Brothers, 336–339
 Leverage of the banking sector, 16
 LIBOR-OIS spread, 431–432, 438
 Mortgage defaults and modifications, 580
 Mortgage OAS, 586
 OTC derivatives, 446
 Regression hedge of 10-year EUR swaps, 183–184
 Repo financing, 19, 327, 329, 333–339
 Special purpose vehicles, 18
 Standardization of CDS, 548
 Subprime mortgages, 564
 Swaps pricing, 69 fn. 1, 326, 437–438

- Financing risk, 327, 333–339
 - Banks, 14–15, 28
 - Broker-Dealers, 19
 - Corporations, 444
 - Fixed income markets
 - United States, 1–3, 4–19
 - Europe, 19–34
 - Money markets, 33–34
 - Eurozone, 1–3
 - Japan, 1–3, 34–45
 - United Kingdom, 1–3, 23, 24–25
 - Flat price, *see Quoted price*
 - Flattening, 83–84, 155
 - Floorlets and floors, 486–487, 500
 - FOMC, *see Federal Open Market Committee*
 - Forward agreements, 351–361, 363–371
 - Arbitrage pricing, 352–355
 - Forward drop, 356–357
 - Interest rate sensitivity, 359–361
 - Value, 352, 355–356
 - Forward bucket '01s, 120, 154, 164–170
 - Payer swaption, 166–169
 - Forward drop, 356–357, 384
 - Forward loan, 69–70, 71, 75
 - Forward measure, 300
 - Forward rate agreements (FRAs), 403–404
 - Forward rates, *see Rates, Forward*
 - FRAs, *see Forward rate agreements*
 - Full price, 52–53, 62, 63, 64, 115, 116
 - Futures contracts and rates, 351, 361–371
 - Daily settlement, 361–363
 - Euribor, 407
 - Options, 495–497, 500, 518–519
 - Early exercise, 525
 - Eurodollar, 47, 368–371, 401, 405–407, 408–411, 415–417
 - Mid-curve, 312, 496
 - Options, 495–497, 500, 517–518
 - Early exercise, 521–524
 - Quarterly, 312, 496
 - Serial, 496
 - Fed fund, 401, 419–424
 - Market expectations at the start of the 2004–2006 tightening cycle, 424–429
 - Martingale with money market
 - account as numeraire, 525–526
 - Note and bond, 124, 373–399
 - Basis trades, 383–386, 386–388
 - Beta, 391
 - Cheapest-to-deliver, 376–378, 387–390
 - Conversion factor, 375, 378–380, 380–382
 - Convexity, 389
 - Cost of delivery, 376–378, 385
 - Daily settlement, 374–375
 - Deliverable basket, 373–4, 378–379
 - Delivery option, *see Quality option*
 - DV01-rate curve, 132, 389
 - End-of-month option, 375, 392–393
 - Europe and Japan, 373, 375
 - Gross basis, 383–386
 - Net basis, 383–386, 386–388, 399
 - Notional coupon, 379, 390
 - Options, 124, 395–396, 498–499
 - DV01-rate curve, 132
 - JGB, 498–499
 - Price-rate curve, 125
 - Valuation with Black-Scholes models, 498–499, 501, 520
 - Valuation with term structure models, 498
 - Price-rate curve, 124–125
 - Quality option, 374, 375, 380–383, 386–390, 390–391
 - Term structure models, 390–391
 - Timing option, 374, 391–392
 - Trading case study, 393–399
 - TIBOR, 407–408
 - Futures-forward difference, 366–371
- G
- Gains process, 510
 - Gaussian models, 253

- Gauss+ model, 203, 287–298, 316–322
 Cascade form, 287–290, 291
 Fitting to an initial term structure, 295–296
 Fitting volatilities, 296–298
 Reduced form, 290
 USD and EUR sample results, 292–295
- Government bonds and debt
 Benchmark status, 48–49
 Europe, 30, 31, 32–33, 52
 Japan, 38, 40, 41, 43–45, 52
 United States, *see Treasury securities*
- Government sponsored entities (GSEs) and GSE-Backed Securities, 8–9
- Commercial banks assets, 13
- Conservatorship, 9
- Federal Home Loan Banks (FHLB), 8
- Federal Home Loan Mortgage Corporation (FHLMC), 8, 9, 563
- Federal National Mortgage Association (FNMA), 8, 9, 563
- Federal reserve bank balance sheet, 17
- Government National Mortgage Association (GNMA), 8, 563
- Guarantee business, 8
- Implicit U.S. government guarantee, 9
- Portfolio business, 8
- Securities, 8
- H**
- Haircuts, 330
- Half-life, 268
- Hedging, 119–121
 Bonds with ED futures, 415–417
 Borrowing and Lending with fed fund futures, 421–424
 Butterfly with principal components, 190–192
 Forward bucket '01s of a payer swaption, 167–169
 Futures option application, 127–130
 Issuance of corporate bonds, 444
 Key rate exposures, 159–162
- Lending or Borrowing with
 Eurodollar futures, 408–411
- Mortgage-backed securities, 584–585, 590
- Mortgage servicing rights, 589
- Number of factors, 189–190
- Principal components, 190–192
- Regression-based, 171–185, 190, 195–196
- Short convexity position, 135–137
- Tails, 371, 386, 399
- 20-year EUR interest rate swaps, 180–184
- Ho-Lee model, 259–260, 281, 282
- Households
 United States, 10–12, 20–21, 35
 Europe, 20–22, 35
 Japan, 20, 35–37
- Hybrid pension plans, *see Pension funds*
- I**
- IMM dates, 405, 449, 548
- Independent variable, 174
- Inflation-linked bonds, 47
 Pension fund investments, 24, 25, 26
 Japan, 45
 UK, 24
 United States, *see Treasury Securities, Inflation protected securities*
- In-sample, 182–183
- Interest rate risk, 14, 111, 119–120, 441–442
 Curve risk, 120, 153, 154
 Forward agreements, 359–361
 Immunization, 153
 With credit risk, 556–557
- Interest rate swaps, *see Swaps, Interest rate*
- Investment banking, 19
- Invoice price, *see Full price*
- ISDA master agreement, 437
- Ito's lemma, 241, 280
- J**
- Japanese Government Bonds (JGBs), *see Government Debt, Japan*
- Japanese Postal Savings System, 36, 39–40

Japanese simple yield, 104
Jensen's inequality, 232, 369

K

Key rate analysis, 120, 153,
154–162

L

Law of one price, 51, 55–58
Least-squares, *see Regression hedging*
Leverage, 15–16
L'Hôpital's rule, 88
Liability hedging, 25
LIBOR, *see London Interbank Offered Rate*
Libor Market Model (LMM), 203–204,
205, 298–316
 Calibrating the volatility and
 correlation functions,
 307–313
 General expression for drift changes,
 306–307
 Pricing an interest rate exotic,
 314–316
LIBOR-OIS spread, 401, 429, 431–432,
438, 458
Life insurance products and companies
 Europe, 27–28, 83
 Household financial assets, 21–22
 Japan, 36–37, 38–39, 39–40
 United States, 18
Lines of credit, 15 fn. 14
Liquidity management, 333
LMM, *see LIBOR Market Model*
Lognormal short-rate model, 278,
280–282, 282–284
London Interbank Offered Rate
 (LIBOR), 72, 163, 401–403,
 450
Long coupon, 491 fn. 2
Long-lived shock, 272

M

Margin for repurchase agreements, 330,
334, 336, 338
Mark-to-market, 362
Martingale property and pricing, 300,
508–513, 525–526

Mean reversion, 262–263, 268–273,
276–277, 282–284
Medium-Term Notes (MTNs), 33
Monetary Policy, 16–17, 31, 35, 42–43
Monte Carlo simulation, 302, 322–323,
581–583
 Path dependence, 581–583
Mortgages, 4–5, 45, 47, 563–568
 Adjustable-rate, 564
 Amortization table, 565–566
 Conforming, 8, 563
 FICO scores, 570
 Fixed rate mortgage payments,
 564–567
 Foreclosure, 4
 Household liability, 11
 Japan, 45
 Jumbos, 563
 Loans, 563–564, 590
 Loan-to-value (LTV) ratio, 570
 Prepayment option, 567–568
 Recourse, 5
 Subprime, 563, 564
Mortgage-backed securities, 8, 47,
568–575
 Agency, 563, 568
 Collateralized mortgage obligation
 (CMO), 575
 Companion bonds, 575
 Conforming, 563
 Constant maturity mortgage (CMM)
 products, 575
 Constant or conditional prepayment
 rate (CPR), 571
 Current coupon, 572
 Current mortgage rate, 572
 Dollar rolls, 573–574
 Guarantors, 568
 Hedge ratios, 584–585
 Hedging risk with swaps, 445
 Interest only (IO), 575, 588–589
 Japan, 45
 Option-adjusted spread (OAS),
 585–587
Options, 575
Pass-through, 568
Planned amortization class (PAC),
575

Mortgage-backed securities

(Continued)

Pools, 568–573

Prepayment rates, 571

Specific, 571–572

Prepayment models, 575–580

Burnout, 578–579

Curtailments, 580

Defaults and modifications, 580

Media effect, 579

Refinancing, 576–579

S-curve, 577–578

Seasoning ramp, 579–580

Turnover, 579–580

Price-rate behavior, 587–589

Principal only (PO), 575, 588–589

Private-label, 563, 568

Securitization, 568, 590

Single monthly mortality rate (SMM),
571

Spread at origination (SATO), 577

Support bonds, 575

TBAs, 571, 572–573

Valuation, 580–584

Weighted-average coupon (WAC),
569Weighted-average loan size (WALS),
576–577Weighted-average maturity (WAM),
569

Zero-volatility spread, 585 fn. 8

Mortgage servicing rights (MSR),
589

Municipal securities, 9, 47

Build America Bonds (BABs), 10

General obligation bonds, 9

Revenue bonds, 9

Tax treatment, 9–10

N

Nominal bonds, 172

Nonfinancial, Nonfarm Businesses,
12–13

Normal models, 253

Numeraire asset, 508–513

Bond futures options, 520

Bond options, 515–517

Caplets, 513–514

Euribor futures options, 518–519

Eurodollar futures options, 517–518

Swaptions, 514–515

O

OAS, *see Option-adjusted spread*OIS, *see Swaps, Overnight Index*Option-adjusted spread (OAS), 207,
221–224, 585–587Original Salomon Brothers model,
280–282

Out-of-sample, 182–183

Overnight index swaps, *see Swaps,*
Overnight Index

P

Par rates, *see Rates, Par*Partial PV01s or '01s, 120, 154,
163–164

Pension systems and funds

Bismarckian, 23

Defined-benefit and defined
contribution plans, 22–23, 24,
26

Europe, 22–27, 83

Austria, 23

Belgium, 23

France, 23

Germany, 23, 26–27

Ireland, 23

Italy, 23

The Netherlands, 23, 25–26

Spain, 23

Switzerland, 23, 26

UK, 23, 24–25

Funded and unfunded plans, 23

Household financial assets, 20–21

Hybrid plans, 22–23, 26

Japan, 36–37, 38–39, 40–41

Pay-as-you-go (PAYG) plans, 23, 40

Solvency ratio, 25

United States, 11, 39

Perpetuity, 102

Premium to face value, 101, 106, 107

Present value, 53, 55, 64, 65, 81–82,
93, 102, 104

DV01 and duration, 143–145

Forward-bucket '01s, 165

- Present value of an '01 (PV01), 164
- Principal component analysis and
hedging, 121, 171, 186–195,
196–200
Gauss+ model, 291, 292–293
Risk weights, 191
- Profit and Loss (P&L) decomposition or
attribution, 95, 105–110,
116–117, 222–224
- Proprietary trading, 19
- Pull-to-par, 101, 106
- Public Debt
United States, 6
Statutory ceiling, 6
- PV01, *see Present value of an '01*
- Q**
- Quantitative easing by the Bank of
Japan, 42–43
- Quoted price, 53, 62, 63, 116
- R**
- Rates, 48, 69–70
Forward, 48, 69, 75–76, 77, 78–81,
81–82, 82–87, 91–92, 93, 99,
368–371
Continuously compounded,
88–89
Curve construction, 592
Decomposition, 241–248
Embedded vs. realized, 95
Realized, 108–109, 110,
111–112
Swap, 357–359
Yield, 357
- Futures, 368–371
Negative, 253–255, 279, 281
Par, 48, 69, 76–77, 78–81, 89,
92–93, 99
- Shadow, 255
- Swap, 71–73, 76, 79–81
Principal components
EUR, GBP, and JPY, 192–195
USD, 187–190
- Spot, 48, 69, 74, 76, 77, 78–81, 89,
91–93, 99
Continuously compounded,
88–89
- Rating agencies, 529–530, 533–534
- Regression hedging, 121, 171
Level vs. change regressions, 184–185
R-squared, 176, 182
Risk weights, 178, 181, 182,
184
Single-variable, 172–180
Standard error of a regression,
176–177, 178, 182
Two-variable, 180–184
- Replacement rate of state pension plans,
21–22
- Replicating portfolio, *see Arbitrage
pricing*
- Repo, *see Repurchase agreements*
- Repurchase agreements (repo), 325,
327–349
Bear Stearns, 334–336
Broker-Dealer liabilities, 19
Commercial banking source of funds,
13
Europe, 33
European Central Bank financing of
banks, 32
Fails, 347–348
Financing risk implications, 14–15,
333–339
Futures basis trades, 384
General collateral and rates, 330,
339–340
Vs. fed funds, 340
Haircuts, 330, 331, 336, 338
Japan, 43
JPMorgan Chase's exposure to
Lehman Brothers, 336–339
Monetary policy implementation,
16–17, 31, 42–43
Money market mutual funds, 329,
331
Mortgage, 573
Municipalities, 329
Open, 329
Reverse, 332–333
Special collateral and rates, 327, 330,
333, 339–349
Treasury auction cycle, 341–347
Vs. level of rates, 347–348
Tri-party, 336, 338

- Return, 48, 95, 96–98
 - Gross, 96–97
 - With realized forwards, 111–112
 - With an unchanged term structure, 113
 - With unchanged yield, 113–114
 - Net, 97–98
 - Real, 172
- Rich securities, 55
- Risk aversion, 209, 210–211, 213, 236
- Risk-neutral pricing, 211–214, 223, 245
- Risk premium, 113, 114, 205, 229,
 - 236–241, 243–245, 247–248, 259, 261, 262
 - Ho-Lee model, 259
 - Negative, 240
 - Model 2, 257–258
 - Vasicek model, 263
- Risk weights, *see Regression hedging*
- Roll-down, *see Carry-Roll-Down*
- R-squares, *see Regression hedging*
- S
- SABR model, 504–508
- Sales and trading, 19
- Serial correlation, 185
- Settlement
 - U.S. Treasuries, 52
- Sharpe ratio, 244, 258
- Shifted lognormal model, 503–504
- Short coupon, 491
- Short-lived shock, 272
- Short-rate models, 202–203, 204, 263
 - fn. 6, 287
 - Black-Karsinski model, 282–284
 - Cox-Ingersoll-Ross model, 277–280, 284
 - Drift, 251–273
 - Gauss+ model, 287–298
 - Ho-Lee model, 259–260
 - Model 1, 251–257, 284
 - Model 2, 257–259, 284
 - Model 3, 275–277, 284
 - Original Salomon Brothers model, 280–282
 - Vasicek model, 262–273, 284
- Short selling, 56
- Simple interest, 69–70
- SONIA, 431
- State variables, 289
- Solvency II, 27, 83
- Sovereign debt crisis, 6, 32, 83, 104–105, 535–536
- Special purpose or investment vehicle (SPV or SIV), 17–18
- Spot loan, 74
- Spot rates, *see Rates, Spot*
- Spread, 48, 95, 98–99
 - Profit and loss decomposition, 105–106, 108–110, 116
- Steepening, 84, 155, 161, 162
- Stop-loss, 85
- Swaps, 435–456
 - Basis, 449–450, 458
 - Benchmark status, 48–49
 - CMS (constant-maturity swaps), 450–456
 - Clearing, 445–449
 - Day-count conventions, 65
 - EUR swap curve, second quarter 2010, 69
 - Forward-starting, 166–168, 358, 403–404
 - Interest rate, 47, 69, 71–73
 - Arbitrage pricing with financing, 465–473, 477–480
 - Cash flows, 435–437
 - Credit risk, 442–444
 - Interest rate risk, 441–442
 - Major uses, 444–445
 - Net present value, 440, 458, 465–473
 - Notional amount, 71–72
 - Valuation, 437–441
 - OIS discounting, 437–438, 450, 458, 473–475
 - Overnight index (OIS), 401, 429–431, 438, 470–473
 - Two-curve pricing, 473, 480–482
- Swaptions, 47, 487–490, 501, 514–515
 - Forward bucket '01s of a payer swaption, 166–167
 - Hedging forward bucket '01s of a payer swaption, 167–169
 - Physical vs. cash settlement, 490
 - Skew, 483, 490, 500–507

T

Tails, 371, 386, 399, 410–411
 Tangent line, 126, 136
 Taylor approximations, 137–139, 222
 TED spread, 401, 411–415, 432–434
 Terminal distribution, 253
 Terminal measure, 301
 Term structure models, 201–205, 226–227
 Forward and futures prices, 363–366
 Term structure of interest rates, 74, 76, 89, 91–93, 102, 104, 114
 Expectations, risk premium, and volatility, 229–248
 EUR, GBP, and USD swap curves, May 28, 2010, 79–81, 82–87
 Fitting models, 260–262
 Gauss+ model, 294–295
 Historical, 240–241
 Interest rate risk, 119–120
 Model 1, 255–256
 Model 2, 258
 Unchanged, 81–82, 107, 110, 112–113
 Vasicek model, 268–270
 Term structure of volatility, 169
 Gauss+ model, 293–294
 Ho-Lee model, 260
 Model 1, 256–257
 Model 2, 259
 Model 3, 277
 Vasicek model, 270–271, 273, 294
 TIBOR, *see Tokyo Interbank Offered Rate*
 Tier 1 Capital, 30
 Time-dependent volatility, 275–277, 283
 Time value of money, 47, 48, 55
 Tokyo Average Overnight Rate (TONAR), 42, 431
 Tokyo Interbank Offered Rate, 407–408
 TONAR, *see Tokyo Average Overnight Rate*
 Treasury securities, 5–8, 51–53, 58
 Amounts outstanding, 4, 6
 Auction cycle, 7–8, 341–347

Bills, 7, 47
 Bonds, 7
 Coupon payments, 7
 Current issues, 341
 DV01-rate curve, 132
 Face amount, 51
 Federal reserve bank balance sheet, 17
 Flight-to-quality trade, 6
 Idiosyncratic pricing, 60–62
 Inflation protected securities (TIPS), 7, 8
 Hedging vs. nominal bonds, 172–180
 Liquidity, 5
 Maturity structure, 6–7
 Notes, 7
 Off-the-run, 341
 On-the-run, 49, 341–347
 Par value, 51
 Price-rate curve, 124
 Principal amount, 51
 Reopening, 8
 Savings bonds, 7
 Separate Trading of Registered Interest and Principal of Securities (STRIPS), 51, 58–59, 60–62, 102
 Key rate exposures, 157–159
 TIPS, *see Treasury securities, Inflation protected securities*
 Tree representation of rates and prices, 207–209, 214–219
 Black-Karasinski model, 282–284
 Ho-Lee Model, 260
 Model 1, 252–253
 Model 2, 258
 Mortgage valuation, 580
 Nonrecombining, 214, 264–265
 Recombining, 215, 265–267, 283
 Time step, 224–225
 Vasicek model, 264–268

U
 Unbiased estimators, 184
 U.S. Treasury securities, *see Treasury securities*

V

Vanilla derivative, 201
Variation margin, 362–363
Vasicek model, 262–273, 277, 282, 284
Volatility
 Basis-point, 251–252, 277–279
 Lognormal model, 281–282
 Multifactor exposures, 169
 Portfolio, 119
 Regression-hedged portfolio, 178
 Short position, 137
 Short-term rates, 229, 231–236
 State-dependent, 214, 215

W

Wings of a butterfly trade,
 190

Y

Yield, 48, 95, 99–105,
 114–116
 Unchanged, 111, 113–114
Yield-volatility, 278

Z

Zero coupon securities, 58, 65,
 102–104