FIXED-INCOME SECURITIES AND DERIVATIVES HANDBOOK

ANALYSIS AND VALUATION

MOORAD CHOUDHRY

CRACKED TRADING SOFTWARE

70+ DVD's FOR SALE & EXCHANGE

www.traders-software.com

www.forex-warez.com

www.trading-software-collection.com

www.tradestation-download-free.com

Contacts

andreybbrv@gmail.com
andreybbrv@yandex.ru
Skype: andreybbrv

FIXED-INCOME
SECURITIES AND DERIVATIVES
HANDBOOK

**Also available
from Bloomberg Press**

Inside the Yield Book:
The Classic That Created the Science of Bond Analysis
by Sidney Homer and Martin L. Liebowitz, Ph.D.

The Securitization Markets Handbook:
Structures and Dynamics of Mortgage- and Asset-Backed Securities
by Charles Austin Stone and Anne Zissu

PIPEs:
A Guide to Private Investments in Public Equity
edited by Steven Dresner with E. Kurt Kim

Hedge Fund of Funds Investing:
An Investor's Guide
by Joseph G. Nicholas

Market-Neutral Investing:
Long/Short Hedge Fund Strategies
by Joseph G. Nicholas

––––––––––––––––

A complete list of our titles is available at
www.bloomberg.com/books

FIXED-INCOME SECURITIES AND DERIVATIVES HANDBOOK

Analysis and Valuation

Moorad Choudhry

*Dedicated to great Americans,
including Geronimo, Muhammad Ali,
and Smokey Robinson*

CONTENTS

PART THREE   SELECTED MARKET TRADING CONSIDERATIONS

16   The Yield Curve, Bond Yield, and Spot Rates   293

17   Approaches to Trading   315

This page is intentionally blank

FOREWORD

I think it's customary when writing a foreword to say "I couldn't put the book down." While one can't necessarily say that about *Fixed-Income Securities and Derivatives Handbook,* this is certainly not a book to read once and then leave to gather dust in the attic. The only fitting place for it is within arms' reach. Given the volume of accessible material about the market, it takes something special for a single book to find a permanent home at the desk and become a first point of reference for both practitioners and students alike.

This book reflects the emerging role of securitization within the bond markets. For instance, the Collateralized Debt Obligation (CDO) market is a constant source of innovation. The low interest rate environment of the past few years and increasing number of downgrades in the corporate bond market has made the rating-resilient securitization issuance an attractive source of investment for investors. Virtually everyone who buys fixed-income products is looking at CDOs.

There is clear evidence of a growing comfort with structured products from an increasing number of asset managers and investors. Spread tightening in the cash bond market has helped draw these parties toward CDOs, but even existing participants are looking for higher yields than those offered by traditional structures. We are now seeing CDO^2, CDOs with other CDOs as the underlying thus increasing leverage, and even beginning to see CDO^3. Further, margins on synthetic CDOs are dependent on the underlying portfolio spread, and so have been fueling demand for the associated credit default swaps. The result is significant product development, of the kind described in this fine book.

The author's refreshingly straightforward style enables him to uniquely bridge the gap between mathematical theory and its current application in the bond markets today. Moreover, given the dynamic, evolutionary nature of structured credit products, Choudhry's comprehensive work on bond market fundamentals is up to date with the most recent structural and product innovations.

SULEMAN BAIG
Research Partner, YieldCurve.com

This page is intentionally blank

PREFACE

The market in bond market securities, also known as the fixed-income market, is incredibly large and diverse, and one that plays an irreplaceable part in global economic development. The vast majority of securities in the world today are debt instruments, with outstanding volume estimated at more than $10 trillion.

Fixed-Income Securities and Derivatives Handbook provides a concise and accessible description of the main elements of the markets, concentrating on the instruments used and their applications. As it has been designed to be both succinct and concise, the major issues are introduced and described, and where appropriate certain applications are also analyzed.

Part One, "Introduction to Bonds," includes a detailed treatment of bond mathematics, including pricing and yield analytics. This includes modified duration and convexity. Chapters also cover the concept of spot (zero-coupon) and forward rates, and the rates implied by market bond prices and yields; yield-curve fitting techniques; an account of spline fitting using regression techniques; and an introductory discussion of term structure models.

Part Two, "Selected Cash and Derivative Instruments," has an analysis of various instruments including callable bonds that feature embedded options. There is also a discussion of mortgage-backed securities, techniques used in the analysis of U.S. Treasury TIPS securities, and a section on the use and applications of credit derivatives by participants in the fixed-income markets.

Finally, Part Three, "Selected Market Trading Considerations," covers the practical uses of redemption yield and duration as well as trading techniques based on the author's personal experience.

This book is designed to be a good starting place for those with little or no previous understanding of or exposure to the bond markets; however, it also investigates the markets to sufficient depth to be of use to the more experienced practitioner. Readers who are part of a front office, middle office, or back office banking and fund management staff involved to some extent with fixed-income securities will also find value here. Corporate

and local authority treasurers, bank auditors and consultants, risk managers, and legal department staff may also find the book useful.

Fixed-Income Securities and Derivatives Handbook builds on the content of the author's earlier book, *Bond Market Securities* (FT Prentice Hall, 2001), and includes an updated treatment of credit derivatives and synthetic structured products, a detailed analysis of bond futures, and case studies specific to the U.S. market. A very detailed treatment of specific markets, exchanges, or trading conventions is left out, as that would result in a very large book; interested readers can access the References section at the back of the book. Where possible these references are indicated for their level of analysis and technical treatment.

Comments on the text are most welcome and should be sent to the author care of Bloomberg Press.

A Word on the Mathematics

Financial subjects such as the debt capital markets are essentially quantitative disciplines, and as such it is not possible to describe them, let alone analyze them, without a certain amount of numerical input. To maintain accessibility of this book, the level of mathematics used has been limited; as a result many topics could not be reviewed in full detail. There are very few derivations, for example, and fewer proofs. This has not, in the author's opinion, impaired the analysis as the reader is still left with an understanding of the techniques required in the context of market instruments.

Website and Further Market Research

For the latest market research on fixed-income securities from Moorad Choudhry visit www.YieldCurve.com, which also contains conference presentations and educational teaching aids written by the author and other YieldCurve.com associates.

Acknowledgments

The author thanks the following for their help during the writing of this book: Serj Walia and Maj Haque at KBC Financial Products and Paul Kerlogue and Andrew Lipton at Moody's.

Thanks also to Suleman Baig for writing the Foreword.

FIXED-INCOME
SECURITIES AND DERIVATIVES
HANDBOOK

This page is intentionally blank

Introduction to Bonds

Part One describes fixed-income market analysis and the basic concepts relating to bond instruments. The analytic building blocks are generic and thus applicable to any market. The analysis is simplest when applied to plain vanilla default-free bonds; as the instruments analyzed become more complex, additional techniques and assumptions are required.

The first two chapters of this section discuss bond pricing and yields, moving on to an explanation of such traditional interest rate risk measures as modified duration and convexity. Chapter 3 looks at spot and forward rates, the derivation of such rates from market yields, and the yield curve. Yield-curve analysis and the modeling of the term structure of interest rates are among the most heavily researched areas of financial economics. The treatment here has been kept as concise as possible. The References section at the end of the book directs interested readers to accessible and readable resources that provide more detail.

This page is intentionally blank

The Bond Instrument

B onds are the basic ingredient of the U.S. debt-capital market, which is the cornerstone of the U.S. economy. All evening television news programs contain a slot during which the newscaster informs viewers where the main stock market indexes closed that day and where key foreign exchange rates ended up. Financial sections of most newspapers also indicate at what yield the Treasury long bond closed. This coverage reflects the fact that bond prices are affected directly by economic and political events, and yield levels on certain government bonds are fundamental economic indicators. The yield level on the U.S. Treasury long bond, for instance, mirrors the market's view on U.S. interest rates, inflation, public-sector debt, and economic growth.

The media report the bond yield level because it is so important to the country's economy—as important as the level of the equity market and more relevant as an indicator of the health and direction of the economy. Because of the size and crucial nature of the debt markets, a large number of market participants, ranging from bond issuers to bond investors and associated intermediaries, are interested in analyzing them. This chapter introduces the building blocks of the analysis.

Bonds are debt instruments that represent cash flows payable during a specified time period. They are essentially loans. The cash flows they represent are the interest payments on the loan and the loan redemption. Unlike commercial bank loans, however, bonds are tradable in a secondary market. Bonds are commonly referred to as *fixed-income* instruments. This term goes back to a time when bonds paid fixed coupons each year. That is

no longer necessarily the case. Asset-backed bonds, for instance, are issued in a number of tranches—related securities from the same issuer—each of which pays a different fixed or floating coupon. Nevertheless, this is still commonly referred to as the fixed-income market.

In the past bond analysis was frequently limited to calculating *gross redemption yield,* or *yield to maturity.* Today basic bond math involves different concepts and calculations. These are described in several of the references for chapter 3, such as Ingersoll (1987), Shiller (1990), Neftci (1996), Jarrow (1996), Van Deventer (1997), and Sundaresan (1997). This chapter reviews the basic elements. Bond pricing, together with the academic approach to it and a review of the term structure of interest rates, are discussed in depth in chapter 3.

In the analysis that follows, bonds are assumed to be *default-free.* This means there is no possibility that the interest payments and principal repayment will not be made. Such an assumption is entirely reasonable for government bonds such as U.S. Treasuries and U.K. gilt-edged securities. It is less so when you are dealing with the debt of corporate and lower-rated sovereign borrowers. The valuation and analysis of bonds carrying default risk, however, are based on those of default-free government securities. Essentially, the yield investors demand from borrowers whose credit standing is not risk-free is the yield on government securities plus some *credit risk* premium.

The Time Value of Money

Bond prices are expressed "per 100 nominal"—that is, as a percentage of the bond's face value. (The convention in certain markets is to quote a price per 1,000 nominal, but this is rare.) For example, if the price of a U.S. dollar–denominated bond is quoted as 98.00, this means that for every $100 of the bond's face value, a buyer would pay $98. The principles of pricing in the bond market are the same as those in other financial markets: the price of a financial instrument is equal to the sum of the *present values* of all the future cash flows from the instrument. The interest rate used to derive the present value of the cash flows, known as the *discount rate,* is key, since it reflects where the bond is trading and how its return is perceived by the market. All the factors that identify the bond—including the nature of the issuer, the maturity date, the coupon, and the currency in which it was issued—influence the bond's discount rate. Comparable bonds have similar discount rates. The following sections explain the traditional approach to bond pricing for plain vanilla instruments, making certain assumptions to keep the analysis simple. After that, a more formal analysis is presented.

Basic Features and Definitions

One of the key identifying features of a bond is its *issuer,* the entity that is borrowing funds by issuing the bond in the market. Issuers generally fall into one of four categories: governments and their agencies; local governments, or municipal authorities; supranational bodies, such as the World Bank; and corporations. Within the municipal and corporate markets there are a wide range of issuers that differ in their ability to make the interest payments on their debt and repay the full loan. An issuer's ability to make these payments is identified by its *credit rating.*

Another key feature of a bond is its *term to maturity:* the number of years over which the issuer has promised to meet the conditions of the debt obligation. The practice in the bond market is to refer to the term to maturity of a bond simply as its *maturity* or *term.* Bonds are debt capital market securities and therefore have maturities longer than one year. This differentiates them from money market securities. Bonds also have more intricate cash flow patterns than money market securities, which usually have just one cash flow at maturity. As a result, bonds are more complex to price than money market instruments, and their prices are more sensitive to changes in the general level of interest rates.

A bond's term to maturity is crucial because it indicates the period during which the bondholder can expect to receive *coupon payments* and the number of years before the *principal* is paid back. The principal of a bond—also referred to as its *redemption value, maturity value, par value,* or *face value*—is the amount that the issuer agrees to repay the bondholder on the maturity, or *redemption,* date, when the debt ceases to exist and the issuer redeems the bond. The coupon rate, or *nominal rate,* is the interest rate that the issuer agrees to pay during the bond's term. The annual interest payment made to bondholders is the bond's *coupon.* The *cash amount* of the coupon is the coupon rate multiplied by the principal of the bond. For example, a bond with a coupon rate of 8 percent and a principal of $1,000 will pay an annual cash amount of $80.

A bond's term to maturity also influences the volatility of its price. All else being equal, the longer the term to maturity of a bond, the greater its price volatility.

There are a large variety of bonds. The most common type is the *plain vanilla,* otherwise known as the *straight, conventional,* or *bullet* bond. A plain vanilla bond pays a regular—annual or semiannual—fixed interest payment over a fixed term. All other types of bonds are variations on this theme.

In the United States, all bonds make periodic coupon payments except for one type: the *zero-coupon.* Zero-coupon bonds do not pay any coupon. Instead investors buy them at a discount to face value and redeem them at

par. Interest on the bond is thus paid at maturity, realized as the difference between the principal value and the discounted purchase price.

Floating-rate bonds, often referred to as *floating-rate notes* (FRNs), also exist. The coupon rates of these bonds are reset periodically according to a predetermined benchmark, such as 3-month or 6-month LIBOR (London interbank offered rate). LIBOR is the official benchmark rate at which commercial banks will lend funds to other banks in the interbank market. It is an average of the offered rates posted by all the main commercial banks, and is reported by the British Bankers Association at 11.00 hours each business day. For this reason, FRNs typically trade more like money market instruments than like conventional bonds.

A bond issue may include a provision that gives either the bondholder or the issuer the option to take some action with respect to the other party. The most common type of option embedded in a bond is a *call feature.* This grants the issuer the right to "call" the bond by repaying the debt, fully or partially, on designated dates before the maturity date. A *put provision* gives bondholders the right to sell the issue back to the issuer at par on designated dates before the maturity date. A *convertible bond* contains a provision giving bondholders the right to exchange the issue for a specified number of stock shares, or equity, in the issuing company. The presence of embedded options makes the valuation of such bonds more complicated than that of plain vanilla bonds.

Present Value and Discounting

Since fixed-income instruments are essentially collections of cash flows, it is useful to begin by reviewing two key concepts of cash flow analysis: discounting and present value. Understanding these concepts is essential. In the following discussion, the interest rates cited are assumed to be the market-determined rates.

Financial arithmetic demonstrates that the value of $1 received today is not the same as that of $1 received in the future. Assuming an interest rate of 10 percent a year, a choice between receiving $1 in a year and receiving the same amount today is really a choice between having $1 a year from now and having $1 plus $0.10—the interest on $1 for one year at 10 percent per annum.

The notion that money has a time value is basic to the analysis of financial instruments. Money has time value because of the opportunity to invest it at a rate of interest. A loan that makes one interest payment at maturity is accruing *simple interest.* Short-term instruments are usually such loans. Hence, the lenders receive simple interest when the instrument expires. The formula for deriving *terminal,* or *future,* value of an investment with simple interest is shown as (1.1).

$$FV = PV(1+r) \tag{1.1}$$

where
   FV = the future value of the instrument
   PV = the initial investment, or the present value, of the instrument
   r = the interest rate

The market convention is to quote *annualized* interest rates: the rate corresponding to the amount of interest that would be earned if the investment term were one year. Consider a three-month deposit of $100 in a bank earning a rate of 6 percent a year. The annual interest gain would be $6. The interest earned for the ninety days of the deposit is proportional to that gain, as calculated below:

$$I_{90} = \$6.00 \times \frac{90}{365} = \$6.00 \times 0.2465 = \$1.479$$

The investor will receive $1.479 in interest at the end of the term. The total value of the deposit after the three months is therefore $100 plus $1.479. To calculate the terminal value of a short-term investment—that is, one with a term of less than a year—accruing simple interest, equation (1.1) is modified as follows:

$$FV = PV \left[1 + r \left(\frac{\text{days}}{\text{year}} \right) \right] \tag{1.2}$$

where FV and PV are defined as above,
   r = the annualized rate of interest
   days = the term of the investment
   year = the number of days in the year

Note that, in the sterling markets, the number of days in the year is taken to be 365, but most other markets—including the dollar and euro markets—use a 360-day year. (These conventions are discussed more fully below.)

Now consider an investment of $100, again at a fixed rate of 6 percent a year, but this time for three years. At the end of the first year, the investor will be credited with interest of $6. Therefore for the second year the interest rate of 6 percent will be accruing on a principal sum of $106. Accordingly, at the end of year two, the interest credited will be $6.36. This illustrates the principle of *compounding:* earning interest on interest. Equation (1.3) computes the future value for a sum deposited at a compounding rate of interest:

$$FV = PV(1+r)^n \tag{1.3}$$

where *FV* and *PV* are defined as before,

    r = the periodic rate of interest (expressed as a decimal)

    n = the number of periods for which the sum is invested

This computation assumes that the interest payments made during the investment term are reinvested at an interest rate equal to the first year's rate. That is why the example above stated that the 6 percent rate was *fixed* for three years. Compounding obviously results in higher returns than those earned with simple interest.

Now consider a deposit of $100 for one year, still at a rate of 6 percent but compounded quarterly. Again assuming that the interest payments will be reinvested at the initial interest rate of 6 percent, the total return at the end of the year will be

$$100 \times \left[(1+0.015) \times (1+0.015) \times (1+0.015) \times (1+0.015)\right]$$

$$= 100 \times \left[(1+0.015)^4\right] = 100 \times 1.6136 = \$106.136$$

The terminal value for quarterly compounding is thus about 13 cents more than that for annual compounded interest.

In general, if compounding takes place *m* times per year, then at the end of *n* years, *mn* interest payments will have been made, and the future value of the principal is computed using the formula (1.4).

$$FV = PV\left(1+\frac{r}{m}\right)^{mn} \tag{1.4}$$

As the example above illustrates, more frequent compounding results in higher total returns. **FIGURE 1.1** shows the interest rate factors corresponding to different frequencies of compounding on a base rate of 6 percent a year.

This shows that the more frequent the compounding, the higher the annualized interest rate. The entire progression indicates that a limit can be defined for continuous compounding, i.e., where *m* = infinity. To define the limit, it is useful to rewrite equation (1.4) as (1.5).

Interest rate factor $= \left(1 + \dfrac{r}{m}\right)^m$

COMPOUNDING FREQUENCY	INTEREST RATE FACTOR FOR 6%
Annual	$(1 + r) = 1.060000$
Semiannual	$\left(1 + \dfrac{r}{2}\right)^2 = 1.060900$
Quarterly	$\left(1 + \dfrac{r}{4}\right)^4 = 1.061364$
Monthly	$\left(1 + \dfrac{r}{12}\right)^{12} = 1.061678$
Daily	$\left(1 + \dfrac{r}{365}\right)^{365} = 1.061831$

$$
\begin{aligned}
FV &= PV \left[\left(1 + \frac{r}{m}\right)^{m/r}\right]^{rn} \\
&= PV \left[\left(1 + \frac{1}{m/r}\right)^{m/r}\right]^{rn} \\
&= PV \left[\left(1 + \frac{1}{w}\right)^{n}\right]^{rn}
\end{aligned}
\tag{1.5}
$$

where
$w = m/r$

As compounding becomes continuous and *m* and hence *w* approach infinity, the expression in the square brackets in (1.5) approaches the mathematical constant *e* (the base of natural logarithmic functions), which is equal to approximately 2.718281.

Substituting *e* into (1.5) gives us

$$
FV = PVe^{rn}
\tag{1.6}
$$

In (1.6) e^{rn} is the *exponential function* of *rn*. It represents the continuously compounded interest rate factor. To compute this factor for an interest rate of 6 percent over a term of one year, set *r* to 6 percent and *n* to 1, giving

$$e^{rn} = e^{0.06 \times 1} = (2.718281)^{0.06} = 1.061837$$

The convention in both wholesale and personal, or retail, markets is to quote an annual interest rate, whatever the term of the investment, whether it be overnight or ten years. Lenders wishing to earn interest at the rate quoted have to place their funds on deposit for one year. For example, if you open a bank account that pays 3.5 percent interest and close it after six months, the interest you actually earn will be equal to 1.75 percent of your deposit. The actual return on a three-year building society bond that pays a 6.75 percent fixed rate compounded annually is 21.65 percent. The quoted rate is the annual one-year equivalent. An overnight deposit in the wholesale, or *interbank,* market is still quoted as an annual rate, even though interest is earned for only one day.

Quoting annualized rates allows deposits and loans of different maturities and involving different instruments to be compared. Be careful when comparing interest rates for products that have different payment frequencies. As shown in the earlier examples, the actual interest earned on a deposit paying 6 percent semiannually will be greater than on one paying 6 percent annually. The convention in the money markets is to quote the applicable interest rate taking into account payment frequency.

The discussion thus far has involved calculating future value given a known present value and rate of interest. For example, $100 invested today for one year at a simple interest rate of 6 percent will generate $100 × (1 + 0.06) = $106 at the end of the year. The future value of $100 in this case is $106. Conversely, $100 is the present value of $106, given the same term and interest rate. This relationship can be stated formally by rearranging equation (1.3) as shown in (1.7).

$$PV = \frac{FV}{(1+r)^n} \tag{1.7}$$

Equation (1.7) applies to investments earning annual interest payments, giving the present value of a known future sum.

To calculate the present value of an investment, you must prorate the interest that would be earned for a whole year over the number of days in the investment period, as was done in (1.2). The result is equation (1.8).

$$PV = \frac{FV}{\left(1 + r \times \frac{\text{days}}{\text{year}}\right)} \tag{1.8}$$

When interest is compounded more than once a year, the formula for calculating present value is modified, as it was in (1.4). The result is shown in equation (1.9).

$$PV = \frac{FV}{\left(1 + \dfrac{r}{m}\right)^{mn}} \tag{1.9}$$

For example, the present value of $100 to be received at the end of five years, assuming an interest rate of 5 percent, with quarterly compounding is

$$PV = \frac{100}{\left(1 + \dfrac{0.05}{4}\right)^{(4)(5)}} = \$78.00$$

Interest rates in the money markets are always quoted for standard maturities, such as overnight, "tom next" (the overnight interest rate starting tomorrow, or "tomorrow to the next"), "spot next" (the overnight rate starting two days forward), one week, one month, two months, and so on, up to one year. If a bank or corporate customer wishes to borrow for a nonstandard period, or "odd date," an interbank desk will calculate the rate chargeable, by interpolating between two standard-period interest rates. Assuming a steady, uniform increase between standard periods, the required rate can be calculated using the formula for *straight line* interpolation, which apportions the difference equally among the stated intervals. This formula is shown as (1.10).

$$r = r_1 + (r_2 - r_1) \times \frac{n - n_1}{n_2 - n_1} \tag{1.10}$$

where
  r = the required odd-date rate for n days
  r_1 = the quoted rate for n_1 days
  r_2 = the quoted rate for n_2 days

Say the 1-month (30-day) interest rate is 5.25 percent and the 2-month (60-day) rate is 5.75 percent. If a customer wishes to borrow money for 40 days, the bank can calculate the required rate using straight line interpolation as follows: the difference between 30 and 40 is one-third that between 30 and 60, so the increase from the 30-day to the 40-day rate is assumed to be one-third the increase between the 30-day and the 60-day rates, giving the following computation

$$5.25 \text{ percent} + \frac{(5.75 \text{ percent} - 5.25 \text{ percent})}{3} = 5.4167 \text{ percent}$$

What about the interest rate for a period that is shorter or longer than the two whose rates are known, rather than lying between them? What if the customer in the example above wished to borrow money for 64 days? In this case, the interbank desk would extrapolate from the relationship between the known 1-month and 2-month rates, again assuming a uniform rate of change in the interest rates along the maturity spectrum. So given the 1-month rate of 5.25 percent and the 2-month rate of 5.75 percent, the 64-day rate would be

$$5.25 + \left[(5.75 - 5.25) \times \frac{34}{30}\right] = 5.8167 \text{ percent}$$

Just as future and present value can be derived from one another, given an investment period and interest rate, so can the interest rate for a period be calculated given a present and a future value. The basic equation is merely rearranged again to solve for r. This, as will be discussed below, is known as the investment's *yield.*

Discount Factors

An *n*-period discount factor is the present value of one unit of currency that is payable at the end of period *n*. Essentially, it is the present value relationship expressed in terms of $1. A discount factor for any term is given by formula (1.11).

$$d_n = \frac{1}{(1+r)^n} \tag{1.11}$$

where n = the period of discount

For instance, the five-year discount factor for a rate of 6 percent compounded annually is

$$d_5 = \frac{1}{(1+0.06)^5} = 0.747258$$

The set of discount factors for every period from one day to thirty years and longer is termed the *discount function.* Since the following discussion is in terms of *PV,* discount factors may be used to value any financial instrument that generates future cash flows. For example, the present value

FIGURE 1.2 *Hypothetical Set of Bonds and Bond Prices*

COUPON	MATURITY DATE	PRICE
7%	7-Jun-01	101.65
8%	7-Dec-01	101.89
6%	7-Jun-02	100.75
6.50%	7-Dec-02	100.37

of an instrument generating a cash flow of $103.50 payable at the end of six months would be determined as follows, given a six-month discount factor of 0.98756:

$$PV = \frac{FV}{(1+r)^n} = FV \times d_n = \$103.50 \times 0.98756 = \$102.212$$

Discount factors can also be used to calculate the future value of a present investment by inverting the formula. In the example above, the six-month discount factor of 0.98756 signifies that $1 receivable in six months has a present value of $0.98756. By the same reasoning, $1 today would in six months be worth

$$\frac{1}{d_{0.5}} = \frac{1}{0.98756} = \$1.0126$$

It is possible to derive discount factors from current bond prices. This process can be illustrated using the set of hypothetical bonds, all assumed to have semiannual coupons, that are shown in **FIGURE 1.2**, together with their prices.

The first bond in figure 1.2 matures in precisely six months. Its final cash flow will be $103.50, comprising the final coupon payment of $3.50 and the redemption payment of $100. The price, or present value, of this bond is $101.65. Using this, the six-month discount factor may be calculated as follows:

$$d_{0.5} = \frac{101.65}{103.50} = 0.98213$$

FIGURE 1.3 *Discount Factors Calculated Using Bootstrapping Technique*

COUPON	MATURITY DATE	TERM (YEARS)	PRICE	D(N)
7%	7-Jun-01	0.5	101.65	0.98213
8%	7-Dec-01	1.0	101.89	0.94194
6%	7-Jun-02	1.5	100.75	0.92211
6.50%	7-Dec-02	2.0	100.37	0.88252

Using this six-month discount factor, the one-year factor can be derived from the second bond in figure 1.2, the 8 percent due 2001. This bond pays a coupon of $4 in six months and, in one year, makes a payment of $104, consisting of another $4 coupon payment plus $100 return of principal.

The price of the one-year bond is $101.89. As with the 6-month bond, the price is also its present value, equal to the sum of the present values of its total cash flows. This relationship can be expressed in the following equation:

$$101.89 = 4 \times d_{0.5} + 104 \times d_1$$

The value of $d_{0.5}$ is known to be 0.98213. That leaves d_1 as the only unknown in the equation, which may be rearranged to solve for it:

$$d_1 = \left[\frac{101.89 - 4(0.98213)}{104} \right] = \frac{97.96148}{104} = 0.94194$$

The same procedure can be repeated for the remaining two bonds, using the discount factors derived in the previous steps to derive the set of discount factors in **FIGURE 1.3**. These factors may also be graphed as a continuous function, as shown in **FIGURE 1.4**.

This technique of calculating discount factors, known as bootstrapping, is conceptually neat, but may not work so well in practice. Problems arise when you do not have a set of bonds that mature at precise six-month intervals. Liquidity issues connected with individual bonds can also cause complications. This is true because the price of the bond, which is still the sum of the present values of the cash flows, may reflect liquidity considerations (e.g., hard to buy or sell the bond, difficult to find) that do not reflect the market as a whole but peculiarities of that

FIGURE **1.4** *Hypothetical Discount Function*
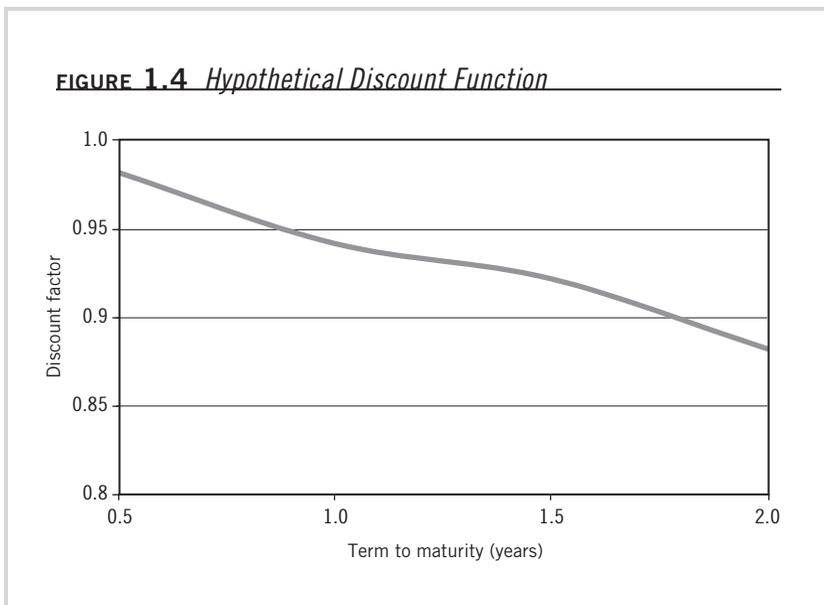
specific bond. The approach, however, is still worth knowing.

Note that the discount factors in figure 1.3 decrease as the bond's maturity increases. This makes intuitive sense, since the present value of something to be received in the future diminishes the farther in the future the date of receipt lies.

Bond Pricing and Yield:
The Traditional Approach

The discount rate used to derive the present value of a bond's cash flows is the interest rate that the bondholders require as compensation for the risk of lending their money to the issuer. The yield investors require on a bond depends on a number of political and economic factors, including what other bonds in the same class are yielding. Yield is always quoted as an annualized interest rate. This means that the rate used to discount the cash flows of a bond paying semiannual coupons is exactly half the bond's yield.

Bond Pricing

The *fair price* of a bond is the sum of the present values of all its cash flows, including both the coupon payments and the redemption payment. The price of a conventional bond that pays annual coupons can therefore be represented by formula (1.12).

$$P = \frac{C}{(1+r)} + \frac{C}{(1+r)^2} + \frac{C}{(1+r)^3} + \dots\dots \frac{C}{(1+r)^N} + \frac{M}{(1+r)^N} \qquad (1.12)$$

$$= \sum_{n=1}^{N} \frac{C}{(1+r)^n} + \frac{M}{(1+r)^N}$$

where

P = the bond's fair price

C = the annual coupon payment

r = the discount rate, or required yield

N = the number of years to maturity, and so the number of interest periods for a bond paying an annual coupon

M = the maturity payment, or par value, which is usually 100 percent of face value

Bonds in the U.S. domestic market—as opposed to international securities denominated in U.S. dollars, such as USD Eurobonds—usually pay semiannual coupons. Such bonds may be priced using the expression in (1.13), which is a modification of (1.12) allowing for twice-yearly discounting.

$$P = \frac{\frac{C}{2}}{\left(1+\frac{1}{2}r\right)} + \frac{\frac{C}{2}}{\left(1+\frac{1}{2}r\right)^2} + \frac{\frac{C}{2}}{\left(1+\frac{1}{2}r\right)^3} + \dots\dots + \frac{\frac{C}{2}}{\left(1+\frac{1}{2}r\right)^{2N}} + \frac{M}{\left(1+\frac{1}{2}r\right)^{2N}}$$

$$= \sum_{n=1}^{2N} \frac{\frac{C}{2}}{\left(1+\frac{1}{2}r\right)^n} + \frac{M}{\left(1+\frac{1}{2}r\right)^{2N}}$$

$$= \frac{C}{r}\left[1 - \frac{1}{\left(1+\frac{1}{2}r\right)^{2N}}\right] + \frac{M}{\left(1+\frac{1}{2}r\right)^{2N}} \qquad (1.13)$$

Note that $2N$ is now the power to which the discount factor is raised. This is because a bond that pays a semiannual coupon makes two interest payments a year. It might therefore be convenient to replace the number of years to maturity with the number of interest periods, which could be represented by the variable n, resulting in formula (1.14).

$$P = \frac{C}{r}\left[1 - \frac{1}{\left(1+\frac{1}{2}r\right)^n}\right] + \frac{M}{\left(1+\frac{1}{2}r\right)^n} \qquad (1.14)$$

This formula calculates the fair price on a coupon payment date, so there is no *accrued interest* incorporated into the price. Accrued interest is an accounting convention that treats coupon interest as accruing every day a bond is held; this accrued amount is added to the discounted present value of the bond (the *clean price*) to obtain the market value of the bond, known as the *dirty price.* The price calculation is made as of the bond's *settlement date,* the date on which it actually changes hands after being traded. For a new bond issue, the settlement date is the day when the investors take delivery of the bond and the issuer receives payment. The settlement date for a bond traded in the *secondary market*—the market where bonds are bought and sold after they are first issued—is the day the buyer transfers payment to the seller of the bond and the seller transfers the bond to the buyer.

Different markets have different settlement conventions. U.K. gilts, for example, normally settle on "T + 1": one business day after the trade date, T. Eurobonds, on the other hand, settle on T + 3. The term *value date* is sometimes used in place of settlement date, however, the two terms are not strictly synonymous. A settlement date can fall only on a business day; a bond traded on a Friday, therefore, will settle on a Monday. A value date, in contrast, can sometimes fall on a non-business day—when accrued interest is being calculated, for example.

Equation (1.14) assumes an even number of coupon payment dates remaining before maturity. If there are an odd number, the formula is modified as shown in (1.15).

$$P = \frac{C}{r}\left[1 - \frac{1}{\left(1 + \frac{1}{2}r\right)^{2N+1}}\right] + \frac{M}{\left(1 + \frac{1}{2}r\right)^{2N+1}} \tag{1.15}$$

Another assumption embodied in the standard formula is that the bond is traded for settlement on a day that is precisely one interest period before the next coupon payment. If the trade takes place between coupon dates, the formula is modified. This is done by adjusting the exponent for the discount factor using ratio *i*, shown in (1.16).

$$i = \frac{\text{Days from value date to next coupon date}}{\text{Days in the interest period}} \tag{1.16}$$

The denominator of this ratio is the number of calendar days between the last coupon date and the next one. This figure depends on the day-count convention (see below) used for that particular bond. Using *i*, the price formula is modified as (1.17) (for annual-coupon-paying bonds; for bonds with semiannual coupons, *r*/2 replaces *r*).

EXAMPLE: *Calculating Consideration for a U.S. Treasury*

The consideration, or actual cash proceeds paid by a buyer for a bond, is the bond's total cash value together with any costs such as commission. In this example, consideration refers only to the cash value of the bond.

What consideration is payable for $5 million nominal of a U.S. Treasury, quoted at a price of 99-16?

The U.S. Treasury price is 99-16, which is equal to 99 and 16/32, or 99.50 per $100. The consideration is therefore:

$$0.9950 \times 5,000,000 = \$4,975,000$$

If the price of a bond is below par, the total consideration is below the nominal amount; if it is priced above par, the consideration will be above the nominal amount.

$$P = \frac{C}{(1+r)^i} + \frac{C}{(1+r)^{1+i}} + \frac{C}{(1+r)^{2+i}} + \dots + \frac{C}{(1+r)^{n-1+i}} + \frac{M}{(1+r)^{n-1+i}}$$

$$(1.17)$$

where the variables *C, M, n,* and *r* are as before

As noted above, the bond market includes securities, known as zero-coupon bonds, or *strips,* that do not pay coupons. These are priced by setting *C* to 0 in the pricing equation. The only cash flow is the maturity payment, resulting in formula (1.18)

$$P = \frac{M}{(1+r)^N}$$

$$(1.18)$$

where *M* and *r* are as before and *N* is the number of years to maturity

Note that, even though these bonds pay no actual coupons, their prices and yields must be calculated on the basis of *quasi-coupon* periods, which are based on the interest periods of bonds denominated in the same currency. A U.S. dollar or a sterling five-year zero-coupon bond, for example, would be assumed to cover ten quasi-coupon peri-

EXAMPLE: *Zero-Coupon Bond Price*

A. Calculate the price of a Treasury strip with a maturity of precisely five years corresponding to a required yield of 5.40 percent.

According to these terms, $N = 5$, so $n = 10$, and $r = 0.054$, so $r/2 = 0.027$. $M = 100$, as usual. Plugging these values into the pricing formula gives

$$P = \frac{100}{(1.027)^{10}} = \$76.611782$$

B. Calculate the price of a French government zero-coupon bond with precisely five years to maturity, with the same required yield of 5.40 percent. Note that French government bonds pay coupon annually.

$$P = \frac{100}{(1.054)^{5}} = €76.877092$$

ods, and the price equation would accordingly be modified as (1.19).

$$P = \frac{M}{\left(1 + \frac{1}{2}r\right)^{n}} \tag{1.19}$$

It is clear from the bond price formula that a bond's yield and its price are closely related. Specifically, the price moves in the opposite direction from the yield. This is because a bond's price is the net present value of its cash flows; if the discount rate—that is, the yield required by investors—increases, the present values of the cash flows decrease. In the same way, if the required yield decreases, the price of the bond rises. The relationship between a bond's price and any required yield level is illustrated by the graph in **FIGURE 1.5**, which plots the yield against the corresponding price to form a convex curve.

Bond Yield

The discussion so far has involved calculating the price of a bond given its yield. This procedure can be reversed to find a bond's yield when its price is known. This is equivalent to calculating the bond's *internal rate of return,* or IRR, also known as its yield to maturity or gross redemption yield

FIGURE 1.5 *The Price/Yield Relationship*
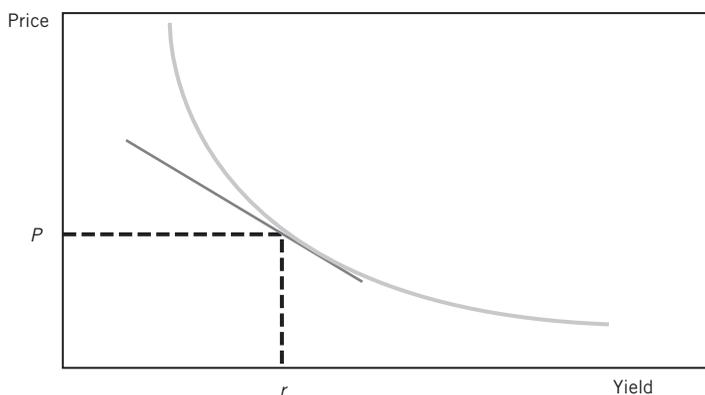
Summary of the Price/Yield Relationship

❏ At issue, if a bond is priced at par, its coupon will equal the yield that the market requires, reflecting factors such as the bond's term to maturity, the issuer's credit rating, and the yield on current bonds of comparable quality.
❏ If the required yield rises above the coupon rate, the bond price will decrease.
❏ If the required yield goes below the coupon rate, the bond price will increase.

(also *yield to workout*). These are among the various measures used in the markets to estimate the return generated from holding a bond.

In most markets, bonds are traded on the basis of their prices. Because different bonds can generate different and complicated cash flow patterns, however, they are generally compared in terms of their yields. For example, market makers usually quote two-way prices at which they will buy or sell particular bonds, but it is the yield at which the bonds are trading that is important to the market makers' customers. This is because a bond's price does not tell buyers anything useful about what they are getting.

Remember that in any market a number of bonds exist with different issuers, coupons, and terms to maturity. It is their yields that are compared, not their prices.

The yield on any investment is the discount rate that will make the present value of its cash flows equal its initial cost or price. Mathematically, an investment's yield, represented by r, is the interest rate that satisfies the bond price equation, repeated here as (1.20).

$$P = \sum_{n=1}^{N} \frac{C_n}{(1+r)^n} + \frac{M}{(1+r)^n} \qquad (1.20)$$

Other types of yield measure, however, are used in the market for different purposes. The simplest is the *current yield,* also know as the *flat, interest,* or *running yield.* These are computed by formula (1.21).

$$rc = \frac{C}{P} \times 100 \qquad (1.21)$$

where *rc* is the current yield

In this equation the percentage for C is not expressed as a decimal. Current yield ignores any capital gain or loss that might arise from holding and trading a bond and does not consider the time value of money. It calculates the coupon income as a proportion of the price paid for the bond. For this to be an accurate representation of return, the bond would have to be more like an annuity than a fixed-term instrument.

Current yield is useful as a "rough and ready" interest rate calculation; it is often used to estimate the cost of or profit from holding a bond for a short term. For example, if short-term interest rates, such as the one-week or three-month, are higher than the current yield, holding the bond is said to involve a *running cost.* This is also known as *negative carry* or *negative funding.* The concept is used by bond traders, market makers, and leveraged investors, but it is useful for all market practitioners, since it represents the investor's short-term cost of holding or funding a bond. The yield to maturity (YTM)—or, as it is known in sterling markets, gross redemption yield—is the most frequently used measure of bond return. Yield to maturity takes into account the pattern of coupon payments, the bond's term to maturity, and the capital gain (or loss) arising over the remaining life of the bond. The bond price formula shows the relationship between these elements and demonstrates their importance in determining price. The YTM calculation discounts the cash flows to maturity, employing the concept of the time value of money.

As noted above, the formula for calculating YTM is essentially that for calculating the price of a bond, repeated as (1.12). (For the YTM of bonds with semiannual coupon, the formula must be modified, as in (1.13).) Note, though, that this equation has two variables, the price P and yield r. It cannot, therefore, be rearranged to solve for yield r explicitly. In fact, the only way to solve for the yield is to use numerical iteration. This involves estimating a value for r and calculating the price associated with it. If the calculated price is higher than the bond's current price, the estimate for r is lower than the actual yield, so it must be raised. This process of calculation and adjustment up or down is repeated until the estimates converge on a level that generates the bond's current price.

To differentiate redemption yield from other yield and interest rate measures described in this book, it will be referred to as *rm*. Note that this section is concerned with the *gross* redemption yield, the yield that results from payment of coupons without deduction of any withholding tax. The *net redemption yield* is what will be received if the bond is traded in a market where bonds pay coupon *net*, without withholding tax. It is obtained by multiplying the coupon rate C by (1 – marginal tax rate). The net redemption yield is always lower than the gross redemption yield.

The key assumption behind the YTM calculation has already been discussed—that the redemption yield *rm* remains stable for the entire life of the bond, so that all coupons are reinvested at this same rate. The assumption is unrealistic, however. It can be predicted with virtual certainty that the interest rates paid by instruments with maturities equal to those of the bond at each coupon date will differ from *rm* at some point, at least, during the life of the bond. In practice, however, investors require a rate of return that is equivalent to the price that they are paying for a bond, and the redemption yield is as good a measurement as any.

A more accurate approach might be the one used to price interest rate swaps: to calculate the present values of future cash flows using discount rates determined by the markets' view on where interest rates will be at those points. These expected rates are known as *forward* interest rates. Forward rates, however, are *implied,* and a YTM derived using them is as speculative as one calculated using the conventional formula. This is because the real market interest rate at any time is invariably different from the one implied earlier in the forward markets. So a YTM calculation made using forward rates would not equal the yield actually realized either. The zero-coupon rate, it will be demonstrated later, is the true interest rate for any term to maturity. Still, despite the limitations imposed by its underlying assumptions, the YTM is the main measure of return used in the markets.

EXAMPLE: *Yield to Maturity for Semiannual-Coupon Bond*

A bond paying a semiannual coupon has a dirty price of \$98.50, an annual coupon of 3 percent, and exactly one year before maturity. The bond therefore has three remaining cash flows: two coupon payments of \$1.50 each and a redemption payment of \$100. Plugging these values into equation (1.20) gives

$$98.50 = \frac{1.50}{\left(1 + \frac{1}{2}rm\right)} + \frac{103.50}{\left(1 + \frac{1}{2}rm\right)^2}$$

Note that the equation uses half of the YTM value rm because this is a semiannual paying bond.

The expression above is a quadratic equation, which can be rearranged as $98.50x^2 - 1.50x - 103.50 = 0$, where x = $(1 + rm/2)$.

The equation may now be solved using the standard solution for equations of the form

$$ax^2 + bx + c = 0$$

$$x = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a}$$

There are two solutions, only one of which gives a positive redemption yield. The positive solution is

$$\frac{rm}{2} = 0.022755, \text{ or } rm = 4.551 \text{ percent}$$

YTM can also be calculated using mathematical iteration. Start with a trial value for rm of $r_1 = 4$ percent and plug this into the right-hand side of equation 1.20. This gives a price P_1 of 99.050, which is higher than the dirty market price P_M of 98.50. The trial value for rm was therefore too low.

Next try $r_2 = 5$ percent. This generates a price P_2 of 98.114, which is lower than the market price. Because the two trial prices lie on either side of the market value, the correct value for rm must lie between 4 and 5 percent. Now use the formula for linear interpolation

$$rm = r_1 + (r_2 - r_1)\frac{P_1 - P_M}{P_1 - P_2}$$

Plugging in the appropriate values gives a linear approximation for the redemption yield of $rm = 4.549$ percent, which is near the solution obtained by solving the quadratic equation.

Calculating the redemption yield of bonds that pay semiannual coupons involves the semiannual discounting of those payments. This approach is appropriate for most U.S. bonds and U.K. gilts. Government bonds in most of continental Europe and most Eurobonds, however, pay annual coupon payments. The appropriate method of calculating their redemption yields is to use annual discounting. The two yield measures are not directly comparable.

It is possible to make a Eurobond directly comparable with a U.K. gilt by using semiannual discounting of the former's annual coupon payments or using annual discounting of the latter's semiannual payments. The formulas for the semiannual and annual calculations appeared above as (1.13) and (1.12), respectively, and are repeated here as (1.22) and (1.23).

$$P_d = \frac{C}{\left(1+\frac{1}{2}rm\right)^2} + \frac{C}{\left(1+\frac{1}{2}rm\right)^4} + \frac{C}{\left(1+\frac{1}{2}rm\right)^6} + \ .... \tag{1.22}$$
$$+ \frac{C}{\left(1+\frac{1}{2}rm\right)^{2N}} + \frac{M}{\left(1+\frac{1}{2}rm\right)^{2N}}$$

$$P_d = \frac{C/2}{(1+rm)^{\frac{1}{2}}} + \frac{C/2}{(1+rm)} + \frac{C/2}{(1+rm)^{\frac{3}{2}}} + \ .... \tag{1.23}$$
$$+ \frac{C/2}{(1+rm)^{N}} + \frac{M}{(1+rm)^{N}}$$

Consider a bond with a dirty price—including the accrued interest the seller is entitled to receive—of \$97.89, a coupon of 6 percent, and five years to maturity. **FIGURE 1.6** shows the gross redemption yields this

FIGURE 1.6 *Yield and Payment Frequency*

DISCOUNTING	PAYMENTS	YIELD TO MATURITY
Semiannual	Semiannual	6.500
Annual	Annual	6.508
Semiannual	Annual	6.428
Annual	Semiannual	6.605

EXAMPLE: *Comparing Yields to Maturity*

A U.S. Treasury paying semiannual coupons, with a maturity of ten years, has a quoted yield of 4.89 percent. A European government bond with a similar maturity is quoted at a yield of 4.96 percent. Which bond has the higher yield to maturity in practice?

The effective annual yield of the Treasury is

$$rm_a = \left(1 + \tfrac{1}{2} \times 0.0489\right)^2 - 1 = 4.9498 \text{ percent}$$

Comparing the securities using the same calculation basis reveals that the European government bond does indeed have the higher yield.

bond would have under the different yield-calculation conventions.

These figures demonstrate the impact that the coupon-payment and discounting frequencies have on a bond's redemption yield calculation. Specifically, increasing the frequency of discounting lowers the calculated yield, while increasing the frequency of payments raises it. When comparing yields for bonds that trade in markets with different conventions, it is important to convert all the yields to the same calculation basis.

It might seem that doubling a semiannual yield figure would produce the annualized equivalent; the real result, however, is an underestimate of the true annualized yield. This is because of the multiplicative effects of discounting. The correct procedure for converting semiannual and quarterly yields into annualized ones is shown in (1.24).

a. General formula

$$rm_a = (1 + \text{interest rate})^m - 1 \qquad (1.24)$$

where *m* = the number of coupon payments per year

b. Formulas for converting between semiannual and annual yields

$$rm_a = \left(1 + \tfrac{1}{2} rm_s\right)^2 - 1$$

$$rm_s = \left[\left(1 + rm_a\right)^{\frac{1}{2}} - 1\right] \times 2$$

c. Formulas for converting between quarterly and annual yields

$$rm_a = \left(1 + \tfrac{1}{4} rm_q\right)^4 - 1$$

$$rm_q = \left[\left(1 + rm_a\right)^{\frac{1}{4}} - 1\right] \times 4$$

where rm_q, rm_s, and rm_a are, respectively, the quarterly, semiannually, and annually discounted yields to maturity.

The market convention is sometimes simply to double the semiannual yield to obtain the annualized yields, despite the fact that this produces an inaccurate result. It is only acceptable to do this for rough calculations. An annualized yield obtained in this manner is known as a *bond equivalent yield.* It was noted earlier that the one disadvantage of the YTM measure is that its calculation incorporates the unrealistic assumption that each coupon payment, as it becomes due, is reinvested at the rate *rm*. Another disadvantage is that it does not deal with the situation in which investors do not hold their bonds to maturity. In these cases, the redemption yield will not be as great. Investors might therefore be interested in other measures of return, such as the equivalent zero-coupon yield, considered a true yield.

To review, the redemption yield measure assumes that

❑ the bond is held to maturity
❑ all coupons during the bond's life are reinvested at the same (redemption yield) rate

Given these assumptions, the YTM can be viewed as an *expected* or *anticipated* yield. It is closest to reality when an investor buys a bond on first issue and holds it to maturity. Even then, however, the actual realized yield at maturity would be different from the YTM because of the unrealistic nature of the second assumption. It is clearly unlikely that all the coupons of any but the shortest-maturity bond will be reinvested at the same rate. As noted earlier, market interest rates are in a state of constant flux, and this would affect money reinvestment rates. Therefore, although yield to maturity is the main market measure of bond levels, it is not a true interest rate. This is an important point. Chapter 2 will explore the concept of a true interest rate.

Another problem with YTM is that it discounts a bond's coupons at the yield specific to that bond. It thus cannot serve as an accurate basis for comparing bonds. Consider a two-year and a five-year bond. These securities will invariably have different YTMs. Accordingly, the coupon cash flows they generate in two years' time will be discounted at different

rates (assuming the yield curve is not flat). This is clearly not correct. The present value calculated today of a cash flow occurring in two years' time should be the same whether that cash flow is generated by a short- or a long-dated bond.

Accrued Interest

All bonds except zero-coupon bonds accrue interest on a daily basis that is then paid out on the coupon date. As mentioned earlier, the formulas discussed so far calculate bonds' prices as of a coupon payment date, so that no accrued interest is incorporated in the price. In all major bond markets, the convention is to quote this so-called clean price.

Clean and Dirty Bond Prices

When investors buy a bond in the market, what they pay is the bond's *all-in* price, also known as the dirty, or *gross price,* which is the clean price of a bond plus accrued interest.

Bonds trade either *ex-dividend* or *cum dividend.* The period between when a coupon is announced and when it is paid is the ex-dividend period. If the bond trades during this time, it is the seller, not the buyer, who receives the next coupon payment. Between the coupon payment date and the next ex-dividend date the bond trades cum dividend, so the buyer gets the next coupon payment.

Accrued interest compensates sellers for giving up all the next coupon payment even though they will have held their bonds for part of the period since the last coupon payment. A bond's clean price moves with market interest rates. If the market rates are constant during a coupon period, the clean price will be constant as well. In contrast, the dirty price for the same bond will increase steadily as the coupon interest accrues from one coupon payment date until the next ex-dividend date, when it falls by the present value of the amount of the coupon payment. The dirty price at this point is below the clean price, reflecting the fact that accrued interest is now negative. This is because if the bond is traded during the ex-dividend period, the seller, not the buyer, receives the next coupon, and the lower price is the buyer's compensation for this loss. On the coupon date, the accrued interest is zero, so the clean and dirty prices are the same.

The net interest accrued since the last ex-dividend date is calculated using formula (1.25).

$$AI = C \ \times \ \left[\frac{N_{xt} - N_{xc}}{\text{Day Base}} \right] \tag{1.25}$$

where

AI = the next accrued interest

C = the bond coupon

N_{xc} = the number of days between the *ex-dividend* date and the coupon payment date

N_{xt} = the number of days between the *ex-dividend* date and the date for the calculation

Day Base = the day-count base (see below)

When a bond is traded, accrued interest is calculated from and including the last coupon date up to and excluding the value date, usually the settlement date. Interest does not accrue on bonds whose issuer has defaulted.

As noted earlier, for bonds that are trading ex-dividend, the accrued coupon is negative and is subtracted from the clean price. The negative accrued interest is calculated using formula (1.26).

$$AI = -C \ \times \ \frac{\text{days to next coupon}}{\text{Day Base}} \qquad (1.26)$$

Certain classes of bonds—U.S. Treasuries and Eurobonds, for example—do not have ex-dividend periods and therefore trade cum dividend right up to the coupon date.

Day-Count Conventions

In calculating the accrued interest on a bond, the market uses the day-count convention appropriate to that bond. These conventions govern both the number of days assumed to be in a calendar year and how the days between two dates are figured. **FIGURE 1.7** shows how the different conventions affect the accrual calculation.

In these conventions, the number of days between two dates includes the first date but not the second. Thus, using actual/365, there are thirty-seven days between August 4 and September 10. The last two conventions assume thirty days in each month, no matter what the calendar says. So, for example, it is assumed that there are thirty days between 10 February and 10 March. Under the 30/360 convention, if the first date is the 31st, it is changed to the 30th; if the second date is the 31st and the first date is either the 30th or the 31st, the second date is changed to the 30th. The 30E/360 convention differs from this in that if the second date is the 31st, it is changed to the 30th regardless of what the first date is.

FIGURE 1.7 *Accrued Interest, Day-Count Conventions*

Actual/365 $AI = C \times$ actual days to next coupon payment/365

Actual/360 $AI = C \times$ actual days to next coupon/360

Actual/actual $AI = C \times$ actual days to next coupon/actual number
                    of days in the interest period

30/360 $AI = C \times$ days to next coupon, assuming 30 days in
                    a month/360

30E/360 $AI = C \times$ days to next coupon, assuming 30 days in
                    a month/360

This page is intentionally blank

Bond Instruments and Interest Rate Risk

Chapter 1 described the basic concepts of bond pricing. This chapter discusses the sensitivity of bond prices to changes in market interest rates and the key related concepts of duration and convexity.

Duration, Modified Duration, and Convexity

Most bonds pay a part of their total return during their lifetimes, in the form of coupon interest. Because of this, a bond's term to maturity does not reflect the true period over which its return is earned. Term to maturity also fails to give an accurate picture of the trading characteristics of a bond or to provide a basis for comparing it with other bonds having similar maturities. Clearly, a more accurate measure is needed.

A bond's maturity gives little indication of how much of its return is paid out during its life or of the timing and size of its cash flows. Maturity is thus inadequate as an indicator of the bond's sensitivity to moves in market interest rates. To see why this is so, consider two bonds with the same maturity date but different coupons: the higher-coupon bond generates a larger proportion of its return in the form of coupon payments than does the lower-coupon bond and so pays out its return at a faster rate. Because of this, the higher-coupon bond's price is theoretically less sensitive to fluctuations in interest rates that occur during its lifetime. A better indication of a bond's payment characteristics and interest rate sensitivity might be the average time to receipt of its cash flows. The cash

flows generated during a bond's life differ in value, however. The average time to receipt would be a more accurate measure, therefore, if it were weighted according to the cash flows' present values. The average maturity of a bond's cash flow stream calculated in this manner provides a measure of the speed at which a bond pays out its return, and hence of its price risk relative to other bonds having the same maturity.

The average time until receipt of a bond's cash flows, weighted according to the present values of these cash flows, measured in years, is known as *duration* or *Macaulay's duration,* referring to the man who introduced the concept in 1938—see Macaulay (1999) in References. Macaulay introduced duration as an alternative for the length of time remaining before a bond reached maturity.

Duration

Duration is a measure of price sensitivity to interest rates—that is, how much a bond's price changes in response to a change in interest rates. In mathematics, change like this is often expressed in terms of differential equations. The price-yield formula for a plain vanilla bond, introduced in chapter 1, is repeated as (2.1) below. It assumes complete years to maturity, annual coupon payments, and no accrued interest at the calculation date.

$$P = \frac{C}{(1+r)} + \frac{C}{(1+r)^2} + \frac{C}{(1+r)^3} + \ldots + \frac{C}{(1+r)^N} + \frac{M}{(1+r)^N} \quad (2.1)$$

where
  P = the bond's fair price
  C = the annual coupon payment
  r = the discount rate, or required yield
  N = the number of years to maturity, and so the number of interest periods for a bond paying an annual coupon
  M = the maturity payment

Chapter 1 showed that the price and yield of a bond are two sides of the same relationship. Because price P is a function of yield r, we can differentiate the price/yield equation at (2.1), as shown in (2.2). Taking the first derivative of this expression gives (2.2).

$$\frac{dP}{dr} = \frac{(-1)C}{(1+r)^2} + \frac{(-2)C}{(1+r)^3} + \ldots + \frac{(-n)C}{(1+r)^{n+1}} + \frac{(-n)M}{(1+r)^{n+1}} \quad (2.2)$$

Rearranging (2.2) gives (2.3). The expression in brackets is the average time to maturity of the cash flows from a bond weighted according to the present value of each cash flow. The whole equation is the formula for calculating the approximate change in price corresponding to a small change in yield.

$$\frac{dP}{dr} = -\frac{1}{(1+r)}\left[\begin{array}{c}\dfrac{1C}{(1+r)} + \dfrac{2C}{(1+r)^2} + ... + \\[2mm] + \dfrac{nC}{(1+r)^n} + \dfrac{nM}{(1+r)^n}\end{array}\right] \tag{2.3}$$

Dividing both sides of (2.3) by P results in expression (2.4).

$$\frac{dP}{dr}\frac{1}{P} = -\frac{1}{(1+r)}\left[\begin{array}{c}\dfrac{1C}{(1+r)} + \dfrac{2C}{(1+r)^2} + ... + \\[2mm] + \dfrac{nC}{(1+r)^n} + + \dfrac{nM}{(1+r)^n}\end{array}\right]\frac{1}{P} \tag{2.4}$$

Dividing the bracketed expression by P gives expression (2.5), which is the definition of Macaulay duration, measured in years.

$$D = \frac{\dfrac{1C}{(1+r)} + \dfrac{2C}{(1+r)^2} + ... + \dfrac{nC}{(1+r)^n} + \dfrac{nM}{(1+r)^n}}{P} \tag{2.5}$$

Equation (2.5) can be simplified using $\sum$, as shown in (2.6).

$$D = \frac{\displaystyle\sum_{n=1}^{N}\frac{nC_n}{(1+r)^n}}{P} \tag{2.6}$$

where C_n = the bond cash flow at time n

If the expression for Macaulay duration, (2.5), is substituted into equation (2.4), which calculates the approximate percentage change in price, (2.7) is obtained. This is the definition of *modified duration*.

$$\frac{dP}{dr}\frac{1}{P} = -\frac{1}{(1+r)}D = \frac{D}{(1+r)} = -D_{\text{mod}} \tag{2.7}$$

or

$$D_{\text{mod}} = \frac{D}{(1+r)} \tag{2.8}$$

EXAMPLE: *Calculating the Macaulay Duration for the 8 Percent*
          2009 Annual Coupon Bond

Issued	30 September 1999
Maturity	30 September 2009
Price	$102.497
Yield	7.634 percent

PERIOD (N)	CASH FLOW	PV AT CURRENT YIELD *	$N \times PV$
1	8	7.43260	7.4326
2	8	6.90543	13.81086
3	8	6.41566	19.24698
4	8	5.96063	23.84252
5	8	5.53787	27.68935
6	8	5.14509	30.87054
7	8	4.78017	33.46119
8	8	4.44114	35.529096
9	8	4.12615	37.13535
10	108	51.75222	517.5222
TOTAL		102.49696	746.540686

*Calculated as $C/(1 + r)^n$

Macaulay duration = 746.540686 / 102.497

            = 7.283539998 years

Modified duration can be used to demonstrate that small changes in yield produce inverse changes in bond price. This relationship is expressed formally in (2.7), repeated as (2.9).

$$\frac{dP}{dr}\frac{1}{P} = -D_{\text{mod}} \tag{2.9}$$

It is possible to shorten the procedure of computing Macaulay duration longhand, by rearranging the bond-price formula (2.1) as shown in (2.10), which, as explained in chapter 1, calculates price as the sum of the present values of its coupons and its redemption payment. The same assumptions apply as for (2.1).

$$P = C\left[\frac{1 - \frac{1}{(1+r)^n}}{r}\right] + \frac{M}{(1+r)^n} \tag{2.10}$$

Taking the first derivative of (2.10) and dividing the result by the current bond price, P, produces an alternative formulation for modified duration, shown as (2.11).

$$D_{\text{mod}} = \frac{\frac{C}{r^2}\left[1 - \frac{1}{(1+r)^n}\right] + \frac{n\left(M - \frac{C}{r}\right)}{(1+r)^{n+1}}}{P} \tag{2.11}$$

Multiplying (2.11) by $(1 + r)$ gives the equation for Macaulay duration. The example on the following page shows how these shorthand formulas can be used to calculate modified and Macaulay durations.

Up to this point the discussion has involved plain vanilla bonds. But duration applies to all bonds, even those that have no conventional maturity date, the so-called perpetual, or irredeemable, bonds (also known as annuity bonds), which pay out interest for an indefinite period. Since these make no redemption payment, the second term on the right side of the duration equation disappears, and since coupon payments can stretch on indefinitely, n approaches infinity. The result is equation (2.12), for Macaulay duration.

$$D = \frac{1}{rc} \tag{2.12}$$

where $rc = (C/P_d)$ is the *running yield* (or *current yield*) of the bond

Equation (2.12) represents the limiting value to duration. For bonds trading at or above par, duration increases with maturity, approaching the value given by (2.12), which acts as a ceiling. For bonds trading at a discount to par, duration increases to a maximum of around twenty years and then declines toward the floor given by (2.12). In general, duration increases with maturity, with an upper bound given by (2.12).

EXAMPLE: *Calculating the Modified and Macaulay Durations as of 1999 of a Hypothetical Bond Having an Annual Coupon of 8 Percent and a Maturity Date of 2009*

Coupon	8 percent, paid annually
Yield	7.634 percent
n	10
Price	$102.497

Plugging these values into the modified-duration equation (2.11) gives

$$D_{\text{mod}} = \frac{\dfrac{8}{(0.07634^2)}\left[1 - \dfrac{1}{(1.07634)^{10}}\right] + \dfrac{10\left(100 - \frac{8}{0.07634}\right)}{(1.07634)^{11}}}{102.497}$$

$$D_{\text{mod}} = 6.76695 \text{ years}$$

To obtain the bond's Macaulay duration, this modified duration is multiplied by $(1 + r)$, or 1.07634, for a value of 7.28354 years.

Properties of Macaulay Duration

Duration varies with maturity, coupon, and yield. Broadly, it increases with maturity. A bond's duration is generally shorter than its maturity. This is because the cash flows received in the early years of the bond's life have the greatest present values and therefore are given the greatest weight. That shortens the average time in which cash flows are received. A zero-coupon bond's cash flows are all received at redemption, so there is no present-value weighting. Therefore, a zero-coupon bond's duration is equal to its term to maturity.

Duration increases as coupon and yield decrease. The lower the coupon, the greater the relative weight of the cash flows received on the maturity date, and this causes duration to rise. Among the non–plain vanilla types of bonds are some whose coupon rate varies according to an index, usually the consumer price index. Index-linked bonds generally have much lower coupons than vanilla bonds with similar maturities. This is true because they are inflation-protected, causing the real yield required to be lower than the nominal yield, but their durations tend to be higher.

Yield's relationship to duration is a function of its role in discounting future cash flows. As yield increases, the present values of all future cash flows fall, but those of the more distant cash flows fall relatively more. This has the effect of increasing the relative weight of the earlier cash flows and hence of reducing duration.

Modified Duration

Although newcomers to the market commonly consider duration, much as Macaulay did, a proxy for a bond's time to maturity, this interpretation misses the main point of duration, which is to measure price volatility, or interest rate risk. Using the Macaulay duration can derive a measure of a bond's interest rate price sensitivity, i.e., how sensitive a bond's price is to changes in its yield. This measure is obtained by applying a mathematical property known as a Taylor expansion to the basic equation.

The relationship between price volatility and duration can be made clearer if the bond price equation, viewed as a function of r, is expanded as a Taylor series (see Butler, pp. 112–114 for an accessible explanation of Taylor expansions). Using the first term of this series, the relationship can be expressed as (2.13).

$$\Delta P = -\left[\frac{1}{(1+r)}\right] \times D \times \text{Change in yield} \tag{2.13}$$

where r = the yield to maturity of an annual-coupon-paying bond

As stated above, Macaulay duration equals modified duration multiplied by $(1+r)$. The first two components of the right-hand side of (2.13) taken together are therefore equivalent to modified duration, and equation (2.13) expresses the approximate percentage change in price as modified duration multiplied by the change in yield.

Modified duration is a measure of the approximate change in bond price for a 1 percent change in yield. The relationship between modified duration and bond prices can therefore be expressed as (2.14). A negative is used in this equation because the price movement is inverse to the interest rate movement, so a rise in yields produces a fall in price, and vice versa.

$$\Delta P = -D_{\text{mod}} \times (\Delta r) \times P \tag{2.14}$$

The example on the following page illustrates how the relationships expressed in these equations work.

Changes in yield are often expressed in terms of *basis points,* which equal hundredths of a percent. For a bond with a modified duration of

EXAMPLE: *Applying the Duration/Price Relationships to a Hypothetical Bond*

Coupon   8 percent, paid annually
Price   par
Duration   2.74 years

If yields rise from 8 percent to 8.50 percent, the fall in the price of the bond can be computed as follows:

$$\Delta P = -D \times \frac{\Delta(r)}{1+r} \times P$$

$$= -\,(2.74) \times \left(\frac{0.005}{1.080}\right) \times 100$$

$$= -\$1.2685$$

That is, the price of the bond will fall to $98.7315.
The modified duration of a bond with a duration of 2.74 years and a yield of 8 percent is

$$D_{\mathrm{mod}} = \frac{2.74}{1.08} = 2.537 \text{ years}$$

If a bond has a duration of 4.31 years and a modified duration of 3.99, a 1 percent move in the yield to maturity produces a move (in the opposite direction) in the price of approximately 3.99 percent.

3.99, priced at par, an increase in yield of 1 basis point leads to a fall in the bond's price of

$$\Delta P = \left(\frac{-3.24}{100}\right) \times (+0.01) \times 100.00$$

$$\Delta P = \$0.0399, \text{ or } 3.99 \text{ cents}$$

In this example, 3.99 cents is the *basis point value* (BPV) of the bond: the change in its price given a 1 basis point change in yield. The general formula for deriving the basis point value of a bond is shown in (2.15).

$$BPV = \frac{D_{\text{mod}}}{100} \times \frac{P}{100} \qquad (2.15)$$

Basis point values are used in hedging bond positions. Hedging is done by taking an opposite position—that is, one that will rise in value under the same conditions that will cause the hedged position to fall, and vice versa. Say you hold a 10-year bond. You might wish to sell short a similar 10-year bond as a hedge against your long position. Similarly, if you hold a short position in a bond, you might hedge it by buying an equivalent amount of a hedging instrument. A variety of hedging instruments are available, for use both on- and off-balance sheet.

For a hedge to be effective, the price change in the primary instrument should be equal to the price change in the hedging instrument. To calculate how much of a hedging instrument is required to get this type of protection, each bond's BPV is used. This is important because different bonds have different BPVs. To hedge a long position in, say, $1 million nominal of a 30-year bond, therefore, you can't simply sell $1 million of

EXAMPLE: *Calculating Hedge Size Using Basis Point Value*

Say a trader holds a long position of $1 million of the 8 percent bond maturing in 2019. The bond's modified duration is 11.14692, and its price is $129.87596. Its basis point value is therefore 0.14477. The trader decides to protect the position against a rise in interest rates by hedging it using the zero-coupon bond maturing in 2009, which has a BPV of 0.05549. Assuming that the yield beta is 1, what nominal value of the zero-coupon bond must the trader sell?

The hedge ratio is

$$\frac{0.14477}{0.05549} \times 1 = 2.60894$$

To hedge $1 million of the 20-year bond, therefore, the trader must sell short $2,608,940 of the zero-coupon bond. Using the two bonds' BPVs, the loss in the long position produced by a 1 basis point rise in yield is approximately equal to the gain in the hedge position.

FIGURE **2.1** *The Modified Duration Approximation of Bond Price Change at Different Yields*

BOND	MATURITY (YEARS)	MODIFIED DURATION	PRICE DURATION OF 1 BASIS POINT	6.00%	6.50%	7.00%
8% 2009	10	6.76695	0.06936	114.72017	110.78325	107.02358

another 30-year bond. There may not be another 30-year bond with the same BPV available. You might have to hedge with a 10-year bond. To calculate how much nominal of the hedging bond is required, you'd use the *hedge ratio* (2.16).

$$\frac{BPV_p}{BPV_h} \times \frac{\text{Change in yield for primary bond position}}{\text{Change in yield for hedge instrument}} \qquad (2.16)$$

where
   BPV_p = the basis point value of the primary bond (the position to be hedged)
   BPV_h = the basis point value of the hedging instrument

The second term in (2.16) is known as the *yield beta.*
   FIGURE 2.1 shows how the price of the 8 percent 2009 bond changes for a selection of yields. For a 1 basis point change in yield, the change in price, indicated as "price duration for 1 basis point," though not completely accurate because it is a straight line or linear approximation of a non-linear relationship, as illustrated with figure 1.5 of the price/yield relationship, is a reasonable estimate of the actual change in price. For a large move—say, 200 basis points—the approximation would be significantly off base, and analysts would accordingly not use it. This is shown in **FIGURE 2.2**.
   Note that the price duration figure, calculated from the modified duration measurement, underestimates the change in price resulting from a fall in yields but overestimates the change from a rise in yields. This reflects the *convexity* of the bond's price-yield relationship, a concept that will be explained in the next section.

7.50%	7.99%	8.00%	8.01%	8.50%	9.00%	10.00%
103.43204	100.0671311	100.00000	99.932929	96.71933	93.58234	87.71087

FIGURE 2.2 *Approximation of the Bond Price Change Using Modified Duration*

YIELD CHANGE	PRICE CHANGE	ESTIMATE USING PRICE DURATION
Down 1 bp	0.06713	0.06936
Up 1 bp	0.06707	0.06936
Down 200 bp	14.72017	13.872
Up 200 bp	12.28913	13.872

Convexity

Duration is a first-order measure of interest rate risk, using first-order derivatives. If the relationship between price and yield is plotted on a graph, it forms a curve. Duration indicates the slope of the tangent at any point on this curve. A tangent, however, is a line and, as such, is only an approximation of the actual curve—an approximation that becomes less accurate the farther the bond yield moves from the original point. The magnitude of the error, moreover, depends on the curvature, or convexity, of the curve. This is a serious drawback, and one that applies to modified as well as to Macaulay duration.

Convexity represents an attempt to remedy the drawbacks of duration. A second-order measure of interest rate risk uses second-order derivatives. It measures the curvature of the price-yield graph and the degree to which this diverges from the straight-line estimation. Convex-

ity can thus be regarded as an indication of the error made when using Macaulay and modified duration. A bond's convexity is positively correlated to the *dispersion* of its cash flows: all else being equal, a bond whose cash flows are more spread out in time—that is, more dispersed—than another's will have a higher convexity. Convexity is also positively correlated with duration.

The second-order differential of the bond price equation with respect to the redemption yield r is

$$\frac{\Delta P}{P} = \frac{1}{P}\frac{\Delta P}{\Delta r}(\Delta r) + \frac{1}{2P}\frac{\Delta^2 P}{\Delta r^2}(\Delta r)^2 \tag{2.17}$$

$$= -D_{\text{mod}}(\Delta r) + \frac{CV}{2}(\Delta r)^2$$

where CV = the convexity

Equation (2.17) shows that convexity is the rate at which price sensitivity to yield changes as yield changes. That is, it describes how much a bond's modified duration changes in response to changes in yield. Formula (2.18) expresses this relationship formally. The convexity term can be seen as an "adjustment" for the error made by duration in approximating the price-yield curve.

$$CV = 10^8\left(\frac{\Delta P'}{P} + \frac{\Delta P''}{P}\right) \tag{2.18}$$

where
$\Delta P'$ = the change in bond price if yield increases by 1 basis point
$\Delta P''$ = the change in bond price if yield decreases by 1 basis point

The unit in which convexity, as defined by (2.18), is measured is the number of interest periods. For annual-coupon bonds, this is equal to the number of years; for bonds with different coupon-payment schedules, formula (2.19) can be used to convert the convexity measure from interest periods to years.

$$CV_{years} = \frac{CV}{C^2} \tag{2.19}$$

The convexity formula for zero-coupon bonds is (2.20).

$$CV = \frac{n(n+1)}{(1+r)^2} \tag{2.20}$$

Convexity is a second-order approximation of the change in price resulting from a change in yield. This relationship is expressed formally in (2.21).

$$\Delta P = \tfrac{1}{2} \times CV \times (\Delta r)$$

(2.21)

The reason the convexity term is multiplied by one-half is because the second term in the Taylor expansion used to derive the convexity equation contains the coefficient 0.5. The formula is the same for a semiannual-coupon bond.

Note that the value for convexity given by the expressions above will always be positive—that is, the approximate price change due to convexity is positive for both yield increases and decreases, except for certain bonds such as callable bonds.

As noted earlier, the price change estimated using modified duration can be quite inaccurate; the convexity measure is the approximation of the size of the inaccuracy. Summing the two values—the price-change estimate using modified duration plus the convexity error adjustment—gives a more accurate picture of the actual magnitude of the price change. The estimated and adjusted values differ significantly, however, only when the change in yield is very large. In the example below, the modified duration of the hypothetical 5 percent 2015 bond is 7.64498. For the specified rise

---

EXAMPLE: *Calculating the Convexity of a Bond*

Coupon	5 percent, paid annually
Maturity	2015, ten years from present
Price	$96.23119
Yield	5.50 percent

If the yield rises to 7.50 percent, a change of 200 basis points, the convexity adjustment that would be made to the price change calculated using modified duration and equation (2.21) is

$$(0.5) \times 96.23119 \times (0.02)^2 \times 100 = 1.92462 \text{ percent}$$

If the yield fell by 200 basis points, the convexity effect would be the same.

in yield of 200 basis points, the approximate price change given by modified duration is

$$\text{Modified duration} = \text{-}7.64498 \times 2 = \text{-}15.28996$$

Note that the modified duration is given as a negative value, because a rise in yields results in a fall in price. Adjusting the estimate by the convexity of 1.92462 derived above results in a net percentage price change of 13.36534. A Hewlett-Packard (HP) calculator gives the price of the bond at the new yield of 7.50 percent as \$82.83980, representing an actual change of 13.92 percent. So using the convexity adjustment produces a noticeably more accurate estimate.

Now assume that yields fall just 1.50 percent, or 150 basis points. The new convexity value is

$$(0.5) \times 96.23119 \times (0.015)^2 \times 100 = 1.0826 \text{ percent}$$

And the price change estimate based on modified duration is

$$\text{Modified duration} = 7.64498 \times 1.5 = 11.46747$$

Adding the two values together results in an adjusted price change estimate of 12.55007 percent. The actual price change according to the HP calculator is 10.98843 percent. In this case, the unadjusted modified duration estimate is closer. This illustrates that the convexity measure is effective for larger yield changes only. The example at right provides an illustration of the greater accuracy produced by combining the modified duration and convexity measures for larger yield shifts.

The convexity measure increases with the square of maturity; it decreases as both coupon and yield rise. It is a function of modified duration, so index-linked bonds, which have greater duration than conventional bonds of similar maturities, also have greater convexity. For a conventional vanilla bond, convexity is almost always positive. Negative convexity occurs most frequently with callable bonds.

In principle, a bond with greater convexity should fall in price less when yields rise than a less-convex one, and rise in price more when yields fall. This is true because convexity is usually positive, so it lessens the price decline produced by rises in yield and increases the price rise produced by falls in yield. Thus, all else being equal, the higher the convexity of a bond the more desirable it should be to investors. The actual premium attached to higher convexity is a function of current yield levels and market volatility. Remember that modified duration and

EXAMPLE: *Convexity Adjustment*

Assume that the yield of the hypothetical 5 percent 2015 bond rises to 8.50 percent, a change of 300 basis points. The percentage convexity adjustment is

$$0.5 \times 96.23119 \times (0.03)^2 \times 100 = 4.3304 \text{ percent}$$

The modified duration of the bond at the initial yield, as seen above, is 7.64498. So the price change estimate using modified duration is

$$7.64498 \times 3.0 = -22.93494$$

Adjusting this by the convexity value derived above results in a price change of 18.60454 percent. Using an HP calculator, the price of the bond is 77.03528, for an actual percentage price change of 19.9477 percent. In this case, the adjusted estimate is closer than that obtained using only the modified duration measure. The continuing error reflects the fact that convexity is a dynamic measure and changes with yield changes; the effect of a large yield movement compounds the inaccuracy of the adjustments.

convexity are both functions of yield level, and their effects are magnified at lower yield levels. In addition, the cash effect of convexity is more noticeable for large moves in yield. So the value investors attach to convexity will vary according to their expectations about the future size of interest rate changes. Hence, convexity is more highly valued when market volatility is high.

This page is intentionally blank

Bond Pricing and Spot and Forward Rates

As discussed in chapter 1, there are two types of fixed-income securities: zero-coupon bonds, also known as discount bonds or strips, and coupon bonds. A zero-coupon bond makes a single payment on its maturity date, while a coupon bond makes interest payments at regular dates up to and including its maturity date. A coupon bond may be regarded as a set of strips, with the payment on each coupon date and at maturity being equivalent to a zero-coupon bond maturing on that date. This equivalence is not purely academic. Before the advent of the formal market in U.S. Treasury strips, a number of investment banks traded the cash flows of Treasury securities as separate zero-coupon securities.

The discussion in this chapter assumes a liquid market of default-free bonds, where both zero-coupon and coupon bonds are freely bought and sold. Prices are determined by the economy-wide supply of and demand for the bonds at any time. The prices are thus *macroeconomic,* rather than being set by individual bond issuers or traders.

Zero-Coupon Bonds

A zero-coupon bond is the simplest fixed-income security. It makes no coupon payments during its lifetime. Instead, it is a *discount* instrument, issued at a price that is below the face, or principal, amount. The rate earned on a zero-coupon bond is also referred to as the *spot interest rate.* The notation $P(t, T)$ denotes the price at time t of a discount bond that matures at time T, where $T \geq t$. The bond's term to maturity, $T - t$, is

denoted by *n*. The strip's price increases over time until the maturity date, when it reaches maturity, or par, value. If the par value of a bond is \$1, then its yield to maturity at time *t* is denoted by $r(t,T)$, where *r* is "1 plus the percentage yield" that is earned by holding the bond from *t* to *T*. The relationship between the zero-coupon bond's price and its yield to maturity at any point in its life may be expressed as equation (3.1).

$$P(t,T) = \frac{1}{[r(t,T)]^n} \tag{3.1}$$

This equation can be rearranged as (3.2) to derive a bond's yield for a given price.

$$r(t,T) = \left[\frac{1}{P(t,T)}\right]^{\frac{1}{n}} = P(t,T)^{-\frac{1}{n}} \tag{3.2}$$

Analysts and researchers frequently work with logarithms of yields and prices, or *continuously compounded* rates. One advantage of the logarithmic approach is that it converts the nonlinear relationship expressed in (3.2) into a linear one. The zero-coupon bond price equation in continuous time is

$$P(t,T) = e^{-r(t,T)(T-t)} \tag{3.3}$$

The price equation for a specific time t_2, where $t \le t_2 \le T$, is

$$P(t_2,T) = P(t,T)e^{(t_2-t)r(t,T)} \tag{3.4}$$

Note that expression (3.4) contains an exponential function; this is why the rate is characterized as continuously compounded. Yield in continuous time is given by (3.5).

$$r(t,T) = -\log\left(\frac{P(t,T)}{n}\right) \tag{3.5}$$

This is sometimes written as (3.6).

$$\log r(t,T) = -\left(\frac{1}{n}\right)\log P(t,T) \tag{3.6}$$

The *term structure of interest rates* is the set of zero-coupon yields at time *t* for terms $(t,t+1)$ to $(t,t+m)$, where the bonds have maturities of $\{0,1,2,.....,m\}$. The *term structure of interest rates* thus describes the relation-

FIGURE 3.1   *U.S. Treasury Zero-Coupon Yield Curve in September 2000*
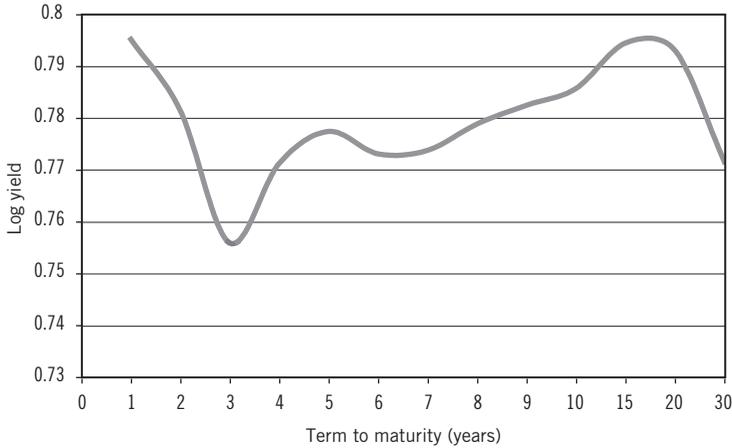
ship between the spot interest rates of credit-risk–free zero-coupon bonds and their maturities.

The *yield curve* is a graph plotting the set of yields $r(t, t+1)$ through $r(t, t+m)$ at time t against m. **FIGURES 3.1, 3.2,** and **3.3** (above and on the following page) show the log zero-coupon yield curves, as of September 27, 2000, for, respectively, U.S. Treasury strips, U.K. gilt strips, and French OAT (Obligations Assimilable du Trésor) strips. The French curve exhibits the most common shape for yield curves: a gentle upward slope. The U.K. curve slopes in the opposite direction, a shape termed *inverted.*

Coupon Bonds

The majority of bonds in the market are coupon bonds. As noted above, such bonds may be viewed as packages of individual strips. The strips corresponding to the coupon payments have face values that equal percentages of the nominal value of the bond itself, with successively longer maturity dates; the strip corresponding to the final redemption payment has the face value and maturity date of the bond.

A bond issued at time i and maturing at time T makes w payments (C_1 … C_w) on w payment dates (t_1, … t_{w-1}, T). In the academic literature, these coupon payments are assumed to be continuous, rather than periodic, so the stream of coupon payments can be represented formally as a positive

FIGURE 3.2 *U.K. Gilt Zero-Coupon Yield Curve in September 2000*
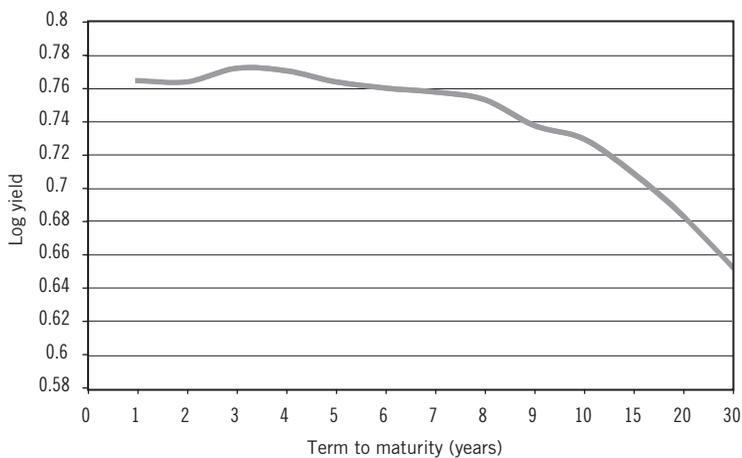
FIGURE 3.3 *French OAT Zero-Coupon Yield Curve in September 2000*
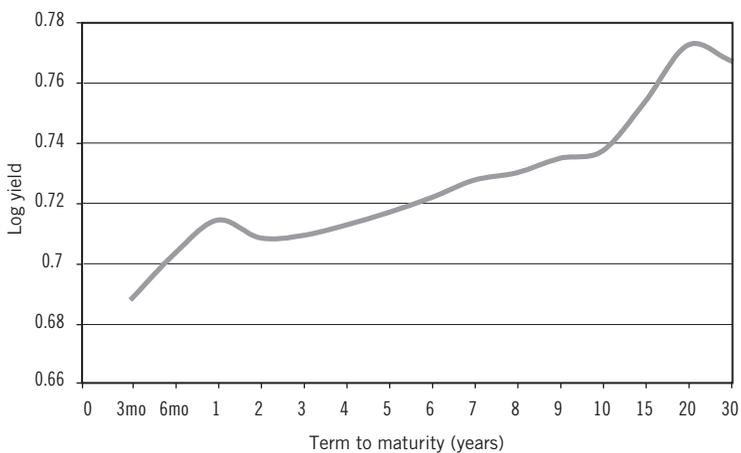
function of time: $C(t), i < t \leq T$. Investors purchasing a bond at time t that matures at time T pay $P(t, T)$ and receive the coupon payments as long as they hold the bond. Note that $P(t, T)$ is the clean price of the bond, as defined in chapter 1; in practice, unless the bond is purchased for settlement on a coupon date, the investor will pay a dirty price, which includes the value of the interest that has accrued since the last coupon date.

As discussed in chapter 1, yield to maturity is the interest rate that relates a bond's price to its future returns. More precisely, using the notation defined above, it is the rate that discounts the bond's cash flow stream C_w to its price $P(t, T)$. This relationship is expressed formally in equation (3.7).

$$P(t,T) = \sum_{t_i > t} C_i e^{-(t_i - t)r(t,T)}$$

(3.7)

Expression (3.7) allows the continuously compounded yield to maturity $r(t, T)$ to be derived. For a zero-coupon, it reduces to (3.5). In the academic literature, $\sum$, which is used in mathematics to calculate sums of a countable number of objects, is replaced by $\int$, the integral sign, which is used for an infinite number of objects. (See Neftci (2000), pages 59–66, for an introduction to integrals and their use in quantitative finance.)

Some texts refer to the graph of coupon-bond yields plotted against maturities as the term structure of interest rates. It is generally accepted, however, that this phrase should be used for zero-coupon rates only and that the graph of coupon-bond yields should be referred to instead as the yield curve. Of course, given the law of one price—which holds that two bonds having the same cash flows should have the same values—the zero-coupon term structure is related to the yield to maturity curve and can be derived from it.

Bond Price in Continuous Time

This section is an introduction to bond pricing in continuous time. Chapter 4 presents a background on price processes.

Fundamental Concepts

Consider a trading environment in which bond prices evolve in a w-dimensional process, represented in (3.8).

$$X(t) = [X_1(t), X_2(t), X_3(t), \ldots\ldots, X_w(t)], t > 0$$

(3.8)

where the *random variables* (variables whose possible values are numerical outcomes of a random process) X_i are *state variables*, representing the state of the economy at times t_i

The markets assume that the state variables evolve through a geometric Brownian motion, or Weiner process. It is therefore possible to model their evolution using a stochastic differential equation. The market also assumes that the cash flow stream of assets such as bonds and equities is a function of the state variables.

A bond is characterized by its coupon process, represented in (3.9).

$$C(t) = \tilde{C}\left[X_1(t), X_2(t), X_3(t), \ldots, X_w(t), t\right] \qquad (3.9)$$

The coupon process represents the cash flow investors receive while they hold the bond. Assume that a bond's term can be divided into very small intervals of length *dt* and that it is possible to buy very short-term discount bonds, such as Treasury strips, maturing at the end of each such interval and paying an annualized rate $r(t)$. This rate is the *short,* or *instantaneous,* rate, which in mathematical bond analysis is defined as the rate of interest charged on a loan taken out at time *t* that matures almost immediately. The short rate is given by formulas (3.10) and (3.11).

$$r(t) = r(t,t) \qquad (3.10)$$

The short rate is the interest rate on a loan that is paid back almost instantaneously; it is a theoretical construct. Equation (3.11) states this mathematical notation in terms of the bond price.

$$r(t) = -\frac{\partial}{\partial T}\log P(t,t) \qquad (3.11)$$

If the principal of the short-term security described above is continuously reinvested at this short rate, the cumulative amount obtained at time *t* is equal to the original investment multiplied by expression (3.12).

$$M(t) = \exp\left[\int_0^t r(s)\,ds\right] \qquad (3.12)$$

where *M* is a money market account that offers a return of the short rate $r(t)$. The bond principal is multiplied by an account and $M(t)$ is the total rate of return of such an account.

If the short rate is constant—that is, $r(t) = r$—then the price of a risk-free bond that pays \$1 on its maturity date *T* is given by expression (3.13).

$$P(t,T) = e^{-r(T-t)} \qquad (3.13)$$

Expression (3.13) states that the bond price is a function of the continuously compounded interest rate. The right-hand side is the discount factor at time t. At $t = T$—that is, on the redemption date—the discount factor is 1, which is the redemption value of the bond and hence the price of the bond at that time.

Consider a scenario in which a market participant can either

❑ invest $e^{-r(T-t)}$ units of cash in a money market account, for a return of $1 at time T, or
❑ purchase a risk-free zero-coupon bond that has a maturity value of $1 at time T

The bond and the money market are both risk-free and have identical payouts at time T, and neither will generate any cash flow between now and time T. Since the interest rates involved are constant, the bond must have a value equal to the initial investment in the money market account: $e_t^{-r(T-t)}$. In other words, equation (3.13) must hold. This is a restriction placed on the zero-coupon bond price by the requirement for markets to be arbitrage-free.

If the bond were not priced at this level, arbitrage opportunities would present themselves. Say the bond was priced higher than $e_t^{-r(T-t)}$. In this case, investors could sell the bond short and invest $et-r(T-t)$ of the sale proceeds in the money market account. At time T, the short position would have a value of $-1 (negative, because the bond position is short); the money market, meanwhile, would have grown to $1, which the investors could use to close their short bond positions. And they would still have funds left over from the short sale, because at time t

$$P(t,T) - e^{-r(T-t)} > 0$$

So they would profit from the transaction at no risk to themselves.

A similar situation obtains if the bond price $P(t, T)$ is less than $e_t^{-r(T-t)}$. In that case, the investors borrow $e_t^{-r(T-t)}$ at the money market rate. They then use $P(t, T)$ of this loan to buy the strip. At maturity the bond pays $1, which the investors use to repay the loan. But they still have surplus borrowed funds, because

$$e^{-r(T-t)} - P(t,T) > 0$$

This demonstrates that the only situation in which no arbitrage profit can be made is when $P(t,T) = e^{-r(T-t)}$. (For texts providing more detail on arbitrage pricing theory, see the References section.)

In the academic literature, the *risk-neutral* price of a zero-coupon bond is expressed in terms of the evolution of the short-term interest rate, $r(t)$—the rate earned on a money market account or on a short-dated risk-free security such as the T-bill—which is assumed to be continuously compounded. These assumptions make the mathematical treatment simpler. Consider a zero-coupon bond that makes one payment, of $1, on its maturity date T. Its value at time t is given by equation (3.14), which is the redemption value of 1 divided by the value of the money market account, given by (3.12).

$$P(t,T) = \exp\left[-\int_t^T r(s)\,ds\right] \qquad (3.14)$$

The price of a zero-coupon bond in terms of its yield is given by equation (3.15).

$$P(t,T) = \exp\left[-(T-t)\,r(T-t)\right] \qquad (3.15)$$

Expression (3.14) is the formula for pricing zero-coupon bonds when the spot rate is the nonconstant instantaneous risk-free rate $r(s)$ described above. This is the rate used in formulas (3.12), for valuing a money market account, and (3.15), for pricing a risk-free zero-coupon bond.

Stochastic Rates

In the academic literature, the bond price given by equation (3.15) evolves as a *martingale* process under the risk-neutral probability measure $\tilde{P}$. This process is the province of advanced fixed-income mathematics and lies outside the scope of this book. An introduction, however, is presented in chapter 4, which can be supplemented by the readings listed in the References section.

Advanced financial analysis produces the bond price formula (3.16) (for the formula's derivation, see Neftci (2000), page 417).

$$P(t,T) = E_t^{\tilde{P}}\left[e^{-\int_t^T r(s)ds}\right] \qquad (3.16)$$

The right-hand side of (3.16) is the randomly evolved discount factor used to obtain the present value of the $1 maturity payment. The expression states that bond prices are dependent on the entire spectrum of short-term interest rates $r(s)$ during the period $t < s < T$. It also implies, given

the view that the short rate evolves as a martingale process, that the term structure at time t contains all the information available on future short rates. From (3.16) the discount curve, or *discount function,* at time t can be derived as $T \rightarrow P_t^T, t < T$. In other words, (3.16) states that the price of the bond is the continuously compounded interest rate as applied to the zero-coupon bond from issue date t to maturity T. The complete set of bond prices assuming a nominal value of $1 is the discount function. Avellaneda (2000) notes that the bond analysts usually replace the term T with a term $(T - t)$, meaning time to maturity, so the function becomes

$$\tau \rightarrow P_t^{t+\tau}, \tau > 0, \text{ where } \tau = (T - t)$$

The relationship between the yield $r(t, T)$ of the zero-coupon bond and the short rate $r(t)$ can be expressed by equating the right-hand sides of equations (3.16) and (3.3) (the formula for deriving the zero-coupon bond price, repeated here as (3.17)). The result is (3.18).

$$P(t,T) = e^{-r(t,T)(T-t)} \tag{3.17}$$

$$e^{-r(T,t)(T-t)} = E_t^{\tilde{P}}\left[e^{-\int_t^T r(s)ds}\right] \tag{3.18}$$

Taking the logarithm of both sides gives (3.19).

$$r(t,T) = \frac{-\log E_t^{\tilde{P}}\left[e^{-\int_t^T r(s)ds}\right]}{T - t} \tag{3.19}$$

Equation (3.19) describes a bond's yield as the average of the spot rates that apply during the bond's life. If the spot rate is constant, the yield equals it.

For a zero-coupon bond, assuming that interest rates are positive, $P(t, T)$ is less than or equal to 1. The yield of this bond is given by (3.20).

$$r(t,T) = -\frac{\log\left(P(t,T)\right)}{T - t} \tag{3.20}$$

Rearranging (3.20) to solve for price results in (3.21).

$$P(t,T) = e^{-(T-t)r(T-t)} \tag{3.21}$$

In practice, this equation means that investors will earn $r(t, T)$ if they purchase the bond at t and hold it to maturity.

Coupon Bonds

The price of coupon bonds can also be derived in terms of a risk-neutral probability measure of the evolution of interest rates. The formula for this derivation is (3.22).

$$P_c = 100 E_t^{\tilde{P}} \left(e^{-\int_t^{tT} r(s)ds} \right) + \sum_{n:t_n > t}^{N} \frac{C}{w} E_t^{\tilde{P}} \left(e^{-\int_t^{t_n} r(s)ds} \right) \tag{3.22}$$

where

P_c = the price of a coupon bond

C = the bond coupon

t_n = the coupon payment dates, with $n \leq N$, and $t = 0$ at the time of valuation

w = the coupon frequency (annual or semiannual for plain vanilla bonds; monthly for certain floating-rate notes and asset-backed securities), expressed in number of times per year

T = the maturity date

Note that "100" on the right-hand side captures the fact that prices are quoted per 100 of the bond's principal, or nominal, value.

Expression (3.22) is written in some texts as (3.23).

$$P_c = 100 e^{-rN} + \int_n^N C e^{-rn} \tag{3.23}$$

Expression (3.22) can be simplified by substituting Df for the part of the expression representing the discount factor. Assuming an annual coupon, the result is (3.24).

$$P_c(t,T) = 100 \times Df_N + \sum_{n:t_n \geq t}^{N} C \times Df_n \tag{3.24}$$

Expression (3.24) states that the market value of a risk-free bond on any date is determined by the discount function on that date.

Forward Rates

An investor can combine positions in bonds of differing maturities to guarantee a rate of return that begins at a point in the future. The trade ticket is written at time t to cover the period T to $T + 1$ where $t < T$. The interest rate earned during this period is known as a *forward rate.* The mechanism by which a forward rate is guaranteed is described below, following Campbell et al (1997) and Jarrow (1996).

Guaranteeing a Forward Rate

Say an investor at time t simultaneously buys one unit of a zero-coupon bond maturing at time T that is priced at $P(t, T)$ and sells $P(t, T)/P(t, T + 1)$ units of zero-coupon bonds maturing at $T + 1$. Together these two transactions generate a zero cash flow: The investor receives a cash flow equal to one unit at time T and pays out $P(t, T)/P(t, T + 1)$ at time $T + 1$. These cash flows are identical to those that would be generated by a loan contracted at time t for the period T to $T + 1$ at an interest rate of $P(t, T)/P(t, T + 1)$. Therefore $P(t, T)/P(t, T + 1)$ is the forward rate. This is expressed formally in (3.25).

$$f(t,T) = \frac{P(t,T)}{P(t,T+1)} \qquad (3.25)$$

Using the relationships between bond price and yield described earlier, (3.25) can be rewritten in terms of yield as shown in (3.26), which represents the rate of return earned during the forward period $(T, T + 1)$. This is illustrated in **FIGURE 3.4**.

FIGURE 3.4 *Forward Rate Mechanics*

TIME

Transactions	t	T	$T+1$
Buy 1 unit of T-period bond	$-P(t,T)$	$+1$	
Sell $P(t,T)/P(t, T+1)$ $T+1$ period bonds	$+[(P(t,T)/P(t, T+1)]P\ m(t, T+1)$		$-P(t,T)/P(t, T+1)$
Net cash flows	0	$+1$	$-P(t,T)/P(t, T+1)$

$$f(t,T,T+1) = \frac{1}{\frac{P(t,T+1)}{P(t,T)}}$$

$$= \frac{r(t,T+1)^{(T+1)}}{r(t,T)^T} \tag{3.26}$$

Expression (3.25) can be rewritten as (3.27), which solves for bond price in terms of the forward rates from t to T. (See Jarrow (1996), chapter 3, for a description of this derivation.)

$$P(t,T) = \frac{1}{\prod_{k=t}^{T-1} f(t,k)}$$

where $\qquad \prod_{k=t}^{T-1} f(t,k) = f(t,t) f(t,t+1) \times \ldots \times f(t,T-1) \tag{3.27}$

Equation (3.27) states that the price of a zero-coupon bond is equal to the nominal value, here assumed to be 1, receivable at time T, after it has been discounted by the set of forward rates that apply from t to T.

Calculating a forward rate is equivalent to estimating what interest rate will be applicable for a loan beginning at some point in the future. This process exploits the law of one price, or no arbitrage. Consider a loan that begins at time T and matures at $T + 1$. The process of calculating the rate for that loan is similar to the one described above. Start with the simultaneous purchase at time t of one unit of a bond with a term of $T + 1$ for price $P(t, T+1)$ and sale of p amount of a bond with a term of T for price $P(t, T)$. The net cash position at t must be zero, so p must be

$$p = \frac{P(t,T+1)}{P(t,T)}$$

To preclude arbitrage opportunities, the value of p amount of one bond price divided by another must be the price of the $T + 1$–term bond at time T. Therefore the forward yield for the period T to $T + 1$ is given by expression (3.28).

$$f(t,T+1) = -\frac{\log P(t,T+1) - \log P(t,T)}{(T+1) - T} \tag{3.28}$$

If the period between T and the maturity of the longer-term bond is progressively reduced, the result is an instantaneous forward rate, which is given by formula (3.29).

$$f(t,T) = -\frac{\partial}{\partial T}\log P(t,T) \tag{3.29}$$

This is the price today of borrowing at time T. When $T = t$, the forward rate is equal to the instantaneous short rate $r(t)$; in other words, the spot and forward rates for the period (t, t) are identical. For other terms, the forward-rate yield curve will lie above the spot-rate curve if the spot curve is positively sloping; below it, if the spot-rate curve is inverted. Campbell et al (1997, pages 400–401) observes that this is a standard property for marginal and average cost curves. That is, when the cost of a marginal unit (say, of production) is above that of an average unit, the addition of a marginal unit increases the average cost. Conversely, the average cost per unit decreases if the marginal cost is below the average cost.

The Spot and Forward Yield Curve

From the preceding discussion of the relationships among bond prices, spot rates, and forward rates, it is clear, given any one of these sets, that it is possible to calculate the other two. As an illustration, consider the set of zero-coupon rates listed in **FIGURE 3.5**, which are assumed to be observed in the market. From these figures, the corresponding forward rates and zero-coupon bond prices may be calculated. **FIGURES 3.6** and **3.7** show the two derived curves plotted against the curve defined by the observed zero-coupon rates.

Note that the zero-coupon–yield curve has a positive, upward slope. The forward-rate curve should, therefore, lie above it, as discussed earlier. This is true until the later maturities, when the forward curve develops a serious kink. A full explanation for why this occurs lies outside the scope of this book. In simplest terms, though, it boils down to this: the forward rate, or *marginal rate of return,* is equal to the spot rate, or *average rate of return,* plus the rate of increase in the spot rate multiplied by the sum of the increases between t and T. If the spot rate is constant (corresponding to a flat curve), the forward-rate curve will equal it. An increasing spot-rate curve, however, does not always generate an increasing forward curve, only one that lies above it; it is possible for the forward curve to be increasing or decreasing while the spot rate is increasing. If the spot rate reaches a maximum level and then levels off or falls, the forward curve will begin to decrease at a maturity point earlier than the spot curve high point. In figure 3.6 the rate of increase in the spot rate in the last period is magnified when converted to the equivalent forward rate; if the last spot rate had been below the previous-period rate, the forward-rate curve would look like that in figure 3.7.

FIGURE 3.5 *Hypothetical Zero-Coupon Yield and Forward Rates*

TERM TO MATURITY (0, T)	SPOT RATE r (0, T)*	FORWARD RATE f (0, T)*	BOND PRICE P (0, T)
0			1
1	1.054	1.054	0.94877
2	1.055	1.056	0.89845
3	1.0563	1.059	0.8484
4	1.0582	1.064	0.79737
5	1.0602	1.068	0.7466
6	1.0628	1.076	0.69386
7	1.06553	1.082	0.64128
8	1.06856	1.0901	0.58833
9	1.07168	1.0972	0.53631
10	1.07526	1.1001	0.48403
11	1.07929	1.1205	0.43198

*Interest rates are given as $(1 + r)$

Calculating Spot Rates

It has been noted that a coupon bond may be regarded as a portfolio of zero-coupon bonds. An implied zero-coupon interest rate structure can therefore be derived from the yields on coupon bonds.

If the actual prices $P_1, P_2, \ldots, P_N$ of zero-coupon bonds with different maturities and \$1 nominal values are known, then the price P_C of a coupon bond of nominal value \$1 and coupon C can be derived using (3.30).

$$P_C = P_1 C + P_2 C + \ldots + P_N (1 + C) \tag{3.30}$$

Conversely if the coupon bond prices $P_{C1}, P_{C2}, \ldots, P_{CN}$ are known, the implied zero-coupon term structure can be derived through an iterative process using the relationship formalized in (3.30), as shown in (3.31) and (3.32).

FIGURE **3.6** *Hypothetical Zero-Coupon and Forward Yield Curves*
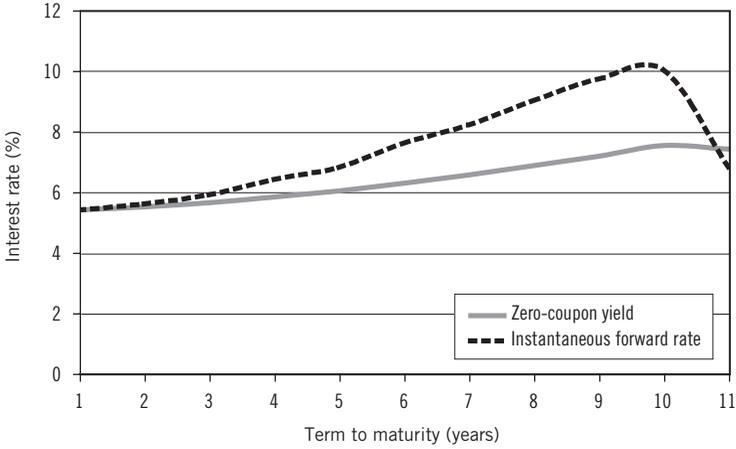
FIGURE **3.7** *Hypothetical Spot and Forward Yield Curves*
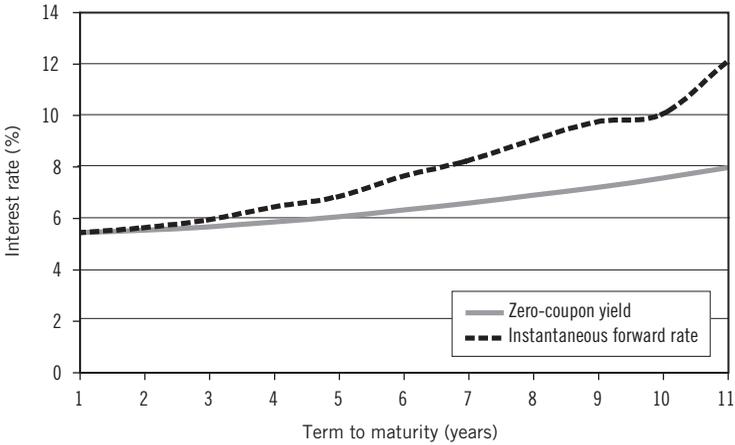
$$P_{C1} = P_1(1+C)$$

so

$$P_1 = \frac{P_{C1}}{(1+C)}$$

and so on, for prices P_1, P_2, ..., P_{N-1} (3.31)

$$P_N = \frac{P_{CN} - P_{N-1}C - ..... - P_1C}{1+C}$$ (3.32)

Finally, a regression technique known as *ordinary least squares,* or OLS (discussed in chapter 5), is applied to fit the term structure. Expression (3.30) restricts coupon-bond prices by requiring them to be precise functions of the prices of other coupon bonds. In practice, this strict relationship is vitiated by the effects of liquidity, taxes, and other factors. For this reason an error term u is added to (3.30), and the price is estimated using cross-sectional regression against all the other bonds in the market, as shown in (3.33).

$$P_{C_iN_i} = P_1C_i + P_2C_i + ... + P_{N_i}(1+C_i) + u_i$$ (3.33)

where $i = 1,2,....,I$
    C_i = the coupon on the ith bond
    N_i = the maturity of the ith bond

In (3.33) the regressor parameters are the coupons paid on each coupon-payment date, and the coefficients are the prices of the zero-coupon bonds P_j where $j = 1, 2, \ldots, N$. The values are obtained using OLS as long as the term structure is complete and $I \geq N$.

In practice, the term structure of coupon bonds is not complete, so the coefficients in (3.33) cannot be identified. To address this problem, McCulloch (1971, 1975) prescribes a *spline estimation* method that assumes zero-coupon bond prices vary smoothly with term to maturity. This approach defines price as a *discount function* of maturity, $P(N)$, which is a given by (3.34).

$$P(N) = 1 + \sum_{j=1}^{J} a_j f_j(N)$$ (3.34)

The function $f_j(N)$ is a known function of maturity N, and the coefficients a_j must be estimated. A regression equation is created by substituting (3.34) into (3.33) to give (3.35), whose value can be estimated using OLS

$$\prod_i = \sum_{j=1}^{J} a_j X_{ij} + u_i, \qquad i = 1, 2, ..., I$$

(3.35)

where

$$\prod_i \equiv P_{C_i N_i} - 1 - C_i N_i$$

$$X_{ij} \equiv f_j(N_i) + C_i \sum_{l=1}^{N_i} f_j(l)$$

The function $f_j(N)$ is usually specified by setting the discount function as a polynomial. In certain texts, including McCulloch, this is done by applying a spline function, which is discussed in the next chapter. (For further information, see the References section, particularly Suits et al (1978).)

Term Structure Hypotheses

As befits a subject that has been extensively researched, the term structure of interest rates has given rise to a number of hypotheses about how maturity terms are related to spot and forward rates and why it assumes certain shapes. This section briefly reviews the hypotheses.

The Expectations Hypothesis

Simply put, the *expectations hypothesis* states that the slope of the yield curve reflects the market's expectations about future interest rates. It was first enunciated in 1896, by Irving Fisher, a Yale economist, and later developed in Hicks (1946) and other texts. Shiller (1990) suggests that the hypothesis derives from the way market participants' view on future interest rates informs their decisions about whether to purchase long- or short-dated bonds. If interest rates are expected to fall, for instance, investors will seek to lock in the current high yield by purchasing long-dated bonds. The resultant demand will cause the prices of long-dated bonds to rise and their yields to decline. The yields will remain low as long as short-dated rates are expected to fall, rising again only when the demand for long-term bonds is reduced. Downward-sloping yield curves, therefore, indicate that the market expects interest rates to fall, while upward-sloping curves reflect expectations of a rise in short-term rates.

There are four distinct and incompatible versions of the expectations hypothesis. The *unbiased* version states that current forward rates are unbiased predictors of future spot rates. Let $f_t(T, T+1)$ be the forward rate at time t for the period from T to $T+1$ and r_T be the one-period spot rate at time T. The unbiased expectations hypothesis states that $f_t(T, T+1)$ is the expected value of r_T. This relationship is expressed in (3.36).

$$f_t(T, T+1) = E_t[r_T] \tag{3.36}$$

The *return-to-maturity* expectations hypothesis states that the return generated by holding a bond for term t to T will equal the expected return generated by continually rolling over a bond whose term is a period evenly divisible into $T - t$. This relationship is expressed formally in (3.37).

$$\frac{1}{P(t,T)} = E_t\left[(1+r_t)(1+r_{t+1})\ldots\ldots(1+r_{T-1})\right] \tag{3.37}$$

The left-hand side of (3.37) represents the return received by an investor holding a zero-coupon bond to maturity; the right-hand side is the expected return from time t to time T generated by rolling over a \$1 investment in one-period maturity bonds, each of which has a yield equal to the future spot rate r_t. In essence, this version represents an equilibrium condition, in which the expected returns for equal holding periods are themselves equal, although it does not state that the equality holds for different bond strategies. Jarrow (1996, page 52) argues for this hypothesis by noting that in an environment of economic equilibrium, the returns on zero-coupon bonds of similar maturity cannot be significantly different, since investors would not hold the bonds with the lower return. A similar argument can be made for coupon bonds of differing maturities. Any difference in yield would therefore not disappear as equilibrium was re-established. There are a number of reasons, however, that investors will hold shorter-dated bonds, irrespective of their yields. So it is possible for the return-to-maturity version of the hypothesis to be inapplicable.

From (3.36) and (3.37), it is clear that these two versions of the expectations hypothesis are incompatible unless no correlation exists between future interest rates. Ingersoll (1987) notes that although such an economic environment would be both possible and interesting to model, it is not related to reality, since interest rates are in fact highly correlated. Given a positive correlation between rates over a period of time, bonds with terms longer than two periods will have higher prices under the unbiased version than under the return-to-maturity version. Bonds with maturities of exactly two periods will have the same price under both versions.

The *yield-to-maturity* expectations hypothesis is stated in terms of yields, as expressed in equation (3.38).

$$\left[\frac{1}{P(t,T)}\right]^{\frac{1}{T-t}} = E_t\left[\left\{(1+r_t)(1+r_{t+1})\ldots(1+r_{T-1})\right\}^{\frac{1}{T-t}}\right] \tag{3.38}$$

The left-hand side of (3.38) specifies the yield-to-maturity at time t of the zero-coupon bond maturing at time T. The equation states that the expected holding-period yield generated by continually rolling over a series of one-period bonds will be equal to the yield guaranteed by holding a long-dated bond until maturity.

The *local* expectations hypothesis states that all bonds will generate the same expected rate of return if held for a small term. It is expressed formally in (3.39).

$$\frac{E_t\left[P(t+1,T)\right]}{P(t,T)} = 1 + r_t \tag{3.39}$$

This version of the hypothesis is the only one that permits no arbitrage, because the expected rates of return on all bonds are equal to the risk-free interest rate. For this reason, the local expectations hypothesis is sometimes referred to as the *risk-neutral* expectations hypothesis.

Liquidity Premium Hypothesis

The liquidity premium hypothesis, which has been described in Hicks (1946), builds on the insight that borrowers prefer to borrow long and lenders to lend short. It states that current forward rates differ from future spot rates by a *liquidity premium*. This is expressed formally as (3.40).

$$f_t(T,T+1) = E_t[r_T] + \pi_t(T,T+1) \tag{3.40}$$

Expression (3.40) states that the forward rate $f_t(T,T+1)$ is the expected one-period spot rate at time T, rate given by r_T, plus the liquidity premium, which is a function of the maturity of the bond (or the term of the loan). This premium reflects the conflicting requirements of borrowers and lenders: traders and speculators will borrow short and lend long in an effort to earn the premium.

Segmented Markets Hypothesis

The *segmented markets hypothesis,* first described in Culbertson (1957), seeks to explain the shape of the yield curve. It states that different types of market participants, with different requirements, invest in different parts of the term structure. For instance, the banking sector needs short-dated bonds, while pension funds require longer-term ones. Regulatory reasons may also affect preferences for particular maturity investments.

A variation on this hypothesis is the *preferred habitat* theory, described in Modigliani and Sutch (1967), which states that although investors have preferred maturities, they may choose other terms if they receive a

premium for so doing. This explains the "humped" shapes of yield curves, because if they have preferred maturities they will buy at those dates, depressing yields there and creating a hump. Cox, Ingersoll, and Ross (1981) describe the preferred habitat theory as a version of the liquidity preference hypothesis, where the preferred habitat is the short end of the yield curve, so that longer-dated bonds must offer a premium to entice investors.

Interest Rate Modeling

Chapter 3 introduced the basic concepts of bond pricing and analysis. This chapter builds on those concepts and reviews the work conducted in those fields. Term-structure modeling is possibly the most heavily covered subject in the financial economics literature. A comprehensive summary is outside the scope of this book. This chapter, however, attempts to give a solid background that should allow interested readers to deepen their understanding by referring to the accessible texts listed in the References section. This chapter reviews the best-known interest rate models. The following one discusses some of the techniques used to fit a smooth yield curve to market-observed bond yields.

Basic Concepts

Term-structure modeling is based on theories concerning the behavior of interest rates. Such models seek to identify elements or factors that may explain the dynamics of interest rates. These factors are random, or stochastic. That means their future levels cannot be predicted with certainty. Interest rate models therefore use statistical processes to describe the factors' stochastic properties and so arrive at reasonably accurate representations of interest rate behavior.

The first term-structure models described in the academic literature explain interest rate behavior in terms of the dynamics of the short rate. This term refers to the interest rate for a period that is infinitesimally small. (Note that *spot rate* and *zero-coupon rate* are terms used often to

mean the same thing.) The short rate is assumed to follow a statistical process, and all other interest rates are functions of the short rate. These are known as *one-factor* models. *Two-factor* and *multifactor* interest rate models have also been proposed. The model described in Brennan and Schwartz (1979), for instance, assumes that both the short rate and a long-term rate are the driving forces, while one presented in Fong and Vasicek (1991) takes the short rate and short-rate volatility as primary factors.

Short-Rate Processes

The original interest rate models describe the dynamics of the short rate; later ones—known as HJM, after Heath, Jarrow, and Morton, who created the first whole yield-curve model—focus on the forward rate.

In a one-factor model of interest rates, the short rate is assumed to be a random, or stochastic, variable—that is, it has more than one possible future value. Random variables are either *discrete* or *continuous.* A discrete variable moves in identifiable breaks or jumps. For example, although time is continuous, the trading hours of an exchange-traded future are discrete, since the exchange is shut outside of business hours. A continuous variable moves without breaks or jumps. Interest rates are treated in academic literature as continuous, although some, such as central bank base rates, actually move in discrete steps. An interest rate that moves in a range from 5 to 10 percent, assuming any value in between—such as 5.671291 percent—is continuous. Assuming that interest rates and the processes they follow are continuous, even when this does not reflect market reality, allows analyses to employ calculus to derive useful results.

The short rate follows a stochastic process, or *probability distribution.* So, although the rate itself can assume a range of possible future values, the process by which it changes from value to value can be modeled. A one-factor model of interest rates specifies the stochastic process that describes the movement of the short rate.

The analysis of stochastic processes employs mathematical techniques originally used in physics. An instantaneous change in value of a random variable x is denoted by dx. Changes in the random variable assume to follow a normal distribution, that is, the bell-shaped curve distribution. The shock, or *noise,* that impels a random variable to change value follows a randomly generated Weiner process, also known as a geometric Brownian motion. A variable following a Weiner process is a random variable, denoted by x or z, whose value alters instantaneously but whose patterns of change follow a normal distribution with mean 0 and standard deviation 1. Consider the zero-coupon bond yield r. Equa-

tion (4.1) states that *r* follows a continuous Weiner process with mean 0 and standard deviation 1.

$$dr = dz \qquad (4.1)$$

Changes or jumps in yield that follow a Weiner process are scaled by the volatility of the stochastic process that drives interest rates, which is denoted by σ. The stochastic process for change in yields is expressed by (4.2).

$$dr = \sigma dz \qquad (4.2)$$

The value of the volatility parameter is user-specified—that is, it is set at a value that the user feels most accurately describes the current interest rate environment. The value used is often the volatility implied by the market price of interest rate derivatives such as caps and floors.

The zero-coupon bond yield has thus far been described as a stochastic process following a geometric Brownian motion that drifts with no discernible trend. This description is incomplete. It implies that the yield will either rise or fall continuously to infinity, which is clearly not true in practice. To be more realistic, the model needs to include a term capturing the fact that interest rates move up and down in a cycle. The short rate's expected direction of change is the second parameter in an interest rate model. This is denoted in some texts by a letter such as *a* or *b,* in others by μ. The short-rate process can therefore be described as function (4.3).

$$dr = adt + \sigma dz \qquad (4.3)$$

where
　　dr = the change in the short rate
　　a = the expected direction of the change, or *drift*
　　dt = the incremental change in time
　　σ = the standard deviation of the price movements
　　dz = the random process

Equation (4.3), which is sometimes written with *dW* or *dx* in place of *dz*, is similar to the models first described in Vasicek (1977), Ho and Lee (1986), and Hull and White (1991). It assumes that, on average, the instantaneous change in interest rates is given by the function *adt*, with random shocks specified by σdz.

Because this process is a geometric Brownian motion, it has two important properties. First, the drift rate is equal to the expected change in

the short rate; if the drift rate is zero, the expected change is also zero, and the expected level of the short rate is equal to its current level. Second, the variance—that is, the square of the standard deviation—of the change in the short rate over a period T is equal to T, and its standard deviation is $\sqrt{T}$.

Equation (4.3) describes a stochastic short-rate process modified to include the direction of change. To be more realistic, it should also include a term describing the tendency of interest rates to drift back to their long-run average level. This process is known as *mean reversion* and is perhaps best captured in the Hull-White model. Adding a general specification of mean reversion to (4.3) results in (4.4).

$$dr = a\left(b - r\right)dt + \sigma dz \qquad (4.4)$$

where
   b = the long-run mean level of interest rates
   a = the speed of mean reversion, also known as the drift rate

Equation (4.4) represents an Ornstein-Uhlenbeck process. When r is above or below b, it will be pulled toward b, although random shocks generated by dz may delay this process.

Ito's Lemma

Market practitioners armed with a term-structure model next need to determine how this relates to the fluctuation of security prices that are related to interest rates. Most commonly, this means determining how the price P of a zero-coupon bond moves as the short rate r varies over time. The formula used for this determination is known as Ito's lemma. It transforms the equation describing the dynamics of the bond price P into the stochastic process (4.5).

$$dP = P_r dr + \frac{1}{2}P_{rr}\left(dr\right)^2 + P_t \qquad (4.5)$$

The subscripts rr in (4.5) indicate partial derivatives—the derivative with respect to one variable of a function involving several variables. The terms dr and $(dr)^2$ are dependent on the stochastic process that is selected for the short rate r. If this is the Ornstein-Uhlenbeck process represented in (4.4), the dynamics of P can be expressed as (4.6), which gives these dynamics in terms of the drift and volatility of the short rate.

$$dP = P_r \left[a \left(b - r \right) dt + \sigma dz \right] + \frac{1}{2} P_{rr} \sigma^2 dt + P_t dt \qquad (4.6)$$

$$= \left[P_r a \left(b - r \right) + \frac{1}{2} P_{rr} \sigma^2 + P_t \right] dt + P_r \sigma dz$$

$$= a \left(r, t \right) dt + \sigma \left(r, t \right) dz$$

Building a term-structure model involves these steps:
❑ Specify the stochastic process followed by the short rate, making certain assumptions about the short rate itself.
❑ Use Ito's lemma to express the dynamics of the bond price in terms of the short rate.
❑ Impose no-arbitrage conditions, based on the principle of hedging a position in one bond with a position in another bond (for a one-factor model; a two-factor model requires two bonds as a hedge) of a different maturity, to derive the partial differential equation of the zero-coupon bond price.
❑ Solve the partial differential equation for the bond price, which is subject to the condition that the price of a zero-coupon bond on maturity is 1.

One-Factor Term-Structure Models

This section briefly discusses some popular term-structure models, summarizing the advantages and disadvantages of each under different conditions and for different user requirements.

Vasicek Model

The Vasicek model was the first term-structure model described in the academic literature, in Vasicek (1977). It is a yield-based, one-factor equilibrium model that assumes the short-rate process follows a normal distribution and incorporates mean reversion. The model is popular with many practitioners as well as academics because it is analytically tractable—that is, it is easily implemented to compute yield curves. Although it has a constant volatility element, the mean reversion feature removes the certainty of a negative interest rate over the long term. Nevertheless, some practitioners do not favor the model because it is not necessarily arbitrage-free with respect to the prices of actual bonds in the market.

The Vasicek model describes the dynamics of the instantaneous short rate as (4.7).

$$dr = a \left(b - r \right) dt + \sigma dz \qquad (4.7)$$

where
  a = the speed of the mean reversion
  b = the mean-reversion level of r
  z = the standard Weiner process with mean 0 and standard deviation 1

Note that some texts use different notation, presenting the formula as

$$dr = \kappa(\theta - r)dt + \sigma dz$$

or

$$dr = \alpha(\mu - r)dt + \sigma dZ$$

The price at time t of a zero-coupon bond that matures at time T is given by (4.8). (For the derivation, see Vasicek (1977) and section 5.3 in Van Deventer and Imai (1997).)

$$P(t,T) = A(t,T)e^{-B(t,T)r(t)} \tag{4.8}$$

where
  $r(t)$ = the short rate at time t

$$B(t,T) = \frac{1 - e^{-a(T-t)}}{a}$$

and

$$A(t,T) = \exp\left[\frac{\left(B(t,T) - T + t\right)\left(a^2 b - \sigma^2/2\right)}{a^2} - \frac{\sigma^2 B(t,T)^2}{4a}\right]$$

In Vasicek's model, the short rate r is normally distributed. It therefore has a positive probability of being negative. Model-generated negative rates are an extreme possibility. Their occurrence depends on the initial interest rate and the parameters chosen for the model. They have been generated, for instance, when the initial rate was very low, like those seen in Japan for some time, and volatility was set at market levels. This possibility, which other interest rate models also allow, is inconsistent with a no-arbitrage market: as Black (1995) states, investors will hold cash rather than invest at a negative interest rate. For most applications, however, the model is robust, and its tractability makes it popular with practitioners.

Hull-White Model

The well-known model described in Hull and White (1993) uses Vasicek's model to obtain a theoretical yield curve and fit it to the observed market curve. It is therefore sometimes referred to as the extended Vasicek

model, with time-dependent drift. (Haug (1998) equates Hull-White to the Ho-Lee model, with mean reversion.) *Time-dependent drift* is a drift rate whose value is dependent on the time period used to calculate it, based on historical movement up until now. The model is popular with practitioners because it enables them to calculate a theoretical curve that is identical to yields observed in the market and that can be used to price bonds and bond derivatives and to calculate hedges.

The model is expressed in equation (4.9).

$$dr = a\left(\frac{b(t)}{a} - r\right)dt + \sigma dz \tag{4.9}$$

where

a = the speed of mean reversion
$\frac{b(t)}{a}$ = a time-dependent mean reversion

Given this description of the short-rate process, the price at time t of a zero-coupon bond with maturity T may be expressed as (4.10).

$$P(t,T) = A(t,T)e^{-B(t,T)r(t)} \tag{4.10}$$

where

$r(t)$ = the short rate at time t

$$B(t,T) = \frac{1 - e^{-a(T-t)}}{a}$$

$$\ln A(t,T) = \ln\left[\frac{P(0,T)}{P(o,t)}\right] - B(t,T)\frac{\partial P(0,t)}{\partial t} - \frac{v(t,T)^2}{2}$$

and

$$v(t,T)^2 = \frac{1}{2a^3}\sigma^2\left(e^{-aT} - e^{-at}\right)^2\left(e^{2at} - 1\right)$$

Further One-Factor Term-Structure Models

The academic literature and market participants have proposed a large number of alternatives to the Vasicek term-structure model and models, such as the Hull-White model, that are based on it. Like those they seek to replace, each of the alternatives has advantages and disadvantages.

The main advantage of Vasicek-type models is their analytic tractability. Their main weakness is that they permit negative interest rates. Negative interest rates are not impossible in the actual market; a bond

sold as part of a repurchase transaction for which excess market demand exists may have a negative rate. Academic researchers, however, often prefer to work with environments where negative rates are impossible (see Black (1995)). Since such environments are functions of more than one variable, Vasicek-type models are considered incapable of modeling them. This limitation does not preclude their use; rather, the choice of using them or not will depend on the conditions obtaining in the economy. In Japan from 1997 to 1998, for example, money market rates were frequently below 0.5 percent. At that level, even a volatility value below 5 percent, when plugged into the Vasicek model, will imply a high probability of negative interest rates. In this environment, therefore, practitioners may wish to use an alternative model, perhaps a two-factor or multifactor one.

Cox-Ingersoll-Ross (CIR) Model

Although published officially in 1985, the Cox-Ingersoll-Ross model was described in academic circles in 1977, or perhaps even earlier, which would make it the first interest rate model. Like Vasicek's it is a one-factor model that defines interest rate movements in terms of the dynamics of the short rate. It differs, however, in incorporating an additional feature, which relates the variation of the short rate to the level of interest rates. This feature precludes negative interest rates. It also reflects the fact that interest rate volatility rises when rates are high and correspondingly decreases when rates are low. The Cox-Ingersoll-Ross model is expressed by equation (4.11).

$$dr = k(b - r)dt + \sigma\sqrt{r}dz \qquad\qquad (4.11)$$

k = the speed of mean reversion

Note that the CIR model is often stated with k used to denote the speed of mean reversion as a was earlier.

Deriving the zero-coupon bond price given this model is formalized in equation (4.12), from Ingersoll (1987), chapter 18.

$$P(r,\tau) = A(\tau)e^{-B(\tau)r} \qquad\qquad (4.12)$$

where
τ = the term to maturity of the bond, or $(T - t)$

$$A(\tau) = \left[\frac{2\gamma e^{(\gamma+\lambda+k)\frac{\tau}{2}}}{g(\tau)} \right]^{\frac{2kb}{\sigma^2}}$$

$$B(\tau) = \frac{-2\left(1-e^{-\gamma\tau}\right)}{g(\tau)}$$

$$g(\tau) = 2\gamma + (k+\lambda+\gamma)\left(e^{\gamma\tau}-1\right)$$

$$\gamma = \sqrt{(k+\lambda)^2 + 2\sigma^2}$$

Again we use k to denote the speed of mean reversion, and λ is a user-defined parameter to adjust the level of this if necessary, connected with the risk premium associated with long-dated bond yields.

Some researchers—Van Deventer and Imai (1997) cite Fleseker (1993), for example—have stated that the difficulties in determining parameters for the CIR model have limited its use among market practitioners. Van Deventer and Imai, however, conclude that it deserves further empirical analysis and remains worthwhile for practical applications.

Two-Factor Interest Rate Models

This section briefly introduces a number of two-factor interest rate models. (The References section indicates sources for further research.) As their name suggests, these models specify the yield curve in terms of two factors, one of which is usually the short rate. A number of factors can be modeled when describing the dynamics of interest rates. Among them are

- ❏ the short-term or instantaneous interest rate
- ❏ the long-term—say, 10-year—interest rate
- ❏ short- and long-term real inflation-adjusted rates of interest
- ❏ the current or expected spread between the short- and long-term interest rates
- ❏ the current or expected corporate credit spread—that is, the difference in yield between the equivalent-maturity Treasury and bonds having the same credit rating as the subject bond
- ❏ the current inflation rate
- ❏ the long-term average expected inflation rate

Which factors the model incorporates depends in part on the purpose it is intended to serve—whether, for example, it is being used for pricing or hedging derivative instruments or for arbitrage trading. Other considerations also apply, such as the ease and readiness with which the parameters involved can be determined.

Brennan-Schwartz Model

The model described in Brennan and Schwartz (1979) uses the short rate and the long-term interest rate to specify the term structure. The long-term rate is defined as the market yield on an irredeemable, or perpetual, bond, also known as an undated or consol bond. Both interest rates are assumed to follow a Gaussian-Markov process. A Gaussian process is one whose marginal distribution, where parameters are random variables, displays normal distribution behavior; a Markov process is one whose future behavior is conditional on its present behavior only, and independent of its past. A later study, Longstaff and Schwartz (1992), found that Brennan-Schwartz modeled market bond yields accurately.

In the model, the dynamics of the logarithm of the short rate are defined by equation (4.13).

$$d\left[\ln(r)\right] = a\left[\ln(l) - \ln(p) - \ln(r)\right]dt + \sigma_1 dz_1 \tag{4.13}$$

In (4.13), p represents the relationship between the short rate, r, and the long-term rate, l. The short rate changes in response to moves in the level of the long rate, which follows the stochastic process described in equation (4.14).

$$dl = l\left[l - r + \sigma_2^2 + \lambda_2\sigma_2\right]dt + l\sigma_2 dz_2 \tag{4.14}$$

where
   λ = the risk premium associated with the long-term interest rate

Extended Cox-Ingersoll-Ross Model

Chen and Scott (1992) transformed the CIR model into a two-factor model specifying the interest rate as a function of two uncorrelated variables, both assumed to follow a stochastic process. The article demonstrated that this modification of the model has a number of advantages and useful applications.

The model specifies the relationship expressed in equation (4.15).

$$dy_i = k_i\left(\theta_i - y_i\right)dt + \sigma_i\sqrt{y_i}\,dz_i \tag{4.15}$$

where

    y_i = the independent variables y_1 and y_2

    k and θ = the parameters that describe the drift rate and the dz or stochastic factor

In this model, the formula for deriving the price of a zero-coupon bond is (4.16).

$$P(y_1, y_2, t, T) = A_1 A_2 e^{-B_1 y_1 - B_2 y_2} \qquad (4.16)$$

where *A* and *B* are defined as before

Heath-Jarrow-Morton (HJM) Model

The approach described in Heath-Jarrow-Morton (1992) represents a radical departure from earlier interest rate models. The previous models take the short rate as the single or (in two- and multifactor models) key state variable in describing interest rate dynamics. The specification of the state variables is the fundamental issue in applying multifactor models. In the HJM model, the entire term structure and not just the short rate is taken to be the state variable. Chapter 3 explained that the term structure can be defined in terms of default-free zero-coupon bond prices or yields, spot rates, or forward rates. The HJM approach uses forward rates.

The single-factor HJM model captures the change in forward rates at time *t*, with a maturity at time *T*, using
- ❏ a volatility function
- ❏ a drift function
- ❏ a geometric Brownian or Weiner process, which describes the shocks, or noise, experienced by the term structure

Consider a forward-rate term structure $f(0, T)$ that is *T*-integrable— for which the integral can be computed in terms of the variable *T*. The dynamics of this structure may be described by the stochastic differential equation (4.17).

$$df(t, T) = a(t, T) \, dt + \sigma \, dz \qquad (4.17)$$

where

    a = the drift rate, denoted in some texts by α

In some texts the term *k* is used for *a*.

    σ = the constant volatility level

dz = the geometric Brownian motion or Weiner process, denoted in some texts by W_t

By applying Ito calculus, i.e., Ito's lemma, (4.17) can be transformed to solve for the price of an asset. Taking the integral of expression (4.17) results in equation (4.18), which derives the forward rate,

$$f(t,T) = f(0,T) + \int_0^t a(s,T)\,ds + \sigma dz \qquad (4.18)$$

where
    s = an incremental move forward in time, so that

$$0 \le t \le T \text{ and } t \le s \le T$$

Equation (4.18) assumes that the forward rate is normally distributed. Crucially, the forward rates with maturities $f(0,1)$, $f(0,2)\ldots.f(0,\,T)$ are assumed to be perfectly correlated. The random element is the Brownian motion dz. The impact of this process is felt over time, rather than over different maturities.

The single-factor HJM model states that, given an initial forward-rate term structure $f(t,\,T)$ at time t, the forward rate for each maturity T is given by (4.20), which is the integral of (4.19).

$$df(t,T) = a(t,T)\,dt + \sigma(t,T)\,dz \qquad (4.19)$$

$$f(t,T) = f(0,T) + \int_0^t a(s,T)\,ds + \int_0^t \sigma(s,T)\,dz(s) \qquad (4.20)$$

Note that the expressions at (4.19) and (4.20) are simplified versions of the formal model being discussed.

Equations (4.19) and (4.20) state that the development of the forward rate for any maturity period T can be described in terms of the drift and volatility parameters $a(t,\,T)$ and $\sigma(t,\,T)$. The HJM model's primary assumption is that, for each T, the drift and volatility processes are dependent only on the histories up to the current time t of the Brownian motion process and of the forward rates themselves.

The Multifactor HJM Model

In the single-factor HJM model, forward rates of all maturities move in perfect correlation. For actual market applications—pricing an interest rate instrument that is dependent on the spread between two points on

the yield curve, for instance—this assumption can be too restrictive. In the multifactor model, each of the state variables is described by its own Brownian motion process. An *m*-factor model, for example, would include Brownian motions $dz_1, dz_2, \ldots, dz_m$. This allows each *T*-maturity forward rate to be described by its own volatility level $\sigma_i(t_i, T)$ and Brownian motion process dz_i. Under this approach, the forward rates derived from the bonds of differing maturities that define the current term structure evolve under more appropriate random processes, and different correlations among forward rates of differing maturities can be accommodated.

The multifactor HJM model is represented by equation (4.21).

$$f(t,T) = f(0,T) + \int_0^t a(s,T)\,ds + \sum_{i=1}^m \int_0^t \sigma_i(s,T)\,dz_i(s) \qquad (4.21)$$

Equation (4.21) states that the dynamics of the forward-rate process, beginning with the initial rate $f(0, T)$, are specified by the set of Brownian motion processes and the drift parameter. For practical applications, the evolution of the forward-rate term structure is usually derived in a binomial-type path-dependent process. Path-independent processes, however, have also been used, as has simulation modeling based on Monte Carlo techniques (see Jarrow (1996)). The HJM approach has become popular in the market, both for yield-curve modeling and for pricing derivative instruments, because it matches yield-curve maturities to different volatility levels realistically and is reasonably tractable when applied using the binomial-tree approach.

Choosing a Term-Structure Model

Selecting the appropriate term-structure model is more of an art than a science, depending on the particular application involved and the user's individual requirements. The Ho-Lee and BDT versions, for example, are *arbitrage*, or arbitrage-free, models, which means that they are designed to match the current term structure. With such models—assuming, of course, that they specify the evolution of the short rate correctly—the law of no-arbitrage can be used to determine the price of interest rate derivatives.

Equilibrium interest rate models also exist. These make the same assumptions about the dynamics of the short rate as arbitrage models do, but they are not designed to match the current term structure. The prices of zero-coupon bonds derived using such models, therefore, do not match prices seen in the market. This means that the prices of bonds and interest rate derivatives are not given purely by the short-rate process. In brief, arbitrage models take as a given the current yield curve described by the

market prices of default-free bonds; equilibrium models do not.

Among the considerations that should be taken into account when deciding which term-structure model to use are the following:

❑ *Ease of application.* In this respect, arbitrage models have the advantage. Their key input is the current spot-rate term structure. This, unlike the input to equilibrium models, can be determined in a straightforward process from the market price of bonds currently trading in the market.

❑ *Desirability of capturing market imperfections.* The term structure generated by an arbitrage model will reflect the current market term structure, which may include pricing irregularities arising from liquidity and other considerations. Equilibrium models do not reflect such irregularities. Selection of a model will depend on whether or not modeling market imperfections is desirable.

❑ *Application in pricing bonds or interest rate derivatives.* Traditional seat-of-the-pants bond pricing often employs a combination of good sense, prices observed in the market (often from interdealer-broker screens), and gut feeling. A more scientific approach may require a yield-curve model, as will relative value trading—trading bonds of different maturities against each other, for example, in order to bet on changes in yield spreads. In such cases, equilibrium models are clearly preferable, since traders will want to compare the theoretical prices given by the model with the prices observed in the market. Arbitrage models, in contrast, assume that the market bond prices are correct. So, an arbitrage model would always suggest that there was no gain to be made from a relative value trade.

Pricing derivative instruments such as interest rate options (or swaptions) requires a different emphasis. The primary consideration of the derivative market maker is the technique and price of hedging the derivative. When writing derivative contracts, market makers simultaneously hedge their exposure using either the underlying assets or a combination of these and other derivatives, such as exchange-traded futures. They profit from the premium they extract and from the difference in price over time between the derivative and the hedge. In this enterprise, only arbitrage models are appropriate, because they price derivatives relative to the actual market. An equilibrium model, in contrast, prices derivatives relative to a theoretical market, which is not appropriate, since those used in the hedge are market instruments.

❑ *Use of models over time.* The parameters in an interest rate model —most notably the drift, volatility, and, if applicable, mean reversion rate—reflect the current state of the economy. This state is not constant; the drift rate used today, for example, may well differ from the value used

tomorrow. Over time, any model must be recalibrated. For arbitrage models, however, this is a constant process, since their parameters change continuously. Equilibrium model parameters, in contrast, are calculated from historical data or using logic, and so may not change as frequently. On the other hand, the accuracy of these models may suffer over time, as current rates diverge from historic average rates. Users must decide whether the greater accuracy of the arbitrage model is worth the constant tweaking that makes it possible.

This list is just a sample. Users must consider a wide range of issues when selecting an interest rate model. It has been observed, for instance, that models that incorporate mean reversion are more accurate than those that do not. Another crucial consideration is the computer processing power available to the user. Single-factor models are often preferred precisely because their processing is more straightforward.

This page is intentionally blank

Fitting the Yield Curve

This chapter considers some of the techniques used to fit the model-derived term structure to the observed one. The Vasicek, Brennan-Schwartz, Cox-Ingersoll-Ross, and other models discussed in chapter 4 made various assumptions about the nature of the stochastic process that drives interest rates in defining the term structure. The zero-coupon curves derived by those models differ from those constructed from observed market rates or the spot rates implied by market yields. In general, market yield curves have more-variable shapes than those derived by term-structure models. The interest rate models described in chapter 4 must thus be calibrated to market yield curves. This is done in two ways: either the model is calibrated to market instruments, such as money market products and interest rate swaps, which are used to construct a yield curve, or it is calibrated to a curve constructed from market-instrument rates. The latter approach may be implemented through a number of *non-parametric* methods.

There has been a good deal of research on the empirical estimation of the term structure, the object of which is to construct a zero-coupon or spot curve (or, equivalently, a forward-rate curve or discount function) that represents both a reasonably accurate fit to market prices and a smooth function—that is, one with a continuous first derivative. Though every approach must make some trade off between these two criteria, both are equally important in deriving a curve that makes economic sense.

This chapter presents an overview of some of the methods used to fit the yield curve. A selection of useful sources for further study is given, as usual, in the References section.

Yield Curve Smoothing

Carleton and Cooper (1976) describes an approach to estimating term structure that assumes default-free bond cash flows, payable on specified discrete dates, to each of which a set of unrelated discount factors are applied. These discount factors are estimated as regression coefficients, with the bond cash flows being the independent variables and the bond price at each payment date the dependent variable. This type of simple linear regression produces a discrete discount function, not a continuous one. The forward-rate curves estimated from this function are accordingly very jagged.

McCulloch (1971) proposes a more practical approach, using polynomial *splines.* This method produces a function that is both continuous and linear, so the *ordinary least squares* regression technique can be employed. A 1981 study by James Langetieg and Wilson Smoot, cited in Vasicek and Fong (1982), describes an extended McCulloch method that fits *cubic splines* to zero-coupon rates instead of the discount function and uses nonlinear methods of estimation.

The term structure can be derived from the complete set of discount factors—the discount function—which can themselves be extracted from the price of default-free bonds trading in the market using the bootstrapping technique described in chapter 1. This approach is problematic, however. For one thing, it is unlikely that the complete set of bonds in the market will pay cash flows at precise six-months intervals from today to thirty years from now or longer, which, as explained in chapter 1, is necessary for the bootstrapping derivation to work. Adjustments must be made for cash flows received at irregular intervals or, in the case of longer maturities, not at all. Another issue is that bootstrapping calculates discount factors for terms that are multiples of six months, but in reality, non-standard periods, such as 4-month or 14.2-year maturities, may be involved, particularly in pricing derivative instruments. A third problem is that bonds' market prices often reflect investor considerations such as the following:

❑ how liquid the bonds are, which is itself a function of issue size, market-maker support, investor demand, whether their maturities are standard or not, and other factors
❑ whether the bonds trade continuously (if they don't, some prices will be "newer" than others)
❑ the tax treatment of the cash flows
❑ the bid-offer spread

These considerations introduce what is known in statistics as *error* or *noise* into market prices. To handle this, *smoothing* techniques are used in the derivation of the discount function.

FIGURE **5.1** *U.S. Treasury Yields to Maturity on December 23, 2003*

YIELD CURVE – US TREASURY ACTIVES
Currency ■   B Bid   RANGE 3M –30
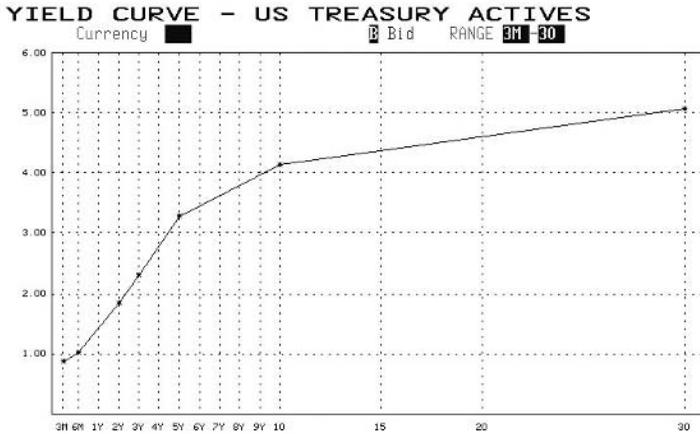
FIGURE **5.2** *Discount Function Derived from U.S. Treasury Prices on December 23, 2003*
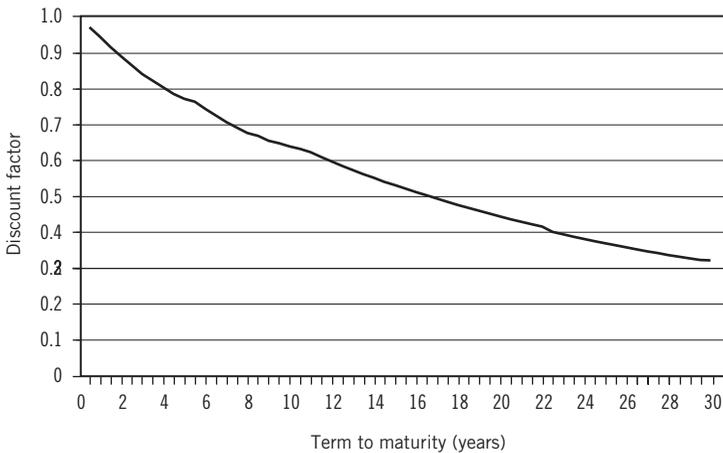
FIGURE **5.2** is the graph of the discount function derived by bootstrapping from the U.S. Treasury prices as of December 23, 2003. FIGURE **5.3** shows the zero-coupon yield and forward-rate curves corresponding to this discount function. Compare these to the yield curve in FIGURE **5.1**

FIGURE 5.3 *Zero-Coupon (Spot) Yield and Forward-Rate Curves Corresponding to the Discount Function*
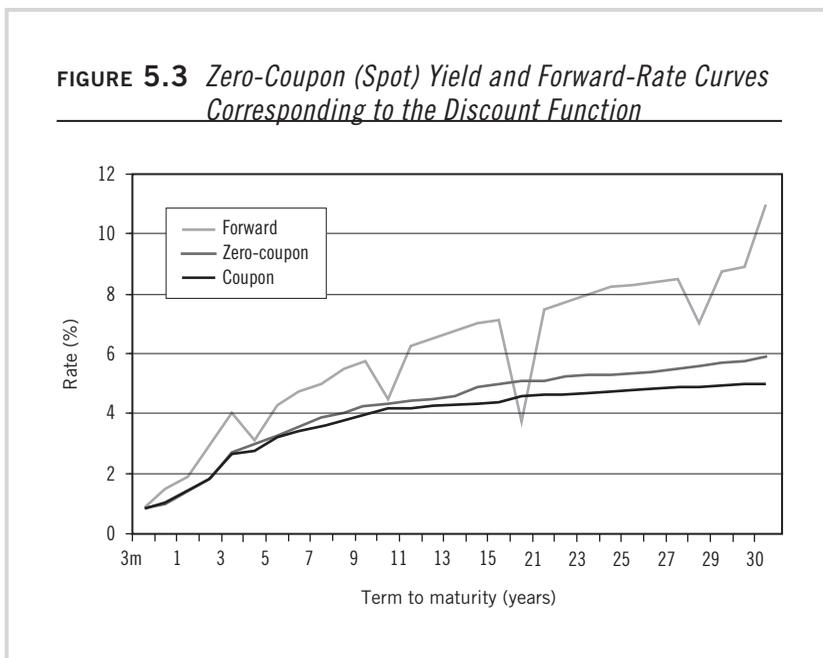
(a Bloomberg screen), which is plotted from Treasury redemption yields using Bloomberg's IYC function.

The zero-coupon curve in figure 5.3 is relatively smooth, though not quite as smooth as the discount function curve in figure 5.2. The forward-rate curve, in contrast, is jagged. Irregularities in implied forward rates indicate to the fixed-income analyst that the discount function and the zero-coupon curve are not as smooth as they appear. The main reason that the forward-rate curve is so jagged is that minor errors in the discount factors, arising from any of the reasons given above, are compounded in calculating spot rates from them and further compounded when these are translated into the forward rates.

Smoothing Techniques

A common smoothing technique is *linear interpolation.* This approach fills in gaps in the market-observed yield curve caused by associated gaps in the set of observed bond prices by interpolating missing yields from actual yields.

Linear interpolation is simple but not accurate enough to be recommended. Market analysts use multiple regression or spline-based methods instead. One technique is to assume that the discount factors represent a functional form—that is, a higher-order function that takes

other functions as its parameters—and estimate its parameters from market bonds prices.

Cubic Polynomials

One simple functional form that can be used in smoothing the discount function is a *cubic polynomial.* This approach approximates the set of discount factors using a cubic function of time, as shown in (5.1).

$$d(t) = a_0 + a_1(t) + a_2(t)^2 + a_3(t)^3 \qquad (5.1)$$

where
$d(t)$ = the discount factor for maturity t

Some texts use a, b, and c in place of a_1 and so on.

The discount factor for $t = 0$, that is for a bond maturing right now, is 1, i.e., the present value of a cash flow received right now is simply the value of the cash flow. Therefore $a_0 = 1$, and (5.1) can then be rewritten as (5.2).

$$\hat{d}(t) - 1 = a_1(t) + a_2(t)^2 + a_3(t)^3 \qquad (5.2)$$

As discussed in chapter 1, the market price of a coupon bond can be expressed in terms of discount factors. Equation (5.3) derives the price of an N-maturity bond paying identical coupons C at regular intervals and a redemption value M at maturity.

$$P = d(t_1)C + d(t_2)C + \dots + d(t_N)(C + M) \qquad (5.3)$$

Replacing the discount factors in (5.3) with their cubic polynomial expansions, given by (5.2), results in expression (5.4).

$$P = C\left[1 + a_1(t_1) + a_2(t_1)^2 + a_3(t_1)^3\right] + \dots + (C + M)\times$$

$$\times\left[1 + a_1(t_N) + a_2(t_N)^2 + a_3(t_N)^3\right] \qquad (5.4)$$

To describe the yield curve, it is necessary to know the value of the coefficients of the cubic function. They can be solved by rearranging (5.4) as (5.5) and further rearranging this expression to give (5.6), which has been simplified by substituting X_i for the appropriate bracketed expressions. This is the form most commonly found in textbooks.

$$P = M + \sum C + a_1\left[C(t_1) + ... + (C + M)(t_N)\right] +$$

$$+ a_2\left[C(t_1)^2 + ... + (C + M)(t_N)^2\right] +$$

$$+ a_3\left[C(t_1)^3 + ... + (C + M)(t_N)^3\right] \qquad (5.5)$$

$$P - \left(M + \sum C\right) = a_1 X_1 + a_2 X_2 + a_3 X_3 \qquad (5.6)$$

The cubic polynomial approach has several drawbacks that limit its practical application. First, equation (5.6) must be solved for each bond in the data set. More important, the result is not a true curve but a set of independent discount factors that have been adjusted with a line of best fit. Third, small changes in the data can have a significant impact at the nonlocal level. A change in a single data point in the early maturities, for example, can result in bad behavior in the longer maturities.

One solution is to use a *piecewise cubic polynomial,* where $d(t)$ is associated with a different cubic polynomial, with different coefficients. A special case of this approach, the cubic spline, is discussed in the next section.

Non-Parametric Methods

Beyond the cubic polynomial, there are two main approaches to fitting the term structure: *parametric* and *non-parametric curves.* Parametric curves are based on term-structure models such as those discussed in chapter 4. As such, they need not be discussed here. Non-parametric curves, which are constructed employing spline-based methods, are not derived from any interest rate models. Instead, they are general approaches, described using sets of parameters. They are fitted using econometric principles rather than stochastic calculus, and are suitable for most purposes.

Spline-Based Methods

A spline is a type of linear interpolation. It takes several forms. The spline function fitted using regression is the most straightforward and easiest to understand. Unfortunately, as illustrated in James and Webber (2000), section 15.3, when applied to yield-curve construction, this method can be overly sensitive to changes in parameters, causing curves to jump wildly.

An nth-order spline is a piecewise polynomial approximation using n-degree polynomials that are differentiable n-1 times, i.e., they have n-1 derivatives. *Piecewise* signifies that the different polynomials are connected

FIGURE 5.4 *Cubic Spline with Knot Points at 0, 2, 5, 10, and 25*
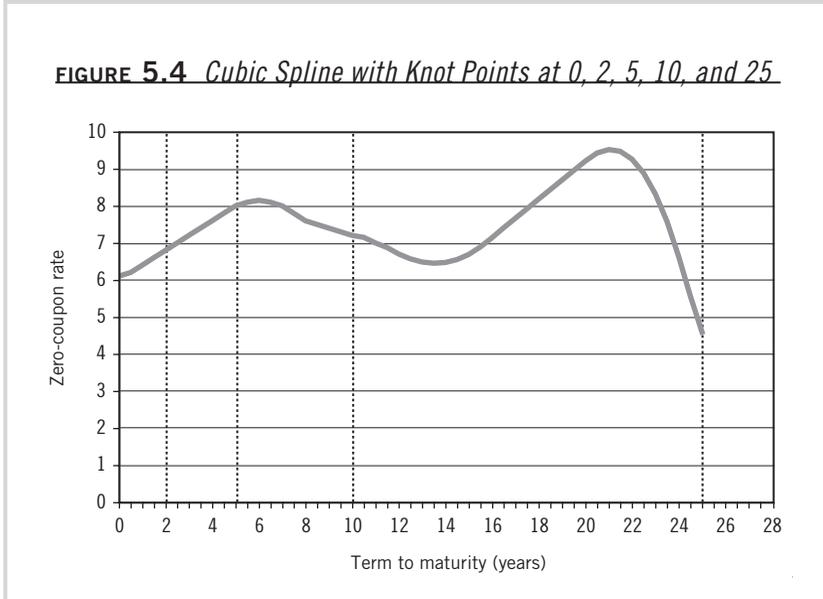
at arbitrarily selected *knot* points. A cubic spline is a piecewise three-degree, or cubic, polynomial that is differentiable twice along all its points.

The *x*-axis in the regression is divided into segments at the knot points, at each of which the slopes of adjoining curves on either side of the point must match, as must the curvatures. **FIGURE 5.4** shows a cubic spline with knot points at 0, 2, 5, 10, and 25 years, at each of which the curve is a cubic polynomial. This function permits a high and low to be accommodated in each space bounded by the knot points. The values of the curve can be adjoined at the knot point in a smooth function.

Cubic spline interpolation assumes that there is a cubic polynomial that can estimate the yield curve at each maturity gap. A spline can be thought of as a number of separate polynomials of the form $y = f(X)$, where X is the complete range of the maturity term divided into user-specified segments that are joined smoothly at the knot points. Given a set of bond yields $r_0, r_1, r_2, \ldots, r_n$ at maturity points $t_0, t_1, t_2, \ldots, t_n$, the cubic spline function can be estimated as follows:

❑ The yield of bond i at time t is expressed as a cubic polynomial of the form $r_i(t) = a_i + b_i t + c_i t^2 + d_i t^3$ for the interval between t_{i-1} and t_i.

❑ The $4n$ unknown coefficients of the cubic polynomial for all n intervals between the $n + 1$ data points are calculated.

❑ These equations are solved, which is possible because they are made to fit the observed data. They are twice differentiable at the knot points, and the two derivatives at these points are equal.

❑ The curve is constrained to be instantaneously straight at the shortest and the longest maturities, that is $r''(0) = 0$, with the double prime notation representing the second-order derivative.

The general formula for a cubic spline is (5.7).

$$s(\tau) = \sum_{i=0}^{3} a_i \tau^i + \frac{1}{3!} \sum_{p=1}^{n-1} b_p (\tau - X_p)^3 \tag{5.7}$$

where
τ = the time of receipt of cash flows
X_p = the knot points, with $\{X_0,\ldots,X_n\}, X_p < X_{p+1}, p = 0,\ldots,n-1$

$$(\tau - X_p) = \max(\tau - X_p, 0)$$

In practice, the spline is expressed as a set of basis functions, with the general spline being a combination of these. This may be arrived at using *B-splines.* The B-spline for a specified number of knot points $\{X_0,\ldots,X_n\}$ is (5.8).

$$B_p(\tau) = \sum_{j=p}^{p+4} \left[\prod_{i=p, i \neq i}^{p+4} \frac{1}{X_i - X_j} \right] (\tau - X_p)^3 \tag{5.8}$$

where $B_p(\tau)$ denotes cubic splines that are approximated on $\{X_0,\ldots,X_n\}$ using function (5.9)

$$\delta(\tau) = \delta(\tau \mid \lambda_{-3},\ldots,\lambda_{n-1}) = \sum_{p=-3}^{n-1} \lambda_p B_p(\tau) \tag{5.9}$$

where
$\lambda = (\lambda_{-3},\ldots,\lambda_{n-1})$ are the required coefficients

The maturity periods $\tau_1,\ldots,\tau_n$ specify the B-splines, so $B = \{B_p(\tau_j)\}_{p=-3,\ldots,n-1, j=1,\ldots,m}$ and $\hat\delta = (\delta(\tau_1),\ldots,\delta(\tau_m))$. From this, the equivalence (5.10) and the regression equation (5.11) follow.

$$\hat\delta = B'\lambda \tag{5.10}$$

$$\lambda^* = \arg\min_{\lambda} \{\varepsilon'\varepsilon \mid \varepsilon = P - D\lambda\} \tag{5.11}$$

where
$D = CB'$
$\varepsilon'\varepsilon$ = the minimum errors

Equation (5.11) is computed using ordinary least squares regression.

Nelson and Siegel Curves

The curve-fitting technique first described in Nelson and Siegel (1985) has since been modified by other authors, resulting in a "family" of curves. This is not a bootstrapping technique but a method for estimating the zero-coupon yield curve from observed T-bill yields, assuming a forward-rate function. The method creates a satisfactory rough fit of the complete term structure, with some loss of accuracy at the very short and very long ends.

The original article specifies four parameters. The implied forward-rate yield curve is modeled along the entire term structure using function (5.12).

$$rf(m,\beta) = \beta_0 + \beta_1 \exp\left(\frac{-m}{t_1}\right) + \beta_2 \left(\frac{m}{t_1}\right) \exp\left(\frac{-m}{t_1}\right) \tag{5.12}$$

where

$\beta = (\beta_0, \beta_1, \beta_2, t_1)$ = the vector of parameters describing the yield curve

m = the maturity at which the forward rate is calculated

The three components on the right side of equation (5.12) are a constant term, a decay term, and a term representing the "humped" nature of the curve. The long end of the curve approaches an asymptote, the value of which is given by β_0 at the long end, with a value of $\beta_0 + \beta_1$ at the short end.

The Svensson model, proposed in Svensson (1994), is a version of the Nelson and Siegel curve with an adjustment for the hump in the yield curve. This is accomplished by expanding (5.12) as (5.13).

$$\tag{5.13}$$

$$rf(m,\beta) = \beta_0 + \beta_1 \exp\left(\frac{-m}{t_1}\right) + \beta_2 \left(\frac{m}{t_1}\right) \exp\left(\frac{-m}{t_1}\right) + \beta_3 \left(\frac{m}{t_2}\right) \exp\left(\frac{-m}{t_2}\right)$$

The Svensson curve is thus modeled using six parameters, with the additional input of β_3 and t_2.

As approximations, Nelson and Siegel curves are appropriate for no-arbitrage applications. They are popular in the market because they are straightforward to calculate. Jordan and Mansi (2000) imputes two further advantages to them: they force the long-date forward curve into a horizontal asymptote, and the user is not required to specify knot points, whose choice determines how effective the cubic spline curves are. The

same article also notes two disadvantages: these curves are less flexible than spline-based ones, and they may not fit the observed data as accurately. James and Webber (2000), pages 444–445, suggests that because of the limited number of parameters involved, Nelson and Siegel curves lack flexibility and states further that they are accurate only for yield curves with one hump, not for those with both a hump and trough.

Comparing Curves

The choice of curve depends on the user's requirements and intended application. It usually represents a trade-off between ease of computation and accuracy. The user must determine how well any curve meets the following criteria, which are all met to a greater or lesser extent by the methodologies discussed:

- ❏ *Accuracy.* Does the curve fit reasonably well?
- ❏ *Flexibility.* Is it flexible enough to accommodate a variety of yield curve shapes?
- ❏ *Model consistency.* Is the fitting method consistent with a model such as Vasicek or Cox-Ingersoll-Ross?
- ❏ *Simplicity.* Is the curve tractable—that is, reasonably simple to compute?

A good summary of the advantages and disadvantages of popular modeling methods can be found in James and Webber (2000), chapter 15.

Selected Cash and Derivative Instruments

Part Two discusses selected instruments traded in the debt capital markets. The products—hybrid securities, mortgage-backed bonds, and callable bonds—have been chosen to give the reader an idea of the variety available in the market. Also described are index-linked bonds and a structured product known as a collateralized debt obligation (CDO). Some of the techniques for analyzing these more complex products are explained.

This part also considers the primary fixed-income derivative instruments. These are not securities in the cash markets and are fixed-income derivatives (or interest rate derivatives) in the synthetic markets.

The products discussed include interest rate swaps, options, and credit derivatives. There is also a chapter on the theory behind forward and futures pricing, with a case study featuring the price history and implied repo rate for the CBOT long bond future.

This page is intentionally blank

Forwards and Futures Valuation

I nterest rate futures will be described in chapter 13. This chapter develops basic valuation concepts. The discussion is adapted, with permission, from section 2.2 of Rubinstein (1999).

Forwards and Futures

A forward is a contract between two parties in which one agrees to purchase from the other a specified asset at a specified price for delivery at a specified future date. The following discussion refers to these variables:

P = the current price of the underlying asset, also known as the *spot* price

P_T = the spot price of the underlying asset at the time of delivery

X = the delivery price specified in the forward contract

T = the term to maturity of the contract, in years, also referred to as the time to delivery

r = the risk-free T-bill interest rate

R = the return of the payout, or its yield

F = the current forward price, that is, the current market prediction of the underlying asset's price on the delivery date

The payoff of a forward contract is given by expression (6.1).

$$Payoff = P_T - X \tag{6.1}$$

The payout yield, R, is the percentage of the spot price that is paid out at contract expiry, i.e., $R = (P_T - X) / P_T$. The forward contract terms are set so that the present value of the payout $(P_T - X)$ is zero. This means that the forward price, F, on day one of the contract equals X. (Note that the forward price is not the same as the value of the contract, which at this point is zero.) From the initiation of the contract until its expiration, the value of X remains fixed. The forward price, F, however, fluctuates, generally rising and falling with the spot price of the underlying asset.

Like forwards, futures contracts also represent agreements to purchase/sell a specified asset at a specified price for delivery on a specified date. While forwards, however, are custom instruments designed to meet individual requirements, futures are standardized contracts that are traded on recognized futures exchanges.

Commodity futures are settled by the physical delivery of the underlying asset; many financial futures are settled in cash. A bond future—which is written on a *notional* bond that can represent any of a set of bonds fitting the contract terms, known as the contract's *delivery basket*—is settled by delivering to the long counterparty one of these bonds. Only a very small percentage of either commodity or financial futures contracts are *delivered into*—that is, involve the actual transfer of the underlying asset to the long counterparty. This is because the majority of futures trading is done to hedge or to speculate. Accordingly, most futures positions are netted out to zero before contract expiry, although, if the position is held into the delivery month, depending on the contract's terms and conditions, the long future may be delivered into.

Cash Flow Differences

Aside from how they are constructed and traded, the most significant difference between forwards and futures, and the feature that influences differences between their prices, concerns their cash flows. The profits or losses from futures trading are realized at the end of each day. Because of this daily settlement, at expiration all that needs to be dealt with is the change in the contract value from the previous day. With forwards, in contrast, the entire payout occurs at contract expiry. (In practice, the situation is somewhat more complex, because the counterparties have usually traded a large number of contracts with each other, across a number of maturity periods and, perhaps, instruments, and as these contracts expire they exchange only the *net* loss or gain on the contract.)

FIGURE 6.1 shows the daily cash flows for a forward and a futures contract having identical terms. The futures contract generates intermediate cash flows; the forward doesn't. As with the forward contract, the delivery price specified in the futures contract is set so that at initiation, the pres-

FIGURE 6.1 *Cash Flows for Forwards and Futures Contracts*
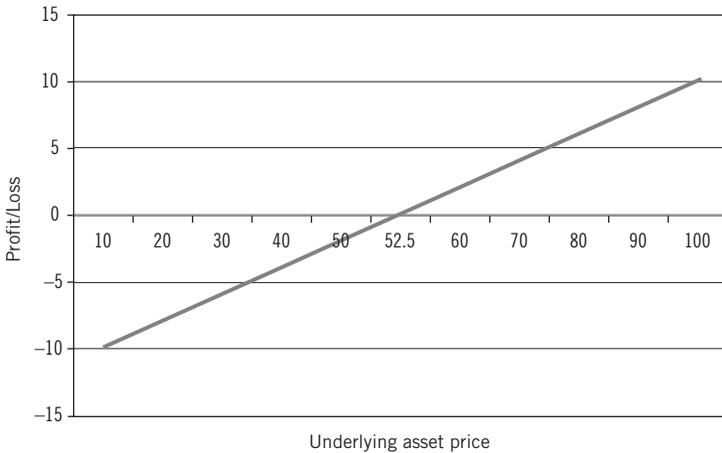
ent value of the futures contract is zero. At the end of each day, the future is marked to market at the close price. This will result in a profit or a gain—or neither, if the closing price is unchanged from the previous day's closing price, which technical traders call a *doji*—which is handed over to the appropriate counterparty. Through this daily settlement, the nominal delivery price is reset each day so that the present value of the contract is always zero. This means that the future and nominal delivery prices of a futures contract are the same at the end of each trading day.

As illustrated in figure 6.1, the process works as follows: After day one, the future price is reset from F to F_1. The amount $(F_1 - F)$, if positive, is handed over by the short counterparty to the long counterparty. If the amount is negative, it is paid by the long counterparty to the short. On the expiry day, T, of the contract, the long counterparty receives a settlement amount equal to $P_T - F_{T-1}$, which is the difference between the future price and the price of the underlying asset. The daily cash flows cancel each other out, so that at expiry the value of the contract is identical to that for a forward, that is $(P_T - F)$.

All market participants in exchange-traded contracts trade with a central counterparty, the exchange's clearing house. This eliminates counterparty risk. The clearing house is able to guarantee each deal, because all participants are required to contribute to its clearing fund through *margining:* each participant deposits an *initial margin* and then every day, as profits and losses are recorded, deposits a further *variation margin* as

needed. Marking the futures contract to market is essential to this margin process.

Daily settlement has both advantages and disadvantages. If a position is profitable, receiving part of this profit daily, as happens with a futures contract, is advantageous because the funds can be reinvested while the position is still maintained. On the other hand, a losing position saddles the holder of a futures contract with daily losses not suffered by the holder of a loss-making forward position.

Relationship Between Forward and Futures Prices

Under certain specified conditions, the prices of futures and forwards with identical terms must be the same. Consider two trading strategies with identical terms to maturity and written on the same underlying asset, one using forward contracts and the other futures. Neither strategy requires an initial investment, and both are *self-financing*—that is, all costs and future funding are paid out of proceeds. Assume the following conditions:

❑ no risk-free arbitrage opportunities
❑ an economist's perfect market
❑ certainty of returns

For the strategy employing forwards, r^T contracts are bought, where r is the daily return (or instantaneous money market rate) and T the maturity term in days. The start forward price is $F = X$, and the payoff on expiry is $r^T (P_T - F)$.

The futures strategy is more involved, because of the margin cash flows that are received or paid daily during the term of the trade. On day one, r contracts are bought, each priced at F. After the close that day, $F_1 - F$ is received. The position is closed out, and the cash received is invested at the daily rate, r, up to the expiry date. The return on this investment is r^{T-1}. Thus, on expiry the counterparty will receive $r(F_1 - F)r^{T-1}$.

The next day, r^2 futures contracts are bought at a price of F_1. At the close, the cash flow of $F_2 - F_1$ is received and invested at r^{T-2}, generating a return on expiry of $r^2(F_2 - F_1)r^{T-2}$. This process is repeated until the expiry date, which is assumed to be the delivery date. The return from following this strategy is expression (6.2).

$$r^T (F_1 - F) + r^T (F_2 - F_1) + r^T (F_3 - F_2) +$$

$$+ ..... + r^T (P_T - F_{T-1}) = r^T (P_T - F) \qquad (6.2)$$

This is also the payoff from the forward contract strategy. The key point is that if equation (6.3) below holds, then so must (6.4).

$$r^T \left(P_T - F\right)_{forward} = r^T \left(P_T - F\right)_{future} \tag{6.3}$$

$$F_{forward} = F_{future} \tag{6.4}$$

Forward-Spot Parity

The forward strategy can be used to imply the forward price, provided that the current price of the underlying and the money market interest rate are known. **FIGURE 6.2** illustrates how this works, using the one-year forward contract whose profit/loss profile is graphed in figure 6.1 and assuming an initial spot price, P, of \$50, a risk-free rate, r, of 1.05 percent, and a payout yield, R, of 1 percent.

FIGURE 6.3 shows that the payoff profile illustrated in figure 6.1 can be replicated by a portfolio composed of one unit of the underlying asset

FIGURE 6.2 *Forward Contract Profit/Loss Profile*

TIME	FORWARD CONTRACT	FUTURES CONTRACT
0	0	0
1	0	$F_1 - F$
2	0	$F_2 - F_1$
3	0	$F_3 - F_2$
4	0	$F_4 - F_3$
5	0	$F_5 - F_4$
…	0	…
…	0	…
…	0	…
$T-1$	0	$F_{T-1} - F_{T-2}$
T	$P_T - F$	$P_T - F_{T-1}$
Total	$P_T - F$	$P_T - F$

FIGURE **6.3** *Forward Strategy*

	CASH FLOWS START DATE	EXPIRY
Buy forward contract	0	$P_T - F$
Buy one unit of the underlying asset	−50	P_T
Borrow zero present-value of forward price	$F / 1.05$	F
Total	$-50 + F / 1.05$	$P_T - F$

Result

Set $-50 + F / 1.05$ equal to zero (no-arbitrage condition)

Therefore $F = 52.5$

whose purchase is financed by borrowing at the risk rate, which equals the discount rate, for the term of the forward contract. At contract expiry the loan is settled by paying a sum equal to $(F/1.05) \times 1.05$, which reduces to F. Since no arbitrage opportunity exists, the cost of creating the portfolio must be identical to the cost of the forward. As noted, the cost of the forward contract at initiation is set at zero. Equation (6.5) must therefore hold.

$$-50 + F / 1.05 = 0$$
$$F = 50 \times 1.05 = 52.50$$

(6.5)

The general relationship between the forward price and the spot prices is demonstrated in **FIGURE 6.4**. As in figure 6.3, the first step is to replicate a forward's payoff profile with a portfolio composed of R^{-T} units of the underlying asset, whose creation is funded by borrowing a sum equal to the present value of the forward price. Again, as in figure 6.3, the loan is settled on the forward expiry date by paying an amount equal to $\left(Fr^{-T}\right) \times r^T = F$.

By design, the portfolio has a payoff that is identical to the forward's: $(P_T - F)$. The cost of setting up the portfolio must therefore equal the current price of the forward: if one were cheaper than the other, a trader could make a risk-free profit by buying the cheaper instrument and short-

FIGURE **6.4** *Proof of Forward-Spot Parity*

	START DATE	EXPIRY
Buy forward contract	0	$P_T - F$
Buy R^{-T} units of the underlying asset	$-PR^{-T}$	P_T
Borrow zero present-value of forward price	Fr^{-T}	$-F$
Total	$-PR^{-T} + Fr^{-T}$	$P_T - F$

Result
Set $-PR^{-T} + Fr^{-T} = 0$
Therefore $F = P(r/R)^T$

ing the more expensive one. The current cost of the forward—its present value—is zero. The cost of constructing the duplicating portfolio must therefore be zero as well. This is expressed in equation (6.6a), which may be solved for F as (6.6b).

$$-PR^{-T} + Fr^{-T} = 0 \qquad (6.6a)$$

$$F = P(r/R)^T \qquad (6.6b)$$

The price of the forward contract is thus a function of the current underlying spot price, the risk-free or money market interest rate, the payoff, and the maturity of the contract. It can be shown that neither $F > P(r/R)^T$ nor $F < P(r/R)^T$ is possible unless arbitrage is admitted. The only possibility is (6.6b), which describes the state in which the futures price is at *fair value*.

The Basis and Implied Repo Rate

This section introduces some terms used in the futures markets. The first is *basis:* the difference between the price of a futures contract and the current underlying spot price. The size of the basis is a function of issues such as *cost of carry,* the net cost of holding the underlying asset from

CASE STUDY: *CBOT September 2003 U.S. Long Bond Futures Contract*

In theory, a futures contract represents the price for forward delivery of the underlying asset. The price of the future and that of the underlying asset should therefore converge as the contract approaches maturity. In actuality, however, this does not occur. For a bond futures contract, convergence is best viewed through the basis. This is illustrated in **FIGURE 6.5**. This shows all of the U.S. Treasury securities whose terms, at that time, make them eligible to deliver into the futures contract. The underlying whose price is used is the U.S. Treasury 6.5 percent due 2023, the cheapest-to-deliver Treasury for most of this contract's life, as demonstrated

FIGURE 6.5 *Bloomberg Screen Shows the Cheapest to Deliver for September 2003 U.S. Long Bond Futures Contract*
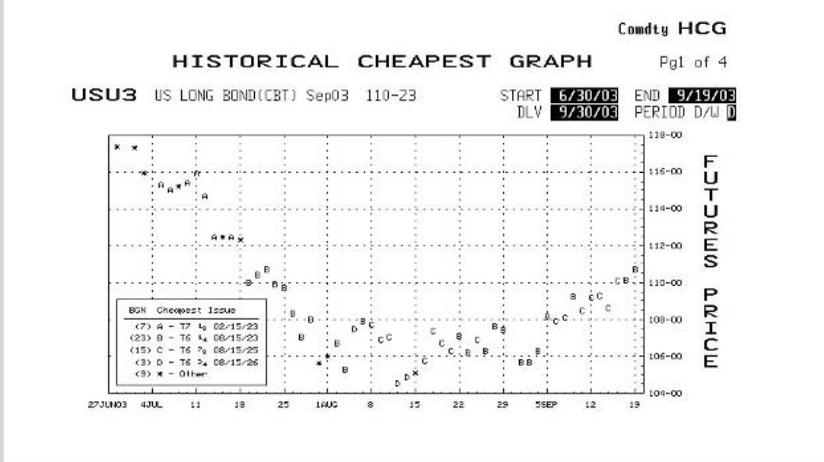
the trade date to expiry or delivery date. The basis is positive or negative depending on the type of market involved. When it is positive—that is, when $F > P$, which is common in precious metals markets—the situation is termed a *contang*. A negative basis, $P < F$, which is common with oil contracts and in foreign currency markets, is known as *backwardation*.

in **FIGURE 6.6**. Figure 6.6 demonstrates the convergence of the future and underlying asset price through the contraction of the basis, as the contract approaches expiry. Note that the implied repo rate remains fairly stable through most of the future's life, confirming the analysis suggested earlier. It does spike towards maturity, illustrating its sensitivity to very small changes in cash or futures price. The rate becomes more sensitive in the last days because there are fewer days to expiry and delivery, so small changes have larger effects.

FIGURE 6.6 *Bloomberg DLV Screen for September 2003 U.S. Long Bond Futures Contract*
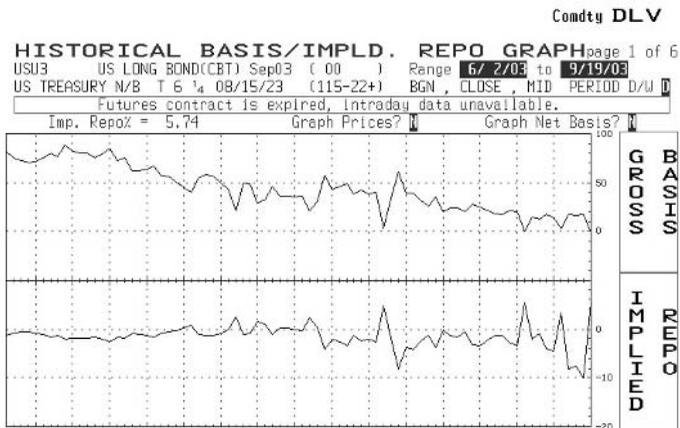
For bond futures, the *gross basis* represents the cost of carry associated with the notional bond from the present to the delivery date. Its size is given by equation (6.7).

$$Basis = P_{bond} - \left(P_{fut} \times CF\right) \qquad (6.7)$$

where

   CF = the *conversion factor* for the bond in question

The conversion factor equalizes each deliverable bond to the futures price. The bond with the lowest gross basis is known as the *cheapest to deliver.*

Generally, the basis declines over time, becoming zero on the contract's expiry date. The size of the basis, however, changes continuously, creating an uncertainty termed *basis risk.* The significance of this risk is greatest for market participants who use futures contracts to hedge positions in the underlying asset. Hedging futures and the underlying asset requires keen observation of the basis. One way to hedge a position in a futures contract is to take an opposite position in the underlying asset. This, however, entails a cost of carry, which, depending on the nature of the asset, may include storage costs, the opportunity cost of forgoing interest on the principal, the funding cost of holding the asset, and other expenses.

The futures price can be analyzed in terms of the forward-spot parity relationship and the risk-free interest rate. Say that the risk-free rate is $r - 1$. The forward-spot parity equation (repeated as (6.8a)) can be rewritten in terms of this rate as (6.8b), which must hold because of the no-arbitrage assumption.

$$F = P(r/R)^T \tag{6.8a}$$

$$r - 1 = R(F/P)^{1/T} - 1 \tag{6.8b}$$

This risk-free rate is known as the *implied repo rate,* because the rate is similar to a repurchase agreement carried out in the futures market. Generally, high implied repo rates indicate high futures prices, low rates imply low prices. The rates can be used to compare contracts that have different terms to maturity and even underlying assets. The implied repo rate for the contract is more stable than its basis, but as maturity approaches it becomes very sensitive to changes in the futures price, spot price, and (by definition) time to maturity.

Swaps

S *waps* are off-balance-sheet transactions involving two or more basic building blocks. Most swaps currently traded in the market involve combinations of cash-market rates and indexes—for example, a fixed-rate security combined with a floating-rate one, with a currency transaction perhaps thrown in. Swaps also exist, however, that have futures, forward, or option components.

The market for swaps is overseen by the International Swaps and Derivatives Association (ISDA). Swaps are among the most important and useful instruments in the debt capital markets. They are used by a wide range of institutions, including commercial banks, mortgage banks and building societies, corporations, and local governments. Demand for them has grown because the continuing volatility of interest and exchange rates has made hedging exposures to these rates ever more critical. As the market has matured, swaps have gained wide acceptance and are now regarded as plain vanilla products. Virtually all commercial and investment banks quote swap prices for their customers. Since they are over-the-counter instruments, transacted over the telephone, it is possible for banks to tailor swaps to match the precise requirements of individual customers. There is a close relationship between the bond and swap markets, and corporate finance teams and underwriting banks watch the government and the swap yield curves for opportunities to issue new debt.

This chapter discusses the uses of interest rate swaps, including as a hedging tool, from the point of view of bond-market participants. The discussion touches on pricing, valuation, and credit risk, but for complete

coverage of these topics, the reader is directed to the works listed in the References section.

Interest Rate Swaps

The market in dollar, euro, and sterling interest rate swaps is very large and very liquid. These are the most important type of swaps in terms of transaction volume. They are used to manage and hedge interest rate exposure or to speculate on the direction of interest rates.

An interest rate swap is an agreement between two counterparties to make periodic interest payments to one another during the life of the swap. These payments take place on a predetermined set of dates and are based on a *notional* principal amount. The principal is notional because it is never physically exchanged—hence the off-balance-sheet status of the transaction—but serves merely as a basis for calculating the interest payments.

In a plain vanilla, or generic, swap, one party pays a fixed rate, agreed upon when the swap is initiated, and the other party pays a floating rate, which is tied to a specified market index. The fixed-rate payer is said to be long, or to have bought, the swap. In essence, the long side of the transaction has purchased a floating-rate note and issued a fixed-coupon bond. The floating-rate payer is said to be short, or to have sold, the swap. This counterparty has, in essence, purchased a coupon bond and issued a floating-rate note.

An interest rate swap is thus an agreement between two parties to exchange a stream of cash flows that are calculated by applying different interest rates to a notional principal. For example, in a trade between Bank A and Bank B, Bank A may agree to pay fixed semiannual coupons of 10 percent on a notional principal of $1 million in return for receiving from Bank B the prevailing 6-month LIBOR rate applied to the same principal. The known cash flow is Bank A's fixed payment of $50,000 every six months to Bank B.

Interest rate swaps trade in a secondary market, where their values move in line with market interest rates, just as bonds' values do. If, for instance, a 5-year interest rate swap is transacted at a fixed rate of 5 percent and 5-year rates subsequently fall to 4.75 percent, the swap's value will decrease for the fixed-rate payer and increase for the floating-rate payer. The opposite would be true if 5-year rates moved to 5.25 percent. To understand why this is, think of fixed-rate payers as borrowers. If interest rates fall after they settle their loan terms, are they better off? No, because they are now paying above the market rate on their loan. For this reason, swap contracts decrease in value to the fixed-rate payers when rates fall. On the other hand, floating-rate payers gain from a fall in rates, because

their payments fall as well, and the value of the contract rises for them.

A bank's swaps desk has an overall net interest rate position arising from all the swaps currently on its books. This position represents an interest rate exposure at all points along the term structure out to the maturity of the longest-dated swap. At the close of business each day, all the swaps on the books are marked to market at the interest rate quote for the day.

A swap can be viewed in two ways. First, it may be seen as a strip of forward or futures contracts that mature every three or six months out to the maturity date. Second, it may be seen as a bundle of cash flows arising from the sale and purchase of cash market instruments—the preferable view in the author's opinion.

Say a bank has only two positions on its books:

❑ A long $100 million position in a 3-year floating-rate note (FRN) that pays 6-month LIBOR semiannually and is trading at par

❑ A short $100 million position in a 3-year Treasury that pays a 6 percent coupon and is also trading at par

Being short a bond is the equivalent to being a borrower of funds. Assuming that these positions are held to maturity, the resulting cash flows are those shown in **FIGURE 7.1**.

There is no net outflow or inflow at the start of these trades, because the $100 million spent on the purchase of the FRN is netted with the receipt of $100 million from the sale of the Treasury. The subsequent net cash flows over the three-year period are shown in the last column. As at the start of the trade, there is no cash inflow or outflow on maturity. The net position is exactly the same as that of a fixed-rate payer in an interest rate swap. For a floating-rate payer, the cash flow would mirror exactly that of a long position in a fixed-rate bond and a short position in an FRN. Therefore, the fixed-rate payer in a swap is said to be short in the bond market—that is, a borrower of funds—and the floating-rate payer is said to be long the bond market.

Market Terminology

Virtually all swaps are traded under the legal terms and conditions stipulated in the ISDA standard documentation. The trade date for a swap is, not surprisingly, the date on which the swap is transacted. The terms of the trade include the fixed interest rate, the maturity and notional amount of the swap, and the payment bases of both legs of the swap. Most swaps tie the floating-rate payments to LIBOR, although other reference rates are used, including the U.S. prime rate, euribor, the Treasury bill rate, and the commercial paper rate. The dates on which the floating rates for a period are determined are the *setting dates,* the first of which may also be the

FIGURE 7.1 *Cash Flows Resulting from a Long Position in a 3-Year FRN and a Short Position in a 3-Year 6 Percent Treasury*

PERIOD (6 MOS)	FRN	GILT	NET CASH FLOW
0	−£100m	+£100m	£0
1	+(LIBOR x 100)/2	−3	+(LIBOR x 100)/2 − 3.0
2	+(LIBOR x 100)/2	−3	+(LIBOR x 100)/2 − 3.0
3	+(LIBOR x 100)/2	−3	+(LIBOR x 100)/2 − 3.0
4	+(LIBOR x 100)/2	−3	+(LIBOR x 100)/2 − 3.0
5	+(LIBOR x 100)/2	−3	+(LIBOR x 100)/2 − 3.0
6	+[(LIBOR x 100)/2] + 100	−103	+(LIBOR x 100)/2 − 3.0

The LIBOR rate is the six-month rate prevailing at the time of the setting, for instance, the LIBOR rate at period 4 will be the rate actually prevailing at period 4.

trade date. As for forward-rate-agreement (FRA) and Eurocurrency deposits, the rate is fixed two business days before the interest period begins. Interest on the swap is calculated from the *effective date,* which is typically two business days after the trade date.

Although for the purposes of explaining swap structures both parties are said to pay and receive interest payments, in practice only the net difference between both payments changes hands at the end of each interest period. This makes administration easier and reduces the number of cash flows for each swap. The final payment date falls on the maturity date of the swap. Interest is calculated using equation (7.1).

$$I = M \times r \times \frac{n}{B} \tag{7.1}$$

where
    I = the payment amount
    M = the swap's notional principal
    B = the day-count base for the swap: actual/360 for dollar and euro-denominated swaps, actual/365 for sterling swaps
    r = the fixed rate in effect for the period
    n = the number of days in the period

FIGURE 7.2 *Cash Flows for a Plain Vanilla Interest Rate Swap*

(i) Cash flows for fixed-rate payer

(ii) Cash flows for floating-rate payer

(iii) Net cash flows

→ Fixed payments   ┈→ Floating payments
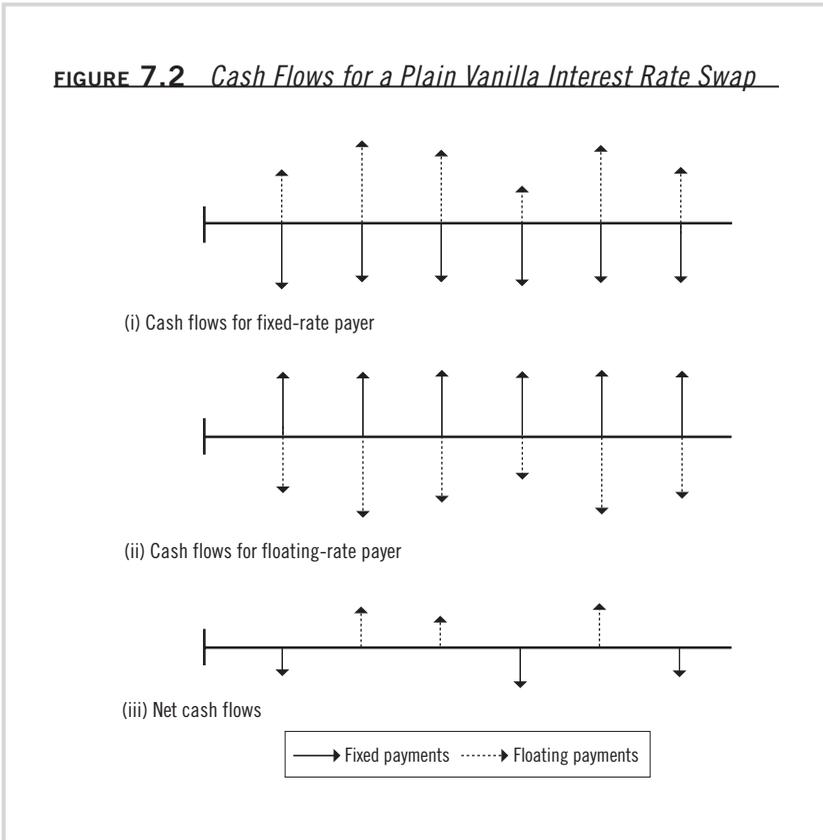
FIGURE 7.2 illustrates the cash flows from a plain vanilla interest rate swap, indicating inflows with arrows pointing up and outflows with downward-pointing ones. The net flows actually paid out are also shown.

Swap Spreads and the Swap Yield Curve

Pricing a newly transacted interest rate swap denotes calculating the swap rate that sets the net present value of the cash flows to zero. Banks quote two-way swap rates on screens or over the telephone or through dealing systems such as Reuters. Brokers also relay prices in the market. The convention is for the swap market maker to set the floating leg at LIBOR and quote the fixed rate that is payable for a particular maturity. For a 5-year swap, for example, a bank's swap desk might quote the following:

Floating-rate payer: pay 6-month LIBOR
               receive a fixed rate of 5.19 percent

Fixed-rate payer: pay a fixed rate of 5.25 percent
                      receive 6-month LIBOR

In this example, the bank is quoting an *offer* rate of 5.25 percent, which is what the fixed-rate payer will pay, and a *bid* rate of 5.19 percent, which is what the floating-rate payer will receive. The *bid-offer spread* is therefore 6 basis points. The fixed rate is always set at a spread over the government bond yield curve and is often quoted that way. Say the 5-year Treasury is trading at a yield of 4.88 percent. The 5-year swap bid and offer rates in the example are 31 basis points and 37 basis points, respectively, above this yield, and the bank's swap trader could quote the swap rates as a *swap spread:* 37–31. This means that the bank would be willing to enter into a swap in which it paid 31 basis points above the benchmark yield and received LIBOR or one in which it received 37 basis points above the yield curve and paid LIBOR.

A bank's swap screen on Bloomberg or Reuters might look something like **FIGURE 7.3**. The first column represents the length of the swap agreement, the next two are its offer and bid quotes for each maturity, and the last is the current bid spread over the government benchmark bond.

The swap spread is a function of the same factors that influence other instruments' spreads over government bonds. For swaps with durations of up to three years, other yield curves can be used in comparisons, such as the cash-market curve or a curve derived from futures prices. The spreads of longer-dated swaps are determined mainly by the credit spreads prevailing in the corporate bond fixed- and floating-rate markets. This is logical, since the swap spread essentially represents a premium compensating the investor for the greater credit risk involved in lending

FIGURE 7.3 *Swap Quotes*

1YR	4.50	4.45	+17
2YR	4.69	4.62	+25
3YR	4.88	4.80	+23
4YR	5.15	5.05	+29
5YR	5.25	5.19	+31
10YR	5.50	5.40	+35

FIGURE **7.4** *Sterling 2-Year and 10-Year Swap Spreads 1998–99*
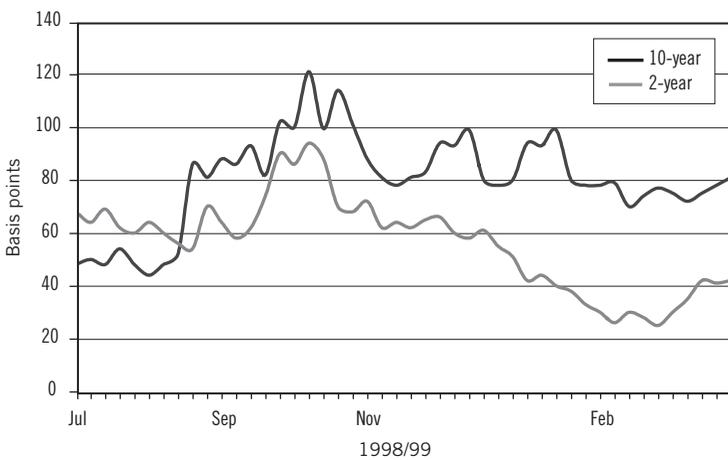
Source: Bank of England *Quarterly Bulletin*

to corporations than to the government. Day-to-day fluctuations in swap rates often result from technical factors, such as the supply of corporate bonds and the level of demand for swaps, plus the cost to swap traders of hedging their swap positions.

In summary, swap spreads over government bonds reflect the supply and demand conditions of both swaps and government bonds, as well as the market's view on the credit quality of swap counterparties. Considerable information is contained in the swap yield curve, as it is in the government bond yield curve. When the market has credit concerns—as it did in 1998, during the corrections in Asian and Latin American markets, and in September 1998, when fears arose about the Russian government's defaulting on its long-dated U.S.-dollar bonds—a "flight to quality" increases the swap spread, particularly at the longer maturities. During the second half of 1998, in reaction to bond market volatility around the world brought about by the concerns and events mentioned, the U.K. swap spread widened, as did the spread between 2- and 10-year swaps, reflecting market worries about credit and counterparty risk. Spreads narrowed in the first quarter of 1999, as credit concerns sparked by the 1998 market corrections declined. The evolution of the 2- and 10-year swap spreads is shown in **FIGURE 7.4**.

Generic Swap Valuation

Banks generally use *par* or *zero-coupon* swap pricing, which is discussed in detail in the next section. This section introduces the subject with a description of intuitive swap valuation.

Intuitive Swap Valuation

Consider a plain vanilla interest rate swap with a notional principal of M that pays n interest payments through its maturity date, T. Payments are made on dates t_i, where $i = 1, \ldots n$. The present value today of a future payment made at time t_i is denoted as $PV(0, t_i)$. If the swap rate is r, the present value of the fixed-leg payments, PV_{fixed}, is given by equation (7.2).

$$PV_{fixed} = N \sum_{i=1}^{n} PV(0,t_i) \times \left[r \times \left(\frac{t_i - t_{i-1}}{B} \right) \right] \tag{7.2}$$

where
  B = the money market day base
  $t_i - t_{i-1}$ = the number of days between the ith and the i-1th payments

The value of an existing swap's floating-leg payments on date t_1 is given by equation (7.3).

$$PV_{float} = N \times \left[rl \times \frac{t_1}{B} \right] + N - \left[N \times PV(t_1, t_n) \right] \tag{7.3}$$

where
  rl = the LIBOR rate for the next interest payment

The present value at time 0 of the floating-rate payment is given by equation (7.4).

$$PV(0,t_1) = \frac{1}{1 + rl(t_1)\left(\frac{t_1}{B} \right)} \tag{7.4}$$

For a new swap, the present value as of t_1 of the floating payments is given by equation (7.5).

$$PV_{float} = N \left[rl \times \frac{t_1}{B} + 1 \right] \times PV(0,t_1) - PV(0,t_n) \tag{7.5}$$

The swap valuation is equal to $PV_{fixed} - PV_{float}$. The swap rate quoted by a market-making bank, known as the par or zero-coupon swap rate, is the rate for which $PV_{fixed} = PV_{float}$.

Zero-Coupon Swap Valuation

As discussed above, vanilla swap rates are often quoted as a spread that is a function mainly of the credit spread required by the market over the risk-free government rate. This convention is logical, because government bonds are the principal instrument banks use to hedge their swap books. It is unwieldy, however, when applied to nonstandard tailor-made swaps, each of which has particular characteristics that call for particular spread calculations. As a result, banks use zero-coupon pricing, a standard method that can be applied to all swaps.

As explained in chapter 3, zero-coupon, or spot, rates are true interest rates for their particular terms to maturity. In zero-coupon swap pricing, a bank views every swap, even the most complex, as a series of cash flows. The zero-coupon rate for the term from the present to a cash flow's payment date can be used to derive the present value of the cash flow. The sum of these present values is the value of the swap.

Calculating the Forward Rate from Spot-Rate Discount Factors

A swap's fixed-rate payments are known in advance, so deriving their present values is a straightforward process. In contrast, the floating rates, by definition, are not known in advance, so the swap bank predicts them using the forward rates applicable at each payment date. The forward rates are those that are implied from current spot rates. These are calculated using equation (7.6).

$$rf_i = \left(\frac{df_i}{df_{i+1}} - 1 \right) N \tag{7.6}$$

where
  rf_i = the one-period forward rate starting at time i
  df_i = the discount factor for the term from the present to time i
  df_{i+1} = the discount factor for the period $i + 1$
  N = the number of times per year that coupons are paid

Although the term *zero-coupon rate* refers to the interest rate on a discount instrument that pays no coupon and has one cash flow at maturity, constructing a zero-coupon yield curve does not require a functioning zero-coupon bond market. Most financial pricing models use a combination of the following instruments to construct zero-coupon yield curves:
  ❑ Money market deposits
  ❑ Interest rate futures
  ❑ Government bonds

Frequently an overlap in the maturity period of all these instruments is used; FRA rates are usually calculated from interest rate futures, so only one or the other is needed.

Once a zero-coupon yield curve is derived, it can be used to derive the forward rates, using equation (7.6), which in turn are used to estimate the floating payments. These, together with their fixed counterparts, can then be present valued using the zero-coupon yield curve. In valuing an interest rate swap, each of the cash flows is present-valued using the zero-coupon rates, and the results are added together. The swap's present value is the difference between the present values of its fixed- and floating-rate legs.

Remember that one way to view a swap is as a long position in a fixed-coupon bond that is funded by taking out a LIBOR loan or shorting a floating-rate bond. The holder of such a position would pay a floating rate and receive the fixed rate. In the arrangement where the long position in the fixed-rate bond is funded with a floating-rate loan, the principal cash flows cancel out, assuming the bond was purchased at par, since they are equal and opposite. That leaves a collection of cash flows that mirror those of an interest rate swap paying floating and receiving fixed. Since the fixed rate on an interest rate swap is the same as the coupon (and yield) of a bond priced at par, calculating the swap rate is the same as calculating the coupon for a bond to be issued at par.

Equation (7.7), used to derive the price of a bond paying semiannual coupons, can be solved for the coupon rate r. The result is equation (7.8), which can be simplified as shown and used to derive the par yield and so the swap rate r.

$$P = \frac{r_n}{2} df_1 + \frac{r_n}{2} df_2 + \ldots + \frac{r_n}{2} df_n + M df_n \qquad (7.7)$$

where

r_n = the coupon rate on an n-period bond with n coupons
M = the maturity payment

And, since P is assumed to be par, $M = P$.

$$r_n = r \times M = \frac{2(P - M df_n)}{(df_1 + \ldots + df_n)} = \frac{M(1 - df_n)}{\left(\frac{df_1}{2} + \ldots + df_n\right)} \qquad (7.8)$$

$$r = \frac{(1 - df_n)}{\left(\frac{df_1}{2} + \ldots + df_n\right)}$$

Expression (7.8) is for bonds paying semiannual coupons. It can be generalized to apply to bonds whose coupon frequency is N, where $N = 1$ (for an annual coupon payment) to 12, and replacing 2 in the discount factor's denominator with N^1. Solving (7.8) thus modified for the nth discount factor results in equation (7.9).

$$df_n = \frac{1 - r_n \sum_{i=1}^{n-1} \frac{df_i}{N}}{1 + \frac{r_n}{N}} \tag{7.9}$$

where
N = the number of coupon payments per year

Expression (7.9) formalizes the bootstrapping process described in chapter 3. Essentially, the n-year discount factor is computed using the discount factors for years one to n-1 and the n-year swap or zero-coupon rate. Given the discount factor for any period, that period's zero-coupon, or spot, rate can be derived using (7.9) rearranged as (7.10).

$$rs_n = {}^{t_n}\!\sqrt{\frac{1}{df_n}} - 1 \tag{7.10}$$

where
rs_n = the spot rate for period n
t_n = time in period

The relationship between discount factors and the spot rates for the same periods can be used to calculate forward rates. Say the spot rate for period 1 is known. The corresponding discount rate can be derived using (7.9), which reduces to (7.11).

$$df_1 = \frac{1}{\left(1 + \frac{rs_1}{N}\right)} \tag{7.11}$$

From this, the discount rate for the next period can be computed, using the forward rate, as shown in (7.12).

$$df_2 = \frac{df_1}{\left(1 + \frac{rf_1}{N}\right)} \tag{7.12}$$

where
rf_1 = the forward rate

This can be generalized to form an expression, (7.13), that calculates the discount factor for any period, $n+1$, given the discount rate for the previous period, n, and the forward rate, rf_n, for the period n to $n+1$. Expression (7.13) can then be rearranged as (7.14), to solve for the forward rate

$$df_{n+1} = \frac{df_n}{\left(1 + \frac{rf_n}{N}\right)} \tag{7.13}$$

$$rf_n = \left(\frac{df_n}{df_{n+1}} - 1\right) N \tag{7.14}$$

The general expression for deriving an n-period discount rate at time n from the previous periods' forward rates is (7.15).

$$df_n = \frac{1}{\left(1 + \frac{rf_{n-1}}{N}\right)} \times \frac{1}{\left(1 + \frac{rf_{n-2}}{N}\right)} \times \ldots \times \frac{1}{\left(1 + \frac{rf_n}{N}\right)}$$

$$df_n = \prod_{i=0}^{n-1} \left[\frac{1}{\left(1 + \frac{rf_i}{N}\right)}\right] \tag{7.15}$$

Equations (7.8) and (7.14) can be combined to obtain (7.16) and (7.17), the general expressions for, respectively, an n-period swap rate and an n-period zero-coupon rate.

$$r_n = \frac{\displaystyle\sum_{i=1}^{n} \frac{rf_{i-1}df_i}{N}}{\displaystyle\sum_{i=1}^{n} \frac{df_i}{F}} \tag{7.16}$$

where
N = the frequency of coupon payments

$$1 + rs_n = \sqrt[t_n]{\prod_{i=0}^{n-1} \left(1 + \frac{rf_i}{N}\right)} \tag{7.17}$$

Equation (7.16) captures the insight that an interest rate swap can be considered as a strip of futures. Since this strip covers the same period as the swap, it makes sense that, as (7.16) states, the swap rate can be com-

puted as the average of the forward rates from rf_0 to rf_n, weighted according to the discount factor for each period.

   Note that although swap rates are derived from forward rates, a swap's interest payments are paid in the normal way, at the end of an interest period, while FRA payments are made at the beginning of the period and must be discounted.

   Equation (7.17) states that the zero-coupon rate is the geometric average of one plus the forward rates. The *n*-period forward rate is obtained using the discount factors for periods *n* and *n*-1. The discount factor for the complete period is obtained by multiplying the individual discount factors together. Exactly the same result would be obtained using the zero-coupon interest rate for the whole period to derive the discount factor.[2]

The Key Principles of an Interest Rate Swap

As noted earlier, *pricing* a newly transacted interest rate swap denotes calculating the swap rate that sets the net present value of the cash flows to zero. *Valuation* signifies the process of calculating the net present value of an *existing* swap by setting its fixed rate at the current market rate. Consider a plain vanilla interest rate swap with the following terms:

Nominal principal	$10,000,000
Day count fixed	Actual/360
Day count floating	Actual/360
Payment frequency fixed	Annual
Payment frequency floating	Annual
Trade date	January 31, 2000
Effective date	February 2, 2000
Maturity date	February 2, 2005
Term	Five years

   Although in practice the fixed payments would differ slightly from year to year, to simplify the pricing, assume that they are identical. Also assume that the relevant set of zero-coupon yields has been derived, as shown in the second column of **FIGURE 7.5**. These rates are used to calculate the discount factors in the third column, which are then plugged into equation (7.14) to derive the forward rates in column four. These forward rates are used to predict what the floating-rate payments will be at each interest period. Both fixed-rate and floating-rate payments are then present-valued at the appropriate zero-coupon rates their present values netted together. The Excel formulae behind figure 7.5 are shown in **FIGURE 7.6**, on pages 124–125. The fixed rate for the swap is calculated as follows, using equation (7.8):

FIGURE 7.5 *Pricing a Plain Vanilla Interest Rate Swap*

PERIOD	ZERO-COUPON RATE %	DISCOUNT FACTOR	FORWARD RATE %
1	5.5	0.947867298	5.5
2	6	0.88999644	6.502369605
3	6.25	0.833706493	6.751770257
4	6.5	0.777323091	7.253534951
5	7	0.712986179	9.023584719
		4.161879501	

$$r = \frac{1 - 0.71298618}{4.16187950} = 6.8963 \text{ percent}$$

It is not surprising that the net present value is zero. The zero-coupon curve is used to derive the discount factors that are then used to derive the forward rates that are used to determine the swap rate. As with any financial instrument, the fair value is its break-even price or hedge cost. The bank that is pricing this swap could hedge it with a series of FRAs transacted at the forward rates shown. This method is used to price any interest rate swap, even exotic ones.

Valuation Using the Final Maturity Discount Factor

The floating-leg payments of an interest rate swap can be valued using just the discount factor for the final maturity period and the notional principal. This short-cut method is based on the fact that the value of the floating-leg interest payments is conceptually the same as that of a strategy that consists of exchanging the notional principal at the beginning and end of the swap and investing it at a floating rate in between. In both cases, the net result is a collection of floating-rate interest payments. The principal plus the payments from investing it for the term of the swap must discount to the value of the principal at the beginning of the swap, and the appropriate discount value for this is the final discount rate.

To understand this principal, consider figure 7.5, which shows the present value of both legs of the 5-year swap to be $2,870,137. The same result is obtained by using the 5-year discount factor, as shown in (7.18).

FIXED PAYMENT	FLOATING PAYMENT	PV FIXED PAYMENT	PV FLOATING PAYMENT
689,625	550,000	653,672.99	521,327.01
689,625	650,236.96	613,763.79	578,708.58
689,625	675,177.03	574,944.84	562,899.47
689,625	725,353.50	536,061.44	563,834.02
689,625	902,358.47	491,693.09	643,369.12
		2,870,137	2,870,137

$$PV_{floating} = (10,000,000 \times 1) - (10,000,000 \times 0.71298618) = 2,870,137$$

(7.18)

The first term in (7.18) represents the notional principal multiplied by the discount factor 1. This reflects the fact that the present value of an amount received immediately is the amount itself.

Non–Plain Vanilla Interest Rate Swaps

The discussion so far has involved plain vanilla swaps. These have been shown to have the following characteristics:

❑ One leg pays a fixed rate of interest; the other pays a floating rate, usually linked to a standard index such as LIBOR
❑ The fixed rate is fixed for the entire life of the swap
❑ The floating rate is set before the start of each payment period and paid in arrears
❑ Both legs have the same payment frequency (quarterly, semiannual, annual)
❑ The maturity is whole years, up to thirty
❑ The notional principal remains constant during the life of the swap

Each of these characteristics can be altered to meet particular customer demands. The resulting swaps are non–plain vanilla, or nongeneric.

A wide variety of swap contracts have been traded in the market. Although six-month LIBOR is the most common reference rate for the

floating-leg of a swap making semiannual payments, for example, three-month LIBOR also has been used, as well as the prime rate (for dollar swaps), the one-month commercial paper rate, the Treasury bill rate, the municipal bond rate (again, for dollar swaps), and others.

Swaps may also be *extendable* or *putable.* In an extendable swap, one of the parties has the right, but not the obligation, to extend the life of the swap beyond the fixed maturity date. In a putable swap, one party has the right to terminate the swap ahead of the specified maturity date. The fixed rate would be adjusted to reflect the cost of the implicit option. For example, if the fixed payer has the right to extend the swap, the fixed rate would be higher than for a plain vanilla swap with similar terms.

A *forward-start* swap's effective date is a considerable period—say, six months—after the trade date, rather than the usual one or two days. A forward start is used when one counterparty, perhaps foreseeing a rise in interest rates, wants to fix the cost of a future hedge or a borrowing now. The swap rate is calculated in the same way as for a vanilla swap.

The floating leg of a *margin* swap pays LIBOR plus or minus a specified number of basis points. The swap's fixed-rate quote is adjusted to allow for this margin. Say a bank finances its fixed-rate lending by borrowing at 25 basis points over LIBOR. It may wish to receive LIBOR plus 25 bps in a swap so that its cash flows match exactly. If the swap rate for the appropriate maturity is 6 percent, the margin swap's fixed leg would be fixed at around 6.25 percent (the margins on the two legs may differ if their day-count conventions or payment frequencies are different). A floating-rate margin might also be dictated by the credit quality of the counterparty. A highly rated counterparty, for example, might pay slightly below LIBOR.

An *off-market* swap is one whose fixed rate is different from the market swap rate. To compensate for this difference, one counterparty pays the other a sum of money. An off-market rate may be required for a particular hedge, or by a bond issuer that wants to cover the issue costs as well as hedge the loan.

In a *basis* swap both legs pay floating rates, but they are linked to different money market indexes. One is normally LIBOR, while the other might be the certificate-of-deposit or the commercial-paper rate. A U.S. bank that had lent funds at prime and financed its loans at LIBOR might hedge the *basis risk* thus created with a swap in which it paid prime and received LIBOR. Both legs of a basis swap may be linked to LIBOR rates, but for different maturities. In such a swap, the payment frequencies of the two legs also differ. One counterparty, for instance, might pay 3-month LIBOR quarterly, while the other pays 6-month LIBOR semiannually.

Note that this situation exposes one party to greater risk that the counterparty will default on payments. For instance, a party paying monthly and receiving semiannual cash flows will have made five interest payments before receiving one in return.

A *constant maturity* swap, or CMS, is a basis swap in which one leg is reset periodically not to LIBOR or some other money market rate but to a long-term rate, such as the current 5-year swap rate or 5-year government bond rate. For example, the counterparties to a CMS might exchange 6-month LIBOR for the 10-year Treasury rate in effect on the reset date. In the U.S. market, a swap one of whose legs is reset to a government bond is referred to as a *constant maturity Treasury,* or CMT, swap. The other leg is usually tied to LIBOR, but may be fixed or use a different long-term rate as its reference.

A *differential* swap is a basis swap in which one of the legs is calculated in a different currency. Typically, this leg is linked to a reference index rate for another currency but is denominated in the domestic currency. For example, one party might pay 6-month sterling LIBOR, in sterling, on a notional principal of $10 million and receive euro-LIBOR minus a margin, in sterling, on the same notional principal. Differential swaps are not very common and are the most difficult for a bank to hedge.

In an *accreting,* or *step-up,* swap, the notional principal increases over the life of the swap; in an *amortizing* swap, the principle decreases. Swaps whose notional principal fluctuates—increasing one year and decreasing the next, for example—are known as *roller coasters.*

An accreting swap might be used by an institution trying to hedge a funding liability expected to grow. An amortizing swap might be employed to hedge an amortizing loan or a bond with a sinking fund feature. It frequently has a forward-start feature, synchronized with the cash flows payable on the loan. In principle, amortizing and accreting swaps are priced and valued in the same way as plain vanilla ones.

In a *LIBOR-in-arrears,* or *back-set,* swap, the floating rate is set just before the end of the payment period, rather than just before the start. Such a swap would be attractive to a counterparty with a view on interest rates that differed from the market consensus. In an upward-sloping yield curve, for instance, forward rates are higher than current market ones, and this is reflected in swap pricing. Floating-rate payers believing that interest rates will rise more slowly than forward rates (and the market) suggest might enter into a LIBOR-in-arrears swap, which would be priced higher than a conventional one.

Swaptions

A bank or corporation may buy or sell an option on a swap, known as a *swaption.* The buyer of a swaption has the right, but not the obligation, to transact an interest rate swap during the life of the option. An option on a swap where the buyer is the fixed-rate payer is termed a *call* swaption; one where the buyer becomes the floating-rate payer is a *put* swaption. The writer of the swaption becomes the buyer's counterparty in underlying the transaction.

Swaptions are similar to forward-start swaps, except that the buyer can choose *not* to commence payments on the effective date. A bank may purchase a call swaption if it expects interest rates to rise; it will exercise only if rates do indeed rise. A company may use swaptions to hedge future interest rate exposures. Say it plans to take out a five-year bank loan in three months. This transaction will make the company liable for floating-rate interest payments, which are a mismatch for the fixed-rate income it earns on the long-term mortgages on its books. To correct this mismatch, the company intends to transact a swap in which it receives LIBOR and pays fixed after getting the loan. To hedge against an unforeseen rise in interest rates in the meantime, which would increase the swap rate it has to pay, it may choose to purchase an option, expiring in three months, on a swap in which it pays a fixed rate of, say, 10 percent.

If the 5-year swap rate is above 10 percent in three months, after the company has taken out its loan, it will exercise the swaption. If the rate is below 10 percent, however, it will transact the swap in the normal way, and the swaption will expire worthless. The swaption thus enables a company to hedge against unfavorable movements in interest rates but also to gain from favorable ones. There is, of course, a cost associated with this benefit: the swaption premium.

Valuation

Since a floating-rate bond is valued on its principal value at the start of a swap, a swaption may be viewed as the value on a fixed-rate bond, with a strike price that is equal to the face value of the floating-rate bond.

Swaptions are typically priced using the Black-Scholes or the Black pricing model. With a European swaption, the appropriate swap rate on the expiry date is assumed to be lognormal. The swaption payoff is given by equation (7.19).

$$\text{Payoff} = \frac{M}{F}\max\left(r - r_n, 0\right)$$

(7.19)

where
  r_n = the strike swap rate
  r = the actual swap rate at expiry
  n = the swap's term
  M = the notional principal
  F = the swap payment frequency

The Black model uses equation (7.20) to derive the price of an interest rate option.

$$c = P(0,T)\big[f_0 N(d_1) - XN(d_2)\big] \tag{7.20}$$

where
  c = the price of the call option
  $P(t, T)$ = the price at time t of a zero-coupon bond maturing at time T
  f = the forward price of the underlying asset with maturity T
  f_t = the forward price at time t
  X = the strike price of the option
  N = normal distribution

and where
$$d_1 = \frac{\ln(f_0 / X) + \sigma^2 T / 2}{\sigma\sqrt{T}}$$

$$d_2 = \frac{\ln(f_0 / X) - \sigma^2 T / 2}{\sigma\sqrt{T}} = d_1 - \sigma\sqrt{T}$$

and
  σ = the volatility of f

Equation (7.20) can be combined with (7.19) to form (7.21), which derives the value of a swap cash flow received at time t_i.

$$\frac{M}{F} P(0,t_i)\big[f_0 N(d_1) - r_n N(d_2)\big] \tag{7.21}$$

where
  f_0 = the forward swap rate at time 0
  r_i = the continuously compounded zero-coupon interest rate for an instrument with maturity t_i

From (7.21), equation (7.22) can be constructed to derive the total value of the swaption.

$$PV = \sum_{i=1}^{Fn} \frac{M}{F} P(0, t_i) \left[f_0 N(d_1) - r_n N(d_2) \right] \tag{7.22}$$

Interest Rate Swap Applications

This section discusses how swaps are used to hedge bond instruments and how swap books are themselves hedged.

Corporate and Investor Applications

As noted earlier, swaps can be tailored to suit a user's requirements. For example, swaps' payment dates, payment frequencies, and LIBOR margins are often specified to match customers' underlying exposures. Because the market is so large, liquid, and competitive, banks are willing to structure swaps to meet the requirements of virtually all customers, although smaller customers may have difficulty obtaining competitive quotes for notional values below $10 million.

FIGURE 7.6 *Excel Formulae for Figure 7.5*

CELL	C	D	E	F
21			10000000	
22				
23	PERIOD	ZERO-COUPON RATE %	DISCOUNT FACTOR	FORWARD RATE %
24	1	5.5	0.947867298	5.5
25	2	6	0.88999644	"((E24/E25)–1)*100
26	3	6.25	0.833706493	"((E25/E26)–1)*100
27	4	6.5	0.777323091	"((E26/E27)–1)*100
28	5	7	0.712986179	"((E27/E28)–1)*100
			"SUM(E24:E28)	

In terms of their applications, swaps generally fall into two categories: those linked to assets and those linked to liabilities. An asset-linked swap is used by investors to change the characteristics of the income stream generated by an asset, such as a bond. Liability-linked swaps are used by borrowers to change the pattern of their cash flows. Of course, the assignment of a swap to one of these categories depends on the point of view of the assignor. What is an asset-linked swap for one party is a liability-linked hedge for the counterparty, except in the case of market-making banks that make two-way quotes on the instruments.

One straightforward application of an interest rate swap is to convert a floating-rate liability into a fixed-rate one, usually in an effort to remove exposure to upward moves in interest rates. Say a company has borrowed money at a floating rate of 100 basis points over 6-month LIBOR. Fearing that interest rates will rise in the three years remaining on the loan, it enters into a 3-year semiannual interest rate swap with a bank, depicted in **FIGURE 7.7** on the following page, in which it pays a fixed rate of 6.75 percent and receives 6-month LIBOR. This fixes the company's borrowing costs for three years at 6.75 percent plus 100 basis points, or 7.75 percent, for an effective annual rate of 7.99 percent.

G	H	I	J
FIXED PAYMENT	**FLOATING PAYMENT**	**PV FIXED PAYMENT**	**PV FLOATING PAYMENT**
689,625	"(F24*10000000)/100	"G24/1.055	"H24/(1.055)
689,625	"(F25*10000000)/100	"G24/(1.06)^2	"H25/(1.06)^2
689,625	"(F26*10000000)/100	"G24/(1.0625)^3	"H26/(1.0625^3)
689,625	"(F27*10000000)/100	"G24/(1.065)^4	"H27/(1.065)^4
689,625	"(F28*10000000)/100	"G24/(1.07)^5	"H28/(1.07)^5
		2,870,137	2,870,137

FIGURE 7.7 *Transforming a Liability from Fixed Rate to Floating and Back to Fixed*

Two-Swap Structure
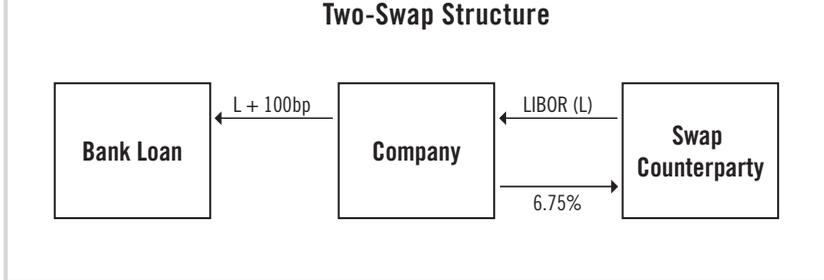
Say a corporation borrows funds for five years at a rate of 6.25 percent. Shortly after taking out the loan, it enters into a swap in which it pays a floating rate of LIBOR and receives 5.85 fixed (see **FIGURE 7.8**). Its net borrowing cost is thus LIBOR plus 40 basis points (6.25 minus 5.85). After one year, interest rates have fallen, and the 4-year swap rate is quoted at 4.90–84 percent—that is, banks are willing to receive 4.90 or pay 4.84 fixed. The company decides to take advantage of the lower interest rates by switching back to a fixed-rate liability. To this end, it enters into a second swap in which it pays 4.90 percent fixed and receives LIBOR. Its borrowing cost is now 5.30 percent (4.90 plus 40 basis points), or 95 basis points—the difference between the two swap rates—below its original borrowing cost.

Investors might use asset-linked swaps if they want fixed-rate securities, and the only assets available with the required credit quality and terms pay floating rates. For instance, a pension fund may have invested in 2-year floating-rate gilts, an asset of the highest quality, that pay 5.5 basis points below the London interbank bid rate, or LIBID (the interest rate at which a bank in the City of London is willing to borrow short term from another City bank). As it is expecting interest rates to fall, however, it prefers to receive a fixed rate. Accordingly, it arranges a tailor-made swap in which it pays LIBID and receives a fixed rate of 5.50 percent. By entering into this swap, the pension fund creates a structure, shown on page 128 in **FIGURE 7.9**, that generates a fixed-income stream of 5.375 percent.

FIGURE **7.8** *Swap Transforming a Floating-Rate Asset to a Fixed-Rate One*
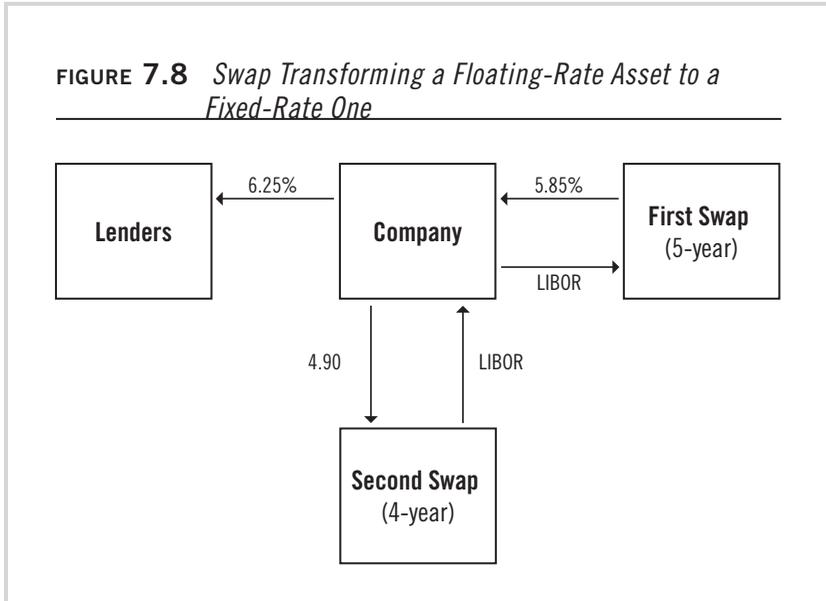
Hedging Bond Instruments Using Interest Rate Swaps

Bond traders wishing to hedge the interest rate risk of their bond positions have several tools to choose from, including other bonds, bond futures, and bond options, as well as swaps. Swaps, however, are particularly efficient hedging instruments, because they display positive convexity. As explained in chapter 2, this means that they increase in value when interest rates fall more than they lose when rates rise by a similar amount—just as plain vanilla bonds do.

The primary risk measure required when using a swap to hedge is the present value of a basis point. PVBP, known in the U.S. market as the dollar value of a basis point, or DVBP, indicates how much a swap's value will move for each basis point change in interest rates and is employed to calculate the hedge ratio. PVBP is derived using equation (7.23).

$$PVBP = \frac{dS}{dr}$$

(7.23)

where

dS = change in swap value
dr = change in market interest rate, in basis points

It was suggested earlier that a swap be seen as a bundle of cash flows arising from the sale and purchase of two cash-market instruments: a

---

---

fixed-rate bond with a coupon equal to the swap rate and a floating-rate bond with the same maturity and paying the same rate as the floating leg of the swap. Considering a swap in this way, equation (7.23) can be rewritten as (7.24).

$$PVBP = \frac{d\text{Fixed bond}}{dr} - \frac{d\text{Floating bond}}{dr} \tag{7.24}$$

Equation (7.24) essentially states that PVBP of the swap equals the difference between the PVBPs of the fixed- and floating-rate bonds. This value is usually calculated for a notional principal of \$1 million, based on the duration and modified duration of the bonds (defined in chapter 2) and assuming a parallel shift in the yield curve.

FIGURE 7.10 illustrates how the PVBP of a 5-year swap may be calculated using the relationships expressed in (7.23) and (7.24). The two derivations are shown in equations (7.25) and (7.26), respectively. (Bonds' PVBPs can be calculated using Bloomberg's YA screen or Microsoft Excel's MDURATION function.)

$$PVBP_{swap} = \frac{dS}{dr} = \frac{4264 - (-4236)}{20} = 425 \tag{7.25}$$

$$PVBP_{swap} = PVBP_{fixed} - PVBP_{floating}$$

$$= \frac{1004940 - 995171}{20} - \frac{1000640 - 999371}{20}$$

$$= 488.45\text{-}63.45$$

$$= 425.00 \tag{7.26}$$

FIGURE 7.10 *PVBPs of a 5-Year Swap and a Fixed-Rate Bond with the Same Maturity Date*

INTEREST RATE SWAP			
Term to maturity	5 years		
Fixed leg	6.50%		
Basis	Semi-annual, act/365		
Floating leg	6-month LIBOR		
Basis	Semi-annual, act/365		
Nominal amount	$1,000,000		

Present value $

	Rate change – 10 bps	0 bps	Rate change + 10 bps
Fixed-coupon bond	1,004,940	1,000,000	995,171
Floating-rate bond	1,000,640	1,000,000	999,371
Swap	4,264	0	4,236

Note that the swap PVBP, $425, is lower than that of the 5-year fixed-coupon bond, which is $488.45. This is because the floating-rate bond PVBP reduces the risk exposure of the swap as a whole by $63.45. As a rough rule of thumb, the PVBP of a swap is approximately the same as that of a fixed-rate bond whose term runs from the swap's next coupon reset date through the swap's termination date. Thus, a 10-year swap making semiannual payments has a PVBP close to that of a 9.5-year fixed-rate bond, and a swap with 5.5 years to maturity has a PVBP similar to that of a 5-year bond.

One corollary of the relationship expressed in (7.25) is that swaps' PVBPs behave differently from those of bonds. Immediately preceding a reset date, when the PVBP of the floating-rate bond corresponding to the swap's floating leg is essentially nil, a swap's PVBP is almost identical to that of the fixed-rate bond maturing on the same day as the swap. For example, if it's a five-year swap, and it's just before the second semiannual payment, the PVBP will be similar to that of a four-year bond. Immediately after the reset date the swap's PVBP will be nearly identical to that of

a bond maturing at the next reset date. Therefore, from a point just before the reset to one just after, the swap's PVBP will decrease by the amount of the floating-rate PVBP. In between reset dates, the swap's PVBP is quite stable, since the effects of changes in the fixed- and floating-rate PVBPs cancel each other out. In contrast, a fixed-rate bond's PVBP decreases steadily over time, assuming that no sudden large-scale yield movements occur. The evolution of the swap and bond PVBPs is illustrated in **FIGURE 7.11**. Note that the graph does not reflect a slight anomaly that occurs in the swap's PVBP, which actually increases by a small amount between reset dates because the floating-rate bond's PVBP decreases at a slightly faster rate than that of the fixed-rate bond.

Hedging a bond with an interest rate swap is conceptually similar to hedging it with another bond or with bond futures. Hedging a long position in a vanilla bond requires a long position in the swap—that is, taking the side that pays fixed and receives floating. Hedging a short position in a bond requires a short swap position, matching the fixed swap income with the pay-fixed liability of the short bond position.

The swap's value will change by approximately the same amount, but in the opposite direction, as the bond's value. The match will not be exact. It is very difficult to establish a precise hedge for a number of reasons, including differences in day count and in maturity, and basis risk. To minimize the mismatch, the swap's maturity should be as close as possible to

FIGURE **7.11** *PVBP of a 5-Year Swap and Fixed-Rate Bond Maturity Period*
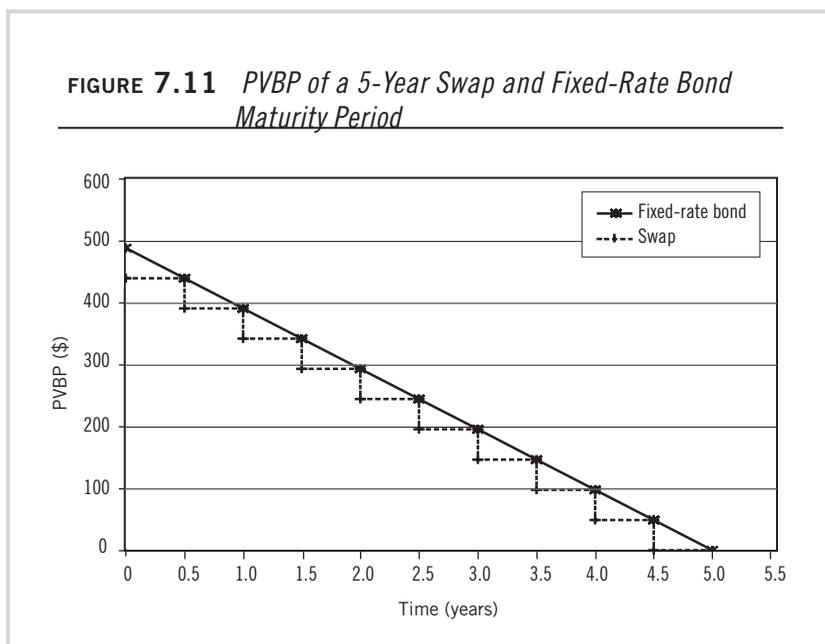
the bond's. Since swaps are OTC contracts, it should be possible to match interest-payment dates as well as maturity dates.

The correct nominal amount of the swap is established using the PVBP hedge ratio, shown as equation (7.27). Though the market still uses this method, its assumption of parallel yield-curve shifts can lead to significant hedging error.

$$\text{Hedge ratio} = \frac{PVBP_{bond}}{PVBP_{swap}}$$

(7.27)

<u>*Chapter Notes*</u>

1. The expression also assumes an actual/365 day count. If any other day-count convention is used, the $1/N$ factor must be replaced by a fraction whose numerator is the actual number of days and whose denominator is the appropriate year base.

2. Zero-coupon and forward rates are also related in another way. If the zero-coupon rate rs_n and the forward rate rf_i are transformed to their continuously compounded equivalent rates, $\ln(1 + rs_n)$ and $\ln(1 + rf_i)$, the result is the following expression, which derives the continuously compounded zero-coupon rate as the simple average of the continuously compounded forward rates:

$$rs_n = \frac{1}{t_n} \sum_{i=0}^{n-1} \frac{rf_i}{F}$$

This page is intentionally blank

Options

Options were originally written on commodities such as wheat and sugar. Today investors can buy or sell options on a wide range of underlying instruments in addition to commodities, including financial products such as foreign exchange rates, bonds, equities, and derivatives such as futures, swaps, and equity indexes. Contracts on commodities are known as *options on physicals;* those on financial assets are known as *financial options.*

The spectrum of trading combinations and structured products involving options is constrained only by imagination and customer requirements. Virtually all participants in capital markets have some requirement that may be met by using options. Market makers, for instance, use options for speculation and arbitrage, to generate returns. One of the most important uses of options, however, is as hedging tools.

Options are unique among hedging instruments in enabling banks and corporations to profit from upside market moves while covering their risk exposures. The contracts also have special characteristics that set them apart from other classes of derivatives. Because they confer the right to conduct a transaction without imposing an obligation to do so, they need be exercised only if the protection they offer is required. Options thus function more like insurance policies than like pure hedging instruments. The option price is in effect an insurance premium paid for peace of mind.

A number of specialized texts are devoted to options. This chapter introduces the basics, including the complex topic of pricing. Chapter

9 looks at the main sensitivity measures employed in running an option book and the uses to which these instruments may be put. For further reading, see the articles and books listed in the References section.

Option Basics

An option is a contract between a buyer and a seller in which the buyer has the right, but not the obligation, to purchase (in the case of a call option) or sell (in the case of a put option) a specified underlying asset at a specified price during or at the end of a specified period. The option seller, or writer, grants this right in return for the option price, or premium. The option buyer is long the contract; the seller is short.

An option's payoff profile is unlike that of any other instrument. Compare, for instance, the profiles of a vanilla call option and of a vanilla bond futures contract. Traders who buy one lot of the bond futures at 114 and hold it for a month before selling it realize a profit if the contract's price at the end of the month is above 114 and a loss if it is below 114. The amount of the gain or loss is $1,000 for each point above or below 114. The same applies in reverse to those with short positions in the futures contract. The futures' payoff profile is thus *linear,* for both the long and the short position. This is illustrated in **FIGURE 8.1**.

Other derivatives, such as forward-rate agreements and swaps, have similar profiles, as, of course, do cash instruments such as bonds and

FIGURE 8.1 *Payoff Profile for a Bond Futures Contract*
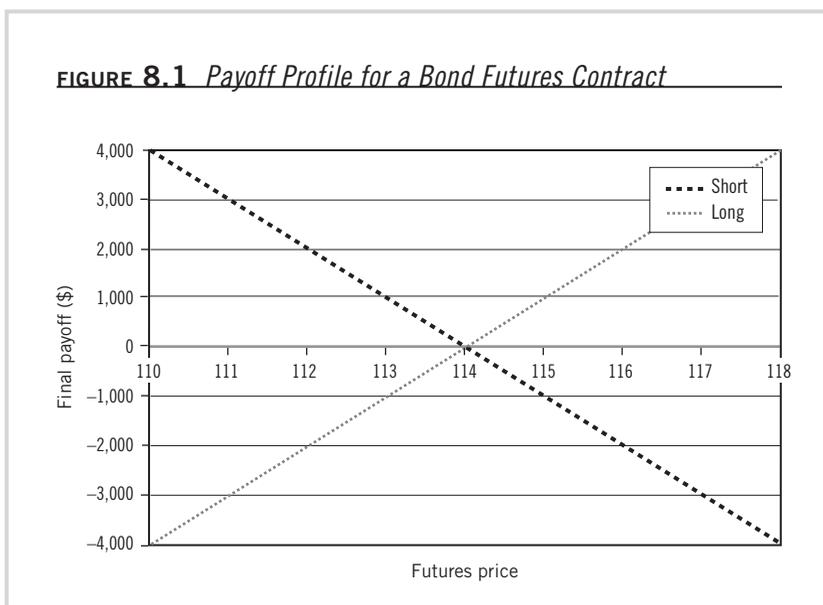
stocks. Options break the pattern. Because these contracts confer a right but impose no obligation on their holders and impose an obligation but confer no right on their sellers, the payoff profiles for the two parties are different. If, instead of the futures contract itself, the traders in the previous example take long and short positions in a call option on the contract at a strike price of 114, their payoff profiles will be those shown in **FIGURE 8.2**.

Should the price of the futures contract rise above 114 during the life of the option, the traders long the call will exercise their right to buy the future. Should the price of the future never rise above 114, these traders will not exercise, and the option will eventually expire worthless and they will suffer a loss in the amount of the premium they paid. In fact, they don't realized a profit until the contract rises above 114 by the amount of the premium. In this respect it is exactly like an equity or bond warrant. The option sellers have a very different payout profile. If the price of the future rises above 114 and the option is exercised, they bear a loss equal to the profit the buyers make. If the option expires without being exercised, the sellers keep the premium income.

This example illustrates that the holders of long and short positions in options, unlike holders of other financial instruments, have asymmetrical payoff profiles. Call option buyers benefit if the price of the underlying asset rises above the strike by at least the amount of the premium but lose

FIGURE 8.2 *Payoff Profiles for Long and Short Positions in a Call Option Contract*
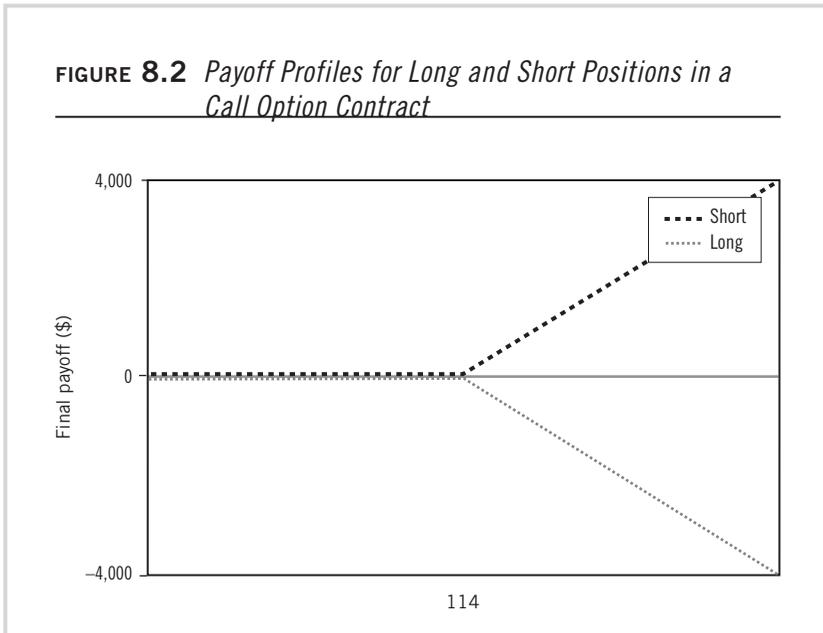
only what they paid for the option if it fails to do so. The option sellers suffer a loss if the price of the underlying asset rises above the strike by more than the premium amount but realize only the funds received for writing the option if it fails to do so.

Traders who wish to benefit from a fall in the market level but don't want to short the market might buy put options. Put options have the same asymmetrical payoff profiles for buyers and sellers as call options, but in the opposite direction. Put buyers profit if the market price of the underlying asset falls below the strike but lose only the premium they paid if the price remains above the strike. Put writers do not profit from moves in the underlying, whatever direction these moves take, and lose if the market falls below the strike by more than the premium amount. The premium they earn on the option sale is their compensation for taking on this risk.

Terminology

When an option is exercised, the option writer will deliver the asset or the cash value of the profit to the buyer. An option trader can take any of the following four positions:

1 Long a call
2 Long a put
3 Short a call
4 Short a put

FIGURE 8.3 shows the payoff profiles for these positions.

The strike price is the price of the underlying asset at which the option is exercised. For example, a call option to buy ordinary shares of a listed company might have a strike price of $10, meaning that the underlying stock can be bought at $10 a share by exercising the contract. Options that
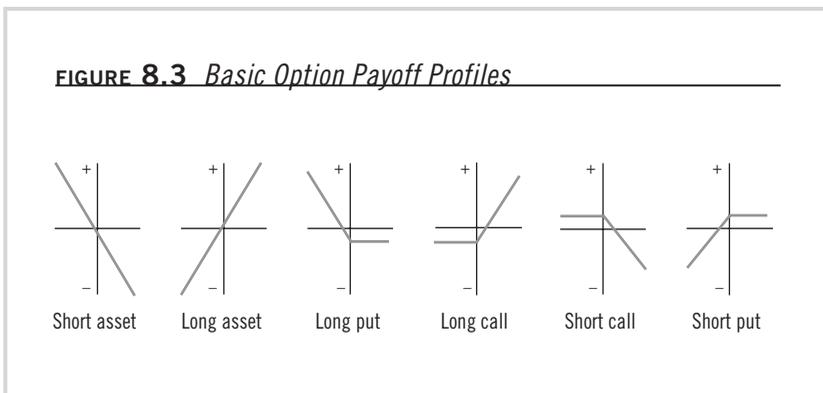
FIGURE **8.3** *Basic Option Payoff Profiles*

Short asset   Long asset   Long put   Long call   Short call   Short put

can be exercised anytime from the day they are struck up to and including the expiry date are called American options. Those that can be exercised only on the expiry date are known as European options. The terms have no geographic relevance; both styles of trade can be used in any market. Bermudan options, which can be exercised on any of a set of specified dates, also exist.[1] It is very rare for an American option to be exercised ahead of its expiry date, so this distinction has little impact in practice. Pricing models developed for European options, however, must be modified to handle American ones.

The option premium has two constituents: intrinsic value and time value. Intrinsic value equals the difference between the strike price and the underlying asset's current price. It is what the option's holders would realize if they were to exercise it immediately. Say a call option on a bond futures contract has a strike price of $100 and the contract is trading at $105. A holder who exercises the option, buying the futures at $100 and selling the contract immediately at $105, earns a profit of $5; that is the option's intrinsic value. A put option's intrinsic value is the amount by which the current underlying asset's price is below the strike. Because an option holder will exercise it only if there is a benefit to so doing, the intrinsic value will never be less than zero. Thus, if the bond future in the example were trading at $95, the intrinsic value of the call option would be zero, not –5.

An option that has intrinsic value is in the money. One with no intrinsic value is out of the money. An option whose strike price is equal to the underlying's current price is at the money. This term is normally used only when the option is first traded.

The time value of an option is the amount by which the option value exceeds the intrinsic value. Because of the risk they are taking on, illustrated in the payoff profiles above, option writers almost always demand premiums that are higher than the contracts' intrinsic value. The value of an option that is out of the money is composed entirely of time value. Time value reflects the potential for an option to move into, or more deeply into, the money before expiry. It diminishes up to the option's expiry date, when it becomes zero. The price of an option on expiry is composed solely of intrinsic value. **FIGURE 8.4** lists basic option market terminology.

Option Instruments

Options are traded both on recognized exchanges and over the counter (OTC). Exchange-traded options are standardized plain vanilla contracts; OTC options can take on virtually any form. Options traded on an exchange are often written on futures contracts. For example, a gilt option

FIGURE **8.4** *Basic Option Terminology*

Call	The right to buy the underlying asset
Put	The right to sell the underlying asset
Buyer	The person who has purchased the option and has the right to exercise it if she wishes
Writer	The person who has sold the option and has the obligation to perform if the option is exercised
Strike price	The price at which the option may be exercised, also known as the *exercise price*
Expiry date	The last date on which the option can be exercised, also known as the maturity date
American	The style of option; an American option can be exercised at any time up to the expiry date
European	An option which may be exercised on the maturity date only, and not before
Premium	The price of the option, paid by the buyer to the seller
Intrinsic value	The value of the option if it was exercised today, which is the difference between the strike price and the underlying asset price
Time value	The difference between the current price of the option and its intrinsic value
In-the-money	The term for an option that has intrinsic value
At-the-money	An option for which the strike price is identical to the underlying asset price
Out-of-the-money	An option that has no intrinsic value

on the London International Financial Futures and Options Exchange (LIFFE) is written on the exchange's gilt futures contract. Exercising a futures option results in a long position in a futures contract being *assigned* to the option holder and a short position in the future being assigned to the option writer. Exchange-traded options on U.S. Treasuries are quoted in option ticks.

Both OTC and exchange-traded options can be either American or European. Exchange-traded options are available on the following instruments:

❑ ***Shares of common stock.*** Major exchanges, including the New York Stock Exchange (NYSE), LIFFE, Eurex, the Chicago Board Options Exchange (CBOE), and the Singapore International Monetary Exchange (SIMEX), trade options on stock shares.

❑ ***Futures.*** Most exchanges trade options on the futures contracts that they trade. These options expire one or two days before the underlying futures do. Some, such as those traded on the Philadelphia Currency Options Exchange, allow cash settlement. This means that when the holders of a futures call exercise it, they are assigned both a long position in the future and the cash value of the difference between the strike price and the futures price.

❑ ***Stock indexes.*** Equity index options, such as the contracts on the Standard & Poor's 500 Index traded on the CBOE and those on the FTSE-100 traded on LIFFE, are popular for both speculating and hedging. Settlement is in cash, not the shares that constitute the underlying index, much like the settlement of an index futures contract.

❑ ***Bonds.*** Exchange-traded options on bonds are invariably written on the bonds' futures contracts. One of the most popular exchange-traded options contracts, for example, is the Treasury bond option, which is written on the Treasury futures contract and traded on the Chicago Board of Trade Options Exchange. Options written on actual bonds must be traded in the OTC market.

❑ ***Interest rates.*** All major exchanges write interest rate options on their 90-day interest rate futures contract.

❑ ***Foreign currency.*** Exchange-traded options on foreign currencies are rare. The major exchange trading them is the one in Philadelphia, which offers, for instance, a sterling option contract on an underlying amount of £31,250. The option gives the holder the right to buy or sell a given amount of the foreign currency at a given price per unit. A sterling call would give the holder the right to buy £31,250 for a certain dollar amount, which would be the strike price.

Option trading on an exchange, like futures trading, involves the daily computation and transfer of margin. Each exchange has its own procedures. On the LIFFE, for example, the option buyer pays no premium on the day the position is put on. Rather, the premium is paid via the daily variation margin, which reflects the daily changes in the option price. The sum of all the variation margin payments made during the life of an option that expires with no intrinsic value equals the difference between the

contract's value on the day it was traded and zero. On other exchanges, option buyers pay premiums on the day of purchase but no variation margin. Some exchanges allow traders to select either method. Margin is compulsory for option writers.

In the OTC market, a large variety of instruments are traded. As with other OTC products, such as swaps, the great advantage of OTC options is that they can be tailored to meet each buyer's requirements. Because of this flexibility, corporations and financial institutions can use them to structure hedges that perfectly match their risk exposures.

Option Pricing: Setting the Scene

An option's price is a function of the following six factors:

1 Its strike price
2 The current price of the underlying asset
3 The option's time to expiry
4 The risk-free interest rate during the option's life
5 The volatility of the underlying asset's price
6 The value of any dividends or cash flows paid by the underlying asset during the option's life

Possibly the two most important of these factors are the current price of the underlying and the option's strike price. As noted above, the relationship between these two determines the option's intrinsic value. The value of a call option thus rises and falls with the price of the underlying. And given several calls on the same asset, the higher the strike, the lower the option price. All this is reversed for a put option.

For all options, the longer the time to maturity, generally, the higher the premium. All other parameters being equal, a longer-dated option is worth at least as much as one with a shorter life. This rule—always true for American options and usually true for European ones—makes intuitive sense: the longer the term to maturity, the more time the underlying asset has to move in a direction and by an amount that increases the option's intrinsic value. Certain factors, however, may cause a longer-dated option to have only a slightly higher value than a shorter-dated one.

One such factor is the coupon payments made by the underlying during the option's life. These payments reduce the price of the underlying asset on the ex-dividend date and so depress the price of a call option and boost that of a put.

A rise in interest rates increases the value of most call options. For stock options, this is because the equity markets view a rate increase as a sign that share price growth will accelerate. Generally, the relationship is the

same for bond options. Not always, however, since in the bond market, rising rates tend to depress prices, because they lower the present value of future cash flows. A rise in interest rates has the opposite effect on put options, causing their value to drop. The risk-free interest rate applicable to a bond option with a term to expiry of, say, three months is a three-month government rate—commonly the government bond repo rate for bond options, usually the T-bill rate for other types.

Limits on Option Prices

Setting option prices' upper and lower limits is relatively straightforward since, like all security prices, they must obey the no-arbitrage rule. A call option grants the buyer the right to buy the underlying asset at the strike price. Clearly, therefore, it cannot have a higher value than the underlying asset. This relationship is expressed formally in (8.1).

$$C \leq S \tag{8.1}$$

where
 C = the price of a call option
 S = the current price of the underlying asset

Similarly, a put option, which grants the buyer the right to sell the underlying at the strike price, can never have a value greater than the strike price. This is expressed formally in (8.2).

$$P \leq X \tag{8.2}$$

where
 P = the price of the put option
 X = the put option strike price

This rule applies to European put options on their expiry date as well as to American puts. This means that a put option cannot have a value greater than the present value of the strike price at expiry. This is expressed formally in (8.3).

$$P \leq Xe^{-rT} \tag{8.3}$$

where
 r = the risk-free interest rate applicable to the option's term
 T = the number of years in the option's life

The lower limit on an option's price depends on whether or not the underlying asset pays dividends. Remembering that intrinsic value can never be less than zero, the lower bound on the price of a call option on a non-dividend-paying security is given by (8.4).

$$C \geq \max\left[S - Xe^{-rT}, 0\right] \tag{8.4}$$

For put options on non-dividend-paying stocks, the lower limit is given by (8.5).

$$P \geq \max\left[Xe^{-rT} - S, 0\right] \tag{8.5}$$

Since, as noted above, payment of a dividend by the underlying asset affects the option's price, the formulas for the lower price bounds must be modified for options on dividend-paying stocks as shown in (8.6) and (8.7).

$$C \geq S - D - Xe^{rT} \tag{8.6}$$

$$P \geq D + Xe^{-rT} - S \tag{8.7}$$

where
D = the present value of the dividend payment made by the underlying asset

Option Pricing

The pricing of other interest rate products, both cash and derivatives, that was described in previous chapters used rigid mathematical principles. This was possible because what happens to these instruments at maturity is known, allowing their fair values to be calculated. With options, however, the outcome at expiry is uncertain, since they may or may not be exercised. This uncertainty about final outcomes makes options more difficult to price than other financial market instruments.

Essentially, the premium represents the buyer's *expected profit*. Like insurance premiums, option premiums depend on how the option writers assess the likelihood of the payout equaling the premium. That, in turn, is a function of the probability of the option being exercised. An option's price is, therefore, a function of the probability that it will be exercised, from which is derived an expected outcome and a fair value.

The following factors influence an option's price.

❑ *The strike price.* Since the deeper in the money the option, the more likely it is to be exercised, the difference between the strike and the underlying asset's price when the option is struck influences the size of the premium.

❑ *The term to maturity.* The longer the term of the option, the more likely it is to move into the money and thus be exercised.

❑ *The level of interest rates.* As noted above, the option premium, in theory, equals the present value of the gain the buyer expects to realize at exercise. The discount rate used therefore affects the premium, although it is less influential than the other factors discussed.

❑ *The price behavior of financial instruments.* One of the key assumptions of option pricing models such as Black-Scholes (B-S), which is discussed below, is that asset prices follow a lognormal distribution—that is, the logarithms of the prices show a normal distribution. This characterization is not strictly accurate: prices are not lognormally distributed. Asset returns, however, are. Returns are defined by formula (8.8).

$$\ln\left(\frac{P_{t+1}}{P_t}\right) \tag{8.8}$$

where

P_t = the asset market price at time t

P_{t+1} = the price one period later

Given this definition, and assuming a lognormal distribution, an asset's expected return may be calculated using equation (8.9).

$$E\left[\ln\left(\frac{P_t}{P_0}\right)\right] = rt \tag{8.9}$$

where

$E[\ ]$ = the expectation operator

r = the annual rate of return

$\left(\frac{P_t}{P_0}\right)$ = the price relative

❑ *Volatility.* The higher the volatility of the underlying asset's price, the greater the probability of exercise, so the asset's volatility when the option is initiated will also influence the premium. The volatility of an asset is the annualized standard deviation of its price returns—that is, of the returns that generate the asset's prices, not the prices themselves. (Using prices would give inconsistent results, because the standard deviation would change as prices increased.) This is expressed formally in (8.10).

$$\sigma = \sqrt{\sum_{i=1}^{R} \frac{(x_i - \mu)^2}{R - 1}} \tag{8.10}$$

where

X_i = the ith price relative

μ = the arithmetic mean of the observed prices

R = the total number of observations

σ = volatility of the price returns

The volatility value derived by (8.10) may be converted to an annualized figure by multiplying it by the square root of the number of days in a year, usually taken to be 250 working days. Using this formula based on market observations, it is possible to calculate the *historical volatility* of an asset.

In pricing an option that expires in the future, however, the relevant factor is not historical but *future* volatility, which, by definition, cannot be measured directly. Market makers get around this problem by reversing the process that derives option prices from volatility and other parameters. Given an option price, they calculate the *implied volatility.* The implied volatilities of options that are either deeply in or deeply out of the money tend to be high.

The Black-Scholes Option Model

Most option pricing models use one of two methodologies, both of which are based on essentially identical assumptions. The first method, used in the Black-Scholes model, resolves the asset-price model's partial differential equation corresponding to the expected payoff of the option. The second is the martingale method, first introduced in Harrison and Kreps (1979) and Harrison and Pliska (1981). This derives the price of an asset at time 0 from its discounted expected future payoffs assuming risk-neutral probability. A third methodology assumes lognormal distribution of asset returns but follows the two-step binomial process described in chapter 11.

Employing pricing models requires the assumption of a *complete market.* First proposed in Arrow and Debreu (1953, 1954), this is a viable financial market where no-arbitrage pricing holds—that is, risk-free profits cannot be generated because of anomalies such as incorrect forward interest rates. This means that a zero-cost investment strategy that is initiated at time t will have a zero value at maturity. The martingale method assumes that an accurate estimate of the future price of an asset may be obtained from current price information. This property of future prices is also incorporated in the semistrong and strong market efficiency scenarios described in Fama (1965).

Assumptions

The Black-Scholes model is neat and intuitive. It describes a process for calculating the fair value of a European call option, but one of its many attractions is that it can easily be modified to handle other types, such as foreign-exchange or interest rate options.

Incorporated in the model are certain assumptions. For instance, apart from the price of the underlying asset, S, and the time, t, all the variables in the model are assumed to be constant, including, most crucially, the volatility variable. In addition, the following assumptions are made:

❑ There are no transaction costs, and the market allows short selling
❑ Trading is continuous
❑ The asset is a non-dividend-paying security
❑ The interest rate during the life of the option is known and constant
❑ The option can only be exercised on expiry

The behavior of underlying asset prices follows a geometric Brownian motion, or Weiner process, with a variance rate proportional to the square root of the price. This is stated formally in (8.11).

$$\frac{\mathrm{d}S}{S} = a\mathrm{d}t + b\mathrm{d}W \tag{8.11}$$

where

S = the underlying asset price
a = the expected return on the underlying asset
b = the standard deviation of the asset's price returns
t = time
W = the Weiner process

The following section presents an intuitive explanation of the B-S model, in terms of the normal distribution of asset price returns.

Pricing Derivative Instruments Using the Black-Scholes Model

To price an option, its fair value at contract initiation must be calculated. This value is a function of the option's expected terminal payoff, discounted to the day the contract was struck. Expression (8.12) describes the expected value of a call option at maturity T.

$$E(C_T) = E\big[\max(S_T - X, 0)\big] \tag{8.12}$$

where

C_T = the price of the call option at maturity T
E = is the expectations operator

S_T = the price of the underlying asset at maturity T
X = the strike price of the option

According to (8.12), only two outcomes are possible at maturity: either the option is in the money and the holder earns $S_T - X$, or it is out of the money and expires worthless. Modifying (8.12) to incorporate probability gives equation (8.13).

$$E(C_T) = p \times \left(E[S_T \mid S_T > X] - X\right) \qquad (8.13)$$

where
p = the probability that on expiry $S_T > X$
$E[S_T \mid S_T > X]$ = the expected value of S_T such that $S_T > X$

Equation (8.13) derives the expected value of a call option on maturity. Equation (8.14) derives the fair price of the option at contract initiation by discounting the value given by (8.13) back to this date.

$$C = p \times e^{-rt} \times \left(E[S_T \mid S_T > X] - X\right) \qquad (8.14)$$

where
r = the continuously compounded risk-free rate of interest
t = the period from today until maturity

Pricing an option therefore requires knowing the value of both p, the probability that the option will expire in the money, and $E[S_T \mid S_T > X] - X$, its expected payoff should this happen. In calculating p, the probability function is modeled. This requires assuming that asset prices follow a stochastic process.

The B-S model is based on the resolution of partial differential equation (8.15), given the appropriate parameters. The parameters refer to the payoff conditions corresponding to a European call option.

$$\tfrac{1}{2}\sigma^2 S^2 \left(\frac{\partial^2 C}{\partial S^2}\right) + rS\left(\frac{\partial C}{\partial S}\right) + \left(\frac{\partial C}{\partial t}\right) - rC = 0 \qquad (8.15)$$

The derivation of and solution to equation (8.15) are too complex to discuss in this book. The interested reader is directed to the works listed in the References section. What follows is a discussion of how the probability and expected-value functions are solved.

The probability, p, that the underlying asset's price at maturity will exceed X equals the probability that its return over the option's holding period will exceed a certain critical value. Since asset price returns are assumed to be lognormally distributed and are themselves defined as the logarithm of price relatives, this equivalence can be expressed as (8.16).

$$p = prob[S_T > X] = prob\left[return > \ln\left(\frac{X}{S_0}\right)\right] \tag{8.16}$$

where
S_0 = the price of the underlying asset at the time the option is initiated

Generally, the probability that a normally distributed variable x will exceed a critical value x_c is given by (8.17).

$$p[x > x_c] = 1 - N\left(\frac{x_c - \mu}{\sigma}\right) \tag{8.17}$$

where
μ and σ = the mean and standard deviation, respectively, of x
$N(\ )$ = the cumulative normal distribution

As discussed earlier, an expansion for μ is the natural logarithm of the asset price returns, and for the standard deviation of returns is $\sigma\sqrt{t}$. Equations (8.16) and (8.17) can therefore be combined as (8.18).

$$\tag{8.18}$$
$$prob[S_T > X] = prob\left[return > \ln\left(\frac{X}{S_0}\right)\right] = 1 - N\left(\frac{\ln\left(\frac{X}{S_0}\right) - \left(r - \frac{\sigma^2}{2}\right)t}{\sigma\sqrt{t}}\right)$$

The symmetrical shape of a normal distribution means that the probability can be obtained by setting $1-N(d)$ equal to $N(-d)$. The result is (8.19).

$$p = prob[S_T > X] = N\left(\frac{\ln\left(\frac{S_0}{X}\right) + \left(r - \frac{\sigma^2}{2}\right)t}{\sigma\sqrt{t}}\right)$$
$$\tag{8.19}$$

Now a formula is required to calculate the second part of the expression at (8.14): the expected value of the option at expiration, T. This involves the integration of the normal distribution curve over the range from X to infinity. The derivation is not shown here, but the result is (8.20).

$$E[S_T \mid S_T > X] = S_0 e^{rt} \frac{N(d_1)}{N(d_2)} \tag{8.20}$$

where

$$d_1 = \frac{\ln\left(\dfrac{S_0}{X}\right) + \left[r + \dfrac{\sigma^2}{2}\right]t}{\sigma\sqrt{t}}$$

and

$$d_2 = \frac{\ln\left(\dfrac{S_0}{X}\right) + \left[r - \dfrac{\sigma^2}{2}\right]t}{\sigma\sqrt{t}} = d_1 - \sigma\sqrt{t}$$

Plugging the probability and expected value expressions, (8.19) and (8.29), into (8.14) results in (8.21).

$$C = N(d_2) \times e^{-rt} \times \left(S_0 e^{rt} \frac{N(d_1)}{N(d_2)} - X \right) \tag{8.21}$$

Equation (8.21) can be simplified as (8.22), the well-known Black-Scholes option pricing model for a European call option. It states that the fair value of a call option is the expected present value of the option on its expiry date, assuming that prices follow a lognormal distribution.

$$C = S_0 N(d_1) - X e^{-rt} N(d_2) \tag{8.22}$$

where
C = the price of a call option
S_0 = the price of the underlying asset at the time the option is struck
X = the strike price
r = the continuously compounded risk-free interest rate
t = the time to option maturity

$N(d_1)$ and $N(d_2)$ are the cumulative probabilities from the normal distribution of obtaining the values d_1 and d_2, defined above. $N(d_1)$ is the *delta* of the option—that is, the change in the option price for a given change in the price of the underlying. $N(d_2)$ represents the probability that

the option will be exercised. The term e^{-rt} is the present value of one unit of cash received t periods from the time the option is struck. Assuming that $N(d_1)$ and $N(d_2)$ both equal 1—that is, assuming complete certainty—(8.22) simplifies to (8.23), Merton's lower bound (the lower bound of call prices) for continuously compounded interest rates, introduced in intuitive fashion previously in this chapter. Assuming complete certainty, therefore, the B-S model reduces to Merton's bound.

$$C = S - Xe^{-rt} \tag{8.23}$$

Put-Call Parity

The prices of call and put options are related via the *put-call parity theorem*. This important relationship obviates the need for a separate model for put options.

Consider two portfolios, Y and Z. Y consists of a call option with a maturity date T and a zero-coupon bond that pays out X on T; Z consists of a put option also maturing on date T and one unit of the underlying asset. The values of portfolios Y and Z on the expiry date are given by equations (8.24) and (8.25), respectively.

$$MV_{Y,T} = \max[S_T - X, 0] + X = \max[X, S_T] \tag{8.24}$$

$$MV_{Z,T} = \max[X - S_T, 0] + S_T = \max[X, S_T] \tag{8.25}$$

Both portfolios have the same value at maturity. Since prices are assumed to be arbitrage free, the two sets of holdings must also have the same initial value at start time t. The put-call relationship expressed in (8.26) must therefore hold.

$$C_t - P_t = S_t - Xe^{-r(T-t)} \tag{8.26}$$

From this relationship it is possible to construct equation (8.27) to derive the value of a European put option.

$$P(S,T) = -SN(-d_1) + Xe^{-rT}N(-d_2) \tag{8.27}$$

Pricing Options on Bonds Using the Black-Scholes Model

The theoretical price of a call option written on a zero-coupon bond is calculated using equation (8.28).

$$C = PN(d_1) - Xe^{-rT}N(d_2) \tag{8.28}$$

EXAMPLES: *Options Pricing Using the Black-Scholes Model*

1   Calculate the price of a call option written with strike price 21 and a maturity of three months written on a non-dividend-paying stock whose current share price is 25 and whose implied volatility is 23 percent, given a short-term risk-free interest rate of 5 percent.

Call price is given by equation $C = S_0 N(d_1) - Xe^{-rt} N(d_2)$

a. Assign values to the relevant variables:

$S_0 = 25$

$X = 21$

$r = 5$ percent

$t = 0.25$

$\sigma = 23$ percent

b. Calculate the discounted value of the strike price:

$$Xe^{-rt} = 21e^{-0.05(0.25)} = 20.73913$$

c. Calculate the values of d_1 and d_2:

$$d_1 = \frac{\ln\left(\frac{S_0}{X}\right) + \left(r + \frac{\sigma^2}{2}\right)t}{\sigma\sqrt{t}}$$

$$d_1 = \frac{\ln(25/21) + \left[0.05 + (0.5)(0.23)^2\right]0.25}{0.23\sqrt{0.25}} = \frac{0.193466}{0.115} = 1.682313$$

$$d_2 = \frac{\ln\left(\frac{S_0}{X}\right) + \left(r - \frac{\sigma^2}{2}\right)t}{\sigma\sqrt{t}} = d_1 - \sigma\sqrt{t}$$

$$d_s = d_1 - 0.23\sqrt{0.25} = 1.567313$$

d. Insert these values into the main price equation:

$$C = S_0 N(d_1) - Xe^{-rt} N(d_2)$$

$$C = 25N(1.682313) - 21e^{-0.05(0.25)} N(1.567313)$$

Using the approximation of the cumulative Normal distribution at the points 1.68 and 1.56, the price of the call option is

$$C = 25(0.9535) - 20.73913(0.9406) = 4.3303$$

2   Calculate the price of a put option on the same stock, given the same risk-free interest rate.

The equation for calculating a put option's price is $P(S,T) = -SN(-d_1) + Xe^{-rT} N(-d_2)$

a. Plugging in the given and derived values, including those for $N(d_1)$ and $N(d_2)$—0.9535 and 0.9406:

$$P = 20.7391\ (1 - 0.9406) - 25\ (1 - 0.9535) = 0.06943$$

3 Use the put-call parity theorem to calculate the price of the put option, plugging in the call price derived above: 4.3303. The put-call parity equation is

$$P = C - S + Xe^{-rt}$$

$$= 4.3303 - 25 + 21e^{-0.05(0.25)}$$

$$= 0.069434$$

Note that this equals the price obtained by applying the put option formula.

4 Calculate the price of a call on the same stock, with the same strike, but with only six months to maturity.

a. All the variable values remain the same except the following: $t = 0.5$
b. Calculate the discounted value of the strike price:
$Xe^{-rt} = 21e^{-0.05(0.5)} = 20.48151$
c. Calculate the values of $N(d_1)$ and $N(d_2)$:

$$d_1 = \frac{\ln\left(\frac{S_0}{X}\right) + \left(r + \frac{\sigma^2}{2}\right)t}{\sigma\sqrt{t}}$$

$$d_2 = \frac{\ln\left(\frac{S_0}{X}\right) + \left(r - \frac{\sigma^2}{2}\right)t}{\sigma\sqrt{t}} = d_1 - \sigma\sqrt{t}$$

$d_1 = 1.3071$, giving $N(d_1) = 0.9049$
$d_2 = 1.1445$, giving $N(d_2) = 0.8740$

d. Plug these values into the call price equation:

$$C = S_0 N(d_1) - Xe^{-rt} N(d_2)$$

$$C = 25(0.9049) - 20.48151(0.8740) = 4.7217$$

This demonstrates the point made earlier that an option's premium increases when the time to expiry, the volatility, or the interest rate (or any combination of these) is increased.

where
  P = the price of the underlying bond
  All other parameters remain the same.

Note that although a key assumption of the model is that interest rates are constant, in the case of bond options, it is applied to an asset price that is essentially an interest rate assumed to follow a stochastic process.

For an underlying coupon-paying bond, the equation must be modified by reducing P by the present value of all coupons paid during the life of the option. This reflects the fact that prices of call options on coupon-paying bonds are often lower than those of similar options on zero-coupon bonds because the coupon payments make holding the bonds themselves more attractive than holding options on them.

Interest Rate Options and the Black Model

In 1976 Fisher Black presented a slightly modified version of the B-S model, using similar assumptions, for pricing forward contracts and interest rate options. Banks today employ this modified version, the Black model, to price swaptions and similar instruments in addition to bond and interest rate options, such as caps and floors. The bond options described in this section are options on bond futures contracts, just as the interest rate options are options on interest rate futures.

The Black model refers to the underlying asset's or commodity's spot price, $S(t)$. This is defined as the price at time t payable for immediate delivery, which, in practice, means delivery up to two days forward. The spot price is assumed to follow a geometric Brownian motion. The theoretical price, $F(t,T)$, of a futures contract on the underlying asset is the price agreed at time t for delivery of the asset at time T and payable on delivery. When $t = T$, the futures price equals the spot price. As explained in chapter 12, futures contracts are cash settled every day through a clearing mechanism, while forward contracts involve neither daily marking to market nor daily cash settlement.

The values of forward, futures, and option contracts are all functions of the futures price $F(t,T)$, as well as of additional variables. So the values at time t of a forward, a futures, and an option can be expressed, respectively, as $f(F,t)$, $u(F,t)$, and $C(F,t)$. Since the value of a forward contract is also a function of the price of the underlying asset S at time T, it can be represented by $f(F,t,S,T)$. Note that the value of the forward contract is not the same as its price. As explained in chapter 12, a forward's price, at any given time, is the delivery price that would result in the contract having a zero present value. When the contract is transacted, the forward value is zero. Over time both the price and the value fluctuate. The futures price

EXAMPLE: *Bond-Option Pricing*

Calculate the price of a European call option with a strike price of $100 and a maturity of one year, written on a bond with the following characteristics:

Price	$98
Semiannual coupon	8.00 percent
Time to maturity	5 years
Bond price volatility	6.02 percent
Coupon payments	$4 each, one payable in three months and another in nine months from the option start date
Three-month risk-free interest rate	5.60 percent
Nine-month risk-free interest rate	5.75 percent
One-year risk-free interest rate	6.25 percent

a. Calculate the present value of the coupon payments made during the life of the option:

$$4e^{-0.056\times0.25} + 4e^{-0.0575\times0.75} = 3.9444 + 3.83117 = 7.77557$$

b. Subtract the present value of the coupon payments from the bond price: $P = 98 - 7.78 = \$90.22$

c. Calculate d_1 and d_2:

$$d_1 = \frac{\ln\left(\dfrac{P_0}{X}\right) + \left(r + \dfrac{\sigma^2}{2}\right)t}{\sigma\sqrt{t}}$$

$$d_2 = \frac{\ln\left(\dfrac{P_0}{X}\right) + \left(r - \dfrac{\sigma^2}{2}\right)t}{\sigma\sqrt{t}} = d_1 - \sigma\sqrt{t}$$

$$d_1 = \left[\ln\left(90.22/100\right) + 0.0625 + 0.001812\right]/0.0602 = -0.6413$$

$$d_2 = d_1 - (0.0602\times1) = -0.7015$$

d. Plug these values into the equation for calculating the bond call option price:

$$C = PN(d_1) - Xe^{-rT}N(d_2)$$

$$C = 90.22N(-0.6413) - 100e^{-0.0625}N(-0.7015)$$

$$= 1.1514$$

The premium of the call option, $1.15, is composed entirely of time value.

is the price at which a forward contract has a zero current value. When a forward is traded, therefore, its price is equal to the futures price F. This equivalence is expressed in equation (8.29), which states that the value of the forward contract is zero when the contract is taken out, and the contract price $S(T)$ is always equal to the current futures price, $F(t, T)$.

$$f(F, t, F, T) = 0 \qquad (8.29)$$

Because futures contracts are repriced each day at the new forward price, their prices imply those of forward contracts. When F rises, so that $F > S$, f is positive; when F falls, f is negative. When the contract expires and delivery takes place, the forward contract value equals the spot price minus either the contract price or the spot price, futures price equals the spot price, and the value of the forward contract equals the spot price minus the contract price or the spot price.

$$f(F, T, S, T) = F - S \qquad (8.30)$$

The value of a bond or commodity option at maturity is either the difference between the spot price of the underlying and the contract price or zero, whichever is larger. Since the futures price on the maturity date equals the spot price, the equivalence expressed in (8.31) holds.

$$C(F, T) = \begin{cases} F - S & \text{if } F \geq S_T \\ 0 & \text{else} \end{cases} \qquad (8.31)$$

The Black model assumes that the prices of futures contracts follow a lognormal distribution with a constant variance that the capital asset pricing model applies in the market, and that no transaction costs or taxes apply. Under these assumptions, a risk-free hedged position can be created that is composed of a long position in the option and a short position in the futures contract. Following the B-S model, the number of options put on against one futures contract is given by $[\partial C(F, t) / \partial F]$, which is the derivative of $C(F, t)$ with respect to F. The change in the hedged position resulting from a change in the price of the underlying is given by expression (8.32).

$$\partial C(F, t) - [\partial C(F, t) / \partial F] \partial F \qquad (8.32)$$

The principle of arbitrage-free pricing requires that the hedged portfolio's return equal the risk-free interest rate. This equivalence plus an expansion of $\partial C(F, t)$ produces partial differential equation (8.33).

$$\left[\frac{\partial C(F,t)}{\partial t}\right] = rC(F,t) - \tfrac{1}{2}\sigma^2 F^2 \left[\frac{\partial^2 C(F,t)}{\partial F^2}\right] \qquad (8.33)$$

Solving this equation (not shown here) involves rearranging it as (8.34).

$$\tfrac{1}{2}\sigma^2 F^2 \left[\frac{\partial^2 C(F,t)}{\partial F^2}\right] - rC(F,t) + \left[\frac{\partial C(F,t)}{\partial t}\right] = 0 \qquad (8.34)$$

The solution to the partial differential equation (8.22) is not presented here.

Setting $T = T\text{-}t$ and using (8.32) and (8.33), an equation can be created for deriving the fair value of a commodity option or option on a forward contract, shown as (8.35).

$$C(F,t) = e^{-rT}\left[FN(d_1) - S_T N(d_2)\right] \qquad (8.35)$$

where

$$d_1 = \frac{1}{\sigma\sqrt{T}}\left[\ln\left(\frac{F}{S_T}\right) + \left(\tfrac{1}{2}\sigma^2\right)T\right]$$

$$d_2 = d_1 - \sigma\sqrt{T}$$

Comments on the Black-Scholes Model

The introduction of the Black-Scholes model paved the way for the rapid development of options as liquid tradable products. B-S is widely used today to price options and other derivatives. Nevertheless, academics have pointed out several weaknesses related to the main assumptions on which it is based. The major criticisms involve the following:

❑ *The assumption of frictionless markets.* This is, at best, only for large markets and then only approximately.

❑ *The assumption of a constant interest rate.* This is possibly the model's most unrealistic assumption. Not only are rates dynamic, but those at the short end of the yield curve often move in the opposite direction from asset prices, particularly the prices of bonds and bond options.

❑ *The volatility data.* Of all the inputs to the B-S model, the variability of the underlying asset—its volatility—is the most problematic. The distribution of asset prices is assumed to be lognormal, meaning that the logarithms of the prices are normally distributed (lognormal rather than normal distribution is used because prices cannot have negative values, which would be allowed in a normal distribution). Though accepted

as a reasonable approximation of reality, this assumption is not completely accurate and fails to account for the extreme moves and market shocks that sometimes occur.

❑ *Limitation to European exercise.* Although American options are rarely exercised early, sometimes the situation requires early exercise, and the B-S model does not price these situations.

❑ *The assumption, for stock options, of a constant dividend yield.*

Stochastic Volatility

The volatility figure used in a B-S computation is constant and derived mathematically, assuming that asset prices move according to a geometric Brownian motion. In reality, however, asset prices that are either very high or very low do not move in this way. Rather, as a price rises, its volatility increases, and as it falls, its variability decreases. As a result, the B-S model tends to undervalue out-of-the-money options and overvalue those that are deeply in the money.

To correct this mispricing, stochastic volatility models, such as the one proposed in Hull and White (1987), have been developed.

Implied Volatility

As noted earlier, although many practitioners use a historical volatility figure in applying the B-S model, the pertinent statistic is really the underlying asset's price volatility going forward. To estimate this future value, banks employ the volatility that is implied by the prices of exchange-traded options. It is not possible, however, to rearrange the B-S model to derive the volatility measure, σ, as a function of the observed price and the other parameters. Generally, therefore, a numerical iteration process, usually the Newton-Raphson method, is used to arrive at the value for σ given the price of the option.

The market uses implied volatilities to gauge the volatility of individual assets relative to the market. The price volatility of an asset is not constant. It fluctuates with the overall volatility of the market, and for reasons specific to the asset itself. When deriving implied volatility from exchange-traded options, market makers compute more than one value, because different options on the same asset will imply different volatilities depending on how close to at the money the option is. The price of an at-the-money option is more sensitive to volatility than that of a deeply in- or out-of-the-money one.

Other Option Models

Other pricing models have been developed drawing on the pioneering work done by Black and Scholes. The B-S model is straightforward and easy to apply, and subsequent work has focused on modifying its assumptions and restrictions to improve its accuracy—by, for instance, incorporating nonconstant volatility—and expand its applicability, to American options and those on dividend-paying stocks, among other products. A number of models have been developed to price specific contracts. The Garman and Kohlhagen (1983) and Grabbe (1983) models are applied to currency options, while the Merton (1973) and the Barone-Adesi Whaley (1987), or BAW, model is used for commodity options. Roll (1977), Geske (1979), and Whaley (1981) developed still another model to value American options on dividend-paying assets. More recently Black, Derman, and Toy (1990) introduced a model to price exotic options.

Some of the newer models refer to parameters that are difficult to observe or measure directly. In practice, this limits their application much as B-S is limited. Usually the problem has to do with calibrating the model properly, which is crucial to implementing it. Calibration entails inputting actual market data to create the parameters for calculating prices. A model for calculating the prices of options in the U.S. market, for example, would use U.S. dollar money market, futures, and swap rates to build the zero-coupon yield curve. Multifactor models in the mold of Heath-Jarrow-Morton employ the correlation coefficients between forward rates and the term structure to calculate the volatility inputs for their price calculations.

Incorrect calibration produces errors in option valuation that may be discovered only after significant losses have been suffered. If the necessary data are not available to calibrate a sophisticated model, a simpler one may need to be used. This is not an issue for products priced in major currencies such as the dollar, sterling, or euro, but it can be a problem for other currencies. That might be why the B-S model is still widely used today, although models such as the Black-Derman-Toy and the one proposed in Brace, Gatarek, and Musiela (1994) are increasingly employed for more exotic option products.

Chapter Notes

1. I'm told this terminology originates in the fact that Bermuda is midway between Europe and America. *Asian* options are defined not in terms of when they can be exercised but how their settlement values are determined—not by the difference between the strike and the current price of the underlying asset but by the difference between the strike and the average of the prices recorded by the underlying at a range of specified dates or over a specified period. Hence these are also known as *average* or *average rate* options. A former colleague informs me that they are termed *Asian* because they originated in Japanese commodity markets.

This page is intentionally blank

Measuring Option Risk

This chapter looks at how options behave in response to changes in market conditions and the main issues that a market maker in options must consider when writing the contracts. It also reviews "the Greeks"—the sensitivity measures applied to option books—and an important set of interest rate options: *caps* and *floors.*

Option Price Behavior

As noted in chapter 8, the value of an option is a function of five factors:
- ❑ The price of the underlying asset
- ❑ The option's strike price
- ❑ The option's time to expiry
- ❑ The volatility of the underlying asset's price returns
- ❑ The risk-free interest rate applicable to the life of the option

The price of an option is composed of intrinsic value and time value. An option's intrinsic value is clear. Valuation models, therefore, essentially price time value.

Assessing Time Value

At-the-money options have the greatest time value; in-the-money contracts have more time value than out-of-the-money ones. These relationships reflect the risk the different options pose to the market makers that write them. Out-of-the-money call options, for instance, have the lowest

probability of being exercised, so market makers may not even hedge them. In leaving these positions uncovered, of course, they run the risk that the underlying asset's price will rise sufficiently to drive the option into the money, in which case the writers must purchase the asset in the market and suffer a loss. This scenario is least likely for options that are deeply out of the money. These contracts thus present market makers with the smallest risk, so their time value is lowest.

In-the-money call options are more likely to be exercised. Market makers writing these contracts therefore generally hedge their positions. They may do so using futures contracts or a *risk reversal*—a long or short position in a call that is reversed to the same as the original position by selling or buying the position for forward settlement (and vice versa). The in-the-money call writer buys a call with the same expiry as the one written but with a slightly higher strike, to cap the possible loss (or a lower strike, to hedge it completely), and simultaneously sell a call with a longer term to offset the cost. Or they may use the underlying asset. Market makers choosing the last alternative run the risk that the asset's price will fall, in which case the option will not be exercised and they will be forced to dispose of the asset at a loss. The more deeply in the money the option, the lower this risk—and, accordingly, the smaller the time value.

At-the-money options—which constitute the majority of OTC contracts—are the riskiest to write. They have 50–50 chances of being exercised, so deciding whether or not to hedge them is less straightforward than with other options. It is this uncertainty about hedging that makes them so risky. Accordingly, at-the-money options have the highest time values.

American Options

In theory, American options should have greater value than equivalent European ones, because they can be exercised before maturity. In practice, however, early exercise rarely happens. This is because option holders realize only their contracts' intrinsic value when they exercise. By selling their contracts in the market, in contrast, holders can realize their full value, including time value.

Since the possibility of early exercise, which represents the chief difference between American and European options, is rarely actualized, the two types of contracts generally have equivalent values. Pricing models, however, calculate the probabilities that different options will be exercised. For American options, this entails determining what circumstances make early exercise more likely and assigning the contracts higher prices in these situations.

One situation conducive to early exercise is when an option has negative time value. This can happen when it is deeply in the money and very near maturity. Although, technically, an option in this situation still has a small positive time value, this value is outweighed by the opportunity cost of deferring the cash flows from the underlying asset that the holder would gain by exercising. Consider a deeply in-the-money option on a futures contract. By deferring exercise, holders lose the opportunity to invest the profits they would realize on the futures contract and forgo potential interest income from the contract's daily cash settlements.

The Greeks

Options' price sensitivity is different from that of other financial market instruments. An option contract's value can be affected by changes in any one or any combination of the five factors considered in option pricing models (of course, strike prices are constant in plain vanilla contracts). In contrast, swaps' values are sensitive to one variable only—the swap rate—and bond futures prices are functions of just the current spot price of the cheapest-to-deliver bond and the current money market repo rate. Even more important, unlike for the other instruments, the relationship between an option's value and a change in a key variable is not linear.

All this makes risk management more complex for option books than for portfolios of other instruments. Each variable must be considered and, in some cases, derivatives of these variables. The latter are often referred to as the "Greeks," because Greek letters are used to denote them all, except volatility sensitivity. This is most commonly represented by *vega,* although the Greek *kappa* is also sometimes used.

Delta

The $N(d_1)$ term in the Black-Scholes equation represents an option's *delta.* Delta indicates how much the contract's value, or premium, changes as the underlying asset's price changes. An option with a delta of zero does not move at all as the price of the underlying changes; one with a delta of 1 behaves the same as the underlying. The value of an option with a delta of 0.6, or 60 percent, increases $60 for each $100 increase in the value of the underlying. The relationship is expressed formally in (9.1).

$$\delta = \frac{\Delta C}{\Delta S} \tag{9.1}$$

where
    C = call option price
    S = the price of the underlying asset

Mathematically, an option's delta is the partial derivative of its premium with respect to the price of the underlying. This is expressed in equation (9.2).

$$\delta = \frac{\partial C}{\partial S}$$

or

$$\delta = \frac{\partial P}{\partial S} \tag{9.2}$$

where
  P = the put price

Delta is closely related, but not equal, to the probability that an option will be exercised. It is very important for option market makers and is also the main hedge measure.

To hedge the options they write, market makers can buy matching options, buy or sell other instruments with the same but opposite values, or buy or sell the underlying assets. The amount of the hedging instrument used is governed by the options' delta. Say a trader writes ten call options, each representing one hundred shares of common stock, with a delta of 0.6. A hedge for this position might consist of six hundred shares of the underlying stock. If the share price rises by $1, the $600 rise in the value of the equity position will offset the trader's $600 loss in the option position. The combined positions are *delta neutral*. This process is known as *delta hedging*. As will be discussed later, such hedges are only approximate. Moreover, an option's delta changes during its term, so delta hedges must be monitored and adjusted, a process known as *dynamic hedging.*

A combined option-underlying position with a positive delta is equiva-

FIGURE 9.1 *Adjustments to Delta-Neutral Hedge in Response to Changes in the Underlying Asset's Price*

OPTION	RISE IN UNDERLYING ASSET PRICE	FALL IN UNDERLYING ASSET PRICE
Long call	Rise in delta: sell underlying	Fall in delta: buy underlying
Long put	Fall in delta: sell underlying	Rise in delta: buy underlying
Short call	Rise in delta: buy underlying	Fall in delta: sell underlying
Short put	Fall in delta: buy underlying	Rise in delta: sell underlying

lent to a long position in the underlying asset alone. A rise in the asset price results in a profit, since the market maker could theoretically sell either the underlying or the call option at a higher price. The opposite is true if the price of the underlying falls. A positive delta means that a market maker trying to maintain a delta-neutral position is edged and must buy or sell delta units of the underlying asset, although in practice futures contracts may be used. **FIGURE 9.1** shows the effect of changes in the underlying price on an option book's delta and the transactions necessary to restore neutrality.

Gamma

Just as modified duration becomes inaccurate as the magnitude of the yield change increases, so do inaccuracies occur in the use of delta to determine option-book hedges. This is because delta itself changes as the price of the underlying changes. Accordingly, a book that is delta-neutral at one asset price may not be when the price rises or falls. To help them guard against this, market makers employ *gamma,* which indicates how much an option's delta changes with movements in the underlying price. Gamma is expressed formally as (9.3).

$$\Gamma = \frac{\Delta \delta}{\Delta S} \tag{9.3}$$

Mathematically, gamma is the second partial derivative of the option price with respect to the underlying asset's price. This is expressed in (9.4).

$$\frac{\partial^2 C}{\partial S^2}$$

or

$$\frac{\partial^2 P}{\partial S^2} \tag{9.4}$$

The equation for gamma is (9.5).

$$\Gamma = \frac{N(d_1)}{S\sigma\sqrt{T}} \tag{9.5}$$

Gamma is the only major Greek that does not measure the sensitivity of the option premium; instead, it measures the delta sensitivity. Since delta determines an option's hedge ratio, gamma indicates how much this ratio must change for changes in the underlying's price. A nonzero gamma

is thus problematic, since it entails continually changing the hedge ratio. This relates to the fact that at-the-money options have the highest value because they present the greatest uncertainty and hence the highest risk.

When an option is deeply in or out of the money, its delta does not change rapidly, so its gamma is insignificant. When it is close to or at the money, however, its delta can change very suddenly, and its gamma is accordingly very large. Long option positions have positive gammas; short ones have negative gammas. Options with large gammas, whether positive or negative, are problematic for market makers, since their hedges must be adjusted constantly to maintain delta neutrality, and that entails high transaction costs. The larger the gamma, the greater the risk that the option book will be affected by sudden moves in the market. A position with a negative gamma is the riskiest and can be hedged only through long positions in other options.

A perfectly hedged book is gamma-neutral, meaning that its delta does not change. When gamma is positive, a rise in the price of the underlying asset increases the delta, necessitating a sale in the underlying asset or in the relevant futures contracts to maintain neutrality. The reverse applies if the underlying falls in price. Note that in this situation, the market maker is selling into a rising market, thus generating a profit, and buying in a falling one.

With a negative gamma, an increase in the price of the underlying depresses the delta and a decrease raises it. To adjust the hedge in the first situation, market makers must buy more of the underlying asset or futures equivalents; in the second situation, they must sell the asset or the futures. In this case, the market makers are either selling into a falling market, and generating losses even as the hedge is being put on, or buying assets in a rising market. Negative gamma, therefore, represents a high risk in a rising market.

When the price volatility of the underlying asset is high, a desk pursuing a delta-neutral strategy with a position having a positive gamma should be able to generate profits. Under the same conditions, a position with negative gamma could sustain losses and be excessively costly to hedge.

To adjust an option book so that it is gamma-neutral, a market maker must buy or sell options on the underlying asset or on the corresponding future, rather than trade either of these instruments themselves, since their gammas are zero. Adding to the book's option position, however, changes its delta. To maintain delta-neutrality, therefore, the market maker has to rebalance the book, using the underlying asset or futures contracts. Since gamma, like delta, changes with the market, the gamma hedge must also be continually rebalanced.

Theta

An option's *theta* indicates how its value changes with its time to maturity. This is expressed formally in (9.6).

$$\Theta = \frac{\Delta C}{\Delta T} \text{ or } -\frac{\partial C}{\partial T} \text{ or } -\frac{\partial P}{\partial T} \tag{9.6}$$

where
T = time to maturity

Using the relationships embodied in the B-S model, the theta of a call option can be expressed mathematically as (9.7).

$$\Theta = -\frac{S\sigma}{2\sqrt{2\pi T}} e^{-\frac{d_1^2}{2}} - Xre^{-rT} N(d_2) \tag{9.7}$$

where
π = a constant
X = the strike price
$N(d_2)$ = the probability that the option will be exercised
σ = the volatility of the underlying asset's price

Theta measures an option's time decay. Time decay hurts the holder of a long option position, because as expiry nears, the contract's value consists increasingly of intrinsic value alone, which may be zero. Option writers, in contrast, benefit from time decay, which reduces their risk. So a high theta should be advantageous for contract writers. High theta, however, entails high gamma, so, in practice, writers do not gain.

Some option strategies exploit theta. Writing a short-term option and simultaneously purchasing a longer-term one with the same strike price, for instance, represents a play on the options' theta: if the time value of the longer option decays at a slower rate than that of the short-dated option, the trade will be profitable.

Vega

An option's vega—also known as its *epsilon* (ε), *eta* (η), or *kappa* (κ)—indicates how much its value changes with changes in the price volatility of the underlying asset. For instance, an option with a vega of 12.75 will increase in price by 0.1275 for a 1 percent increase in volatility. Vega is expressed formally as (9.8).

$$v = \frac{\Delta C}{\Delta \sigma} \tag{9.8}$$

or

$$v = \frac{\partial C}{\partial \sigma} \text{ or } \frac{\partial P}{\partial \sigma}$$

Using the relationships embodied in the B-S formula, vega is defined as (9.9) for a call or put.

$$v = \frac{S\sqrt{\frac{T}{2\pi}}}{e^{\frac{d_1^2}{2}}} \tag{9.9}$$

or

$$v = S\sqrt{\Delta T} N(d_1)$$

Vegas are highest for at-the-money options, decreasing as the underlying's price and the strike prices diverge. Options with short terms to expiry have lower vegas than longer-dated ones. Positive vegas generally imply positive gammas. Long call and put positions usually have positive vegas, meaning that an increase in volatility increases their value.

Buying options is equivalent to buying volatility, while selling options is equivalent to selling volatility. Market makers generally like volatility and set up their books so that they have positive vega. In making trades, they calculate the volatilities implied by the option prices, then compare these values with their own estimates. If the implied volatilities appear too high, they short calls and puts, reversing their positions when the implied volatilities decline.

FIGURE 9.2 shows the modifications to a delta hedge in response to changes in volatility.

Managing an option book involves trade-offs between gamma and vega much like those between gamma and theta. A long options position is long vega and long gamma. This is not difficult to manage. If volatility falls, the market makers may choose to maintain positive gamma if they believe that the decrease in volatility can be offset by adjusting gamma in the direction of the market. On the other hand, they may prefer to set up a position with negative gamma by writing options, thus selling volatility. In either case, the costs associated with rebalancing the delta must compensate for the reduction in volatility.

Rho

An option's *rho* indicates how much its value changes with changes in interest rates. It is the least used of the sensitivity measures, because market

FIGURE **9.2** *Dynamic Hedging Responding to Changes in Volatility*

OPTION POSITION	RISE IN VOLATILITY	FALL IN VOLATILITY
LONG CALL		
ATM	No adjustment to delta	No adjustment to delta
ITM	Rise in delta, buy underlying	Rise in delta, sell underlying
OTM	Fall in delta, sell underlying	Fall in delta, buy underlying
LONG PUT		
ATM	No adjustment to delta	No adjustment to delta
ITM	Fall in delta, sell underlying	Rise in delta, buy underlying
OTM	Rise in delta, buy underlying	Fall in delta, sell underlying
SHORT CALL		
ATM	No adjustment to delta	No adjustment to delta
ITM	Fall in delta, sell underlying	Rise in delta, buy underlying
OTM	Rise in delta, buy underlying	Fall in delta, sell underlying
SHORT PUT		
ATM	No adjustment to delta	No adjustment to delta
ITM	Rise in delta, buy underlying	Rise in delta, sell underlying
OTM	Fall in delta, sell underlying	Fall in delta, buy underlying

interest rates are probably the least variable of all the parameters used in option pricing. Longer-dated options tend to have higher rhos. The measure is defined formally as (9.10).

$$\frac{\partial C}{\partial r} \ \text{ or } \ \frac{\partial P}{\partial r} \tag{9.10}$$

Using the relationships expressed in the B-S model formula, the rho of a call option is expressed as (9.11).

$$\rho = Xte^{-rT}N(d_2) \tag{9.11}$$

where

t = time to expiry

Lambda

An option's *lambda* is similar to its delta in that it indicates the change in option premium for a change in the underlying asset's price. Lambda, however, measures this sensitivity against the percentage change in the underlying's price. It thus indicates the option's gearing, or leverage, which, in turn, indicates the expected profit or loss for changes in the price of the underlying. **FIGURE 9.3** shows that in-the-money options have a gearing of at least five and sometimes considerably higher. This means that if the underlying asset rises in price, holders of a long in-the-money call could realize at least five times as great a profit as if they had invested the same cash amount in the asset.

The sensitivity measures enable market makers and portfolio managers to determine the effect of changes in prices and volatility levels on their entire books, without having to reprice all the options in them. All they need to do is compute the weighted sum of the options' deltas, gammas, vegas, and thetas. The Greeks are also important for risk managers and those implementing value-at-risk systems.

FIGURE 9.3 *Option Lambda, Nine-Month Bond Option*
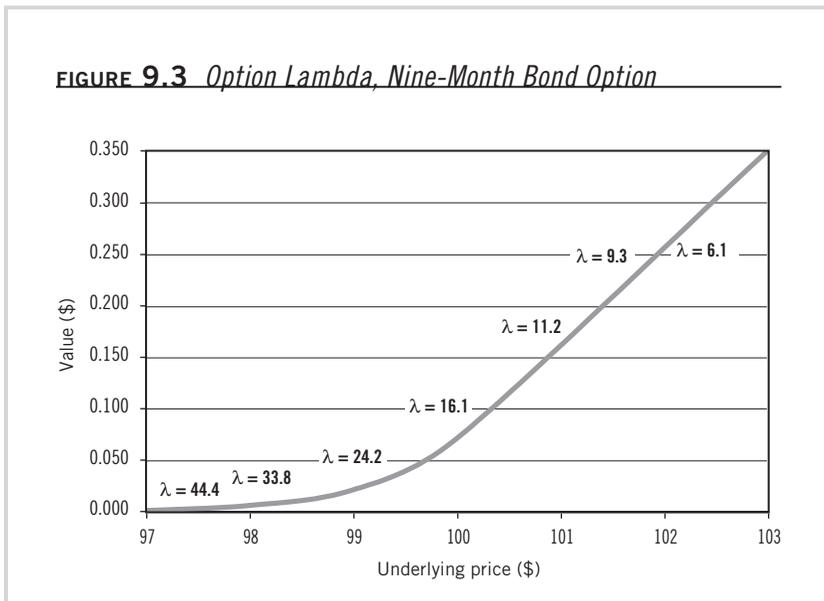
The Option Smile

According to the B-S model, the implied volatility of all options with the same underlying asset and expiry date should be the same, regardless of strike price. In reality, this is not true. The implied volatility observed in the market is a convex function of exercise price. When graphed, this function forms the *volatility,* or *options, smile* shown in generalized form as **FIGURE 9.4**. (Curves graphing actual implied volatilities against strike price are not smooth and often not even true smiles.)

FIGURE 9.4 *The Volatility Smile Curves for Bond and Equity Options*
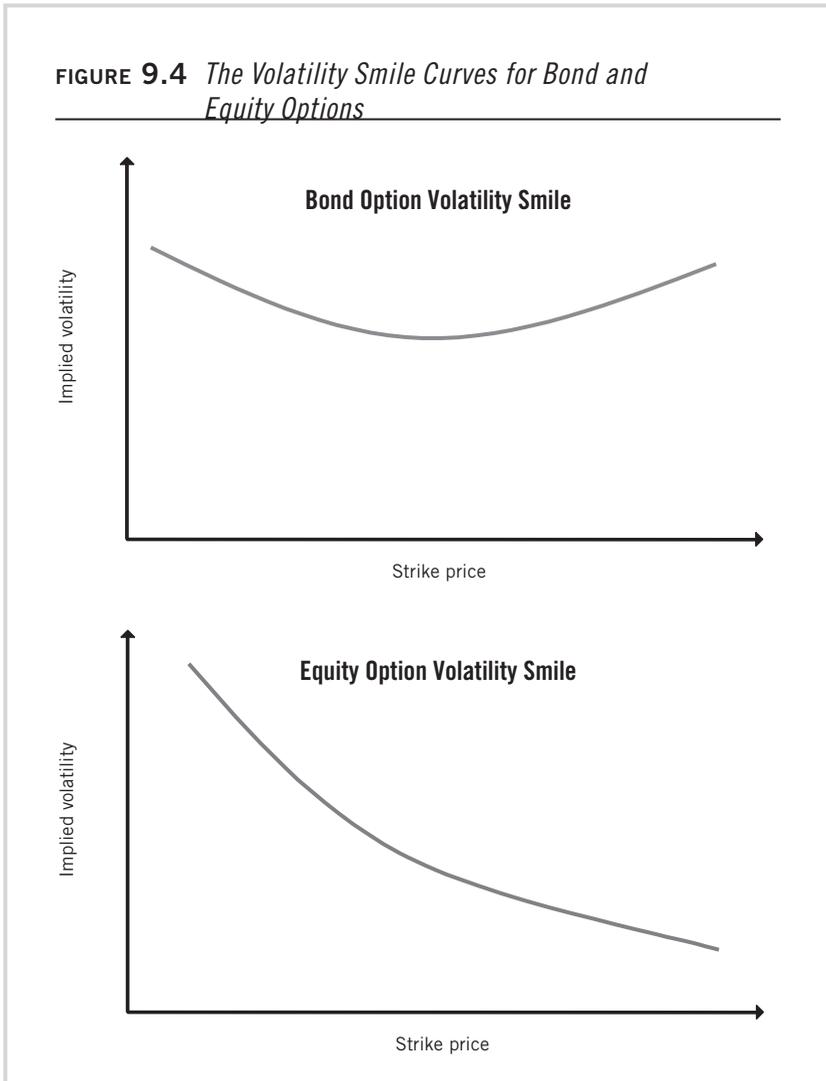
The existence of the smile indicates that market makers price in-the-money options (calls having strikes below the forward price of the underlying asset, *S*, and puts whose strikes are above it) and out-of-the-money ones (calls having strikes higher than *S* and puts having strikes below it) with higher volatilities than they do at-the-money options (those whose strikes equal *S*). This, in turn, suggests that market makers make more complex assumptions about the behavior of asset prices than can be fully explained by the geometric Brownian motion model. Specifically, they attach probabilities to terminal values of the underlying asset price that are inconsistent with a lognormal distribution.

The degree of a smile's convexity indicates how far the market price process differs from the lognormal function contained in the B-S model. The more convex the smile curve, the greater the probability the market attaches to extreme outcomes for the price of the asset on expiry, S_T. In fact, observed asset price returns do follow a distribution with "fatter tails," i.e., with more occurrences at the extremes, than are found in a lognormal distribution. In addition, the direction in which the smile curve slopes reflects the skew of the price-returns function, and a curve with a positive slope indicates a function that is more positively skewed than the lognormal distribution; the opposite is true for a curve with a negative slope.

All this suggests that asset-price behavior is more accurately described by nonstandard price processes, such as the jump diffusion model or a stochastic volatility, than by a model assuming constant volatility. For more-detailed discussion of the volatility smile and its implications, interested readers may consult the works listed in the References section.

Caps and Floors

Caps and floors are options on interest rates such as LIBOR, euribor, U.S. prime, and the commercial paper rate. A cap is a call option on interest rates, written by a market-making bank and sold to the borrower of a cash loan. In return for the premium, the bank agrees that if the reference rate rises above the *cap level,* it will pay the buyer the difference between the two. The cap thus places an upper limit on the rate the borrower must pay. The loan may precede the cap transaction and be with a third party. Alternatively, the cap may be transacted alongside the loan, or as part of it, as a form of interest rate risk management. In that case, the cap's notional amount will equal the loan amount. The cap term can be fairly long to match the loan term—commonly as much as ten years.

Typically, the cap rate is compared with the indexed interest rate on the rate-reset dates—semiannually, for instance, if the reference rate is six-month LIBOR. The cap actually consists of a strip of individual contracts,

called *caplets,* corresponding to each of the reset periods. When the index interest rate is below the cap level, no payment changes hands, except for the interest payment the borrower owes on the loan, computed at the market rate. When the index rate is fixed above the cap level, the cap seller will make a payment computed by applying the difference between the two rates, prorated for the reset period—quarterly, semiannual, and so forth—to the notional amount. Since the payments, like those of FRAs, are made at the beginning of the period covered, they are discounted at the index rate. Equation (9.12) is the payment calculation.

$$Int = \frac{\max[r - rX, 0] \times (N / B) \times M}{1 + r(N / B)} \tag{9.12}$$

where
  r = the index interest rate on the reset date
  rX = the cap level
  M = the notional amount of the cap
  B = the day base (360 or 365)
  N = the number of days in the interest period—that is, the days to the next rate fix

The cap's premium is a function of the probability that it will be exercised, based on the volatility of the forward interest rate. Caps are frequently priced using the Black (1976) model, taking the cap level as the strike and the forward rate as the underlying asset's "price." The resulting equation for computing the premium is (9.13).

$$C = \frac{\phi M}{1 + \phi rf} \times e^{-rT} \times \left[rfN(d_1) - rXN(d_2) \right] \tag{9.13}$$

where
$$d_1 = \frac{\ln(rf / rX)}{\sigma_f \sqrt{T}} + \frac{\sigma_f}{2} \sqrt{T}$$

$$d_2 = d_1 - \sigma_f \sqrt{T}$$

  rf is the forward rate for the relevant term
  ϕ is the number of times a year the rate is fixed—semiannual or quarterly, for instance
  σ_f is the forward-rate volatility
  T is the period from the start of the cap to the next caplet payment date

Each caplet can be priced individually and the results summed to give the total cap premium. The Black model assumes constant volatility, so banks use later models to price products for which this assumption is considered materially unrealistic.

Like caps, floors are series of individual contracts. These are called *floorlets,* and they function essentially as put options on interest rates. Lenders may buy floors to limit their income losses should interest rates fall. A long call cap position combined with a short floor position is a *collar,* so called because it bounds the interest rate payable on the upside at the cap level and on the downside at the floor level. *Zero-cost collars,* in which the cap and floor premiums are identical, are very popular with corporations seeking to manage their interest rate risk.

CHAPTER 10

Credit Derivatives

Unless purchasing what are considered default-free instruments, such as U.S. Treasuries, German bunds, or U.K. gilts, bond investors are exposed to *credit risk.* This is the risk that the debt issuer will default either on servicing the loan—delaying or failing to make the coupon payments, known as a *technical* default—or on paying back the principal, an *actual* default. To hedge this risk, investors may use credit derivatives. These instruments, which were introduced in significant volume only in the mid-1990s, were originally designed to protect banks and other institutions against losses arising from credit events. Today they are used to trade credit and to speculate, as well as for hedging.

Gup and Brooks (1993) noted that swaps' credit risk, unlike their interest rate risk, could not be hedged. That was true in 1993. The situation changed quickly, however, in years following. By 1996 a liquid market existed in instruments designed for just such hedging. Credit derivatives are, in essence, insurance policies against a deterioration in the credit quality of borrowers. The simplest ones even require regular premiums, paid by the protection buyer to the protection seller, and make payouts should a specified credit event occur.

As noted, credit derivatives may be used by investors to manage the extra risk they take on by opting for the higher returns of non–default-free debt. The instruments can also be used, however, to synthesize the exposure itself, if, for instance, compelling reasons exist for not putting on the cash-market position. Since credit derivatives are OTC products, they can be tailored to meet specific requirements.

FIGURE **10.1** *Corporate Bond Credit Ratings*

FITCH	MOODY'S	S&P	SUMMARY DESCRIPTION
Investment Grade			
AAA	Aaa	AAA	Gilt edged, prime, maximum safety, lowest risk, and when sovereign borrower considered "default-free"
AA+	Aa1	AA+	
AA	Aa2	AA	High-grade, high credit quality
AA–	Aa3	AA–	
A+	A1	A+	
A	A2	A	Upper-medium grade
A–	A3	A–	
BBB+	Baa1	BBB+	
BBB	Baa2	BBB	Lower-medium grade
BBB–	Baa3	BBB–	
Speculative Grade			
BB+	Ba1	BB+	
BB	Ba2	BB	Low grade, speculative
BB–	Ba3	BB–	
B+	B1		
B	B	B	Highly speculative
B–	B3		
Predominantly Speculative, Substantial Risk, or in Default			
CCC+		CCC+	
CCC	Caa	CCC	Substantial risk, in poor standing
CC	Ca	CC	May be in default, very speculative
C	C	C	Extremely speculative
		CI	Income bonds—no interest being paid
DDD			
DD			Default
D		D	

Credit Risk

When a technical or actual default occurs, bondholders suffer losses as the value of their assets declines or, in the worst case, disappears entirely. The extent of credit risk varies with the fiscal condition of the borrowers and the health of the greater economy. The magnitude of the risk is usually encapsulated in a formal credit rating, assigned by agencies such as Standard & Poor's, Moody's, and Fitch. The agencies arrive at these ratings, shown in **FIGURE 10.1**, after analyzing a company's business and circumstances. Among the issues they consider are the following:

1 The financial position of the company, as determined by its balance sheet and anticipated cash flows and revenues
   —Other issues, such as the quality of the management and succession planning
   —The company's ability to meet scheduled interest and principal payments both in its domestic and in foreign currencies
2 The outlook for the company's industry as whole and the competition within it
3 The health of the domestic economy

Another indicator of credit risk is the credit risk premium: the spread between the yields on corporate bonds and those of government bonds in the same currency. This spread is the compensation required by investors for holding bonds that are not default-free. The size of the credit premium changes with the market's perception of the financial health of individual companies and sectors and of the economy in general. The variability of the premium is illustrated in **FIGURES 10.2** and **10.3** on the following page, which show the spreads between the U.S.-dollar-swap and Treasury yield curves in, respectively, February 2001 and February 2004.

The credit spread reflects a number of macroeconomic and microeconomic factors, and at any one time is a good snapshot indicator of the perceived health and future prospect of the economy. For example, the reduction in spread from 2001 to 2004 reflects a general increase in favorable perception on the health of the U.S. economy, following the technology and dot-com boom and bust of the previous years and the market disruption following September 11, 2001.

Credit Risk and Credit Derivatives

Credit derivatives are financial contracts designed to reduce or eliminate credit risk exposure by providing insurance against losses suffered because of credit events. The loan or bond carrying the credit risk in question is the ref-

FIGURE **10.2** *Yield Curves for Active U.S. Treasury Securities and U.S. Dollar Swaps on February 2, 2001*
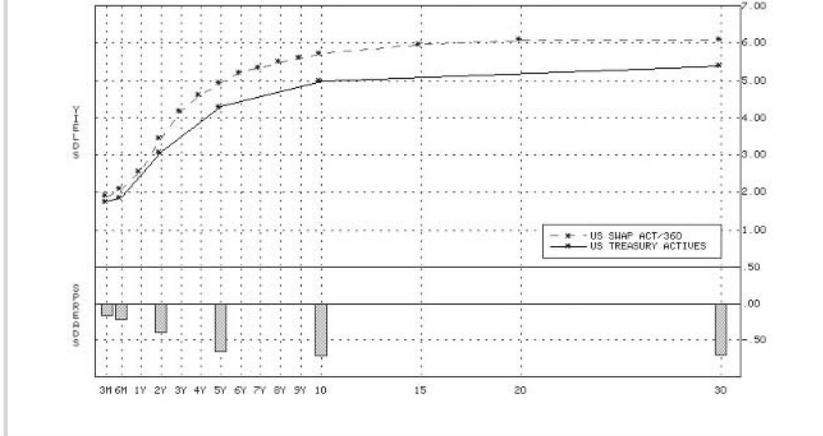
FIGURE **10.3** *Yield Curves for Active U.S. Treasury Securities and U.S. Dollar Swaps on February 3, 2004*
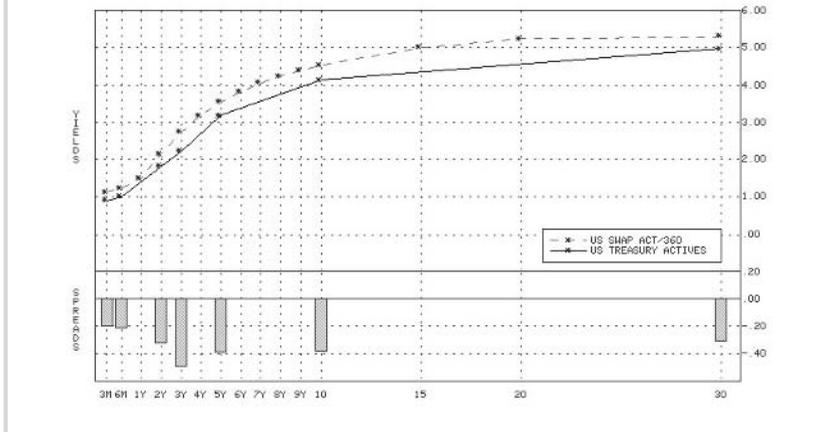
erent, or underlying, asset. The referent asset issuer or borrower is the obligor. The terms of a credit derivative usually include specifications of the type of events that will trigger payouts. These typically include the following:

❑ Bankruptcy or insolvency of the reference asset obligor
❑ Financial restructuring required under bankruptcy protection
❑ Technical default
❑ The bond's credit spread over Treasuries widens beyond a specified level
❑ A downgrade in credit rating below a specified level

Applications of Credit Derivatives

Any institution—including investment and commercial banks, insurance companies, corporations, fund managers, and hedge funds—that is or desires to be exposed to credit risk may use credit derivatives to manage or create it. The credit risk managed may be the one inherent in corporate- or non–AAA-sovereign-bond portfolios or the risk associated with commercial loan books. Indeed, commercial loans were credit derivatives' first area of application. Intense competition to make loans, combined with rapid disintermediation—the lessening of the role of banks as intermediaries—forced commercial banks to reevaluate their lending policies with a view to improving profitability and return on capital. Credit derivatives helped them restructure their businesses by enabling them to repackage and parcel out credit risk while retaining assets on their balance sheets (when required) and thus maintaining client relationships. Credit derivatives are particularly well suited for these and similar purposes for the following reasons:

❑ The issuer of the original debt or borrower of the original loan, known as the *reference entity,* need not be a party to the credit risk transfer and in fact is usually not even aware of the transaction, thus allowing the client relationship between the lending bank and the borrower to be maintained.

❑ The credit derivative can be tailored to the requirements of the protection buyer, as opposed to the liquidity or term of the underlying loan.

❑ Credit derivatives can be used to synthesize the economic effect of selling a bank loan short—a transaction not possible in the cash market—and do so theoretically without the risk of a liquidity or delivery squeeze, since a specific credit risk is being traded.

❑ Because the derivatives isolate credit risk from other factors, such as client relationships and interest rate risk, they offer a formal mechanism for pricing credit issues only. A market in credit alone can thus evolve, allowing still more efficient pricing and the modeling of a term structure of credit rates.

❑ Unless embedded in fixed-income products, such as structured or credit-linked notes, the derivatives are off the balance sheet. This status endows them with tremendous flexibility and leverage, characteristics they

share with other financial derivatives. For instance, bank loans are often deemed unattractive as investments because of the administration that managing and servicing a loan portfolio requires. Investors can acquire exposure to bank loans' returns while avoiding the administrative costs through, say, a total return swap. The same transaction allows banks to distribute their loan credit risk.

Credit derivatives are important tools not only for commercial banks but also for bond portfolio managers, who can use them to increase the liquidity of their portfolios, profit from credit-pricing anomalies, and enhance returns. Credit derivatives separate the ownership and management of the credit risk associated with the assets in question from ownership of their other features. This means that banks can transfer the credit risk exposures of illiquid assets, such as bank loans and illiquid bonds, thus protecting themselves against credit loss, even if they cannot transfer the assets themselves. Among the reasons banks use credit derivatives are the following:

To diversify their credit portfolios. Banks may wish to take on additional credit exposure by selling credit protection on assets they own to other banks or investors, thus enhancing their portfolio returns. They may sell derivatives to non-bank clients who don't want to buy the associated assets directly but do want exposure to the credit risk of the assets. In such a transaction, the bank acts as a credit intermediary.

To reduce their credit exposure. By buying credit default swaps, banks can reduce their credit exposure on individual assets or sector concentrations. This tactic is especially useful when the positions in question cannot be sold because of relationship or tax issues.

To act as market makers or traders in credit derivatives. Credit derivative traders may or may not hold the reference assets directly, depending on their appetite for risk and the liquidity of the market they would need to use to hedge their derivative contracts.

Credit Derivative Instruments

Credit Default Swap

The most common credit derivative, and possibly the simplest, is the *credit default swap*—also known as the *credit* or *default* swap. As diagrammed in **FIGURE 10.4**, it is a bilateral contract in which a *protection seller,* or *guarantor,* in return for a periodic fixed fee or a onetime premium agrees to pay the *beneficiary* counterparty in case any of a list of specified credit events occurs. The fee is usually quoted as a percentage of the nominal value of the reference asset or basket of assets. The swap term does not have to

FIGURE 10.4 *Credit Default Swap*

Bank A
"Beneficiary"

Bank B
"Guarantor"

Fee or premium

Default payment on triggering event
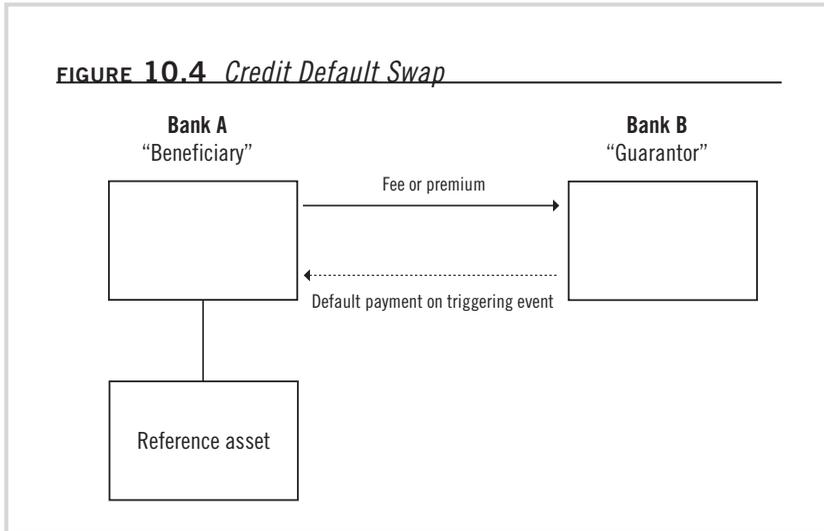
Reference asset

match the maturity of the reference asset and in most cases does not.

In a default, the swap is terminated, and the default payment is calculated and handed over. The amount of this payment may be linked to the change in price of the reference asset or another specified asset or fixed at a predetermined recovery rate. Alternatively, it may involve actual delivery of the reference asset at a specified price.

Banks may use default swaps to trade sovereign and corporate credit spreads without trading the actual assets. The party long the swap—the protection buyer—need not own or ever have owned the reference asset. Nevertheless, he or she profits should the obligor suffer a rating downgrade or perceived reduction in credit quality. This is because such an event increases the cost of protecting the asset, so the swap buyer can sell the position at a profit, assuming a purchaser can be found.

Credit Options

Credit options are also bilateral OTC financial contracts. Like other options, they may be designed to meet specific hedging or speculative requirements. Also like other options, they may be plain vanilla or exotic. The buyer of a vanilla, or *standard,* credit call option has the right, but no obligation, to purchase the underlying credit-sensitive asset or credit spread at a specified price and at a specified time or during a specified period. The buyer of a vanilla credit put option has the right, but no obligation, to sell the underlying credit-sensitive asset or credit spread.

In exotic credit options, one or more parameters differ from the vanilla norm. A *barrier* credit option, for example, specifies a credit event

that would trigger the option or inactivate it. A *digital* credit option has a binary payout: if it is at or out of the money at expiration, it pays zero; otherwise, it pays a fixed amount, no matter how far in the money it is.

Bond investors can use credit options to hedge against rating downgrades and similar events that would depress the value of their holdings. To ensure that any loss resulting from such events will be offset by a profit on their options, they purchase contracts whose payoff profiles reflect their bonds' credit quality. The options also enable banks and other institutions to take positions on credit spread movements without taking ownership of the related loans or bonds. The writer of credit options earns fee income.

Credit options allow market participants to express their views on credit alone, without reference to other factors, such as interest rates, with no cost beyond the premium. For example, investors who believe that the credit spread associated with an individual entity or a sector (such as all AA-rated sterling corporates) will widen over the next six months can buy six-month call options on the relevant spread. If the spread widens beyond the strike during the six months, the options will be in the money, and the investors will gain. If not, the investors' loss will be limited to the premium paid.

Credit-Linked Notes

Credit-linked notes, or CLNs, are known as *funded* credit derivatives, because the protection seller pays the entire notional value of the contract up front. In contrast, credit default swaps pay only in case of default and are therefore referred to as *unfunded.* CLNs are often used by borrowers to hedge against credit risk and by investors to enhance their holdings' yields.

Credit-linked notes are hybrid securities, generally issued by an investment-grade entity, that combine a credit derivative with a vanilla bond. Like a vanilla bond, a standard CLN has a fixed maturity structure and pays regular coupons. Unlike bonds, all CLNs, standard or not, link their returns to an underlying asset's credit-related performance, as well as to the performance of the issuing entity. The issuer, for instance, is usually permitted to decrease the principal amount if a credit event occurs. Say a credit card issuer wants to fund its credit card loan portfolio by issuing debt. To reduce its credit risk, it floats a 2-year credit-linked note. The note has a face value of 100 and pays a coupon of 7.50 percent, which is 200 basis points above the 2-year benchmark. If more than 10 percent of its cardholders are delinquent in making payments, however, the note's redemption payment will be reduced to $85 for every $100 of face value. The credit card issuer has in effect purchased a credit option that lowers its liability should it suffer a specified credit event—in this case, an above-expected incidence of bad debts.

Why would investors purchase such a note? Because its coupon is higher than the one the credit card bank would pay on a vanilla bond and, presumably, higher than the rates for many other investments in the market. In addition, such notes are usually issued below par, so if they are redeemed at par, investors realize a substantial capital gain.

Total Return Swaps

A *total return swap,* or TRS—also known as a *total rate of return swap*—is one of the principal instruments used by banks and other financial instruments to manage their credit risk exposure. As defined in Francis et al (1999), page 29, it is an agreement between two parties to exchange the total return from the reference asset or basket of assets—bank loans or credit-sensitive securities such as corporate loans or sovereign or corporate bonds—for some other fixed or floating cash flow, usually tied to LIBOR or other loans or credit-sensitive securities, in the process transferring the credit risk from one party to the other. The TRS differs from other credit derivatives in that the payments between its counterparties are connected to changes in the market value of the underlying asset as well as to changes resulting from a credit event.

In some versions of a TRS, the underlying asset is sold to the counterparty, with a corresponding swap transacted on the side; in other versions the underlying asset remains unsold. The term of the TRS need not match the maturity of the underlying security and, in fact, rarely does.

FIGURE 10.5 diagrams a generic TRS. The two counterparties are la-

FIGURE 10.5 *Total Return Swap*

Bank A
Total return payer
or "Beneficiary"

Bank B
Total return receiver
or "Guarantor"

Bank A

Total return (interest and appreciation) →

Bank B

← LIBOR + spread, plus depreciation

Underlying asset
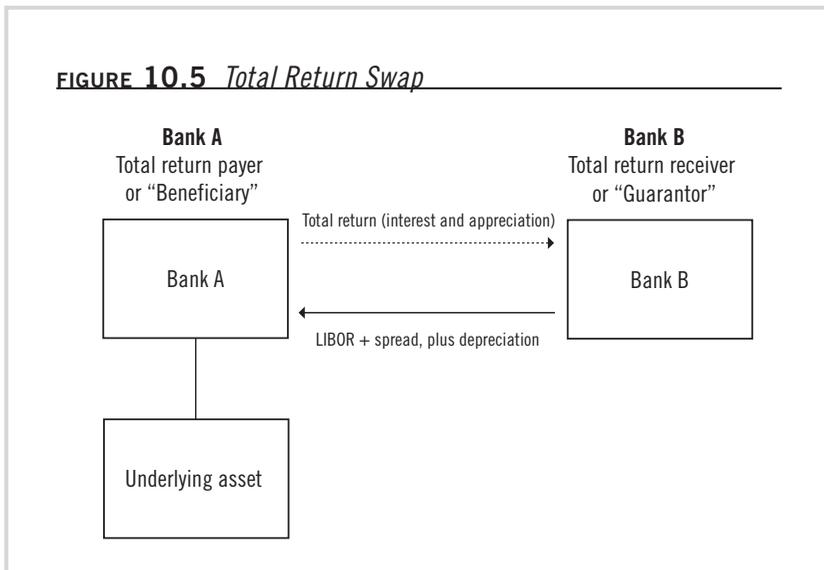
beled as banks, but the total return payer, or beneficiary, may be any financial institution, including an insurance company or hedge fund. In figure 10.5, Bank A, the beneficiary, has contracted to pay the total return—interest payments plus any appreciation in market value—on the reference asset. The appreciation may be cash settled, or Bank A may take physical delivery of the reference asset at swap maturity, paying Bank B, the total return receiver, the initial asset value. Bank B pays Bank A a margin over LIBOR and makes up any depreciation that occurs in the price of the asset—hence the label "guarantor."

The economic effect for Bank B is that of owning the underlying asset. TR swaps are thus synthetic loans or securities. Significantly, the beneficiary usually (though not always) holds the underlying asset on its balance sheet. The TR swap can thus be a mechanism for removing an asset from the guarantor's balance sheet for the term of the agreement.

The swap payments are usually quarterly or semiannual. On the interest-reset dates, the underlying asset is marked to market, either using an independent source, such as Bloomberg or Reuters, or as the average of a range of market quotes. If the reference asset obligor defaults, the swap may be terminated immediately, with a net present value payment changing hands and each counterparty liable to the other for accrued interest plus any appreciation or depreciation in the asset value. Alternatively, the swap may continue, with each party making appreciation or depreciation payments as appropriate. The second option is available only if a market exists for the asset, an unlikely condition in the case of a bank loan. The terms of the agreement typically give the guarantor the option of purchasing the underlying asset from the beneficiary and then dealing directly with the loan defaulter.

Banks and other financial institutions may have a number of reasons for entering into TR swap arrangements. One is to gain off-balance-sheet exposure to the reference asset without having to pay out the cash that would be required to purchase it. Because the swap maturity rarely matches that of the asset, moreover, the swap receiver may benefit, if the yield curve is positive, from positive *carry*—that is, the ability to roll over the short-term funding for a longer-term asset. Higher-rated banks that can borrow at Libid can benefit by funding on-balance-sheet credit-protected assets through a TR swap, assuming the net spread of asset income over credit-protection premium is positive.

The swap payer can reduce or remove credit risk without selling the relevant asset. In a vanilla TR swap, the total return payer retains rights to the reference asset, although, in some cases, servicing and voting rights may be transferred. At swap maturity, the swap payer can reinvest the asset, if it still owns it, or sell it in the open market. The swap can thus be

considered a synthetic repo—that is, an arrangement in which the holder of a bond (usually a government issue) sells it to a lender, promising to buy it back a short time later at an agreed-upon price that gives the lender a low-risk rate of return, termed the *repo rate.*

Total return swaps are increasingly used as synthetic repo instruments, most commonly by investors who wish to purchase the credit exposure of an asset without purchasing the asset itself. This is similar to what happened when interest rate swaps were introduced, enabling banks and other financial institutions to trade interest rate risk without borrowing or lending cash funds. Banks usually enter into synthetic repos to remove assets from their balance sheets temporarily. The reason may be that they are due to be analyzed by credit-rating agencies, or their annual external audit is imminent, or they are in danger of breaching capital limits between quarterly return periods. In the last case, as the return period approaches, lower-quality assets may be removed from the balance sheet by means of a TR swap with, say, a two-week term that straddles the reporting date. Bonds sold as part of a TR swap transaction are removed from the seller's balance sheet because the bank selling the assets is not legally required to repurchase them from the swap counterparty, nor is the total return payer obliged to sell them back—or indeed to sell them at all—at swap maturity.

TR swaps may also be used for speculation. Bond traders who believe that a particular bond not currently on their books is about to decline in price have a couple of ways to profit from this view. One method is to sell the bond short and cover their position through a repo. The cash flow to the traders from this transaction consists of the coupon on the bond that they owe as a result of the short sale and, if the shorted bond falls in price as expected, the capital gain from the short sale plus the repo rate—say, LIBOR plus a spread. The danger in this transaction is that if the shorted bond must be covered through a repo at the *special* rate instead of the higher *general collateral* rate—the one applicable to Treasury securities— the traders will be funding it at a loss. The yield on the bond must also be lower than the repo rate.

Alternatively, the traders can enter into a TR swap in which they pay the total return on the bond and receive LIBOR plus a spread. If the bond yield exceeds the LIBOR payment, the funding will be negative, but the trade will still gain if the bond falls in price by a sufficient amount. The traders will choose this alternative if the swap's break-even point— the price to which the bond must decline for a gain from the short sale to offset the trade's funding cost—is higher than in the repo approach. This is more likely if the bond is special.

Investment Applications

This section explores the ways bond investment managers typically use credit derivatives.

Capital Structure Arbitrage

In capital structure arbitrage, investors exploit yield mismatches between two loans from the same reference entity. Say an issuer has two debt instruments outstanding: a commercial bank loan paying 330 basis points over LIBOR and a subordinated bond issue paying LIBOR plus 230 basis points. This yield anomaly can be exploited with a total return swap in which the arbitrageur effectively purchases the bank loan and sells the bond short.

The swap is diagrammed in **FIGURE 10.6**. The arbitrageur receives the total return on the bank loan and pays the counterparty bank the bond return plus an additional 30 basis points, the price of the swap. These rates are applied to notional amounts of the loan and bond set at a ratio of 2 to 1, since the bond's price is more sensitive to changes in credit status than that of the loan.

The swap generates a net spread of 200 basis points as shown below: + [(100 bps × ½) + (250 bps × ½)]. That is,

$$\text{Receive Loan } CF = (+\text{Libor}+330)\times 1(\text{Loan Notional Principal})$$
$$\text{Bond } CF = (-\text{Libor}+230+30)\times 0.5(\text{Bond Notional Principal})$$

Exposure to Market Sectors

To gain exposure to sectors where, for various reasons, they do not wish to make actual purchases, investors can use a variation on a TR swap called an *index* swap, in which one of the counterparties pays a total return tied to an external reference index and the other pays a LIBOR-linked coupon or the total return of another index. Indexes used include those for government bonds, high-yield bonds, and technology stocks. Investors who believe that the bank loan market will outperform the mortgage-backed bond sector, for instance, might enter into an index swap in which they pay the total return of the mortgage index and receive the total return of the bank-loan index.

Credit Spreads

Credit default swaps can be used to trade credit spreads. Say investors believe the credit spread between certain emerging-market government bonds and U.S. Treasuries is going to widen. The simplest way to exploit

FIGURE 10.6 *Total Return Swap Used in Capital Structure Arbitrage*
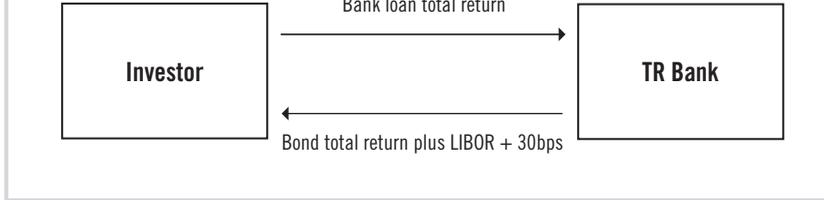
this view would be to go long a credit default swap on the emerging-market bonds paying 600 basis points. If the investors' view is correct and the bonds' credit spread widens, depressing their price, the premium payable on the credit swap will increase. The investors will then be able to sell their swap in the market at the higher premium.

Funding Positions

Investment banks and hedge funds often use TRS contracts to pay for positions in securities that they cannot—for operational, credit, or other reasons—fund using the interbank market or a classic repo. The TRS counterparty that is long the security swaps it with a counterparty that provides the money to pay for the asset in the market. This money is in effect a loan to the asset seller, at a cost of LIBOR plus a spread. During the swap term, the funds provider pays the asset seller the coupon/interest on the asset. On the swap-reset or maturity date, the asset is marked to market. If it has increased in value, the funds provider will pay the asset seller the difference; if it has fallen, the asset seller will pay the difference. The asset seller also pays the LIBOR-plus interest on the initial swap proceeds. In addition to funding the asset, this transaction removes it from the original holder's balance sheet, transferring it to the counterparty's for the term of the swap. As an illustration of how this works, consider the hypothetical transaction in **FIGURE 10.7** on the following page.

Note that the "haircut" is the amount of the bond value that is not handed over in the loan proceeds—it acts as a credit protection to the provider of funds in the event that the bond, which is in effect the collateral for the loan, drops in value during the term of the swap. For ease of illustration the haircut in this example is 0 percent so the loan is the full value of the bond collateral.

FIGURE 10.7 *Using a Total Return Swap to Fund a Security*

Party A: Asset seller, a hedge fund

Party B: Funds provider, an investment bank

Asset: $50 million face value of an A-rated asset-backed security

Price: 99.50

Asset value: $49,750,000

Date when swap is transacted: February 3, 2004

Value date: February 6, 2004

Swap term: 14 days

Maturity date: February 20, 2004

Haircut: 0 percent

Initial swap proceeds: $49,750,000

Floating rate paid by the asset seller: LIBOR plus 35 basis points

2-week LIBOR at start of swap: 1.08125 percent

Swap floating rate: 1.43125 percent

Assume that, at swap maturity, the stock has risen in price to 99.70, making its market value $49,850,000. The investment bank would therefore owe the hedge fund a "performance payment" of $100,000. The hedge fund, meanwhile, owes the investment bank $27,690.71, which is 1.43125 percent interest on $49,750,000 for fourteen days. The bank pays the fund a net payment of $72,309. If there was a coupon payment during the term, it would be paid by the investment bank as part of the performance payment to the hedge fund.

Credit-Derivative Pricing

Banks employ a number of methods to price credit derivatives. This section presents a quick overview. Readers wishing a more in-depth discussion should consult the references listed for this chapter in the References section.

Credit-derivative pricing is similar to the pricing of other off-balance-sheet products, such as equity, currency, and bond derivatives. The main difference is that the latter can be priced and hedged with reference to the underlying asset, and credit derivatives cannot. The pricing model for credit products incorporates statistical data concerning the likelihood of default, the probability of payout, and market level of risk tolerance.

Pricing Total Return Swaps

The guarantor in a TR swap usually pays the beneficiary a spread over LIBOR. Pricing in this case means determining the size of the LIBOR spread. This spread is a function of the following factors:

❑ The credit rating of the beneficiary
❑ The credit quality of the reference asset
❑ The face amount and value of the reference asset
❑ The funding costs of the beneficiary bank
❑ The required profit margin
❑ The capital charge—the amount of capital that must be held against the risk represented by the swap—associated with the TR swap

Related to these factors are several risks that the guarantor must take into account. One crucial consideration is the likelihood of the TR swap receiver defaulting at a time when the reference asset has declined in value. This risk is a function both of the financial health of the swap receiver and of the market volatility of the reference asset. A second important consideration is the probability of the reference asset obligor defaulting, triggering a default by the swap receiver before the swap payer receives the depreciation payment.

Asset-Swap Pricing

Asset-swap pricing is commonly applied to credit-default swaps, especially by risk management departments seeking to price the transactions held on credit traders' book. A par asset swap typically combines an interest rate swap with the sale of an asset, such as a fixed-rate corporate bond, at par and with no interest accrued. The coupon on the bond is paid in return for LIBOR plus, if necessary, a spread, known as the *asset-swap spread.* This spread is the price of the asset swap. It is a function of the credit risk of the underlying asset. That makes it suitable as the basis for the price payable on a credit default swap written on that asset.

The asset swap spread is equal to the underlying asset's redemption yield spread over the government benchmark, minus the spread on the associated interest rate swap. The latter, which reflects the cost of convert-

ing the fixed-rate coupons of benchmark bonds to a floating rate during the life of the asset, or the default swap, is based on the swap rate for the relevant term.

Credit-Spread Pricing Models

Practitioners increasingly model credit risk as they do interest rates and use spread models to price associated derivatives. One such model is the Heath-Jarrow-Morton (HJM) model described in chapter 4. This analyzes interest rate risk, default risk, and recovery risk—that is, the rate of recovery on a defaulted loan, which is always assumed to retain some residual value.

The models analyze spreads as wholes, rather than splitting them into default risk and recovery risk. Das (1999), for example, notes that equation (10.1) can be used to model credit spreads. Credit options can thus be analyzed in the same way as other types of options, modeling the credit spread rather than, say, the interest rate.

$$ds = k(\vartheta - s)\,dt + \sigma \sqrt{s}\,dZ \qquad (10.1)$$

where

s = the credit spread over the government benchmark
ds = change in the spread over an infinitesimal change in time
k = the mean reversion rate of the credit spread
θ = the mean of the spread over t
σ = the volatility of the spread
dt = change in time
dZ = standard Brownian motion or Weiner process

For more detail on modeling credit spreads to price credit derivatives, see Choudhry (2004).

The Analysis of Bonds with Embedded Options

The yield calculation for conventional bonds is relatively straight-forward. This is because their redemption dates are fixed, so their total cash flows—the data required to calculate yield to maturity—are known with certainty. Less straightforward to analyze are bonds with *embedded options*—calls, puts, or sinking funds—so called because the option element cannot be separated from the bond itself. The difficulty in analyzing these bonds lies in the fact that some aspects of their cash flows, such as the timing or value of their future payments, are uncertain.

Because a callable bond has more than one possible redemption date, its future cash flows are not clearly defined. To calculate the yield to maturity for such a bond, it is necessary to assume a particular redemption date. The market convention is to use the earliest possible one if the bond is priced above par and the latest possible one if it is priced below par. Yield calculated in this way is sometimes referred to as *yield to worst* (the Bloomberg term).

If a bond's actual redemption date differs from the assumed one, its return computed this way is meaningless. The market, therefore, prefers to use other methods to calculate the return of callable bonds. The most common method is *option-adjusted spread,* or OAS, *analysis.* Although the discussion in this chapter centers on callable bonds, the principles enunci-ated apply to all bonds with embedded options.

Understanding Option Elements Embedded in a Bond

Consider a callable U.S.-dollar corporate bond issued on December 1, 1999, by the hypothetical ABC Corp. with a fixed semiannual coupon of

FIGURE **11.1** *Call Schedule for the ABC Corp. 6 Percent Bond Due December 2019*

DATE	CALL PRICE
01–Dec–2004	103.00
01–Dec–2005	102.85
01–Dec–2006	102.65
01–Dec–2007	102.50
01–Dec–2008	102.00
01–Dec–2009	101.75
01–Dec–2010	101.25
01–Dec–2011	100.85
01–Dec–2012	100.45
01–Dec–2013	100.25
01–Dec–2014	100.00

6 percent and a maturity date of December 1, 2019. **FIGURE 11.1** shows the bond's call schedule, which follows a form common in the debt market. According to this schedule, the bond is first callable after five years, at a price of $103; after that it is callable every year at a price that falls progressively, reaching par on December 1, 2014, and staying there until maturity.

The call schedule works like this. If market interest rates rise after the bonds are issued, ABC Corp. gains, because it is incurring below-market financing costs on its debt. If rates decline, investors gain, because the value of their investment rises. Their upside, however, is capped at the applicable call price by the call provisions, since the issuer will redeem the bond if it can reduce its funding costs by doing so.

Basic Options Features

An option is a contract between two parties: the option buyer and the option seller. The buyer has the right, but not the obligation, to buy or sell an underlying asset at a specified price during a specified period or at a specified time (usually the expiry date of the contract). The price of an option

is known as its *premium,* which is paid by the buyer to the seller, or *writer.* An option that grants the holder the right to buy the underlying asset is known as a *call* option; one that grants the right to sell the underlying asset is a *put* option. The option writer is short the contract; the buyer is long.

If the owner of an option elects to *exercise* it and enter into the underlying trade, the option writer is obliged to execute under the terms of the contract. The price at which an option specifies that the underlying asset may be bought or sold is the exercise, or *strike,* price. The expiry date of an option is the last day on which it may be exercised. Options that can be exercised anytime from the day they are struck up to and including the expiry date are called *American* options. Those that can be exercised only on the expiry date are known as *European* options.

The profit-loss profile for option buyers is quite different from that for option sellers. Buyers' potential losses are limited to the option premium, while their potential profits are, in theory, unlimited. Sellers' potential profits are limited to the option premium, while their potential losses are, in theory, unlimited; at the least, they can be very substantial. (For a more in-depth discussion of options' profit-loss profile, see chapter 8.)

Option Valuation

The References section contains several works on the technical aspects of option pricing. This section introduces the basic principles.

An option's value, or price, is composed of two elements: its *intrinsic value* and its *time value.* The intrinsic value is what the holder would realize if the option were exercised immediately—that is, the difference between the strike price and the current price of the underlying asset. To illustrate, if a call option on a bond has a strike price of $100 and the underlying bond is currently trading at $103, the option has an intrinsic value of $3. The holder of an option will exercise it only if it has intrinsic value. The intrinsic value is never less than zero. An option with intrinsic value greater than zero is *in the money.* An option whose strike price is equal to the price of the underlying is *at the money;* one whose strike price is above (in the case of a call) or below (in the case of a put) the underlying's price is *out of the money.*

An option's time value is the difference between its intrinsic value and its premium. Stated formally,

$$\text{Time value} = \text{Premium} - \text{Intrinsic value}$$

The premium of an option with zero intrinsic value is composed solely of time value. Time value reflects the potential for an option to move more deeply into the money before expiry. It diminishes up to the option's

expiry date, when it becomes zero. The price of an option on expiry is composed solely of intrinsic value.

The main factors determining the price of an option on an interest rate instrument such as a bond are listed below. (Their effects will differ depending on whether the option in question is a call or a put and whether it is American or European.)

- ❏ the option's strike price
- ❏ the underlying bond's current price and its coupon rate
- ❏ the time to expiry
- ❏ the short-term risk-free rate of interest during the life of the option
- ❏ the expected volatility of interest rates during the life of the option

A number of option-pricing models exist. Market participants often use variations on these models that they developed themselves or that were developed by their firms. The best-known of the pricing models is probably the Black-Scholes, whose fundamental principle is that a synthetic option can be created and valued by taking a position in the underlying asset and borrowing or lending funds in the market at the risk-free rate of interest. Although Black-Scholes is the basis for many other option models and is still used widely in the market, it is not necessarily appropriate for some interest rate instruments. Fabozzi (1997), for instance, states that the Black-Scholes model's assumptions make it unsuitable for certain bond options. As a result a number of alternatives have been developed to analyze callable bonds.

The Call Provision

A bond with early redemption provisions is essentially a portfolio consisting of a conventional bond having the same coupon and maturity and a put or call option on this bond. The value of the bond is the sum of the values of these "portfolio" elements. This is expressed formally as (11.1).

$$P_{bond} = P_{underlying} \pm P_{option} \qquad (11.1)$$

For a conventional bond, the value of the option component is zero. For a putable one, the option has a positive value. The portfolio represented by a putable bond contains a long position in a put, which, by acting as a floor on the bond's price, increases the bond's attractiveness to investors. Thus the greater the value of the put, the greater the value of the bond. This is expressed in (11.2).

$$P_{pbond} = P_{underlying} + P_{put} \qquad (11.2)$$

A callable bond is essentially a conventional bond plus a short position in a call option, which acts as a cap on the bond's price and so reduces its value. If the value of the call option were to increase because of a fall in interest rates, therefore, the value of the callable bond would decrease. This is expressed in (11.3).

$$P_{cbond} = P_{underlying} - P_{call} \qquad (11.3)$$

The difference between the price of the option-free bond and the callable bond at any time is the price of the embedded call option. The behavior of the option element depends on the terms of the callable issue.

If the issuer of a callable bond is entitled to call it at any time after the first call date, the bondholder has effectively sold the issuer an American call option. However, as figure 11.1 illustrates, the redemption value may vary with the call date. This is because the value of the underlying bond at the time the call is exercised is composed of the sum of the present values of the remaining coupon payments that the bondholder would have received had the issue not been called. Of course, the embedded option does not trade on its own. Nevertheless, it is clear that embedded options influence significantly not only a bond's behavior but its valuation as well.

The Binomial Tree of Short-Term Interest Rates

Chapter 3 discussed how a coupon-bond yield curve could be used to derive spot (zero-coupon) and implied forward rates. A forward rate is the interest rate for a term beginning at a future date and maturing one period later. Forward rates form the basis of binomial interest rate trees.

Any models using implied forward rates to generate future prices for options' underlying bonds would be assuming that the future interest rates implied by the current yield curve will actually occur. An analysis built on this assumption would, like yield-to-worst analysis, be inaccurate, because the yield curve does not remain static and neither do the rates implied by it; therefore future rates can never be known with certainty. To avoid this inaccuracy, a binomial tree model assumes that interest rates fluctuate over time. These models treat implied forward rates, sometimes referred to as *short rates,* as outcomes of a binomial process, resulting in a binomial tree of possible short rates for each future period. A binomial tree is constructed by starting from a known interest rate at period 0, and assuming that during the following period, rates can travel along two possible paths, each resulting in a different rate one period forward; those two future rates each serve as the origin for two more paths, resulting in four possible rates at the end

of the next period forward, and so on. The tree is called *binomial* because at each future state, or node, there are precisely two possible paths ending in two possible interest rates for the next period forward.

Arbitrage-Free Pricing

Assume that the current six-month and one-year rates are 5.00 and 5.15 percent, respectively. Assume further that six months from now the six-month rate will be either 5.01 or 5.50 percent, and that each rate has a 50 percent probability of occurring. Bonds in this hypothetical market pay semiannual coupons, as they do in the U.S. and U.K. domestic markets. This situation is illustrated in **FIGURE 11.2**.

Figure 11.2 is a one-period binomial interest rate tree, or lattice, for the six-month interest rate. From this lattice, the prices of six-month and 1-year zero-coupon bonds can be calculated. As discussed in chapter 3, the current price of a bond is equal to the sum of the present values of its future cash flows. The six-month bond has only one future cash flow: its redemption payment at face value, or 100. The discount rate to derive the present value of this cash flow is the six-month rate in effect at point 0. This is known to be 5 percent, so the current six-month zero-coupon bond price is $100/(1 + [0.05/2])$, or 97.56098. The price tree for the six-month zero-coupon bond is shown in **FIGURE 11.3**.

For the six-month zero-coupon bond, all the factors necessary for pricing—the cash flow and the discount rate—are known. In other words, only one "world state" has to be considered. The situation is different for the one-year zero coupon, whose binomial price lattice is shown in **FIGURE 11.4**.

Deriving the one-year bond's price at period 0 is straightforward. Once again, there is only one future cash flow—the period 2 redemption payment at face value, or 100—and one possible discount rate: the one-year interest rate at period 0, or 5.15 percent. Accordingly, the price of the one-year zero-coupon bond at point 0 is $100/(1 + [0.0515/2]^2)$, or 95.0423. At period 1, when the same bond is a six-month piece of paper, it has two possible prices, as shown in figure 11.4, which correspond to the two possible six-month rates at the time: 5.50 and 5.01 percent. Since each interest rate, and so each price, has a 50 percent probability of occurring, the average, or *expected* value, of the one-year bond at period 1 is $[(0.5 \times 97.3236) + (0.5 \times 97.5562)]$, or 97.4399.

Using this expected price at period 1 and a discount rate of 5 percent (the six-month rate at point 0), the bond's present value at period 0 is $97.4399/(1 + 0.05/2)$, or 95.06332. As shown above, however, the market price is 95.0423. This demonstrates a very important principle in financial economics: markets do not price derivative instruments based on their expected future value. At period 0, the one-year zero-coupon bond is a

FIGURE **11.2** *The Binomial Price Tree for the Six-Month Interest Rate*
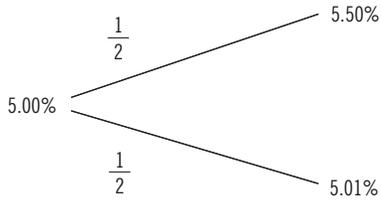
FIGURE **11.3** *The Binomial Price Tree for the Six-Month Zero-Coupon*
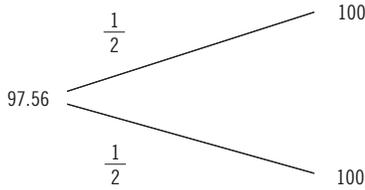
FIGURE **11.4** *The Binomial Price Tree for the One-Year Zero-Coupon*
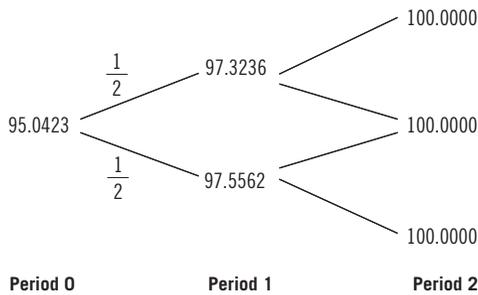
riskier investment than the shorter-dated six-month zero-coupon bond. The reason it is risky is the uncertainty about the bond's value in the last six months of its life, which will be either 97.32 or 97.55, depending on the direction of six-month rates between periods 0 and 1. Investors prefer certainty. That is why the period 0 present value associated with the single estimated period 1 price of 97.4399 is higher than the one-year bond's actual price at point 0. The difference between the two figures is the *risk premium* that the market places on the bond.

Options Pricing

Assume now that the one-year zero-coupon bond in the example has a call option written on it that matures in six months (at period 1) and has a strike price of 97.40. **FIGURE 11.5** is the binomial tree for this option, based on the binomial lattice for the one-year bond in figure 11.4. The figure shows that at period 1, if the six-month rate is 5.50 percent, the call option has no value, because the bond's price is below the strike price. If, on the other hand, the six-month rate is at the lower level, the option has a value of 97.5562 – 94.40, or 0.1562.

What is the value of this option at point 0? Option pricing theory states that to calculate this, you must compute the value of a *replicating portfolio.* In this case, the replicating portfolio would consist of six-month and one-year zero-coupon bonds whose combined value at period 1 will be zero if the six-month rate rises to 5.50 percent and 0.1562 if the rate at that time is 5.01 percent. It is the return that is being replicated. These conditions are stated formally in equations (11.4) and (11.5), respectively.

$$C_1 + 0.973236C_2 = 0 \qquad\qquad (11.4)$$

$$C_1 + 0.975562 = 0.1562 \qquad\qquad (11.5)$$

FIGURE 11.5 *The Binomial Tree for a Six-Month Call Option Written on the One-Year Zero-Coupon Bond*
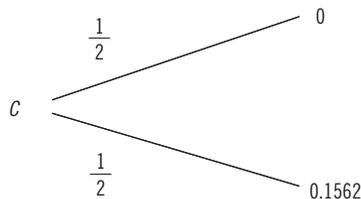
where
   C_1 = the face amount of the six-month bond at period 1
   C_2 = the face amount of the one-year bond at period 1

Since the six-month zero-coupon bond in the replicating portfolio matures at period 1, it is worth 100 percent of face value at point 1, no matter where interest rates stand. The value of the one-year at period 1, when it is a six-month bond, depends on the interest rate at the time. At the higher rate, it is 97.3236 percent of face; at the lower rate, it is 97.5562 percent of face. The two equations state that the total value of the portfolio must equal that of the option, which at the higher interest rate is zero and at the lower 0.1562.

Solving the two equations gives C_1 = –65.3566 and C_2 = 67.1539. This means that to construct the replicating portfolio, you must purchase 67.15 of one-year zero-coupon bonds and sell short 65.36 of the six-month zero-coupon bond. The reason for constructing the portfolio, however, was to price the option. The portfolio and the option have equal values. The portfolio value is known: it is the price of the six-month bond at period 0 multiplied by C_1, plus the price of the one-year bond multiplied by C_2, or

$$(0.9756 \times -65.3566) + (0.950423 \times 67.1539) = 0.0627 \qquad (11.6)$$

The result of this calculation, 0.06, is the *arbitrage-free* price of the option: if the option were priced below this, a market participant could earn a guaranteed profit by buying it and simultaneously selling short the replicating portfolio; if it were priced above this, a trader could profit by writing the option and buying the portfolio. Note that the probabilities of the two six-month rates at point 1 played no part in the analysis. This reflects the arbitrage pricing logic: the value of the replicating portfolio must equal that of the option whatever path interest rates take.

That is not to say that probabilities do not have an impact on the option price. Far from it. If there is a very high probability that rates will increase, as in the example, an option's value to an investor will fall. This is reflected in the market value of the option or callable bond. When probabilities change, the market price changes as well.

Risk-Neutral Pricing

Although, as noted, the market does not price instruments using expected values, it is possible to derive *risk-neutral* probabilities that generate expected values whose discounted present values correspond to actual prices at period 0. The risk-neutral probabilities for the example above are derived in (11.7).

$$\frac{97.3236p + 97.5562(1-p)}{1 + \frac{1}{2}0.05} = 95.0423 \tag{11.7}$$

where
   p = the risk-neutral probability of an interest rate increase
   $1 - p$ = the probability of a rate decrease

Solving equation (11.6) gives p = 0.5926 and $1 - p$ = 0.4074. These are the two probabilities for which the probability-weighted average, or expected, value of the bond discounts to the true market price. These risk-neutral probabilities can be used to derive a probability-weighted expected value for the option in figure 11.5 at point 1, which can be discounted at the six-month rate to give the option's price at point 0. The process is shown in (11.8).

$$\frac{(0.5926 \times 0) + (0.4074 \times 0.1562)}{1 + \frac{1}{2}0.05} = 0.0621 \tag{11.8}$$

The option price derived in (11.8) is virtually identical to the 0.062 price calculated in (11.6). Put very simply, risk-neutral pricing works by first finding the probabilities that result in an expected value for the underlying security or replicating portfolio that discounts to the actual present value, then using those probabilities to generate an expected value for the option and discounting this to its present value.

Recombining and Nonrecombining Trees
The interest rate lattice in figure 11.2 is a one-period binomial tree. Expanding it to show possible rates for period 2 results in a structure like that shown in **FIGURE 11.6**.

The binomial tree in Figure 11.5 is termed *nonrecombining,* because each node branches out to two further nodes. This seems a logical process, and such trees are used in the market. Analyses incorporating them, however, require a considerable amount of computer processing power.

In period 1 there are two possible levels for the interest rate; at period 2 there are four possible levels. After N periods, there will be 2^N possible values for the interest rate. Calculating the current price of a 10-year callable bond that pays semiannual coupons involves generating more than one million possible values for the last period's set of nodes. For a 20-year bond, the number jumps to one trillion. (Note that the binomial models actually used in analyses have much shorter periods than six months, increasing the number of nodes.)

FIGURE **11.6** *A Nonrecombining Two-Period Binomial Tree for the Six-Month Interest Rate*
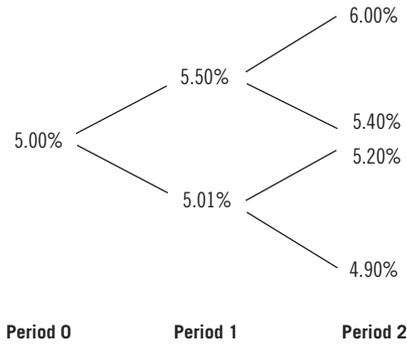
FIGURE **11.7** *A Recombining Two-Period Binomial Tree for the Six-Month Interest Rate*
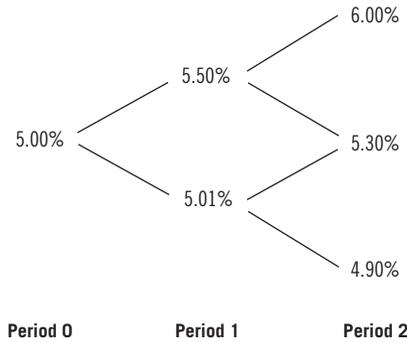
For this reason some market practitioners prefer to use *recombining* binomial trees, in which the branch sloping downward from an upper node ends at the same interest rate state as the one sloping upward from a lower node. This is illustrated in **FIGURE 11.7**.

The number of terminal nodes and possible values is much reduced in a recombining tree. A recombining tree with one-week periods used to price a 20-year bond, for example, has only 52 × 20 + 1, or 1,041, terminal values.

Pricing Callable Bonds

The tools discussed so far in this chapter are the building blocks of a simple model for pricing callable bonds. To illustrate how this model works, consider a hypothetical bond maturing in three years, with a 6 percent semiannual coupon and the call schedule shown in **FIGURE 11.8**.

The first step in the analysis is to create a risk-neutral recombining binomial lattice tracking the evolution of the six-month interest rate. The tree's nodes occur at six-month intervals and at each node the probability of an upward move in the rate is equal to that of a downward move. The tree is shown in **FIGURE 11.9**.

The next step is to use this tree to describe the bond's price evolution, ignoring its call feature. The tree is constructed from the final date backwards, using the bond's ex-coupon values. At each node, the ex-coupon bond price is equal to the sum of the expected value plus the coupon six months forward, discounted at the appropriate six-month yield. At year 3, the bond's price at all the nodes is 100.00, its ex-coupon par value. At year 2.5, the bond's price at the highest yield, 7.782 percent, is calculated by using this rate to discount the bond's expected price six months forward. The price in six months in both the "up" and the "down" state is 103.00—the ex-coupon value plus the final coupon payment. The bond's price at this node, therefore, is derived using the risk-neutral pricing formula as follows:

FIGURE **11.8** *Call Schedule of a Hypothetical 3-Year 6 Percent Bond*

CALL SCHEDULE	
Year 1	103.00
Year 1.5	102.00
Year 2	101.50
Year 2.5	101.00
Year 3	100.00

FIGURE 11.9 *Risk-Neutral Recombining Binomial Tree for the Six-Month Interest Rate*
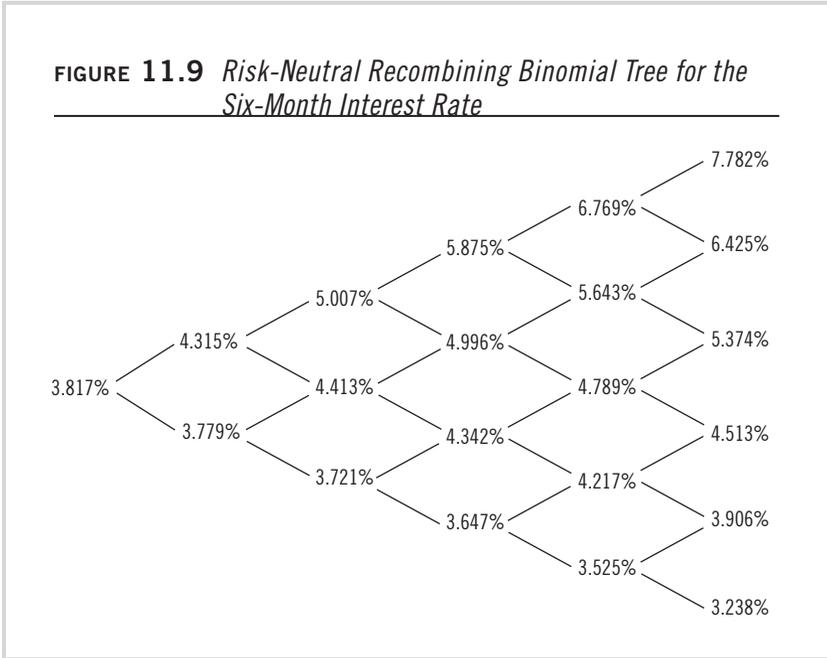
$$P_{bond} = \frac{0.5 \times 103 + 0.5 \times 103}{1 + \dfrac{0.07782}{2}} = 99.14237$$

The same process is used to obtain the prices for every node at year 2.5 and then repeated for each node in year 2. At the highest yield for year 2, 6.769 percent, the two possible future values are

$$99.14237 + 3.0 = 102.14237$$

and

$$99.79411 + 3.0 = 102.79411$$

Therefore the price of the bond in this state is given by

$$P_{bond} = \frac{0.5 \times 102.14237 + 0.5 \times 102.79411}{1 + \dfrac{0.06769}{2}} = 99.11374$$

The procedure is repeated until every node in the lattice is associated with a price. The completed lattice is shown in **FIGURE 11.10**.

FIGURE 11.10 *The Binomial Price Tree for the Hypothetical 3-Year 6 Percent Bond, Ignoring the Call Provision*
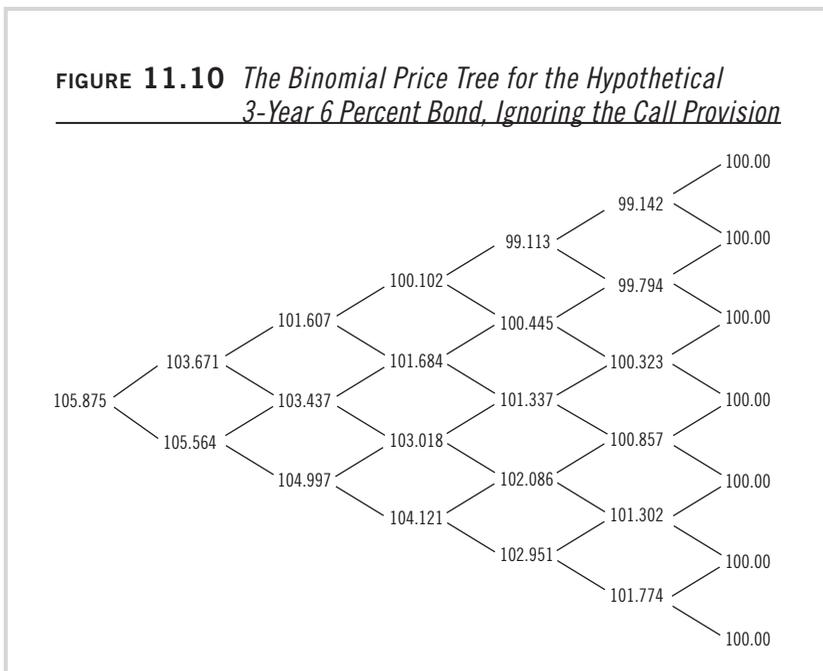
After calculating the prices for the conventional element of the callable bond, the next step is to compute the value of the option element. On the bond's maturity date, the option is worthless, because its "strike" is 100, which is the price the bond is redeemed at in any case. The option needs to be valued, however, at all the other node points.

The holder of the option in the case of a callable bond is the issuing company. At each call date during the life of the bond, the option holder will elect either to exercise it or to wait till the next date. In making this decision, the option holder must consider the following factors:

❑ the value of holding the option for an extra period, denoted by P_{Ct}
❑ the value of exercising the option straight away, P_C

If P_{Ct} is greater than P_C, the holder will not exercise; if P_C is greater than P_{Ct}, the holder will exercise immediately. The general rule determining whether P_C or P_{Ct} is greater is that options have more value "alive than dead." The same is true for callable bonds. It is sometimes more advantageous to run an in-the-money option rather than exercise straight away. On the other hand, at the year 2.5 call date, there is no value in holding the option for another period because it will be worthless at year 3. Therefore, if the option is in the money, the holder will exercise.

A number of factors dictate whether an option is exercised or not. The first is the asymmetric profit-loss profile of option holders: their potential gain is theoretically unlimited when the price of the "underlying" asset rises, but they lose only their initial investment if the price falls. This asymmetry favors running an option position. Another consideration favoring holding is the fact that the option's time value is lost if it is exercised early. In callable bonds, the call price often decreases as the bond approaches maturity. This creates an incentive to delay exercise until a lower strike price is available. Coupon payments, on the other hand, may favor earlier exercise, since, in the case of a normal, nonembedded option, this allows the holder to earn interest sooner.

The general process of valuing embedded options works as follows: say the value of the option for immediate exercise is V_t, the value of the option held for a further period is V_T, and the value of the option at any node is V. These values are defined by equations (11.9), (11.10), and (11.11).

$$V_T = \frac{0.5V_h + 0.5V_l}{1 + \frac{1}{2}r} \tag{11.9}$$

where

V_h = the value of the option in the up state
$V_l t$ = the value of the option in the down state
r = the six-month interest rate at the specified node

$$V_t = \max(0, P - S) \tag{11.10}$$

where

P = the bond's value at the specified node
S = the call option strike price, determined by the call schedule

$$V = \max(V_T, V_t). \tag{11.11}$$

The option value binomial tree, shown in **FIGURE 11.11**, is constructed by applying the appropriate expressions at each node, starting at the final period and working backward in time.

It is now possible to complete the price tree for the callable bond. Remember that the option in the case of a callable bond is held by the issuer. Its value, given by the tree in figure 11.11, must therefore be subtracted from the conventional bond price, given by the tree in figure 11.10, to obtain the callable bond value. For instance, the current price of the callable bond is 105.875 − 0.76, or 105.115. **FIGURE 11.12** shows the tree that results from this process. A tree constructed in this way, which is programmable into a spreadsheet or as a front-end application, can be used to price either a callable or a putable bond.

FIGURE **11.11** *Binomial Price Tree for the Option Embedded in the Callable 3-Year 6 Percent Bond*
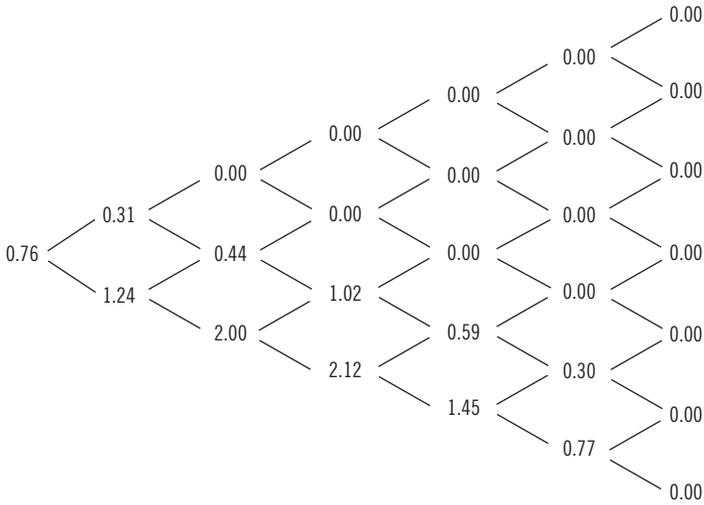
FIGURE **11.12** *Binomial Price Tree for the Hypothetical Callable 3-Year 6 Percent Bond*
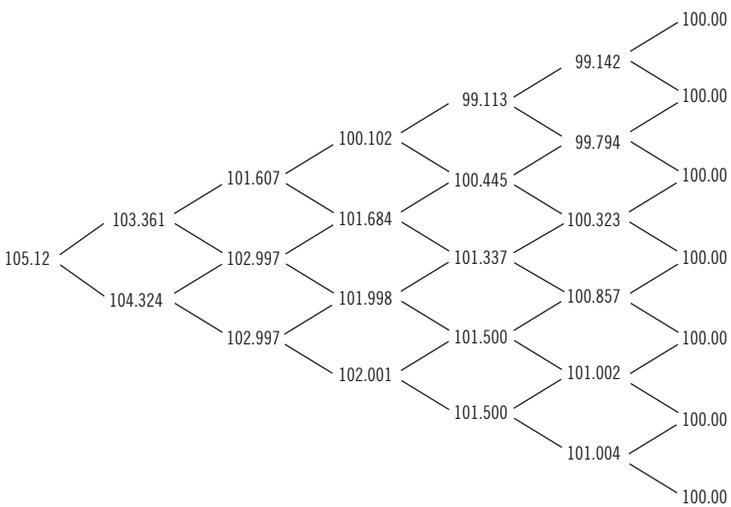
Price and Yield Sensitivity

As explained in chapter 1, the curve representing a plain vanilla bond's price-yield relationship is essentially convex. The price-yield curve for a bond with an embedded option changes shape as the bond's price approaches par, at which point the bond is said to exhibit negative convexity. This means that its price will rise by a smaller amount for a decline in yield than it will fall for a rise in yield of the same magnitude. **FIGURE 11.13** summarizes the price-yield relationships for both negatively and positively convex bonds.

Callable bonds exhibit negative convexity as interest rates fall. Option-adjusted spread analysis highlights this relationship by effecting a parallel shift in the benchmark yield curve, holding the spread constant, and calculating the theoretical prices along the nodes of the binomial price tree. The average present value is the projected price for the bond. **FIGURE 11.14** shows how this process affects the price-yield relationship of a hypothetical callable bond by comparing it with that of a conventional bond having the same coupon and maturity. Note that once the market rate falls below 10 percent, the callable bond exhibits negative convexity. This is because the embedded option is exercisable at that point, effectively capping the bond's price.

The market quotes bonds with embedded options in terms of yield spreads. A "cheap" bond trades at a high spread, a "dear" one at a low spread. The usual convention is to quote the spread between the redemption yield of the bond being analyzed and that of a government bond having an equivalent maturity. This is not an accurate measure of the actual difference in value between the two bonds, however. The reason is that, as explained in chapter 1, the redemption yield computation unrealistically discounts all a bond's cash flows at a single rate.

A better measure of the relative value of a bond with an embedded

FIGURE 11.13 *Price-Yield Relationships Associated with Negative and Positive Convexity*

| | PERCENT PRICE CHANGE FOR | |
CHANGE IN YIELD	POSITIVE CONVEXITY	NEGATIVE CONVEXITY
Fall of 100 bp	X%	Lower than Y%
Rise of 100 bp	Lower than X%	Y%

FIGURE 11.14 *Projected Prices for Callable and Conventional Bonds with Identical Coupons and Final Maturity Dates*
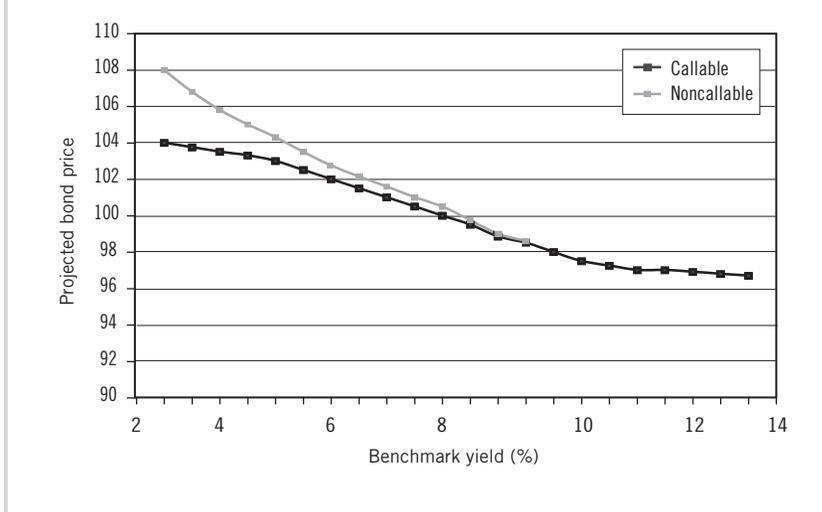
option is the constant spread that, when added to all the short-rates in the binomial tree, makes the bond's theoretical (model-derived) price equal to its observed market price. The constant spread that satisfies this requirement is the option-adjusted spread. It is "option-adjusted" because it reflects the option feature attached to the bond.

The OAS depends on the volatility level assumed in applying the model. For a given price, the higher the specified volatility, the lower the spread for a callable bond and the higher the spread for a putable one. Since the OAS is usually calculated relative to a government spot- or forward-rate curve, it reflects the credit and liquidity premiums over the government bond that are assigned to the corporate bond. OAS analysis depends on the valuation model being used and is only as accurate as the model itself.

Measuring Bond Yield Spreads

The binomial model evaluates a bond's return by measuring the extent to which it exceeds those determined by the risk-free short rates in the tree. The spread between these returns is the bond's *incremental return* at a specified price. Determining the spread involves the following steps:

❑ the binomial tree is used to derive a theoretical price for the bond

❏ this theoretical price is compared with the bond's observed market price

❏ if the two prices differ, the rates in the binomial model are adjusted by a user-specified amount, which is the estimated spread

❏ a new theoretical price is derived using the adjusted rates and compared with the observed one

❏ the previous two steps are repeated until the two prices are the same

Price Volatility of Bonds with Embedded Options

As explained in chapter 2, the duration for any bond may be calculated using equation (11.12) (assuming annualized yields).

$$D = \frac{\sum_{t=1}^{n} \frac{tC_t}{(1+rm)^t}}{P} \tag{11.12}$$

where
C_t = the bond cash flow at time t
P = the bond's fair price
C = the annual coupon payment
rm = the redemption yield
n = the number of years to maturity

To calculate the modified duration of a bond with an embedded option, the bondholder must assume a fixed maturity date based on the bond's current price. When it is unclear what redemption date to use, modified duration may be calculated to both the first call date and the final maturity date. This is an unsatisfactory compromise, however, since neither date, and so neither measure, may be appropriate. The problem is more acute for bonds that are continuously callable or putable from the first call or put date until maturity.

Effective Duration

It is possible to overcome some of the drawbacks of traditional duration by using OAS analysis to derive a bond's effective duration. As will be discussed in chapter 14, effective duration is based on approximate duration. Following Fabozzi (1997), approximate duration is derived using equation (11.13).

$$D_{\text{approx}} = \frac{P_- - P_+}{2P_0\,(\Delta rm)} \tag{11.13}$$

where

  P_0 = the initial price of the bond

  P_- = the estimated price of the bond if the yield falls by Δrm

  P_+ = the estimated price of the bond if the yield rises by Δrm

  Δrm = the change in the yield of the bond

Effective duration recognizes that yield changes may effect the future cash flow of a bond and so its price. For bonds with embedded options the difference between traditional duration and effective duration can be significant. The effective duration of a callable bond, for example, is sometimes half its traditional duration. As noted in chapter 14, for mortgage-backed securities, the difference is sometimes greater still.

Effective duration may be calculated using the binomial model and equation (11.13), as follows:

❑ calculate the bond's OAS spread

❑ apply a downward parallel shift to the benchmark yield

❑ construct an adjusted binomial tree using the new yield curve

❑ add the OAS adjustment to the short rate at each of the tree's nodes

❑ use the modified binomial tree—shown in Figure 11.9—to calculate the new value of the bond

❑ substitute this new price for P_+ in equation (11.13)

The same steps are used to derive P_-, except that the yield curve is shifted upward instead of downward. The effective duration of bonds containing embedded options is often referred to as *option-adjusted spread duration.* This measure has two advantages. The first is that it takes into account the interest-rate-dependent behavior of the embedded option and thus of the bond's cash flows. This is done by incorporating the binomial tree and holding the bond's OAS constant over the specified interest rate shifts, in effect maintaining the credit spread demanded by the market. The second, and possibly more significant, advantage is that OAS duration is based on a parallel shift in the benchmark yield curve and so links changes in a bond's price to changes in market interest rates rather than to shifts in its own yield.

Effective Convexity

Just as standard duration is not appropriate for bonds with embedded options, neither is traditional convexity. This is because traditional convexity, like traditional duration, fails to take into account the impact on a bond's future cash flows of a change in market interest rates. As discussed in chapter 14, the approximate convexity of any bond may be derived, following in Fabozzi (1997), using equation (11.14).

$$CV = \frac{P_+ + P_- - 2P_0}{P_0 (\Delta rm)^2}$$

(11.14)

If the prices used for P_+, P_-, and P_0 are calculated assuming that the bond's remaining cash flows will not change when market rates do, the convexity computed is for an option-free bond. For bonds with embedded options, the prices used in the equation should be derived using a binomial model, in which the cash flows do change with interest rates. The result is *effective* or *option-adjusted convexity.*

Sinking Funds

In some markets, most prominently in the United States, corporate bond issuers set up *sinking fund* provisions. A sinking fund allows the issuer to redeem the principal using one of two methods: by purchasing the stipulated amount of bonds in the open market and delivering them to the trustee for cancellation or by calling the required amount of the bonds at par. The second option is, in effect, a *partial call*—that is, a call involving only a fraction of the issue. The bonds called are generally selected randomly, by certificate serial numbers.

The method the issuer chooses to fulfill the sinking fund requirement is a function of the interest rate level. If interest rates have risen since the bond was issued, depressing the bond's price, the issuer will purchase the required amount of bonds in the open market. If interest rates have fallen, it will call the specified amount at par. As an illustration, consider the hypothetical ABC bond whose terms are listed in **FIGURE 11.15**.

According to the terms in figure 11.15, the ABC bond pays an 8 percent coupon and is set to mature in 2019. It also contains a provision stating that the issuer will redeem $50 million of the bond's $100 million

FIGURE 11.15 *A Hypothetical Bond with a Sinking Fund*

Issuer	ABC plc
Issue date	1–Dec–99
Maturity date	1–Dec–19
Nominal	$5 million
Coupon	8%
Sinking fund provision	$5 million December 1, 2009 to 2018

face value over ten years. This is the formal provision. The actual payments made may differ.

ABC has in effect embedded ten European options in the bond, each relating to $5 million nominal of the bonds and each expiring on December 1 of a different year, starting in 2009 and ending with 2018. The decision to exercise the options as they mature is made using the binomial-tree method discussed earlier.

Inflation-Indexed Bonds

Certain countries have markets in bonds whose coupon or final redemption payment, or both, are linked to their consumer price indexes. Generally, the most liquid markets in these *inflation-indexed,* or *index-linked,* debt instruments are the ones for government issues. Investors' experiences with the bonds differ, since the securities were introduced at different times in different markets and so are designed differently. In some markets, for instance, only the coupon payment, and not the redemption value, is index-linked. This makes comparisons in terms of factors such as yield difficult and has in the past hindered arbitrageurs seeking to exploit real yield differences. This chapter highlights the basic concepts behind indexed bonds and how their structures may differ from market to market.

Basic Concepts

The features considered in the design of index-linked bonds are the type of index, the indexation lag, the coupon frequency, and the type of indexation.

Choice of Index

In principle, bonds can be linked to almost any variable, including various price indexes, earnings measures, GDP output, specific commodities, and the exchange rate of foreign currencies against another currency. Ideally, the chosen index should reflect the hedging requirements of both

FIGURE 12.1 *Bloomberg DES Screen for the July 2013 TIPS*

Govt **DES**

SECURITY DISPLAY

TSY INFL IX N/B   TII1 ⅞ 07/15/13 101-12 /101-14 (1.72 /71) BGN @ 9:35

SECURITY INFORMATION				
CPN FREQ	2			
CPN TYPE	FIXED			
MTY/REFUND TYP	NORMAL			
CALC TYP (648)U.S. I/L REAL YLD				
DAY COUNT(1)ACT/ACT				
MARKET ISS	US GOVT			
COUNTRY/CURR	USA/ DOL			
SECURITY TYPE	USN			
AMT ISSUED	20000(MM)			
AMT OUTSTAND	20000(MM)			
MIN PIECE	1000			

ISSUER INFO
NAME   TSY INFL IX N/B
TYPE   US GOVT NATIONAL

IDENTIFICATION #'s
CUSIP   912828BD1
MLNUM   H2577
SEDOL 1   2914206
WERTPAP   900783
ISIN   US912828BD18
EURO COM 017286196

REDEMPTION INFO
MATURITY DT   7/15/13
NEXT CALL DT
WORKOUT DT   7/15/13
RISK FACTOR   4.37

ISSUANCE INFO
ISSUE DATE   7/15/03
INT ACCRUES   7/15/03
1ST CPN DT   1/15/04
PRC @ ISSUE   97.201

TENDERS ACCEPTED: $11000MM, SEE CPIRJL03 Index FOR INDEX RATIO.
$9000MM ISS'D AS A REOPENING EFF 10/15/03.
MIN RDMPTN VALUE AT MATURITY = PAR.

FIGURE 12.2 *Bloomberg DES Screen for the July 2003 Consumer Price Index*

Index **DES**

Related Functions	News	Price History

US Inflation Indexed CPI Ratio 10 Yr Notes Issued July 2003

Ticker	CPIRJL03 Index
Quoted	Price
Frequency	Daily
Country	UNITED STATES
Currency	US DOLLAR
Last Update	05:00
Source	US Treasury
History	Close
Begins	07/15/03 Daily
End of Week	Friday

Last   1.0054
High on 12/01/03 1.0094
Average 1.0055
Low on 10/16/03 1.0030

5 & 10 year ratios from the CPURNSA Index for revised historical
ratios and CPURNSBR Index for non-revised historical ratios.

the issuer and the investor. Their needs, however, may not coincide. For instance, retail investors overwhelmingly favor indexation to consumer prices, to hedge against inflation, which erodes bond earnings. Pension funds, on the other hand, prefer linking to earnings levels, to offset their earnings-linked pension liabilities. In practice, most bonds have been tied to inflation indexes, since these are usually widely circulated and well understood and issued on a regular basis. U.S. TIIS (Treasury inflation-indexed securities), or TIPS (Treasury inflation-protected securities), for instance, are linked to the U.S. Consumer Price Index (CPI-U), the non–seasonally adjusted average of prices for urban consumers. The securities' daily interest accrual is based on straight-line interpolation, and there is a three-month lag. So, for example, the October 2003 index level is used to determine the adjustment for January 1, 2004. **FIGURE 12.1** is the Bloomberg DES ("description") screen for the TIPS maturing in July 2013. **FIGURE 12.2** is the DES screen for the CPI index in July 2003, which was the base for this security when it was issued.

Indexation Lag

To provide precise protection against inflation, interest payments for a given period would need to be corrected for actual inflation over the same period. Lags, however, exist between the movements in the price index and the adjustment to the bond cash flows. According to Deacon and Derry (1998), such lags are unavoidable for two reasons. First, inflation statistics for one month are usually not known until well into the following month and are published some time after that. This causes a lag of at least one month, as shown in **FIGURE 12.3**. Second, in some markets the size of a coupon payment must be known before the start of the coupon period in

FIGURE **12.3** *The Indexation Lag*
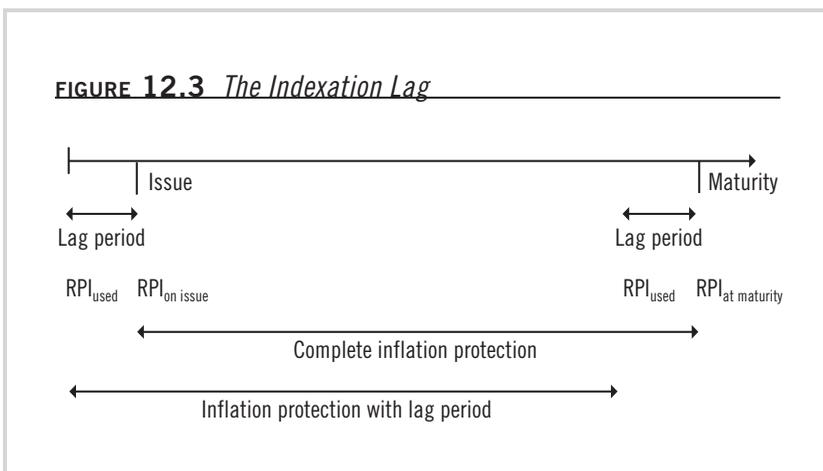
order to calculate the accrued interest. There is thus a delay between the date the coupon amount is fixed and the time the inflation rate for the period affecting that payment is known that is equal to the length of time between coupon payments. Deacon and Derry (1998) also notes that the lag can be minimized—for example, by basing the accrued interest calculation on cumulative movements in the consumer price index since the last coupon date, as is done for Canadian Real Return Bonds.

Coupon Frequency

Index-linked bonds often pay interest semiannually. Certain long-dated investors, such as fund managers whose liabilities include inflation-indexed annuities, may be interested in indexed bonds that pay on a quarterly or even monthly basis.

Type of Indexation

There are five basic methods of linking the cash flows from a bond to an inflation index: interest indexation, capital indexation, zero-coupon indexation, annuity indexation, and current pay. Which method is chosen depends on the requirements of the issuers and of the investors they wish to attract. The principal factors considered in making this choice, according to Deacon and Derry (1998), are duration, reinvestment risk, and tax treatment.

Interest indexation. Interest-indexed bonds have been issued in Australia, although not since 1987. They pay a coupon fixed rate at a real—inflation-adjusted—interest rate. They also pay a principal adjustment (equal to the percentage change in the CPI from the issue date times the principal amount) every period. The inflation adjustment is thus fully paid out as it occurs, and no adjustment to the principal repayment at maturity is needed.

Capital indexation. Capital-indexed bonds have been issued in the United States, Australia, Canada, New Zealand, and the United Kingdom. Their coupon rates are specified in real terms, meaning that the coupon paid guarantees the real amount. For example, if the coupon is stated as 2 percent, what the buyer really gets is 2 percent after adjustment for inflation. Each period, this rate is applied to the inflation-adjusted principal amount to produce the coupon payment amount. At maturity, the principal repayment is the product of the bond's nominal value times the cumulative change in the index since issuance. Compared with interest-indexed bonds of similar maturity, these bonds have longer durations and lower reinvestment risk.

Zero-coupon indexation. Zero-coupon indexed bonds have been issued in Sweden. As their name implies, they pay no coupons; the entire inflation adjustment occurs at maturity, applied to their redemption value.

These bonds have the longest duration of all indexed securities and no reinvestment risk.

In the United States, Canada, and New Zealand, indexed bonds can be stripped, allowing coupon and principal cash flows to be traded separately. This obviates the need for specific issues of zero-coupon indexed securities, since the market can create products such as deferred-payment indexed bonds in response to specific investor demand. In markets allowing stripping of indexed government bonds, a strip is simply a single cash flow with an inflation adjustment. An exception to this is in New Zealand, where the cash flows are separated into three components: the principal, the principal inflation adjustment, and the inflation-linked coupons—the latter being an indexed annuity.

Annuity indexation. Indexed-annuity bonds have been issued in Australia, although not by the central government. They pay a fixed annuity payment plus a varying element that compensates for inflation. These bonds have the shortest duration and highest reinvestment risk of all index-linked debt securities.

Current pay. Current-pay bonds have been issued in Turkey. They are similar to interest-indexed bonds in that their redemption payments at maturity are not adjusted for inflation. They differ, however, in their term cash flows. Current-pay bonds pay an inflation-adjusted coupon plus an indexed amount that is related to the principal. In effect, they are inflation-indexed floating-rate notes.

Duration. Duration measures something slightly different for an indexed bond than it does for a conventional bond, indicating price sensitivity to changes in real, inflation-adjusted interest rates, instead of in nominal, unadjusted ones. As with conventional bonds, however, the duration of zero-coupon indexed bonds is longer than that of equivalent coupon bonds. As noted above, indexed annuities will have the shortest duration of the inflation-linked securities. Investors with long-dated liabilities should theoretically prefer hedging instruments with long durations.

Reinvestment risk. Like holders of a conventional bond, investors in a coupon indexed bond are exposed to reinvestment risk: because they cannot know in advance what rates will be in effect when the bond's coupon payments are made, investors cannot be sure when they purchase their bond what yield they will earn by holding it to maturity. Bonds, such as indexed annuities, that pay more of their return in the form of coupons carry more reinvestment risk. Indexed zero-coupon bonds, like their conventional counterparts, carry none.

Tax treatment. Tax treatment differs from market to market and from product to product. Some jurisdictions, for example, treat the yearly capi-

tal gain on zero-coupon bonds as current income for tax purposes. This is a serious drawback, since the actual gain is not available until maturity, and it reduces institutional demand for these instruments.

Index-Linked Bond Cash Flows and Yields

As noted above, index bonds differ in whether their principal payments or their coupons or both are linked to the index. When the principal alone is linked, each coupon and the final principal payment are determined by the ratio of two values of the relevant index. U.S. TIPS' coupon payments, for instance, are calculated using an accretion factor based on the ratio between two CPI-U levels, defined by equation (12.1).

$$IR_{SetDate} = \frac{CPI_{Settlement}}{CPI_{Issue}} \qquad (12.1)$$

where

$IR_{SetDate}$ = index ratio
Settlement = the bond's settlement date
Issue = the bond's issue date
CPI_{M-3} = the CPI level three months before the bond's redemption date

$CPI_{Settlement}$ and CPI_{Issue} are the consumer price index levels recorded three months before the relevant dates. For a settlement or issue date of May 1, for instance, the relevant CPI level would be the one recorded on February 1. For a settlement or issue occurring on any day besides the first of the month, linear interpolation is used to calculate the appropriate CPI level. This is done by subtracting the reference month's CPI-U level from the following month's level, then dividing the difference by the number of days between the readings and multiplying the result by the number of days in the month leading up to the reference date. As an illustration, consider an issue date of April 7. The relevant index level would be the one for January 7. Say the January 1 CPI-U level is 160.5 and the February 1 level 160.6. The difference between these two values is

$$160.6 - 160.5 = 0.10$$

Dividing this difference by the number of days between January 1 and February 1 gives

$$0.10/31 = 0.00322581$$

And multiplying the result by the number of days in January before the reference date gives

$$0.00322581 \times 6 = 0.19355$$

So the CPI-U for January 7 is 160.5 + 0.19, or 160.69.

TIPS Cash Flow Calculations

TIPS' periodic coupon payments and their final redemption payments are both calculated using an inflation adjustment. Known as the *inflation compensation,* or IC, this is defined as in expression (12.2).

$$IC_{Set\ Date} = (P \times IR_{SetDate}) - P \qquad (12.2)$$

where
P = the bond's principal

The semiannual coupon payment, or interest, on a particular dividend date is calculated using equation (12.3).

$$Interest_{DivDate} = \frac{C}{2} \times (P + IC_{DivDate}) \qquad (12.3)$$

where
C = the annual coupon rate

The principal repayment is computed as in expression (12.4). Note that the redemption value of a TIPS is guaranteed by the Treasury to be a minimum of 100 percent of the face value.

$$\text{Principal repayment} = 100 \times \frac{CPI_{M-3}}{CPI_0} \qquad (12.4)$$

where
CPI_0 = the *base* CPI level—that is, the level three months before the bond's issue date

TIPS Price and Yield Calculations

The price of a TIPS comprises its real price plus any accrued interest, both of which are adjusted for inflation by multiplying them times the index ratio for the settlement date. The bond's unadjusted accrued interest, as explained in chapter 1, is calculated using expression (12.5).

$$\frac{C}{2} \times \frac{(d-f)}{d} \tag{12.5}$$

where

f = the number of days from the settlement date to the next coupon date

d = the number of days in the regular semiannual coupon period ending on the next coupon date

C = the unadjusted coupon payment

The TIP security's real price is given by equation (12.6).

$$\text{Real Price} = \left[\frac{1}{1+\frac{f}{d}\frac{r}{2}}\right]\left(\frac{C}{2} + \frac{C}{2}\sum_{j=1}^{n}\phi^j + 100\phi^n\right) - RAI \tag{12.6}$$

where

$$\phi = \left(\frac{1}{1+\frac{r}{2}}\right)$$

r = the TIPS' real annual yield

RAI = the unadjusted accrued interest

n = the number of full semiannual coupon periods between the next coupon date and the maturity date

EXAMPLE: *TIPS Coupon and Redemption Payment Calculation*

Consider a TIPS issued on January 15, 1998, with coupon of 3.625 percent and a maturity date of January 15, 2008. The base CPI-U level for the bond is the one registered in October 1997. Say this is 150.30. Assume that the CPI for October 2007, the relevant computing level for the January 2008 cash flows, is 160.5. Using these values, the final coupon payment and principal repayment per $100 face value will be:

$$\text{Coupon payment} = \frac{3.625}{2} \times \frac{160.5}{150.3} = \$1.9355$$

$$\text{Principal repayment} = 100 \times \frac{160.5}{150.3} = \$106.786$$

The markets use two main yield measures for all index-linked bonds: the *money,* or *nominal, yield,* and the *real yield.* Both are varieties of yield to maturity.

To calculate a money yield for an indexed bond, it is necessary to forecast all its future cash flows. This requires forecasting all the relevant future CPI-U levels. The market convention is to take the latest available CPI reading and assume a constant future inflation rate, usually 2.5 or 5 percent. The first relevant future CPI level is computed using equation (12.7).

$$CPI_1 = CPI_0 \times (1+\tau)^{m/12} \qquad (12.7)$$

where
   CPI_1 = the forecast CPI level
   CPI_0 = the latest available CPI
   τ = the assumed future annual inflation rate
   m = the number of months between CPI_0 and CPI_1

Consider an indexed bond that pays coupons every June and December. To compute its yield, it is necessary to forecast the CPI levels registered three months before June and eight months before December—that is, the October and April levels. Say this computation takes place in February. The first CPI level that must be forecast is thus next April's. This means that in equation (12.7), $m = 2$. Say the February CPI is 163.7. Assuming an annual inflation rate of 2.5 percent, the CPI for the following April is computed as follows.

$$CPI_1 = 163.7 \times (1.025)^{2/12}$$
$$= 164.4$$

Equation (12.8) is used to forecast the subsequent relevant CPI levels

$$CPI_{j+1} = CPI_1 \times (1+\tau)^{(j+1)/2} \qquad (12.8)$$

where
   j = the number of semiannual forecasts after CPI_1

The forecast CPI level for the following October is calculated as follows.

$$CPI_2 = 164.4 \times (1.025)^{1/2} = 168.5$$

Once the CPIs have been forecast, the bond's yield can be calculated. Assuming that the analysis is carried out on a coupon date so that accrued interest is zero, the money yield of a bond paying semiannual coupons is calculated by solving equation (12.9) for ri.

$$P_d = \frac{\left(\frac{C}{2}\right)\left(\frac{CPI_1}{CPI_0}\right)}{\left(1+\frac{1}{2}ri\right)} + \frac{\left(\frac{C}{2}\right)\left(\frac{CPI_2}{CPI_0}\right)}{\left(1+\frac{1}{2}ri\right)^2} + \ .... \ + \frac{\left(\left[\frac{C}{2}\right]+M\right)\left(\frac{CPI_N}{CPI_0}\right)}{\left(1+\frac{1}{2}ri\right)^N} \qquad (12.9)$$

where
ri = the semiannual money yield
N = the number of coupon payments (interest periods) up to maturity
M = the bond principal
C = the unadjusted coupon payment

The equation for indexed bonds paying annual coupons is (12.10).

$$P_d = \frac{C\left[\frac{CPI_1}{CPI_0}\right]}{(1+ri)} + \frac{C\left[\frac{CPI_2}{CPI_0}\right]}{(1+ri)^2} + \ .... \ + \frac{(C+M)\left[\frac{CPI_N}{CPI_0}\right]}{(1+ri)^N} \qquad (12.10)$$

The real yield, ry, first described by Fisher in *Theory of Interest* (1930), is related to the money yield through equation (12.11) (for bonds paying semiannual coupons).

$$\left(1+\tfrac{1}{2}ry\right) = \left(1+\tfrac{1}{2}ri\right)/(1+\tau)^{\frac{1}{2}} \qquad (12.11)$$

To illustrate this relationship, say the money yield is 5.5 percent and the forecast inflation rate is 2.5 percent. The real yield would then be

$$ry = \left\{ \frac{\left[1+\frac{1}{2}(0.055)\right]}{\left[1+(0.025)\right]^{\frac{1}{2}}} - 1 \right\} \times 2 = 0.0297 \text{ or } 2.97 \text{ percent}$$

Rearranging equation (12.11) to express ri in terms of ry and substituting the resulting expression for ri in equation (12.9) gives equation (12.12), which can be solved to give the real yield, calculated on a coupon date, of index bonds paying semiannual coupons.

$$P_d = \frac{CPI_a}{CPI_0}\left[\frac{\left[\frac{C}{2}\right](1+\tau)^{\frac{1}{2}}}{\left(1+\frac{1}{2}ri\right)} + \frac{\left[\frac{C}{2}\right](1+\tau)}{\left(1+\frac{1}{2}ri\right)^2} + ... + \frac{\left[\left\{\frac{C}{2}\right\}+M\right](1+\tau)^{\frac{N}{2}}}{\left(1+\frac{1}{2}ri\right)^N}\right]$$

$$= \frac{CPI_a}{CPI_0}\left[\frac{\left(\frac{C}{2}\right)}{\left(1+\frac{1}{2}ry\right)} + ... + \frac{\left(\frac{C}{2}\right)+M}{\left(1+\frac{1}{2}ry\right)^N}\right]$$

(12.12)

where

$$CPI_a = \frac{CPI_1}{(1+\tau)^{\frac{1}{2}}}$$

CPI_0 = the base index level

$\dfrac{CPI_a}{CPI_o}$ = the rate of inflation between the bond's issue date and the date the yield calculation is carried out

The equations for money yield and real yield can be interpreted as indicating what redemption yield to employ as the discount rate in calculating the present value of an index bond's future cash flows. From this perspective, equation (12.9) shows that the money yield is the appropriate rate for discounting money, or nominal, cash flows. Equation (12.12) shows that the real yield is the appropriate rate for discounting real cash flows.

Assessing Yields on Index-Linked Bonds

Index-linked bonds do not offer complete protection against a fall in the real value of an investment. These bonds, including TIPS, do not have guaranteed real returns, despite having their cash flows linked to a price index such as the CPI. The reason for this is the lag in indexation, which for TIPS is three months. The time lag means that an indexed bond is not protected against inflation for the last interest period of its life. Any inflation occurring during the final interest period will not be reflected in the bond's cash flows and will reduce the real value of the redemption payment and hence the bond's real yield. This may not be a major consideration when the inflation rate is low, but it can be a worry for investors when the rate is high. The only way to effectively eliminate inflation risk is to reduce the time lag in indexation of payments to one or two months.

Bond analysts frequently compare the yields on index-linked bonds with those on conventional bonds of the same maturity to determine the market's expectation with regard to inflation rates. Of course, many fac-

tors can influence the gap between conventional and indexed bond yields, including supply and demand and liquidity (conventional bonds are generally more liquid than indexed ones). A large part of the difference, however, is the inflation *premium,* which reflects the market's expectations about inflation during the life of the bond. To determine the implied expectation, analysts calculate the *break-even inflation rate:* the rate for which the money yield on an index-linked bond equals the redemption yield on a conventional bond of the same maturity.

As an illustration, say the August 1999 redemption yield on the 5 percent Treasury maturing in 2009 was 5.17 percent and the money yield on the 2 percent TIPS with the same maturity, assuming a constant inflation rate of 3 percent, was 2.23 percent. Plugging these values into equations (12.9) and (12.10), the implied break-even inflation rate is computed as follows.

$$\tau = \left\{ \frac{\left[1 + \frac{1}{2}(0.0517)\right]}{\left[1 + \frac{1}{2}(0.0223)\right]} \right\}^2 - 1 = 0.029287 \text{ or } 2.9 \text{ percent}$$

On the same date, also according to Bloomberg, conventional Treasury securities maturing in 2014 had yields to maturity ranging from 6.03 to 6.54 percent. The lowest conventional yield reflects expected inflation of approximately 4.27 percent over the ten years to maturity.

Which to Hold: Indexed or Conventional Bonds?

Accepting that developed, liquid markets, such as that for Treasuries, are efficient, with near-perfect information available to most if not all participants, then the inflation expectation is built into the conventional Treasury yield. If the inflation premium understates what certain market participants expect, investors will start buying more of the index-linked bond in preference to the conventional bond. This activity will force the indexed yield down (or the conventional yield up). If, on the other hand, investors think that the implied inflation rate overstates expectations, they will buy more of the conventional bond.

The higher yields of the conventional bonds compared with those of the index-linked bonds represent compensation for the effects of inflation. Bondholders will choose to hold index-linked bonds instead of conventional ones if they are worried about unexpected inflation. An individual's view on future inflation will depend on several factors, including the current macroeconomic environment and the credibility of the monetary authorities, be they the central bank or the government. Fund managers take their views of inflation, among other factors, into account in deciding how much of the TIPS and how much of the conventional Treasury to hold. Investment managers often hold indexed bonds in a portfolio against

specific index-linked liabilities, such as pension contracts that increase their payouts in line with inflation each year.

In certain countries, such as the United Kingdom and New Zealand, the central bank has explicit inflation targets, and investors may believe that over the long term those targets will be met. If the monetary authorities have good track records, investors may further believe that inflation is not a significant issue. In such situations, the case for holding index-linked bonds is weakened.

Indexed bonds' real yields in other markets are also a factor in investors' decisions. The integration of markets around the world in the past twenty years has increased global capital mobility, enabling investors to shun markets where inflation is high. Over time, therefore, expected returns should be roughly equal around the world, at least in developed and liquid markets, and so should real yields. Accordingly, index-linked bonds should have roughly similar real yields, whatever market they are traded in.

The yields on indexed bonds in the United States, for example, should be close to those in the U.K. indexed market. In May 1999, however, long-dated indexed bonds in the United States were trading at a real yield of 3.8 percent, compared with just 2 percent for long-dated index-linked gilts. Analysts interpreted this difference as a reflection of the fact that international capital was not as mobile as had been thought and that productivity gains and technological progress in the United States had boosted demand for capital there to such an extent that the real yield had had to rise.

Analysis of Real Interest Rates

Observing the trading patterns of a liquid market in inflation-indexed bonds enables analysts to draw conclusions about nominal versus real interest rates and to construct an inflation term structure. Such analysis is problematic, since conventional and indexed bonds typically differ considerably in liquidity. Nevertheless, as explained above, it is usually possible to infer market estimates of inflation expectations from the difference between the yields of the two types of bonds.

Indexation Lags and Inflation Expectations

As noted earlier, indexation lags prevent indexed bonds' returns from being completely inflation-proof. According to Deacon and Derry (1998), this suggests that an indexed bond can be regarded as a combination of a true indexed instrument (with no lag) and an nonindexed bond. Equation (12.13) expresses the price-yield relationship for a bond whose indexation lag is exactly one coupon period.

$$P = \sum_{j=1}^{n} \frac{C \prod_{i=0}^{j-1}(1+ri_i)}{(1+rm_j)^j \prod_{i=1}^{j}(1+ri_i)} + \frac{M \prod_{i=0}^{n-1}(1+ri_i)}{(1+rm_n)^n \prod_{i=1}^{n}(1+ri_i)} \tag{12.13}$$

where

ri_i = the rate of inflation between dates i-1 and i
rm = the redemption yield
C = the coupon payment
M = the unadjusted redemption payment
j = the specified interest period
n = term to maturity

If the bond has just paid the last coupon before its redemption date, (12.13) reduces to (12.14).

$$P = \frac{C}{(1+rm)(1+ri)} + \frac{M}{(1+rm)(1+ri)} \tag{12.14}$$

In this situation, the final cash flows are not indexed, and the price-yield relationship is identical to that for a conventional bond. This, then, represents the nonindexed component of the indexed bond. Its yield can be compared with those of conventional bonds, making it possible to quantify the indexation element. This implies a true real yield measure for the indexed bond.

To estimate the true real yield, analysts use the Fisher identity, one variant of which is shown as equation (12.15).

$$1 + y = (1+r)(1+i)(1+\rho) \tag{12.15}$$

where

y = the nominal interest rate
r = the real interest rate
i = the expected rate of inflation
ρ = the inflation premium

Essentially, the Fisher identity describes the relationship between nominal and real interest rates. Assuming a value for the risk premium ρ, the two bond price equations—one for a conventional bond and one for an indexed bond—can be linked using (12.15) and solved as a set of simultaneous equations to obtain values for the real interest rate and the expected inflation rate.

One approach is to determine the expected inflation rate using the difference between the yield on a conventional and that on an indexed bond having the same maturity date, if such bonds exist, ignoring any lag effects. This, however, is a flawed measure of inflation, because the calculation of the indexed bond's redemption yield already assumes an expected inflation rate.

As Deacon and Derry (1998, page 91) states, this problem is exacerbated if the maturity of both bonds is relatively short, because the less time to an indexed bond's maturity date, the greater the impact of its nonindexed component. To overcome this flaw, the break-even rate of inflation is used. This is derived by using the Fisher identity, with the risk premium ρ set to an assumed figure, such as 0, to relate the yield on the conventional bond to a yield on the indexed bond derived using an assumed initial inflation rate. The result is a new estimate of the expected inflation rate i, which is then used to recalculate the indexed bond's yield. The new yield, in turn, is used to produce a new estimate of the expected inflation rate. The process is repeated until a consistent value for i is obtained.

The main drawback with this basic technique is that it requires conventional and index-linked bonds of identical maturity. Using bonds with merely similar maturities compromises the results. In addition, the bonds' yields will be influenced not only by inflation expectations but by liquidity, taxation, indexation, and other considerations as well. There is also no equivalent benchmark (or *on-the-run*) indexed security.

An Inflation Term Structure

Where a liquid market in indexed bonds exists across a reasonable maturity term structure, it is possible to construct a term structure of inflation rates. In essence, the process involves constructing the nominal and real interest rate term structures, then using them to infer an inflation term structure. This, in turn, can be used to calculate a forward expected inflation rate for any term or a forward inflation curve in the same way that a forward interest rate curve is constructed.

The U.S. Federal Reserve uses an iterative technique to construct a term structure of expected inflation rates. First the nominal interest rate term structure is constructed using a version of the model described in Waggoner (1997) and discussed in James and Webber (2000). An initial assumed inflation term structure is then used to infer a term structure of real interest rates. This assumed inflation curve is usually set at a flat 3 or 5 percent. The real interest rate curve is then used to calculate an implied real interest rate forward curve. Next, the Fisher identity is applied at each point along the nominal and real interest rate forward

curves, which produces a new estimate of the inflation term structure. A new real interest rate curve is calculated from this. The process is repeated until a single consistent inflation term structure is produced.

Hybrid Securities

U p to now, the discussion has centered essentially on plain vanilla securities, bonds with a fixed-coupon and maturity date. However, much of the fixed-income markets revolves around more complex instruments, arranged to meet specific investor requirements. These include:

❑ *Hybrid securities.* Bonds that are made up of a combination of one or more securities, or which link their performance to the performance of another bond or asset, or an index

❑ *Structured notes.* Bonds that combine a bond with a derivative such as an interest rate swap, or bonds that are created from securitization transactions (see chapter 14)

❑ *Swaps.* Derivative contracts that exchange payments based on a specified notional amount, either as a percent rate (such as an interest rate swap) or other quoted cash flow

The motivations behind the development and use of more exotic, structured notes are varied. They include the desire for increased yield without additional credit risk, as well as the need to alter, transform, hedge, or transfer risk exposure and modify risk-return profiles. These instruments have been issued by banks, corporate institutions, and sovereign authorities. They can be tailored to particular risk profiles and enable investors to gain exposure to different markets, sometimes synthetically, that they have previously been unable to access. For instance, by purchasing structured notes, investors can take positions reflecting their views on exchange rates

or the changes they anticipate in the yield curves in different markets.

This chapter describes a number of structured notes that are currently available but by no means covers all the possible variations. That would, in fact, be impossible, since, if no existing security meets a particular investor or issuer requirement, an investment bank can structure a note that does.

Floating-Rate Notes

Floating-rate notes, or FRNs, are not structured notes. They are described here as a prelude to a discussion of *inverse floating-rate notes,* which *are* structured notes. As explained in chapter 1, an FRN is a bond that has a variable rate of interest: the coupon rate is linked to a specified index and changes periodically to reflect the current index reading. The notes usually pay a fixed spread over their reference index—for example, 50 basis points over the 6-month interbank rate. An FRN whose spread over the reference rate is not fixed is known as a *variable-rate note.*

Generally, the reference rate for FRNs is LIBOR, the London interbank offered rate—that is, the rate at which one bank will lend funds to another. The interest rate is fixed for a three- or six-month period, at the end of which it is reset. If, say, LIBOR is 7.6875 percent at the coupon reset date for a sterling FRN paying six-month LIBOR plus 0.50 percent, the FRN will pay 8.1875 percent for the following period, and interest will accrue at a daily rate of £0.0224315.

FRNs can have additional features, such as *floors,* which specify minimum levels below which the coupon cannot fall; *caps,* which specify maximum rates; and *calls,* which specify possible redemption dates before maturity. Perpetual FRNs also exist. As in other markets, borrowers frequently issue floating notes with specific, even esoteric, terms to meet particular requirements or customer demands. For example, Citibank issued a series of U.S. dollar–denominated FRNs indexed to the Euribor rate and another set of notes whose day count was linked to a specified LIBOR range.

Because the future values for the reference index are not known, it is not possible to calculate the redemption yield of an FRN. On the coupon-reset dates, the note will be priced precisely at par. Between these dates, it will trade very close to par, because of the way the coupon resets. If market rates rise between reset dates, the note will trade slightly below par; if rates fall, it will trade slightly above par. This makes FRNs' behavior very similar to that of money market instruments traded on a yield basis, although, of course, the notes have much longer maturities. FRNs can thus be viewed either as money market instruments or as alternatives to conventional bonds. Similarly, they can be analyzed using two approaches,

corresponding to these two views. The first approach is the *margin method,* which calculates the difference between an FRN's return and that of an equivalent money market security. The second is the *yield-to-maturity spread* approach, which compares the notes with fixed-rate bonds.

The margin method has two forms: *simple* margin and *discounted* margin. The simple margin method compares an FRN's average return throughout its life with the reference interest rate (i.e., LIBOR). This method is sometimes preferred because it does not require the forecasting of future interest index rates and coupon values. It has two components: a *quoted* margin, which is either above or below the reference rate, and a capital gain or loss, calculated assuming that the difference between the current price of the FRN and its maturity value is spread evenly over the remaining life of the bond. The formula for computing simple margin is (13.1).

$$\text{Simple margin} = \frac{(M - P_d)}{(100 \times T)} + M_q \tag{13.1}$$

where
  P_d = the dirty price, or P plus the accrued interest AI
  M = the par value
  T = the number of years from settlement to maturity
  M_q = the quoted margin

A positive simple margin signifies that the FRN's yield is higher than that of a comparable money market security.

The simple margin formula may be adjusted to take into account changes in the reference index rate since the last reset date. This is done by replacing the price in (13.1) with an adjusted price, defined using either (13.2a) or (13.2b) which assume semiannual coupons.

$$AP_d = P_d + (re + M) \times \frac{N_{sc}}{365} \times 100 - \frac{C}{2} \times 100 \tag{13.2a}$$

or

$$AP_d = P_d + (re + M) \times \frac{N_{sc}}{365} \times P_d - \frac{C}{2} \times 100 \tag{13.2b}$$

where
  AP_d = the adjusted dirty price
  re = the current value of the reference interest rate (such as LIBOR)
  $C/2$ = the next coupon payment (that is, C is the reference interest rate on the last coupon reset date plus M_q)

N_{sc} = the number of days between settlement and the next coupon date

Equation (13.2a) differs from (13.2b) in ignoring the current yield effect: all payments are assumed to be received on the basis of par, thus understating the coupon's value for FRNs trading below par and overstating it for those trading above par.

The simple margin method amortizes the discount or premium relative to the money market on the FRN in a straight line over its remaining life. The discounted margin method amortizes at a constantly compounded rate. The discounted margin method has the disadvantage of requiring that the reference index rate be forecast over the remaining life of the bond.

The discounted margin for an FRN paying semiannual coupons can be solved for from equation (13.3). (Equation (13.3) may also be stated in terms of discount factors instead of the reference rate.)

$$P_d = \left\{ \frac{1}{\left[1+\frac{1}{2}(re+DM)\right]^{\frac{days}{year}}} \right\} \times$$

$$\times \left\{ \frac{C}{2} + \sum_{t=1}^{N-1} \frac{(re*+M_q)\times\frac{100}{2}}{\left[1+\frac{1}{2}(re*+DM)\right]^t} + \frac{M_q}{\left[1+\frac{1}{2}(re*+DM)\right]^{N-1}} \right\} \quad (13.3)$$

where

DM = the discounted margin

re = the current value of the reference index rate

$re*$ = the assumed (or forecast) value of the reference rate over the remaining life the bond

M_q = the quoted margin

N = the number of coupon payments before redemption

The spread in the second, yield-to-maturity approach is defined as *rmf* − *rmb,* where *rmf* is the yield of the subject note and *rmb* is the yield of a reference bond. An FRN's *rmf* is calculated using equation (13.3) with both $(re + DM)$ and $(re* + DM)$ replaced by *rmf.* The reference bond yield is calculated using (13.4). If the yield-to-maturity spread is positive, the FRN offers a higher yield than the reference bond.

$$P = \frac{C}{(1+r)} + \frac{C}{(1+r)^2} + \frac{C}{(1+r)^3} + \cdots\cdots \frac{C}{(1+r)^N} + \frac{M}{(1+r)^N}$$

$$= \sum_{n=1}^{N} \frac{C}{(1+r)^n} + \frac{M}{(1+r)^N} \tag{13.4}$$

The discussion so far has involved plain vanilla FRNs. Other types of floaters that have traded include the following:

❏ *Collared FRNs*, which have both caps and floors. Effectively, these notes contain two embedded options—the issuer buys a cap from and sells a floor to the investor.

❏ *Step-up recovery FRNs*, whose coupons are fixed against comparable longer-maturity bonds, i.e., bonds with longer maturities than those of the FRNs in question. These notes enable investors to maintain exposure to short-term assets while capitalizing on a yield curved with positive slope.

❏ *Corridor FRNs*, which accrue daily interest only when the reference index falls within a specified range. Introduced to capitalize on expectations of comparative interest rate inactivity, these are high-risk, high-reward instruments. They offer investors very substantial margins over a chosen reference rate. But if the reference rate does not remain within a relatively narrow corridor, the interest payment is forfeited entirely.

Inverse Floating-Rate Notes

An inverse floating-rate note, or *inverse floater,* pays a coupon that increases as general market rates decline. It offers enhanced returns to investors who, in contrast to the market consensus, believe the outlook for bonds is generally positive. These notes are suitable when inflation is low and the yield curve positive, both conditions that would, in a conventional analysis, suggest rising interest rates in the medium term. Inverse floaters may also be appropriate when the yield curve is negative, i.e., inverted, should the investor agree with the market consensus, which would be for lower rates in the medium term.

The coupon on an inverse floater may be determined in a number of ways. The most common is to subtract from a specified fixed interest rate a variable that is linked to a reference index. Coupons have a floor, which, if unspecified, is 0 percent.

Issuers of inverse floaters are usually corporations. The notes may also, however, be issued to meet specific client requirements, by *specialized*

FIGURE 13.1 *Terms of a Hypothetical Inverse Floater*

Nominal value	$100,000,000
Issue date	5-Jan-2000
Maturity date	5-Jan-2003
Note coupon	15.75% – (2 x LIBOR)
Day-count basis	actual/365
Index	6-month LIBOR
Current LIBOR rate	5.15%
Rate fixing	Semiannual
Initial coupon	5.45%
Minimum coupon	0%

USD LIBOR	COUPON PAYABLE
1.00%	13.75%
1.50%	12.75%
2.00%	11.75%
2.50%	10.75%
3.00%	9.75%
3.50%	8.75%
4.00%	7.75%
4.50%	6.75%
5.00%	5.75%
5.50%	4.75%
6.00%	3.75%
6.50%	2.75%
7.00%	1.75%

investment vehicles—funds, such as wholly owned Citigroup subsidiaries Centauri and Dorada Corp, that are set up to invest in particular areas or sectors. **FIGURE 13.1** illustrates the coupon calculation on a typical inverse floater and how changes in the LIBOR rate affect the payment.

FIGURE **13.2** *Duration of a Year Inverse Floater*

Duration of 3-year note with 5.30% coupon	2.218 years
Duration of 3-year inverse floater (x 3)	6.654 years

The inverse floater in figure 13.1 pays a slightly above-market initial coupon given a positive yield curve. Investors benefit because they get a coupon whose sensitivity is equal to two times the changes in LIBOR.

FRNs in general are highly interest rate sensitive. This is because the leverage involved in the coupon calculation endows them with the highest duration of any instrument traded in the fixed-income market. The note in figure 13.1, for example, has a calendar maturity of three years. As shown in **FIGURE 13.2**, however, its modified duration is much higher.

Inverse floaters are very flexible instruments. They can, for instance, be linked to any reference index and thus serve as vehicles for different views concerning various short- or long-term interest rates, such as the central bank repo rate, 10-year swap rates, or a government benchmark. The notes' leverage can also be altered according to investors' risk preferences. The fixed element can be chosen for the same reason. Equally, the fixed element can be set to move upward or downward as required at next term. Inverse floaters can also give investors exposure to markets they wouldn't otherwise have access to. Someone who has a particular view on a specific foreign interest rate, for example, but who can't invest in that market's securities for some reason can buy an inverse FRN that is linked to the foreign index but pays coupon in the domestic currency.

Hedging Inverse Floaters

Borrowers often issue inverse floating notes in one currency and swap the proceeds into another that fits their funding needs better using a currency swap. They may then hedge their interest rate exposure on the note through an interest rate swap. The counterparty swap bank, in turn, will hedge its own exposure. The interest rate swap structure used to hedge the note in figure 13.1 is shown in **FIGURE 13.3**.

In figure 13.3, the note issuer enters into a swap in which it
- ❏ pays 6-month LIBOR
- ❏ receives the note coupon rate

FIGURE **13.3** *Structure of Swaps Hedging an Inverse Floater*

	15.75% − (2 x LIBOR)		3 x 5.30% (inc cap cost)	
Note Holder	**Note Issuer**	**Swap Bank**	**The Market**	
	15.75% − (2 x LIBOR)	LIBOR	3 x LIBOR	
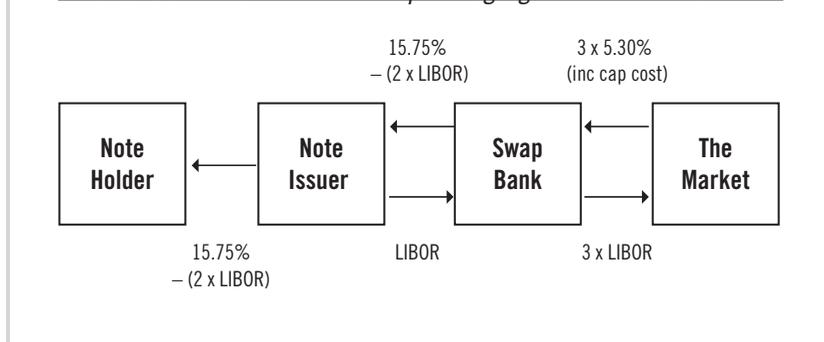
The swap bank, which takes the other side of the transaction, receiving LIBOR and paying the coupon rate, must now hedge its own exposure with another swap. To understand the structure of this second swap, it helps to consider the note coupon of 15.75 percent minus two times LIBOR as composed of the following two transactions:

❑ the note holder receives 15.75 percent
❑ the note holder pays two times LIBOR

In its arrangement with the note issuer, the bank is thus receiving three times LIBOR and paying 15.75 percent. The hedging swap it enters into consists of a payment of three times LIBOR in return for three times the fixed swap rate of 5.30, or 15.90 percent. This total is higher than the fixed component of the coupon by 15 basis points. This difference is the cost of fixing a cap, to hedge against the exposure presented by the floor on the note.

Figure 13.1 specifies that the inverse floater has a minimum coupon on 0 percent. The floor is passed on from the note issuer to the swap bank via the swap. This, in effect, caps the note holders' LIBOR exposure at 7.875 percent (15.75 divided by two). The bank's swap leaves it exposed to a rise in LIBOR above this level. To be fully hedged, the bank must buy an interest rate cap on LIBOR with a strike rate of 7.875 percent. The cap costs 15 basis points, which explains the spread over the coupon rate in the swap structure.

Indexed Amortizing Notes

Another type of hybrid security is the *indexed amortizing note,* or IAN. IANs were introduced in the U.S. market in the early 1990s in response to demand by investors in asset-backed notes known as collateralized mort-

FIGURE 13.4 *Terms of a Hypothetical IAN*

Issuer	Mortgage agency
Nominal value	$250,000,000
Legal maturity	Six years
Coupon	2-year Treasury plus 100 bps
Interest basis	Monthly
"Lock-out period"	Three years
Reference index	6-month LIBOR
6m LIBOR fixing on issue	5.15%
Minimum level of note	20%

AVERAGE LIFE SENSITIVITY LIBOR RATE	AMORTIZATION RATE	AVERAGE LIFE (YEARS)
5.15%	100%	3
6.00%	100%	3
7.00%	21%	4.1
8.00%	7%	5.6
9.00%	0.00%	6

gage obligations, or CMOs. IANs have been issued by banks, corporations, and, in large volumes, by U.S. government agencies.

IANs are short- to medium-term unsecured notes, most commonly with five-year maturities. They have fixed-coupons and variable nominal values—that is, after a "lock-out period," during which it remains stable, their principal is paid down according to a schedule determined by the level of a reference index, such as 6-month LIBOR, relative to a specified base rate. In this respect, IANs are similar to mortgage-backed notes, which also amortize. Mortgage-backed notes' amortization, however, is determined by the principal payments and prepayments of their underlying pools of mortgages. Because mortgage payments follow less clearly defined patterns than the reference indexes, IANs are considered to have an advantage over mortgage-backed securities. They also typically pay yields that are higher than those of conventional debt securities of similar credit quality. **FIGURE 13.4** shows the terms of a hypothetical IAN.

FIGURE 13.5 *IAN Hedge Arrangement*
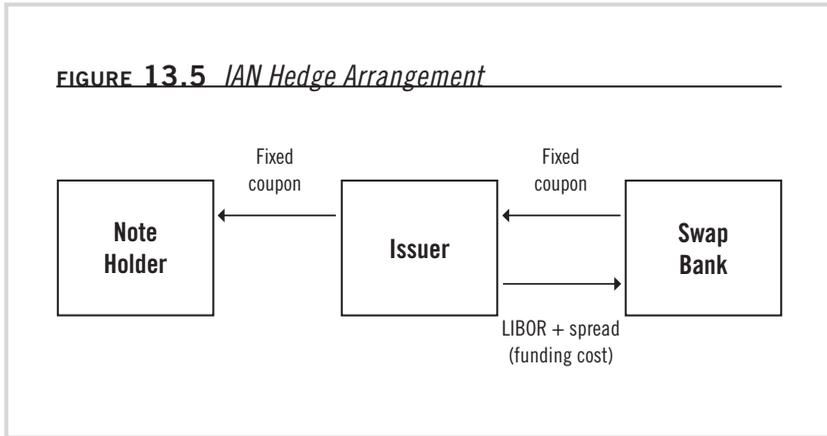
The note in figure 13.3 pays a coupon equal to the current 2-year government benchmark plus a fixed spread of 1 percent. It has a legal maturity of six years, but it will mature in three years if, two years from the issue date, 6-month LIBOR stands at 6 percent or below. Amortization takes place on subsequent rate-fixing dates according to the specified schedule. If at any time less than 20 percent of the nominal value remains, the note is canceled.

IAN issuers typically hedge their exposure through swap arrangements that mirror the notes' structure. A simple hedge is shown in **FIGURE 13.5**. It is more common, however, for the arrangement to involve a series of options on swaps, or *swaptions*. When volatility in the fixed-income market is high and the yield curve steeply positive, swaptions tend to have greater value, so the IAN coupon in those circumstances might be especially attractive.

Advantages for Investors

It has already been noted that IANs offer relatively high yields for relatively short maturities and that their amortization structures are easier to understand than those of mortgage-backed securities. In addition to these advantages, the notes, like the other instruments described here, can be tailored to meet individual investor requirements. The features most frequently subjected to this tailoring are the legal maturity and the lockout period, with the yield premium decreasing as the latter approaches the former. Although the most common reference index is LIBOR, it can also be a government benchmark or an interbank rate such as the swap rate.

Synthetic Convertible Notes

Synthetic convertible notes are securities with fixed coupons, typically set at a relatively low level, whose total return is linked to an external source, such as the level of an equity index or the price of a specific security. In one common structure, the note is redeemable above par if the reference index or security value exceeds a stated minimum. The notes thus give investors the opportunity to profit from the benchmark's performance while providing the safety net of redemption at par should this performance fall short. Another typical synthetic convertible structure is the zero-coupon note. These notes are issued at par and redeemable at par, or higher, if a specified equity index performs better than a stated level.

FIGURE 13.6 shows the terms of a hypothetical sterling synthetic convertible note linked to the FTSE 100 equity index. This note will pay par on maturity unless the FTSE 100 has risen by more than 10 percent from its level on the issue date. In that case, the redemption value will be par plus the amount of the index rise. The note also pays a coupon of 0.5

FIGURE **13.6** *Terms of a Hypothetical Synthetic Convertible Note*

Nominal value	$50,000,000
Term to maturity	Two years
Issue date	17-Jun-99
Maturity date	17-Jun-01
Issue price	$100
Coupon	0.50%
Interest basis	Semiannual
Redemption proceeds	Min [100, Formula level]
Formula level	100 + [100 x (R(I) – (1.1 x R(II)) / R(II)]
Index	FTSE-100
R(I)	Index level on maturity
R(II)	Index level on issue
Hedge terms	
Issuer pays	LIBOR
Swap bank pays	Redemption proceeds in accordance with formula

percent, which is roughly five percentage points less than the 2-year sterling yield at the time. The synthetic convertible described is thus suitable only for investors who are very bullish on the prospects of the FTSE 100.

Investor Benefits

Like convertible bonds, synthetic convertible notes provide investors with fixed coupons plus the possibility of profiting from the performance of the reference index. Unlike the ordinary convertible, however, the synthetic pays in cash rather than in shares of the associated equity.

The reference for a synthetic convertible can be virtually any publicly quoted financial instrument or relationship. Payouts have been linked, for instance, to the exchange rate of two currencies, the days on which LIBOR falls within a specified range, and the performance of a selected basket of stocks, such as technology shares.

Interest Differential Notes

Interest differential notes, or IDNs, are hybrid securities that enable investors to take a position based on their views about interest rates in two different currencies. Notes in the U.S. market are usually denominated in U.S. dollars; Euromarket notes have been issued in a wide range of currencies.

IDNs have a number of variations. Some pay a variable coupon and a fixed redemption amount; others pay a fixed coupon and a redemption amount that is determined by the level or performance of a reference index. Still other IDNs have payoff profiles linked to the difference between interest rates in two specified currencies or between rates for two different maturities in one currency.

FIGURE 13.7 shows the terms and potential returns of a 5-year IDN. The terms specify that the note's coupon will increase as the difference between U.S. dollar LIBOR and euro LIBOR widens, and vice versa. Below the list of terms are various returns that are possible in different interest rate scenarios. These are stated as spreads over the 5-year government benchmark yield. For instance, at the initial LIBOR differential of 2.65, it is stated that the IDN will return 95 basis points over the benchmark. The returns given are not guaranteed, of course, since they are based on the unrealistic assumption the interest differential will remain at the level indicated through to the final coupon-setting date. The listed returns do, however, demonstrate that the note's spread over the benchmark widens as the difference between the two rates increases. And the note will continue to offer a premium over the government yield even if the rate difference declines, as long as this decline doesn't exceed 100 basis points each year.

FIGURE **13.7** *Terms of a Hypothetical Interest Differential Note*

Term to maturity	Five years
Coupon	[(2 x USD LIBOR) – (2 x EUR LIBOR) – 0.50%]
Current USD LIBOR	6.15%
Current EUR LIBOR	3.05%
Rate differential	3.10%
First coupon fix	5.70%
Current 5-year benchmark rate	4.75%
Yield spread over benchmark	0.95%

CHANGE IN LIBOR SPREAD (BPS P.A.)	LIBOR SPREAD AT RATE RESET	SPREAD OVER BENCHMARK
75	4.78%	2.34%
50	3.90%	1.88%
25	3.15%	1.21%
0	2.65%	0.95%
−25	1.97%	0.56%
−50	1.32%	0.34%
−75	0.89%	0.12%
−100	0.32%	−0.28%

For purposes of analysis, this IDN can be regarded as a fixed-coupon bond plus the double indexation of an interest rate differential. Indexation here refers to a reference rate based on an index. The double indexation creates two long positions in a 5-year dollar-denominated fixed-rate note and two short positions in a euro-denominated fixed-rate note. The short positions remove the euro exchange-rate risk, so investors are exposed only to the euro interest rate risk, which is the desired exposure.

As with other hybrid securities, issuers of IDNs hedge their exposure with swaps. For the note in figure 13.7, the hedge would involve both dollar and euro interest rate swaps.

Benefits for Investors

IDNs and similar instruments enable investors to put on positions that reflect their views on the direction or level of foreign interest rates without taking on currency, or exchange-rate, risk. The IDN in figure 13.7, for example, would appeal to investors with a particular view on the U.S.-dollar and euro yield curves. Say the dollar curve is inverted, and the euro curve positively sloping. Investors buying this IDN would be earning a high yield while expressing a view—for example, is the market going up or down—that is different from the market consensus.

IDNs can be structured to enable investors to take positions on the yields in different currencies at the same maturity. A note's coupon, for example, could be determined by the difference between the 10-year government benchmark yields in two specified countries. The notes can also be linked to spreads between yields at different maturities in the same currency. This would be a straight yield-curve, or relative-value, trade in a domestic or foreign currency.

IDNs do expose investors to interest rate risk. If the rate differential moves in the opposite direction from the one desired, the coupon is reduced and the yield may fall below that available on the benchmark bond.

Securitization and Mortgage-Backed Securities

Perhaps the best illustration of the flexibility, innovation, and user-friendliness of the debt capital markets is the rise in use and importance of *securitization.* As defined in Sundaresan (1997, page 359), securitization is "a framework in which some illiquid assets of a corporation or a financial institution are transformed into a package of securities backed by these assets, through careful packaging, credit enhancements, liquidity enhancements and structuring."

The flexibility of securitization is a key advantage for both issuers and investors. Financial-engineering techniques employed by investment banks today enable bonds to be created from any type of cash flow. The most typical such flows are those generated by high-volume loans such as residential mortgages and car and credit card loans, which are recorded as assets on bank or financial-house balance sheets. In a securitization, the loan assets are packaged together, and their interest payments are used to service the new bond issue.

In addition to the more traditional cash flows from mortgages and loan assets, investment banks underwrite bonds secured with flows received by leisure and recreational facilities, such as health clubs, and other entities, such as nursing homes. Bonds securitizing mortgages are usually treated as a separate class, termed *mortgage-backed securities,* or MBSs. Those with other underlying assets are known as *asset-backed securities,* or ABSs. The type of asset class backing a securitized bond issue determines the method used to analyze and value it.

The asset-backed market represents a large and diverse group of securities suited to a varied group of investors. Often these instruments are the

only way for institutional investors to pick up yield while retaining assets with high credit ratings. They are popular with issuers because they represent a cost-effective means of removing assets from their balance sheets, thus freeing up lines of credit and enabling them to access lower-cost funding. Although asset-backed securities were developed in the United States, liquid markets for them also exist in the United Kingdom, Europe, Asia, and Latin America.

Instruments are available backed by a variety of assets, covering the entire yield curve, with either fixed or floating coupons. In the United Kingdom, for example, it is common for mortgage-backed bonds to have floating coupons, mirroring the interest basis of the country's mortgages. To suit investor requirements, however, some of these structures have been modified, through swap arrangements, to pay fixed coupons.

As noted, securitization is used in a large number of markets involving many different currencies. Readers interested in particular areas should consult the References section for specialized texts.

Reasons for Undertaking Securitization

The driving force behind the growth in securitization has been the need for banks to reduce the size of their balance sheets. Reducing assets in this manner has the following benefits:

❏ The return on equity increases, since the revenues from the assets remain roughly unchanged even as the size of the assets decreases.

❏ The level of capital required to support the balance sheet is reduced, leading to cost savings or permitting the institution to allocate the capital to other, perhaps more profitable, businesses.

❏ The interest payable on ABSs is frequently considerably below that earned on the underlying loans, creating a cash surplus.

Securitization also enables an institution to access debt markets its credit rating would otherwise be too low for. The growth in the United States of the "credit card banks," such as MBNA International, would have been severely restricted if these firms had not had a market for their securitized debt.

Market Participants

Securitization involves several participants. The *originator* owns the securitized assets. These are typically acquired by an *issuer,* which is usually a *special purpose vehicle,* or SPV, set up specifically for this purpose and domiciled offshore. Establishing an SPV ensures that the underlying asset pool is separate from the originator's other assets so that a bankruptcy or insolvency suf-

fered by the originator will have minimal impact on the securitized assets.

In addition to the originator and issuer, there are two trustees. The *issuer trustee* has the following responsibilities:
- ❏ representing the interests of the bondholders
- ❏ monitoring the transaction and the issuer for violations of the deal covenants
- ❏ enforcing the rights of the bondholders in the event of bankruptcy

The *security trustee* is responsible for the following:
- ❏ holding the security interest in the underlying asset pool
- ❏ communicating with the pool's manager
- ❏ acting under the direction of the note trustee, who is responsible for representing the noteholder interests—often the same firm acts as security and note trustee—in the event of default

Because the assets are held within an SPV framework, defined in formal legal terms, the originator's financial status and credit rating are almost irrelevant to the bondholders. There may also be a third-party guarantee of credit quality, which enables the securitized notes to be offered with an investment-grade credit rating up to AAA. **FIGURE 14.1** shows a simplified process of securitization.

FIGURE 14.1 *Securitization Process*
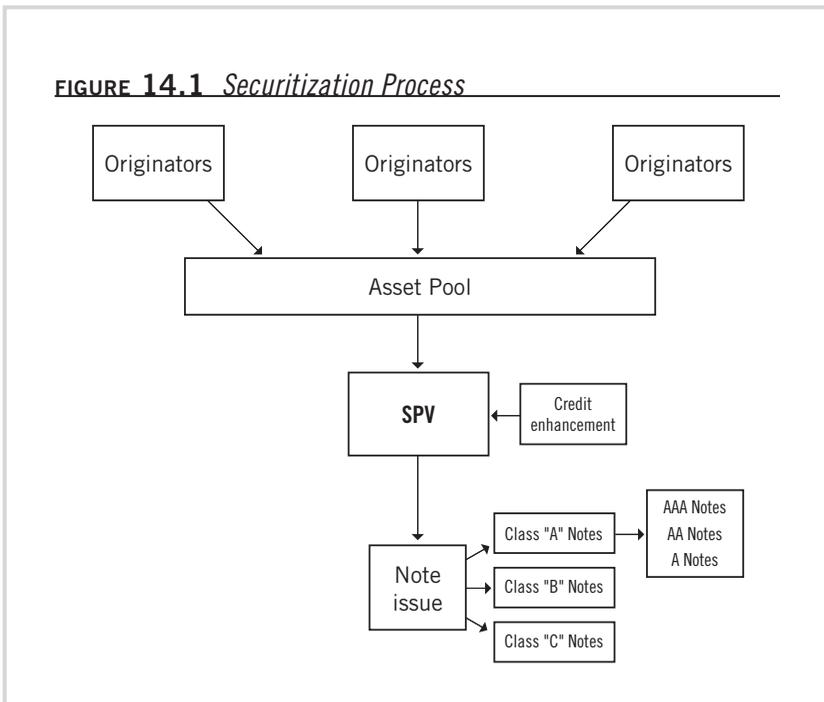
Securitizing Mortgages

A mortgage is a long-term loan taken out to purchase residential or commercial property, which itself serves as security for the loan. The term of the loan is usually twenty-five to thirty years, but a shorter period is possible if the borrower, or *mortgagor,* wishes one. In exchange for the right to use the property during the term of the mortgage, the borrower provides the lender, or *mortgagee,* with a *lien,* or claim, against the property and agrees to make regular payments of both principal and interest. If the borrower defaults on the interest payments, the lender has the right to take over and sell the property, recovering the loan from the proceeds of the sale. The lien is removed when the debt is paid off.

A lending institution may have many hundreds of thousands of individual residential and commercial mortgages on its books. When these are pooled together and used as collateral for a bond issue, the result is a mortgage-backed security. In the U.S. market, certain mortgage-backed securities are backed, either implicitly or explicitly, by the government. A government agency, the Government National Mortgage Association (GNMA, known as "Ginnie Mae"), and two government-sponsored agencies, the Federal Home Loan Corporation and the Federal National Mortgage Association ("Freddie Mac" and "Fannie Mae," respectively), purchase mortgages to pool and hold in their portfolios and, possibly, securitize. The MBSs created by these agencies trade essentially as risk-free instruments and are not rated by the credit agencies.

Mortgage-backed bonds not issued by government agencies are rated in the same way as other corporates. Some nongovernment agencies obtain mortgage insurance for their issues, to boost their credit quality. The credit rating of the insurer then becomes an important factor in the bond's credit rating.

Growth of the Market

Hayre, Mohebbi, and Zimmermann (1998) list the following advantages of mortgage-backed bonds:

❑ Their yields are usually higher than those of corporate bonds with the same credit rating. In the mid-1990s, mortgage-backed bonds traded around 100 to 200 basis points above Treasury bonds; by comparison, corporates traded at a spread of around 80 to 150 for bonds of similar credit quality. This yield gap stems from the mortgage bonds' complexity and the uncertainty of mortgage cash flows.

❑ They offer investors a wide range of maturities, cash flows, and security collateral to choose from.

❑ Agency mortgage-backed bonds are implicitly backed by the government and therefore represent a better credit risk than triple-A-rated corporate bonds; nonagency bonds are often triple-A or double-A rated.

❑ The market is large and thus very liquid; agency mortgage-backed bonds have the same liquidity as Treasury bonds.

❑ Unlike most other bonds, mortgage-backed securities pay monthly coupons, an advantage for investors who require frequent income payments.

Types of Mortgages and Their Cash Flows

Mortgages can have either fixed or floating rates of interest. Borrowers in the United States generally take out fixed-rate *repayment* mortgages, which amortize the principal. There are also *interest-only* mortgages, where the borrower's regular payments consist only of interest; the principal is paid off on maturity using the proceeds of an investment contract taken out at the same time, and for the same term, as the mortgage. In the United Kingdom, these are known as *endowment* mortgages and are popular with homebuyers, although their popularity has waned in recent years.

Level-payment fixed-rate mortgages, the conventional form in the United States, have fixed terms to maturity and specify monthly payments of fixed-rate interest. The monthly interest payment on a conventional fixed-rate mortgage is given by formula (14.3), which is derived from the conventional present-value equation via the steps shown in (14.1) and (14.2).

$$M_{m0} = I \left[\frac{1 - \left[\frac{1}{(1+r)^n} \right]}{r} \right] \tag{14.1}$$

where

M_{m0} = the original mortgage balance (the loan cash amount)

I = the monthly cash mortgage payment

r = the simple monthly interest rate, equal to the annual interest rate divided by 12

n = the term of the mortgage in months

Equation (14.1) is rearranged as (14.2), which may then be simplified as (14.3).

$$I = \cfrac{M_{m0}}{\left[\cfrac{1 - \left[\cfrac{1}{(1+r)^n}\right]}{r}\right]}$$

$$(14.2)$$

$$I = M_{m0}\left[\frac{r(1+r)^n}{\left[(1+r)^n - 1\right]}\right]$$

$$(14.3)$$

The monthly payment includes both the interest servicing and a repayment of part of the principal. In equation (14.3), after the 264th interest payment, the balance will be zero and the mortgage will have been paid off. As the principal is paid off, the base on which the interest is calculated diminishes and the monthly interest payment is reduced. So, assuming a constant monthly payment amount, the proportion of it dedicated to repaying the principal steadily increases. The remaining mortgage balance for any month during the term of the mortgage may be calculated using (14.4).

$$M_{mt} = M_{m0}\left[\frac{\left[(1+r)^n - (1+r)^t\right]}{\left[(1+r)^n - 1\right]}\right]$$

$$(14.4)$$

where
  M_{mt} = the mortgage cash balance after t months

The principal repayment and the interest payment in any month during the mortgage term can be calculated using equations (14.5) and (14.6), respectively.

$$p_t = M_{m0}\left[\frac{\left[r(1+r)^{t-1}\right]}{\left[(1+r)^n - 1\right]}\right]$$

$$(14.5)$$

where
  P_t = the scheduled principal repayment amount for month t

$$i_t = M_{m0}\left[\frac{r\left[(1+r)^n - (1+r)^{t-1}\right]}{\left[(1+r)^n - 1\right]}\right]$$

$$(14.6)$$

where
  i_t = the interest payment in month t

Mortgages may be serviced by the original lender or by a third-party institution that has agreed to service it in return for the fee. Some mortgage contracts specify a *servicing fee* to cover the administrative costs associated with collecting interest payments, sending regular statements and other information to borrowers, chasing overdue payments, maintaining the records and processing systems, and other activities. When applicable, the servicing charge incorporated into the monthly payment is usually stated in the form of a percentage, say 0.25 percent, added to the mortgage rate.

The U.S. market also includes *adjustable-rate mortgages,* or ARMs, which reset interest payments at specified intervals to a specified short-term interest rate index. The interval between resets can be a month, six months, a year, or longer. The interest rate is usually set at a spread over the reference rate, which can be market-determined—the prime rate, for instance—or calculated based on the funding costs for U.S. savings and loan institutions, or thrifts, as indicated by one of the thrift indexes. The two most commonly consulted thrift indexes are the Eleventh Federal Home Loan Bank Board District Cost of Funds Index, or COFI, and the National Cost of Funds Index.

According to Sundaresan (1997, page 366), ARMs account for more than half the U.S. domestic mortgage business. Most borrowers prefer to reduce uncertainty by fixing their mortgage rates. To entice borrowers away from fixed-rate mortgages, ARM lenders often offer below-market interest rates for an introductory period, usually of two to five years although it can be longer. ARMs typically also have interest-rate caps, which limit the maximum rate borrowers will have to pay should market rates rise dramatically.

Balloon mortgages, like many ARMs, offer a fairly low fixed rate for the first five to seven years of their terms, after which the rate is reset; unlike with ARMs, however, this reset occurs only once. Balloon mortgages amortize their principal over a long term, usually thirty years, but require that a large "balloon" payment, equal to the original loan minus amortization, be made before maturity. This effectively transforms a long-dated loan into short-term borrowing. Balloon loans are best suited to borrowers who expect to sell their property soon; bonds securitized with them therefore have actual maturities that are shorter than the stated ones.

Graduated payment mortgages, or GPMs, are aimed at low-income buyers who expect their earnings to grow. They have fixed interest rates and terms, but their monthly payments rise according to a specified schedule. The payments start below those for level-paying mortgages with identical interest rates and terms. Each year, the payment amounts increase by a set percentage over the previous year's until they reach a specified level, at which they then remain fixed.

Because a GPM's initial payments are below the market level for the specified fixed rate, they include little or no repayment of principal. In fact, they may not even cover the amount of interest due. In that case, the interest owed is added to the principal of the loan. The outstanding balance may thus increase during the early years of the mortgage, a process known as *negative amortization.* The higher payments in the remainder of the mortgage term are designed to pay off the entire balance by maturity.

Growing equity mortgages, or GEMs, also have fixed rates and payments that increase over time, but these payments don't start at below-market levels, so there is no negative amortization. In fact, the increasing payments result in faster amortization than with a level-pay mortgage and thus a shorter term to maturity.

Mortgage Bond Risk

Although mortgages are typically long-term contracts, running for twenty to thirty years or even longer, their actual terms may be considerably shorter. This is because borrowers can elect to repay principal at a faster rate than is specified in the contract. There are number of reasons for such *prepayments.* The most common is the sale of the property securing the mortgage. Other possible causes of prepayment include default by the borrower, resulting in repossession of the securing property; a change in interest rates that makes refinancing the mortgage (usually with another lender) attractive; and destruction of the property through accident or natural disaster.

Some lending institutions penalize borrowers who retire their loans early. In the United States, prepayment penalties are levied only for commercial mortgages, not for residential ones (both are penalized in the United Kingdom). Residential lenders, therefore, cannot be certain of the cash flows they will receive. This is known as *prepayment risk.* The uncertainty of mortgages' cash flows, and the risk associated with it, is passed on to the securities backed by the loans. In this way, an MBS is similar to a callable bond, with the "call" exercisable at the discretion of the borrowers, and, as will be explained later, it can be valued using a similar pricing model.

An investor acquiring a pool of mortgages from a lender measures the amount of associated prepayment risk by using a financial model to project the level of expected future payments. Although it is impossible to evaluate with any accuracy the prepayment potential of an individual mortgage, such analysis is reasonable for a large pool of loans. This is similar to what actuaries do when they assess the future liability of an insurer that has written personal pension contracts. The level of prepayment risk for a pool of loans is lower than that for an individual mortgage.

The other significant risk of a mortgage book is that a borrower will fall into arrears or be unable to repay the loan at maturity. This is known as *default risk.* Lenders take steps to minimize default risk by assessing the credit quality of each borrower, as well as the quality of the property itself. The study described in Brown, S., et al. (1990) found that the higher the deposit paid by the borrower, the lower the incidence of default. Lenders prefer that borrowers' equity in the securing property be high enough to protect against a fall in the property's value. The typical deposit required is between 10 and 25 percent, although certain lenders will advance funds against smaller deposits, such 5 percent.

Types of Mortgage-Backed Securities

Mortgage-backed securities—also known as *mortgage pass-through securities,* because the income from the underlying pool of loans is "passed through" to the bondholders—may be formed from residential or commercial loans or from a mixture of both. Bonds backed by commercial mortgages are known as *commercial mortgage-backed securities,* CMBSs. Those created from mortgage pools that have been purchased by government agencies are known as *agency mortgage-backed securities,* or AMBSs, and are regarded as risk-free in the same way that Treasuries are. *Collateralized mortgage securities* and *stripped mortgage-backed securities* are related to pass-throughs but differ in important respects. The following discussion concerns plain vanilla mortgage-backed bonds. More complicated instruments also trade in the market.

A collateralized mortgage obligation, or CMO, differs from a pass-through in that the underlying mortgage pool is separated into different maturity classes, or *tranches,* and the cash flow distributions to investors are prioritized based on the class of the tranche they hold. It is possible to have two tranches of the same rating, but one more senior than the other. Typically, each tranche also pays a different interest rate and appeals to a different class of investors.

All classes of a single CMO receive an equal share of the interest payments; it is the principal repayments received that differ. Consider an issue with a nominal value of $100 million, $60 million of which is allocated to the class A tranche, $25 million to the class B, and the rest to class C. Holders of class A bonds receive all the principal repayments until the bonds retire, after which class B holders get all the repayments, and so on. The class A bonds thus have the shortest maturity and the highest credit rating, and the class C bonds the longest and usually the lowest. A level of uncertainty is still associated with the maturity of each bond, but it is lower than that associated with a pass-through security. CMOs are discussed in more depth below.

As its name suggests, the stripped mortgage-backed security, or *stripped bond,* is created by separating the underlying loan pool's interest and principal payments and assigning each cash flow to a different class of bonds—respectively, the *interest-only,* or IO, class, and *the principal only,* or PO, class. Stripped mortgage-backed bonds are potentially less advantageous to the issuer than a pass-through or a CMO. They are, however, liquid instruments and are often traded to hedge a conventional mortgage bond book. IOs and POs are discussed more fully below.

Cash Flow Patterns

As already noted, the exact term of a mortgage-backed bond cannot be stated with assurance at the time of issue, because of the uncertainty connected with the speed of mortgage prepayments. As a result, it is not possible to analyze these bonds using the same methods as for fixed-coupon bonds. The most common approach is to assume a fixed prepayment rate—recognizing that, in reality, it will fluctuate with changes in mortgage rates and the economic cycle—and use this to project the bond's cash flows and thus its life span. The prepayment rate selected obviously is very important. This section considers some of the ways the rate is arrived at.

Prepayment Analysis

Some market analysts base their estimates of the terms of mortgage pass-through bonds on the average life of a mortgage. Market data have suggested that the average mortgage is paid off after its twelfth year, leading to the traditional assumption of a "12-year prepaid life." This approach, however, does not take into account the effect of mortgage rates and other factors and so is not generally favored. A more common method is to use a *constant prepayment rate,* or CPR, an annualized figure based on the number of mortgages in a pool expected to be prepaid in a selected period. The *constant monthly repayment* figure, also known as the *single monthly mortality rate,* or SMM, is the percentage of the outstanding balance, minus the scheduled principal payment, expected to be repaid each month. The equation for calculating SMM is (14.7).

$$SMM = 1 - (1 - CPR)^{1/12} \qquad (14.7)$$

The convention is to estimate CPR using the prepayment standard developed by the Public Securities Association, or PSA, the domestic bond market trade association now named the Bond Market Association. The PSA benchmark—100 percent PSA—assumes a starting prepayment rate

EXAMPLE: *Constant Prepayment Rate*

Assume that the constant monthly prepayment rate for a pool of mortgages is 2 percent, the outstanding principal balance at the start of the month is $72,200, and the scheduled principal payment is $223. To estimate the amount of principal prepayment for that month, the scheduled payment is subtracted from the balance, giving a total of $71,977, which is multiplied by the constant monthly prepayment rate, giving a prepayment amount for that month of $1,439.

of 0.2 percent. This increases by 0.2 points each month until the thirtieth month, when it levels off at a constant rate of 6 percent. Stated formally, if t is the number of months from the start of the mortgage, then

$$\text{for } t < 30, \text{ CPR} = \left(\frac{6 \times t}{30}\right)\text{percent}$$

$$\text{for } t >= 30, \text{ CPR} = 6 \text{ percent}$$

This benchmark can be altered to suit changing market conditions, e.g., when interest rates are dropping and prepayments are accordingly expected to be faster. For example, the benchmark might be 200 percent PSA, which doubles the starting prepayment rate, to 0.4 percent; the monthly increase, to 0.4 percentage points; and the constant level it reaches in the thirtieth month, to 12 percent. The 50 percent PSA benchmark—for when rates are rising and prepayments are expected to slow—starts at 0.1 percent and increases by 0.1 percent, leveling off at 3 percent.

A mortgage pool's prepayment rate affects its cash flows. If the rate is zero, meaning no prepayments, the cash flows are constant during the mortgages' lifetimes. In a fixed-rate mortgage, the ratio of principal to interest payment changes each month as more and more of the loan amortizes. For a pass-through security issued today, and whose coupon, therefore, reflects the current market rate, the payment pattern at a zero percent prepayment rate will resemble the bar chart in **FIGURE 14.2**.

When the CPR is greater than zero, as in the 100 percent PSA and 200 percent PSA models, the principal payments will increase during the early years of the mortgages, then level off before declining for the

FIGURE **14.2** *Ratio of Interest to Principal Payments for a Mortgage Pass-Through Security with 0 Percent CPR*
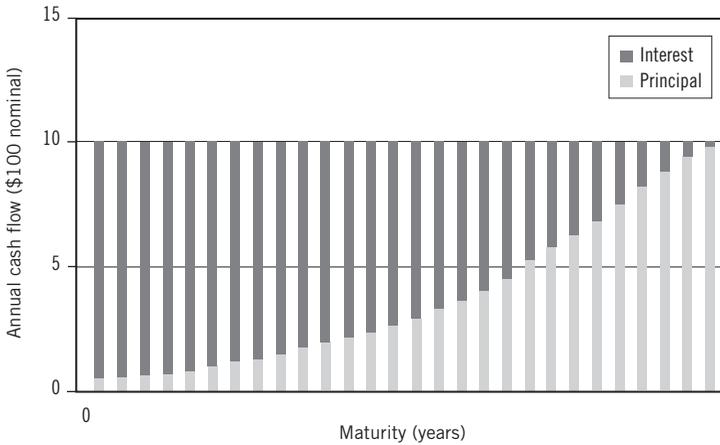
FIGURE **14.3** *Payment Ratio in a 100 Percent PSA Model*
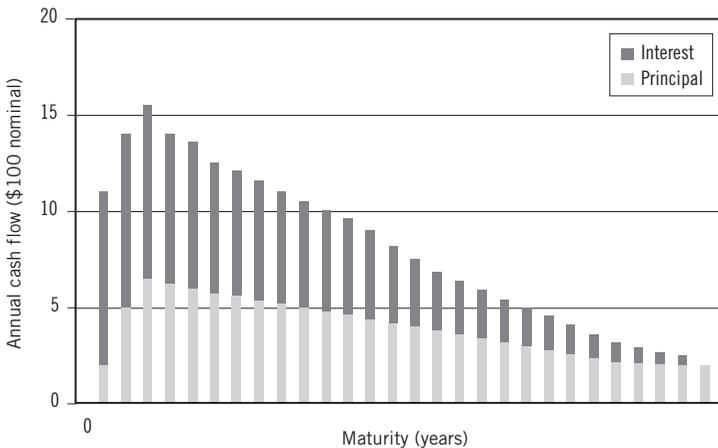
remainder of the term, as the outstanding balances become so small that the scheduled principal payments are insignificant. **FIGURE 14.3** shows the payment pattern for a single 9 percent, 30-year loan at 100 percent PSA.

The prepayment volatility of a mortgage-backed bond depends on the interest rates of the underlying mortgages. Volatility has been observed to be highest when the mortgages' rates are between 100 and 300 basis points above current rates: at the bottom of this range, any fall in interest rates tends to produce a sudden increase in refinancings and thus prepayments, while at the top of the range, a rise in rates leads to a decrease in refinancings and prepayments.

The actual cash flows of a mortgage pass-through depend, of course, on the cash flow patterns of the underlying mortgages. The projected monthly payment for a level-paying fixed-rate mortgage is given by formula (14.8).

$$\overline{I_t} = \overline{M}_{mt-1} \left[\frac{r(1+r)^{n-t+1}}{(1+r)^{n-t+1} - 1} \right] \tag{14.8}$$

where
 $\overline{I_t}$ = the projected monthly mortgage payment for month t
 $\overline{M}_{mt-1}$ = the projected mortgage balance at the end of month $t - 1$, assuming that prepayments have occurred in the past
 n = mortgage term, in months
 r = annualized interest rate

The interest portion of the projected monthly mortgage payment, $\overline{i_t}$, is calculated using equation (14.9).

$$\overline{i_t} = M_{mt-1}i \tag{14.9}$$

where
 i = the interest amount

Equation (14.9) states that the projected monthly interest payment can be obtained by multiplying the mortgage balance at the end of the previous month by the monthly interest rate. The expression for calculating the projected scheduled monthly principal payment for any month is (14.10).

$$\overline{p_t} = \overline{I_t} - \overline{i_t} \tag{14.10}$$

where
 $\overline{p_t}$ = the projected scheduled principal payment for the month t

The projected monthly principal prepayment, which is an expected figure only and not a model forecast, is given by (14.11).

$$\overline{pp_t} = SMM_t\left(\bar{M}_{mt-1} - \overline{p_t}\right) \tag{14.11}$$

where
 $\overline{pp_t}$ = the projected monthly principal prepayment for month t.

Equations (14.9) through (14.11) calculate values for
❑ the projected monthly interest payment
❑ the projected scheduled monthly principal payment
❑ the projected monthly principal prepayment

Combining these values, as in (14.12), gives the total cash flow that a holder of a mortgage-backed bond receives in any month.

$$cf_t = \overline{i_t} + \overline{p_t} + \overline{pp_t} \tag{14.12}$$

where
 cf_t = the cash flow received in month t

Using a projected prepayment rate enables analysts to evaluate mortgage-backed bonds. The original PSA benchmarks were based on the observation that prepayment rates tend to stabilize after the first thirty months of a mortgage and assumed a linear increase in these rates. They do not reflect seasonal variations in prepayment patterns nor the different behavior patterns of different types of mortgages.

The PSA benchmarks can be applied to assumptions about defaults to produce the PSA *standard default assumption* (SDA) benchmark, which is used to assess the potential default rate for a mortgage pool. This benchmark is used only for nonagency mortgage-backed bonds, since agency securities are guaranteed by one of the three government or government-sponsored agencies. The standard benchmark, 100 SDA, assumes that the default rate in the first month of a mortgage is 0.02 percent and that the rate increases in a linear fashion by 0.02 percentage points each month until the thirtieth one, when it levels off at 0.60 percent. It remains at 0.60 percent until month 60, when it begins to fall to 0.03 percent until it reaches 0.03 percent in month 120, remaining at that level for the rest of the mortgage term. The other benchmarks have similar patterns.

Prepayment Models

The PSA standard benchmark discussed in the previous section assumes certain prepayment rates and can be used to calculate the prepayment proceeds of a mortgage. It is not, strictly speaking, a prepayment *model,* because it

cannot be used to estimate actual prepayments of a mortgage pool. A pre-payment model does attempt to do this, by modeling the statistical relationships between the various factors that affect the level of prepayment. These factors are the current mortgage rate, the characteristics of the mortgages in the pool, seasonal influences, and the general business cycle.

The prevailing mortgage interest rate and its spread above or below the original contract rate are probably the most important factors in determining the level of prepayment, since they influence borrowers' decisions about refinancing their mortgages. If the current rate is materially below the original one, for instance, borrowers will prepay their loans. Because the mortgage rate reflects the general bank base rate at the time, the level of market interest rates has the greatest impact on mortgage prepayment levels. The current mortgage rate also affects housing prices: when mortgages are seen as "cheap," people are more likely to purchase homes, driving prices up. This rate has an impact on how likely it is that people will repay early. The pattern followed by mortgage rates since the original loan also has an impact on prepayments, a phenomenon known as *refinancing burnout*. Presumably, the point is that if mortgage rates have been trending consistently downward for a long period, everyone who can refinance will have done so, so the prepayment rate slows.

The housing market and mortgage activity also appear to follow seasonal patterns. Spring and summer see the strongest action, winter the weakest.

Generally, mortgage activity reflects the economic cycle. An economy that is performing strongly will see increasing levels of borrowing. However, as the economy picks up, this brings the risk of higher inflation, such that interest rates need to be raised by the Federal Reserve to cool the economy. This influences mortgage activity and also paradoxically sees higher repayments, as people seek to come out of debt. However, if the slowdown is not graduated, but turns into a recession, prepayment activity falls off as people are not able to repay early. The various factors have been combined to calculate expected prepayment levels. Fabozzi (1997), for example, cites equation (14.13), which a U.S. investment bank uses to calculate expected prepayments.

$$(14.13)$$

$$\text{Monthly prepayment rate} = (\text{Refinancing Incentive}) \times (\text{Season multiplier}) \times \\ \times (\text{Month multiplier}) \times (\text{Burnout})$$

Collateralized Mortgage Securities

Collateralized mortgage obligations, which were described earlier, account for a large segment of the U.S. debt capital market. The majority of CMOs are issued by government-sponsored agencies. They thus have Treasury

bond credit quality, but they offer significantly higher yields—a combination that makes the instruments attractive to a range of institutional investors, as does the ability to tailor their characteristics to specific investment needs. This section reviews some of the newer CMO structures.

The CMO market in the United States experienced rapid growth during the 1990s. According to the securitization newsletter *Asset-Backed Alert,* a high of $324 billion in the securities was issued in 1993; by 1998, this figure had fallen to just under $100 billion. The growth of the market brought with it a range of new structures. To meet the demands of bondholders desiring lower exposure to prepayment risk, for example, *planned amortization classes,* or PACs, and *targeted amortization classes,* TACs, were introduced. *Very accurately defined maturity,* or VADM, bonds, which are guaranteed not to extend beyond a stated date, were created to remove the uncertainty concerning the term of mortgage-backed bonds. In the United Kingdom and certain overseas markets, mortgage-backed bonds pay floating-rate coupons, and the desire of foreign investors to have something similar in the U.S. domestic market led to the creation of bonds with coupons linked to LIBOR.

Originally, mortgage-backed bonds were created from individual underlying mortgages. CMOs created in this manner are known as *whole loan.* In contrast, the mortgages underlying agency-issued CMOs have already been pooled and securitized, usually as pass-throughs. Whole-loan CMOs are thus based on cash flows from an entire pool of individual mortgages rather than on a pass-through security formed from this pool. As with agency CMOs, however, the underlying mortgages in a whole-loan pool generally have the same risk, maturity, and interest rate.

Whole-loan CMOs also differ from agency bonds in the size of the underlying mortgages: those backing agency bonds are limited to a stated maximum size and so tend to be smaller than the ones backing whole-loan CMOs, which may include jumbo loans. Another difference between whole-loan and agency CMOs concerns *compensating interest.* Virtually all mortgage-backed securities pay principal and interest monthly, on a fixed coupon date. The underlying mortgages, however, may be paid off on any day of the month. Agency-issued securities guarantee their bondholders interest payments for the complete month, even if the underlying mortgage has been paid off ahead of the coupon date and so has not generated any interest for that month. Whole-loan CMOs do not offer this guarantee. Holders of these bonds may thus receive less than one month's interest on the coupon date. Some issuers, though not all, will make a compensating interest payment to bondholders to cover the shortfall.

Following Sundaresan (1997, page 389), the primary features of CMOs in the U.S. market may be summarized as follows:

❏ *Credit quality.* CMOs issued by U.S. government agencies have the same guarantee against default as agency pass-through securities and so may be considered risk-free. They therefore do not require any form of credit insurance or credit enhancement. Whole-loan CMOs do not carry a government guarantee and are rated by credit rating agencies. Most of them have triple-A ratings, either because their mortgage pools generate cash flows well in excess of what is required to service the interest obligations of all tranches, or because of the high quality of the issuing vehicles or through some credit enhancement. CMOs also tend to be significantly overcollateralized.

❏ *Interest frequency.* The notes usually pay semiannual or quarterly coupons. The underlying mortgages, however, pay interest monthly or almost daily. The cash generated between coupon dates is reinvested at money market rates. The preferred reinvestment vehicle is a guaranteed investment contract (GIC), but only a few banks and insurance companies offer GICs, so most issuers settle for money market accounts. However, the providers of GICs will usually accept much smaller deposits, sometimes as little as $100,000, than are required for interbank deposits, which are another option for investing the excess cash.

❏ *Cash flow profile.* CMOs' profiles are based on an assumed prepayment rate, known as the *pricing speed,* which reflects current market expectations for prepayment levels and market interest rates.

❏ *Maturity.* Originally, virtually all CMOs were created from underlying mortgage collateral with 30-year stated maturities. During the 1990s, issues were formed from shorter-dated collateral, including 5- to 7-year and 15- to 20-year mortgages. Still, most CMOs have long terms.

❏ *Trading conventions.* CMOs trade on a yield basis, as opposed to a price basis, and are usually quoted as a spread over the yield of the Treasury security with the most similar maturity. Like their cash flows, CMOs' yields are based on assumed prepayment rates. CMOs are settled on a T+3 basis. Agency issues are cleared via an electronic book-entry system run by the Federal Reserve, known as Fedwire; whole-loan issues are cleared using either physical delivery or by electronic transfer. New issues of CMOs settle from one to three months after the initial offer date.

CMOs have two basic structures: *sequential pay* and *planned amortization class,* or PAC.

Sequential Pay

One of the demands that drove the design of CMOs was for mortgage-backed bonds with a wider range of maturities. Most CMO structures redirect principal payments sequentially to individual tranches, according

FIGURE **14.4** *Generic CMO Sequential Structure*

TRANCHE	PRINCIPAL	COUPONS	AVERAGE LIFE (YEARS)	YIELD
A	100	7.00%	2.5	2-year benchmark plus 80 bps
B	250	7.00%	5	5-year benchmark plus 100 bps
C	75	7.00%	10	10-year benchmark plus 120 bps
Z	75	7.00%	20	30-year benchmark plus 150 bps

to each one's stated maturity. That is, principal payments go to the tranche with the shortest stated maturity until it is completely redeemed and are then allocated to the next maturity class, continuing in this manner until all the bonds in the structure are retired. Sequential-pay CMOs are attractive to a wide range of investors, particularly those with shorter investment horizons, because they can purchase only the CMO class whose maturity terms meet their requirements. Investors with longer investment horizons are protected from prepayment risk in the early years of the issue, when principal payments are used to pay off the shorter-dated tranches.

FIGURE 14.4 shows a typical generic CMO sequential structure. The collateral cash flows are allocated to each tranche in specified order. The first tranche is allotted both its coupon and any prepayments. The remaining tranches receive only their coupon payments until the first one is fully retired.

Planned Amortization Class

The first CMO with a PAC structure was issued in 1986, in response to a demand for less interest-rate-sensitive mortgage-backed structures that was sparked by a period of sustained falls in market rates. PAC structures are designed to reduce prepayment risk and the volatility of the weighted average life measure, which is related to the prepayment rate. PAC CMOs have principal-payment schedules that are unaffected by changes in prepayment rates. These schedules, similar to those of corporate-bond sinking funds, are based on the minimum amounts of principal cash flow that will be produced by the underlying mortgage pool at two different prepayment rates. These rates together define a range known as a *PAC band.*

PAC bands are defined by prepayment rates set at a low and a high PSA standard—for example, 50 percent PSA and 250 percent PSA. This

constrains the amount of principal repayment: in the early years of the issue, the lower standard serves as a floor for the minimum principal received, while later in the bond's life the payment schedule is constrained by the upper PSA standard. The total principal cash flows under the PAC schedule determine the value of PACs in an issue.

In the PAC structure, the uncertainty of principal payments is directed to another class of security, i.e., another tranche in the CMO, known as the *companion,* or *support,* class. When prepayment rates are high, companion issues support the main PACs by absorbing any principal prepayments that are in excess of the PAC schedule; when the rate falls, the companion amortization is delayed if principal prepayments are not sufficient to reach the minimum stipulated by the PAC band. Accordingly, when prepayment rates are high, the companions' average life shortens; when rates are low, their average life lengthens. Within the set of PACs and the set of companions, the principal cash flows can be allotted sequentially, as in the sequential-pay structure.

PACs exhibit lower price volatility than other mortgage securities. When the prepayment rates are within the PAC band, their prices are fairly stable; when rates move outside the band, volatility increases by a smaller amount than for non-PAC bonds, because the prepayment risk is transferred to the companions. For this reason, PAC issues trade at lower spreads to the Treasury yield curve than do other issues with similar maturities. The companion bonds are always priced at a wider spread than the PACs, reflecting their higher prepayment risk.

Within the CMO structure may be some PAC bonds with less prepayment risk than others, known as *Type II* and *Type III* PACs. A Type II PAC has a narrower band than a standard PAC, thus reducing prepayment risk. If prepayment rates remain within their narrower bands, Type II PACs trade like standard PACs; if rates move outside their bands, the extra cash flow is redirected to the companions only if the rates move *above* the range. Otherwise, there is no excess and the companion amortization is delayed. Type II PACs are second in priority to the standard PACs and so trade at a higher yield. If prepayment rates remain high for an extended period and all the companions are redeemed, the Type II PACs take over the function of companion and, with it, the higher prepayment risk. Type III PACs function like Type II PACs, but their band ranges are even tighter.

The upper and lower limits of a PAC band may "drift" during the life of the CMO, regardless of actual prepayment rates. This drift results from the interaction of actual cash prepayments with the bands and from changes in the collateral balance and in the ratio between the nominal values of the PACs and the companions. The type of drift and its impact depend on the prepayment rate, as illustrated in the following scenarios:

❏ *The prepayment rate lies within the current band.* If this situation continues, the upper limit will rise, since prepayments have been below the maximum level, leaving more companion issues to receive future prepayments. The lower limit rises as well, because prepayments have been above the minimum level, resulting in less collateral to generate future principal payments. The upper limit, though, tends to rise more quickly than the lower one, widening the range.

❏ *The prepayment rate lies above the band.* If this continues, the number of companions available to receive faster prepayments will fall, and the band will narrow, its upper and lower limits converging completely when all the companion bonds have been redeemed. At that point, the PAC will trade as a conventional sequential-pay security until it is redeemed.

❏ *The prepayment rate lies below the band.* The upper limit will drift upward, because more companion bonds are available to receive a greater level of prepayments in the future; the lower band may also rise by a small amount. This type of drift is relatively rare, however, since PACs have the highest priority of all classes in a CMO structure until the payment schedule is back on track.

Band drift occurs over time and is sometimes not noticeable. Significant changes in the band levels only take place if the prepayment rate is outside the band for a prolonged period. Prepayment rates that move outside the bands for short spells do not have any effect on the bands.

Targeted Amortization Class

Targeted amortization class, or TAC, bonds were created to cater to investors who require prepayment protection but at a higher yield than would be available with a PAC. Essentially, a TAC is a PAC whose "band" consists of only one standard prepayment rate. Like PACs, TACs amortize principal according to a schedule when the actual prepayment rate accords with this standard and, when the rate moves above it, use the extra principal amounts to pay off companion bonds. They differ from PACs mainly in taking on extra prepayment risk when prepayments fall below the rate required to maintain the payment schedule, extending the issue's average life. Because one element of the PAC band is removed, TACs trade at a higher yield.

The preference for TACs over PACs is a function of the prevailing interest rate environment. When current rates are low or are expected to fall, there is a risk of prepayments increasing, reducing the average life of the bond. In this scenario, investors may be willing to forgo the protection against an extension in the bond's average life provided by PACs, deeming it unlikely to be required, and take the extra yield offered by the TAC.

Z-Class Bonds

The *Z-class,* or *Z,* bond ranks below all other classes in the CMO's structure and pays no cash flows for part of its life, functioning essentially like a zero-coupon bond (hence its name). When the CMO is issued, the Z bond, also known as an *accrual,* or *accretion,* bond, has a relatively small nominal value. At the start of its life, it pays out cash flows on a monthly basis, as determined by its coupon. However, when the Z bond itself is not receiving principal payments, its cash flows are used to retire some of the principal of the other classes in the structure. In their place, the bond receives credits, which increase its face value by the amount of the forgone coupon. As a result, its principal amount is higher at the end of its life than at the start. When all classes of bond ahead of the Z bond have been retired, the Z bond itself starts to pay out principal and interest cash flows.

In a conventional sequential-pay structure, the other classes in the CMO receive some of their principal prepayments from the Z bond, which lowers their average-life volatility. Z bonds are an alternative for investors who might otherwise purchase Treasury zero-coupon bonds. Like zero coupons, these bonds have no reinvestment risk, but they have higher yields than Treasury strips with similar average lives.

Interest-Only and Principal-Only Classes

As noted earlier, stripped mortgage-backed securities, or stripped bonds, are created by splitting the cash flows payable by a pool of mortgages into interest and principal payments and assigning the two different streams to two different classes of bonds: interest-only, or IO, and principal-only, or PO, bonds.

The PO bond is similar to a zero-coupon in that it is issued at a discount to par value. The PO bondholder's return is a function of the rapidity at which prepayments are made; the quicker the prepayment, the higher the return. This is like the buyer of a zero-coupon bond receiving the maturity payment ahead of the redemption date. The highest possible return for the bondholder would occur if all the mortgages were prepaid the instant after the PO bond was bought. A low return occurs if all the mortgages are held until maturity, so that there are no prepayments.

The price of a PO bond fluctuates with mortgage interest rates. As noted earlier, the majority of mortgages are fixed-rate loans. If mortgage rates fall below the PO bond's coupon rate, the volume of prepayments should increase as the individuals holding the underlying loans refinance them, speeding the stream of payments to the bondholder. The PO's price will rise both because of the faster cash flows and because the flows are now discounted at a lower rate. The opposite happens when mortgage rates rise.

An IO bond has no par value, since it is essentially just a stream of cash flows, consisting of the interest payments on the underlying mortgage principal outstanding. These cash flows cease once the principal is redeemed, so a higher rate of prepayment depresses the IO price. The risk for investors is that prepayments occur so quickly that they don't receive enough interest payments to cover what they paid for their bond.

An IO's price, like that of a PO, is a function of mortgage rates. The relationship is more complicated, though. When rates fall below the bond coupon, increasing the expected prepayment rate and so reducing expected cash flows to the IO, the IO price falls as well, even though the cash flows themselves are discounted at a lower interest rate. When mortgage rates rise, the outlook for IO cash flows improves, but they will be discounted at a higher rate, so the bond's price may move in either direction. Generally, though, IOs' prices move in the same direction as interest rates—a curious characteristic for a bond.

PO and IO issues have much greater interest rate sensitivity than the pass-through security from which they are created, exhibiting extreme price volatility when the mortgage rate is moving in either direction. Note that POs and IOs can both be created from the mortgage pool underlying one pass-through security, so their combined modified durations must equal that of the original bond.

FIGURE 14.5 compares the price sensitivities of a 7 percent pass-through security and of the IO and PO created from it. Note that the pass-through's price is not particularly sensitive to a fall in the mortgage rate below its coupon rate of 7 percent. This illustrates the *negative convexity* of pass-through securities (discussed further below). The price sensitivities of the two strip issues are very different. As the mortgage rate rises above the coupon rate, the PO's price falls dramatically, while the IO's rises. On the other hand, the IO's price drops significantly when mortgage rates fall below the coupon rate.

Early strip issues were created with an unequal amount of coupon and principal, resulting in a synthetic coupon rate that was different from the coupon on the underlying bond. The early strips were not IOs or POs. Instead principal was distributed unequally among different classes of the pass-through, which were associated with correspondingly different coupons, all of which were different from the interest rate on the underlying mortgages. These instruments were known as *synthetic-coupon pass-throughs.* Nowadays it is more typical to allocate all the interest to one class, the IO, and all of the principal to the PO.

The most common CMO structures have a portion of their principal split into IO and PO bonds. Some CMOs, though, are made up entirely of IO and PO bonds. The amount of principal used to create stripped securities depends on investor demand.

FIGURE **14.5** *Sensitivity to Changes in the Mortgage Rate of Pass-Through, IO, and PO Prices*
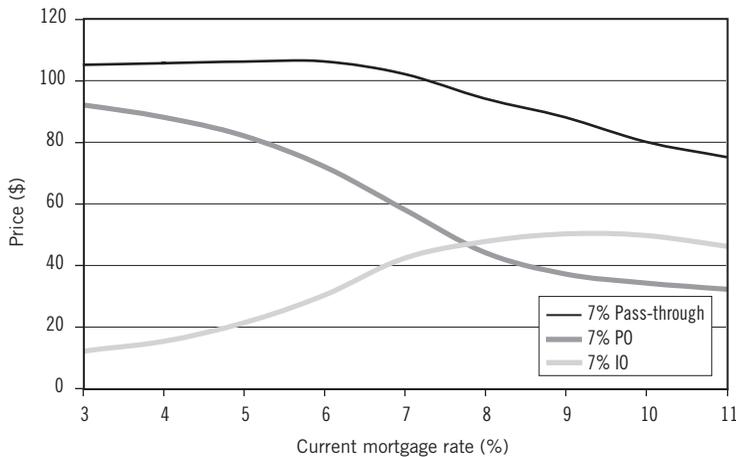
IO issues created from a class of CMOs known as *real estate mortgage investment conduits,* or REMICs, can have quite esoteric terms. For example, the IO classes might be issued with a small amount of principal, known as the nominal balance. The cash flows for IOs with this structure are created by amortizing and prepaying the nominal balance. Because the nominal balance is small, the IO has a multidigit coupon and very high price. Ames (1997) cites the example of an IO paying a 1,183 percent coupon and priced at $3,626-12, which implies a bond whose price has risen by more than thirty times its original face value.

Strips created from whole-loan CMOs trade differently from those issued out of agency CMOs, due to the nature of the underlying collateral: they are viewed differently by investors and so their secondary market characteristics are less liquid.

IOs and POs are particularly useful in whole-loan CMOs, because of the way their coupons are calculated. Agency CMOs pay fixed coupons, but a whole-loan CMO's coupons are based on the weighted average of the underlying mortgages' coupons. During the life of a whole-loan issue, its coupon will change as prepayments alter the amount of principal. Stripping a portion of the principal and interest cash flows from the underlying mortgages leaves collateral with a more stable average life, preserving the coupon payments of all issues within the structure.

Nonagency CMO Bonds

The structure and terms of nonagency CMOs do not differ significantly from those of agency CMOs. The key feature of nonagency CMOs is that they are not guaranteed by government agencies and so carry an element of credit risk, in the same way that corporate bonds expose investors to credit risk. To attract investors, most nonagency CMOs incorporate an element of credit enhancement. This usually results in a triple-A rating. Indeed, a large majority of nonagency CMOs are triple-A rated, with very few falling below double-A. All four major private credit rating agencies analyze and rate nonagency CMOs. The rating granted a particular CMO depends on several factors, including

- ❑ the term of the underlying loans
- ❑ the loans' size, whether conforming or jumbo
- ❑ the loans' interest basis, whether level-pay fixed-rate, variable, or other
- ❑ the type of property securing the loans
- ❑ the geographical area within which the loans were made
- ❑ the loans' purpose, whether for a first purchase or a refinancing

Credit Enhancements

Nonagency CMOs may have either an external or an internal credit enhancement. An external enhancement is a guarantee by a third party to cover losses on the issue. Usually, a set percentage, such as 25 percent, of the issue of the face value is guaranteed, rather than the entire issue. The guarantee can take the form of a letter of credit, bond insurance, or *pool insurance.* Pool insurance policies are issued by specialized agencies to cover losses arising from a default or foreclosure. Usually the coverage is for a cash amount that remains the same during the life of the pool. Some policies, however, are set up so that the coverage falls over time. Since only defaults and foreclosures are covered, investors wishing to be protected against other types of loss, such as sickness or death, must arrange their own insurance.

External credit enhancement still leaves a CMO exposed to credit risk—the risk associated with the insurance provider. That is, the issue can be downgraded because of a deterioration in the credit quality of the provider. Investors who purchase nonagency CMOs must ensure that they are satisfied with the credit quality of the third-party guarantor, as well as with the quality of the underlying mortgage pool.

An external credit enhancement has no impact on the cash flow structure of the CMO. Internal credit enhancements, which are generally more complex, sometimes do affect the cash flow. The two most common types

of internal credit enhancement are *reserve funds* and *senior/subordinated structures.*

There are two types of reserve funds: *cash reserve funds* and *excess servicing spread accounts.* A cash reserve fund is a separate fund into which a portion of the profits from the bond's issuance have been deposited and invested in short-term bank securities. In a default, the cash in the fund is used to compensate investors who have suffered capital losses. A cash reserve fund is often set up in conjunction with another type of credit enhancement, such as a letter of credit.

An excess servicing spread account is a separate account into which is deposited the excess spread or cash left after the mortgages' coupon, servicing fees, and other expenses have been paid. For instance, if an issue's gross weighted average coupon—that is, the average of the interest rates of the underlying mortgages, before adjustment for service fees—is 7.50 percent, the service fee is 0.10 percent, and the net weighted average coupon (the average of the pass-through rates of the mortgage securities backing the CMO) is 7.25 percent, then the excess is 0.15 percent. This amount is paid into the spread account, which will grow steadily during the bond's life. As with the cash reserve fund, the cash in the account can be used to compensate for losses on the bond that affect investors.

Senior/subordinated structures are the most common type of internal credit enhancement encountered in the market. Essentially, the CMO is divided into two classes of bonds, one senior and the other subordinated. The latter absorbs all the losses arising from default or other cause, leaving the senior class unaffected. The subordinated bonds clearly have higher risk than the senior class and so trade at a higher yield. Most senior/subordinated arrangements incorporate a *shifting interest structure,* which redirects prepayments from the subordinated to the senior class. This alters the cash flow characteristics of the senior notes, whether or not defaults or similar events occur.

Commercial Mortgage-Backed Securities

The loans underlying commercial mortgage-backed securities are, as the name implies, for commercial, as opposed to residential, properties. CMBSs trade like other mortgage securities but differ in structure.

Issuing a CMBS

Commercial mortgages are loans made against commercial property. A CMBS is created from a pool, or *trust,* of commercial mortgages, whose interest and principal payments back the bond's cash. It is rated in the same way as a residential mortgage security and usually includes a credit

enhancement. The issuers of CMBS bonds are generally the same as those of RMBS, although in the United States there is of course the agency market for RMBS. However, banks are large issuers of CMBS. Commercial mortgage-backed issues typically have sequential-pay structures in which the bonds are retired in the order of their credit ratings. That is, the triple-A-rated bonds will be retired ahead of the double-A-rated bonds, and so on. This is similar to different tranches that have different ratings.

Borrowers who repay their commercial mortgages early are penalized, usually with an interest charge levied on the final principal. Early payments may also be prohibited by a prepayment "lockout" in the mortgage contract. Commercial mortgages' early-prepayment protection is repeated in the bonds created from them, sometimes in the form of call protection. The ratings of individual issues in the structure already provide some protection, since the highest-rated ones are paid off first. The highest-rated classes also have the most protection from loss of principal because of default of an underlying mortgage, since this will affect the lowest-rated bonds first.

In addition to their early-retirement protection, commercial mortgages differ from residential ones in that many of them are balloon loans. As explained above, a large part of the principal of a balloon loan is paid off on a single date. This makes CMBSs similar to plain vanilla, or *bullet,* bonds, an attraction for investors who prefer more certainty about terms to maturity.

Types of CMBS Structures

In the U.S. market, there are five types of CMBS structures: *liquidating trusts, multiproperty single borrower, multiproperty conduit, multiproperty nonconduit,* and *single-property/single-borrower.*

Single-property/single-borrower loans are self-explanatory and refer to a loan taken out by one obligor on one property. Liquidating trusts account for a small part of the market by face value. They are issued against nonperforming loans and are therefore often referred to as *nonperforming CMBSs.* Among their distinctive features are a *fast-pay structure,* which signifies that all cash flows from the mortgage pool are used to redeem the most senior bond first, and *overcollateralization,* meaning that the collective face value of the bonds created is significantly lower than that of the underlying loans. Because of their overcollateralization, CMBSs are paid off earlier than other mortgage securities and receive cash flows from only a portion of the underlying loans. They are usually issued with relatively short average lives. A target date is set for paying off all the CMBS classes. The bonds usually have a provision to raise the coupon rate should this target not be met.

The single-borrower, multiproperty structure accounts for a large part of the CMBS market. Bonds having this structure are *cross-collateralized*—that is, the property used as collateral for each individual loan is also pledged against every other loan in the underlying pool. Another provision in the structure, known as *cross-default,* allows the lender to call every loan in the pool if any one of them defaults. Together, cross-collateralization and cross-default assure that sufficient cash flow is available to meet the collective debt obligations of all of the loans. These properties boost the issue's credit rating. By the same token, the properties of the liquidating trust enhance the rating of the associated CMO. A *property release provision* in the structure prohibits the lender from removing or prepaying the stronger loans in the book, making it difficult to cherry-pick good loans. Another provision protecting investors against this risk that is commonly included in the structure prevents the issuer from substituting one property for another, which means that the credit quality of original underlying loans is guaranteed. These two provisions also enhance the credit of the CMBS.

Conduits are commercial lending entities set up solely to generate collateral to be used in securitization. They are required by more-frequent issuers. The major investment banks have all established conduit arms. Conduits are responsible for originating collateral that meets the investor's requirements on loan type (whether amortizing or balloon, and so on), loan term, geographic spread of the properties, and the time that the loans were struck. Generally, *pool diversification* in terms of size and location is desirable, since this reduces the default risk for the investor. After it has generated the collateral, the conduit structures the deal with terms similar to those of CMOs but with the additional features described in this section.

The multiproperty nonconduit structure is useful for originators with large pools of assets that do not expect to tap the market frequently. This differs from the single-borrower, multiproperty structure only in terms of how often it issues. It has more in common with a single-issuer, residential MBS transaction than with the conduit structure, with the added features of cross-collateralization and cross-default.

Evaluation and Analysis of Mortgage-Backed Bonds

Pricing and hedging mortgage-backed securities are complex subjects. The following sections present a brief introduction. Readers interested in more-in-depth discussions should consult the Reference section for appropriate texts.

Term to Maturity

Term to maturity is a very important measure. It represents the period during which a bond's return is being generated and helps determine its sensitivity to changes in market interest rates. It is also a basis for comparing bonds. Such analyses and comparisons cannot be made for mortgage pass-throughs using their stated maturities, because these can be reduced by prepayments. Instead, for evaluating and comparing these securities, the market uses estimated values—specifically, *average life* and duration.

The average life of a mortgage pass-through security, also known as its *weighted-average life,* is the weighted-average time to return of a unit of principal payment that comprises both scheduled payments and prepayments. The time from the end of the term measured by the average life to the final scheduled principal payment is the bond's *tail.* Average life is derived using equation (14.14).

$$\text{Average life} = \frac{1}{12} \sum_{t=1}^{n} \frac{t\,(\text{Principal received at } t)}{\text{Total principal received}} \qquad (14.14)$$

where

n = the number of months remaining in the bond's stated term to maturity

As explained in chapter 2, a bond's modified or Macaulay's duration is the average time to receipt of its cash flows, weighted according to their present values. To compute a mortgage-backed bond's duration, it is necessary to project its cash flows using an assumed prepayment rate. These projections, together with the bond price and the periodic interest rate, derived from the yield, may then be used to arrive at the bond's periodic duration, which is divided by twelve (or four, in the case of a bond that pays quarterly) to arrive at its duration in years.

Calculating Yield and Price: Static Cash Flow Model

There are a number of ways to calculate the yield on a mortgage-backed bond. One of the most common is the *static cash flow model.* This assumes a single prepayment rate to estimate the cash flows for the bond and does not take into account how changes in market conditions might affect the prepayment pattern.

As explained in chapter 1, yield is conventionally defined as the rate at which a bond's expected cash flows must be discounted so that the sum of their present values will equal the bond's clean price—that is, the price excluding any accrued interest. This is known as the bond's redemption

yield or yield to maturity. For mortgage-backed bonds, it is the *cash flow* or *mortgage yield.* This is calculated by plugging projected cash flows based on an assumed prepayment rate into formula (14.15).

$$P = \sum_{n=1}^{N} \frac{C(t)}{(1 + ri/1200)^{t-1}} \qquad (14.15)$$

where
    N = number of interest periods
    t = each interest payment date
    ri = the mortgage yield

Equation (14.15) computes the yield for a bond making monthly coupon payments (as the majority of mortgage-backed bonds do, although some pay quarterly). For purposes of comparison, this figure must be converted to a bond-equivalent yield. Equation (14.16) derives the annualized bond-equivalent yield for a mortgage-backed bond paying monthly coupons.

$$rm = 2\left[(1 + ri_M)^6 - 1\right] \qquad (14.16)$$

where
    rm = the bond-equivalent yield
    ri_M = the interest rate that will equate the present value of the projected monthly cash flows for the mortgage-backed bond to its current price

In the U.S. and U.K. markets, the basis for comparison is the relevant government bond yield, which is semiannual. Equation (14.17) derives the equivalent semiannual yield. (See chapter 1 for a discussion of converting from one payment basis to another.)

$$rm_{s/a} = (1 + ri_M)^6 - 1 \qquad (14.17)$$

Cash flow yield calculated in this way is essentially a redemption yield calculated assuming a prepayment rate to project the cash flows. As such, it has the same drawbacks as the redemption yield for a plain vanilla bond: it assumes that all the cash flows will be reinvested at the same interest rate and that the bond will be held to maturity. In fact, the potential inaccuracy is even greater for a mortgage-backed bond because the frequency of interest payments is higher, which makes the reinvestment risk greater. The final yield of a mortgage-backed bond depends on the performance of the mortgages in the pool—specifically, their prepayment pattern.

Given the nature of a mortgage-backed bond's cash flows, its exact yield cannot be calculated. Market participants, however, commonly compare an MBS's cash flow yield to the redemption yield of a government bond with a similar duration or a term to maturity similar to the MBS's average life. The usual convention is to quote the spread over the government bond.

As noted in chapter 3, it is possible to calculate a bond's price given its yield and vice versa. As with a plain vanilla bond, a mortgage-backed bond's price is the sum of the present values of its projected cash flows. The discount rate used to derive the present values is the bond-equivalent yield converted to a monthly basis.

The cash flows of IO and PO bonds are dependent on the cash flows of the underlying pass-through security, which is itself dependent on the cash flows of the underlying mortgage pool. To calculate the prices of these strips, their cash flows must be estimated using a prepayment rate. The price of an IO is the present value of the projected interest payments; the price of the PO is the present values of the projected principal payments, comprising the scheduled principal payments and the projected principal prepayments.

Bond Price and Option-Adjusted Spread

The prepayment option of the holders of mortgages underlying a mortgage security is essentially a call option. Not surprisingly, then, mortgage securities often behave like callable bonds.

The optionality of a mortgage-backed bond, and the volatility of its yield, frequently have a negative impact on the bondholders. This is for two reasons: the actual yield realized during the holding period has a high probability of being lower than the anticipated yield, which was calculated on the basis of an assumed prepayment level, and mortgages are frequently prepaid when the bondholders will suffer the most—that is, when rates have fallen, leaving them to reinvest repaid principal at a lower rate.

For investors these features represent the biggest risk of holding a mortgage security, and market analysts attempt to measure and quantify it. They usually do so using a form of option-adjusted spread analysis. In this approach, the value of the mortgagors' prepayment option is expressed as a basis-point penalty that is subtracted from the bond's expected yield spread. The penalty is calculated using a binomial model or a simulation model to generate a range of future interest rate paths only some of which will cause mortgagors to prepay. Then the paths resulting in prepayment are evaluated with respect to their impacts on the mortgage bond's expected yield spread over a government bond.

The OAS-derived yield spread is based on the present values of expected cash flows discounted using government bond–derived forward rates. The spread between the cash flow yield and the government bond yield is based on yields to maturity. The OAS spread is added to the entire yield curve, whereas a yield spread is over a single point on the government bond yield curve. For these reasons, the two spreads are not strictly comparable.

Because OAS analysis takes into account a mortgage-backed bond's option feature, it is less affected by a change in interest rates or the yield curve, which affect prepayments, than the bond's yield spread. Assuming a flat yield curve, the relationship between the OAS and the yield spread is expressed in equation (14.18).

$$\text{OAS} = \text{Yield spread} - \text{Cost of option feature} \qquad (14.18)$$

This relationship can be observed when yield spreads on current-coupon mortgages widen during declines in interest rates. As the possibility of prepayment increases, the cost of the bonds' option feature rises; put another way, the option feature gets closer to being in the money. To adjust for the increased value of the option, traders price higher spreads into the bond, which keeps the OAS more or less unchanged.

Effective Duration and Convexity

The modified duration of a bond measures its price sensitivity to a change in yield. It is essentially a snapshot of one point in time. It assumes that no change in expected cash flows will result from a change in market interest rates and is thus inappropriate as a measure of the interest rate risk borne by a mortgage-backed bond, whose cash flows are affected by rate changes because of the prepayment effect.

Mortgage-backed bonds react differently from conventional bonds to interest rate changes. When rates fall, prepayments rise (and vice versa), shortening (lengthening) the MBS's duration. This is the opposite of what happens with a conventional bond. Mortgage-backed bonds thus exhibit negative convexity similar to that displayed by callable bonds. The prices of both securities react differently from those of conventional bonds to interest rate changes. For these reasons, a more accurate measure of mortgage-based bonds' interest rate sensitivity is *effective duration*. Effective duration is based on *approximate duration,* which can be thought of as the *median* time to receipt of a bond's cash flows, rather than the average calculated by Macaulay duration. Approximate duration is derived using equation (14.19).

$$D_{app} = \frac{P_- - P_+}{2P_0\,(\Delta rm)} \qquad (14.19)$$

where

P_0 = the initial price of the bond

Δrm = the change in the yield of the bond

P_- = the estimated price of the bond if the yield decreases by Δrm

P_+ = the estimated price of the bond if the yield increases by Δrm

Effective duration is essentially approximate duration where P_- and P_+ are obtained using a valuation model—such as a static cash flow model, a binomial model, or a simulation model—that incorporates the effect of a change in interest rates on the expected cash flows. The values of P_- and P_+ depend on the assumed prepayment rate. Generally analysts assume a higher prepayment rate when the interest rate is at the lower level of the two rates—interest and prepayment.

FIGURE 14.6 illustrates how effective duration—calculated using a 20 basis point change in rates—differs from modified duration for agency mortgage pass-through securities with a range of coupons. It shows that modified duration overestimates the price sensitivity of lower-coupon bonds. This difference has a significant effect when hedging a mortgage-backed bond position: using modified duration to calculate the needed nominal value of a hedging instrument will be accurate for only very small changes in yield.

The formula for calculating *approximate convexity* is (14.20). If P_- and P_+ are obtained using a valuation model that incorporates the effect of a change in interest rates on the expected cash flows, the equation derives *effective convexity*. The effective convexity of a mortgage pass-through security is invariably negative.

$$CV_{app} = \frac{P_+ + P_- - 2P_0}{P_0\,(\Delta rm)^2} \qquad (14.20)$$

Total Return

To assess the value of a mortgage-backed bond over a given investment horizon, it is necessary to measure the return generated during the holding period from the bond's cash flows. This is done using the *total return* framework.

Computing total return starts with calculating total cash flows. A mortgage-backed bond's cash flows comprise:

❑ its projected interest payments and principal repayments and prepayments

FIGURE 14.6 *Modified and Effective Duration of Agency Mortgage-Backed Bonds*
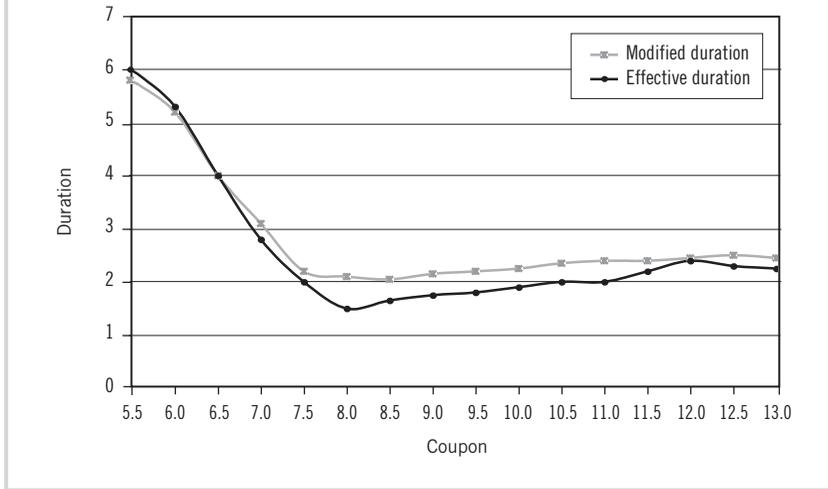
❑ the interest earned by reinvesting all the payments
❑ the bond's projected price at the end of the holding period

The first component can be estimated by assuming a prepayment rate during the holding period; the second entails assuming a reinvestment rate. For the third, two assumptions are necessary: one concerning the bond's bond-equivalent yield at the end of the holding period, and another about the prepayment rate projected by the market at this point, which is a function of the projected yield.

Plugging the total cash flow figure into equation (14.21) gives the bond's total return for the holding period, on a monthly basis.

$$TR = \left[\frac{\text{Total future cash flow amount}}{P_m} \right]^{1/n} - 1 \qquad (14.21)$$

where

P_m = initial investment
n = number of months in the holding period

The monthly return can be converted to an annualized bond-equivalent yield using formulas (1.24a) or (b), as discussed in chapter 1.

The return calculated using (14.21) is based on several assumptions. The best way to obtain an idea of the return likely to be generated over the holding period is to compute a range of returns by using a range of values for each assumption.

Price-Yield Curves of Mortgage Pass-Through, PO, and IO Securities

When interest rates are high, holders of mortgage-backed bonds want prepayments to occur. This is because the rate paid by the underlying mortgages, and thus by their bonds, are lower than those available in the market and the likelihood of mortgage prepayment at par boosts their bonds' value. Conversely, when interest rates are low, bondholders prefer no prepayments, since their bonds' interest rate is higher than that available in the market and their value correspondingly high. **FIGURE 14.7** illustrates how the price of a pass-through security with a nominal coupon of 7 percent behaves under different prepayment scenarios at different market yields.

When no prepayments are made, cash flows are certain and the pass-through's price and yield behave like those of a conventional bond. At an optimal prepayment rate—that is, one based on the assumption that homeowners act rationally and refinance whenever they can reduce their mortgage costs by an amount greater than the refinancing transaction's cost—the bond acts like a callable bond: when interest rates are high, it resembles a plain vanilla bond; when rates are lower, its price is capped at par. Under what Tuckman (1996) calls "realistic payment" conditions, the price behavior is somewhat different. First, when rates are very low, the bond's price is higher than in the other two scenarios. This is because a number of mortgage borrowers do not act "optimally," repaying their loans irrespective of the level of interest rates—even when they're high; since prepayments at high rates are good for bondholders, the bond prices in the realistic scenario are higher at this end of the yield spectrum than are those for the other two models, which predict no prepayments under these conditions.

Second, when interest rates are very low, the bond's price is higher under the realistic scenario than under the optimal one, though not as high as in the no-prepayment model. The reason is that, in this environment, many borrowers will behave "optimally" and prepay their loans, but by no means all will. Since prepayments decrease the value of a mortgage bond when rates are low, the fact that not all borrowers prepay in the "realistic" scenario results in the realistic-prepaid value of a mortgage bond being somewhat greater than its optimal-prepaid value. This nonprepayment behavior can lead to the bond being valued
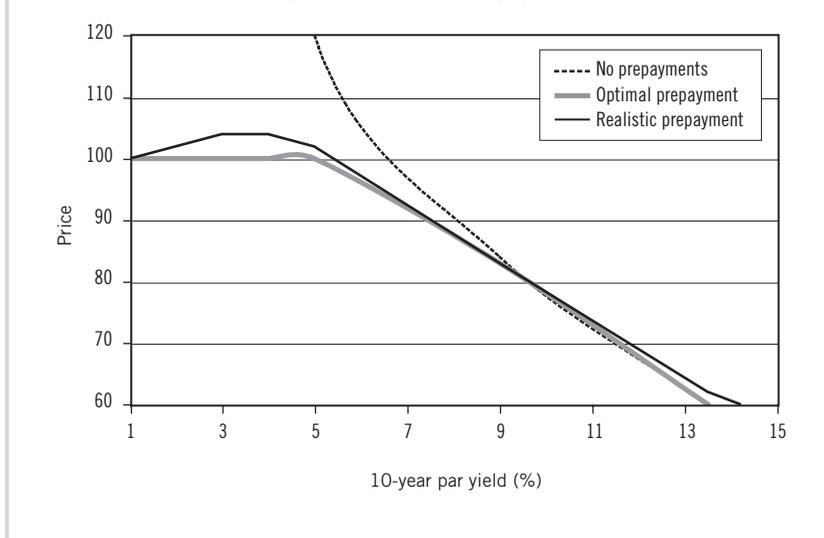
above par. This is something of an anomaly, considering that the bond is then priced above the level at which it can theoretically be called. Eventually, though, rates fall far enough to convince all borrowers to redeem their loans, and the realistic-prepayments curve moves down to par.

Figure 14.7 demonstrates the negative convexity of mortgage bonds through the fact that their prices fall as interest rates decline. This does not mean that investors should avoid mortgage-backed bonds in this environment. As Tuckman (1996) notes, mortgage bonds in this situation are paying rates higher than those available elsewhere in the market, particularly the debt market. The relevant consideration is total return over the holding period, not price. Making investment decisions based on price behavior alone, Tuckman writes (page 256), is "as bad as concluding that premium Treasuries should never be purchased because they will eventually decline in price to par."

As already discussed, IOs, which receive the interest payments of the underlying collateral, and POs, which receive principal payments, exhibit different price behavior from pass-throughs and from each other. Figure 14.5 (page 263) showed that when interest rates are very high and prepayments, accordingly, unlikely, POs act as if repayable at par on maturity, like zero-coupon bonds. When interest rates decline and prepayments

CASE STUDY: *ACE Securities Corp. Home Equity Loan Trust, Series 2004[1]*

Residential MBSs are characterized as prime and sub-prime, depending on the credit quality of the underlying mortgages. A credit quality score known as FICO measures whether the loan is prime or sub-prime. Home equity, while previously referring to a different type of RMBS, now refers to a sub-prime RMBS transaction.

ACE Securities series 2004 is a sub-prime RMBS transaction that closed in the U.S. market in January 2004. It is a securitization of a pool of sub-prime mortgages originally on the balance sheet of Fremont Investment and Loan. Fremont is a commercial banking institution that had been engaged in sub-prime mortgage lending for more than ten years prior to the transaction, and also originated previous home equity securitization deals.

Transaction Summary

Originator	Fremont Investment & Loan
Type	Senior subordinated residential MBS
Amount	$751,303,000
Credit support	Note tranching, overcollateralization, excess spread
Servicer	The Provident Bank
Trustee	HSBC Bank USA
Underwriter	Deutsche Bank Securities

The tranche structure for ACE Securities HELT series 2004 is shown in **FIGURE 14.8.** The transaction was undertaken to provide a diversified funding source for Fremont, with a size of more than $751 million.

The deal is structured as a senior-subordinated overcollateralization, with the first three notes all rated as AAA. These are ranked further into a super-senior and junior-senior tranche. The note tranching is the principal form of credit enhancement, in addition to the overcollateralization of 0.85 percent. There is also a reserve account to trap excess spread, which is a further credit enhancement.

The Class A-1 notes also have credit enhancement from Class A-3. This works as follows: where the subordinated notes are reduced to zero, any losses on the underlying pool of mortgages

FIGURE 14.8 *ACE Securities Corp. HELT Series 2004-FM1*

CLASS	DESCRIPTION	AMOUNT $000	COUPON	RATING
A-1	Super senior principal & interest	571,643	LIBOR + 0.30	Aaa
A-2A	Senior principal & interest	37,604	LIBOR + 0.32	Aaa
A-2B	Senior principal & interest	39,000	LIBOR + 0.19	Aaa
A-2C	Senior principal & interest	19,127	LIBOR + 0.46	Aaa
A-3	Junior senior principal & interest	63,516	LIBOR + 0.40	Aaa
M-1	Subordinate principal & interest	69,547	LIBOR + 0.60	Aa2
M-2	Subordinate principal & interest	57,128	LIBOR + 1.25	A2
M-3	Subordinate principal & interest	17,387	LIBOR + 1.45	A3
M-4	Subordinate principal & interest	17,387	LIBOR + 1.80	Baa1
M-5	Subordinate principal & interest	14,903	LIBOR + 1.95	Baa2
M-6	Subordinate principal & interest	9,935	LIBOR + 3.50	Baa3
B-1A	Subordinate principal & interest	6,955	LIBOR + 3.50	Ba2
B-1B	Subordinate principal & interest	6,955	6.00%	Ba2
CE	Residual	–	Not rated	
P	Prepayment penalties	–	Not rated	
R	Residual	–	Not rated	

supporting the notes that are not covered by the overcollateralization and the excess spread will be borne by the A-3 notes ahead of the A-1 notes.

This transaction features an unusual feature in that the underlying pool of mortgages is split into two groups, Loan Groups 1 and 2. Classes A-1 and A-3 are supported by Loan Group 1, and classes A-2A, A-2B, and A-2C are supported by Loan Group 2; however, there is also cross-collateralization for the senior notes.

This is an interesting structure but nevertheless represents a routine transaction in the highly developed U.S. MBS market.

increase, the POs' price increases. Other factors are at work, however, that make PO prices highly volatile. These are:

- ❑ the hypersensitivity of POs to the conventional price/yield effect, which states that lower interest rates cause higher prices and vice versa
- ❑ the effect on POs' maturity of prepayment rates—specifically, the higher the actual and expected rates, the lower the effective maturity and, so, the higher the POs' price

IOs' price/yield relationship is a function of that for POs, obtained by subtracting the value of the latter from that of the underlying mortgage pass-through. IOs' prices are very volatile when interest rates are low and falling. This may be explained as follows: when rates are high and prepayments very low, IOs' cash flows are known with virtual certainty, so they act like plain vanilla bonds. When rates fall and prepayments rise, diminishing the nominal amount of the mortgages on which interest is charged, IOs' cash flows effectively disappear because, unlike pass-throughs and other mortgage securities, they don't receive any principal payments. Their prices in these circumstances decline dramatically. Such negative duration makes IOs attractive to market makers in mortgage-backed securities as interest-rate hedging instruments.

Chapter Notes

Collateralized Debt Obligations

Collateralized bond obligations (CBOs) and collateralized loan obligations (CLOs), which together make up collateralized debt obligations (CDOs), are among the newest developments in securitization. The instruments are generally held to have originated in the repackaging of high-yield debt or loans into higher-rated bonds that began in the late 1980s. Today many types of CDOs exist, and the market has expanded from the United States into Europe and Asia.

Both CBOs and CLOs are securities issued against an underlying collateral of assets. These assets almost invariably are diverse corporate bonds or loans or both. CBOs are backed by corporate or sovereign bonds; CLOs, by secured and/or unsecured corporate and commercial bank loans. There are two types of CDOs: *arbitrage* and *balance sheet.* Some analysts also recognize a third category: *emerging market* CDOs, which are CDOs securitized from a portfolio of emerging market bonds (or loans).

A typical CDO structure involves the transfer of the credit risk associated with an underlying asset pool from the originating institution to a special purpose vehicle, or SPV, created specifically to make this transfer possible. The SPV—typically bankruptcy remote and isolated from the originator's credit risk, often in a tax haven—then transfers the risk to investors by issuing CDO notes. The return to investors in the issued notes depends on the performance of the underlying asset pool. The manager, who is responsible for managing the portfolio of underlying assets and bonds, would be expected to manage the portfolio after the CDO transaction is brought to market. As the bonds in the underlying portfolio

might need to be hedged (to remove risk exposure arising from issuing a series of notes that have different interest pay dates and also possibly different currencies), an interest rate and currency swap is entered into with a hedge counterparty.

Among institutions' objectives in originating CDO transactions are the following:

❑ Optimizing their returns on regulatory capital, by reducing the need for capital to support assets on the balance sheet. Regulatory capital is the capital needed to be put up by a financial institution in accordance with the "Basel" rules, issued by the Bank for International Settlement.

❑ Improving their returns on economic capital, the actual capital used by the bank to support its operations, by managing risk effectively

❑ Managing their credit risk and balance sheets
❑ Issuing securities as a means of funding
❑ Gaining funding for acquiring assets
❑ Increasing funds under management

FIGURE 15.1 shows a typical conventional CDO structure.

As noted above, CLOs are backed by pools of bank loans and CBOs by portfolios of bonds. The two types of underlying assets differ in ways that affect the analyses of the securities they collateralize. Among the differences are the following:

❑ Loans have less uniform terms than bonds, varying widely in their interest dates, amortization schedules, reference indexes, reset dates, maturities, and so on. How their terms are defined affects the analysis of cash flows.

❑ In part because of this lack of uniformity, the legal documentation for loans is less standardized than that for bonds. Securities backed by loans, therefore, require more in-depth legal review.

FIGURE 15.1 *A Typical Conventional CDO Structure*
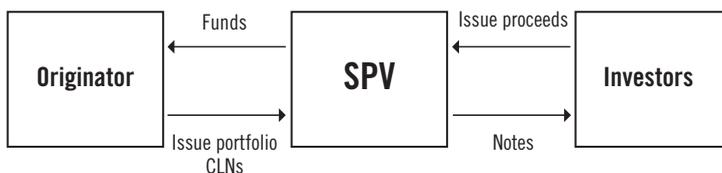
❑ It is often possible to restructure a loan portfolio to reflect the changed or changing status of the borrowers—for example, their ability to service the debt. This provides participants in a CLO with more flexibility than they usually have with a CBO.

❑ The market in bank loans is far less liquid than that in bonds, which has the effect of making overlying notes sometimes less liquid in the secondary market for CLOs.

CDO Structures

CDO structures may be either conventional or synthetic. The conventional structures were the first to be widely used, but synthetic ones have become increasingly common since the late 1990s. The difference between the two structures lies in how they transfer credit risk from the originator to the SPV: in conventional CDO structures, this is achieved by transferring assets; in synthetic structures, credit derivative instruments are used.

CDOs of both types are also categorized by the motivation behind their creation. The two main categories are issuer- or *balance sheet*–driven transactions and investor-driven or *market value* arbitrage transactions.

Conventional CDO Structures

In a conventional structure, such as the one illustrated in **FIGURE 15.2**, the creation of an SPV usually involves the transfer from the originator of a nominal amount of equity. The main funding comes from issuing CDO notes. The proceeds from the issuance are used to acquire the pool of underlying assets (bonds or loans) from the originator in what is known as a *true sale.* If performed and structured properly, this asset transfer removes assets from the regulatory balance sheet of a bank originator. As a result, the securitized assets are not included in the calculation of the bank's capital ratios. This provides regulatory capital relief, which is the main motivation for many of the CDO transactions in the market today.

Because the SPV now owns the assets, it has an asset-and-liability profile that must be managed during the term of the CDO. The typical liability structure includes a senior tranche rated Aaa/Aa, a junior tranche rated Ba, and an unrated equity tranche. The equity tranche is the riskiest, since it is the first to absorb any losses in the underlying portfolio. For this reason, it is often referred to as the *first-loss* tranche.

In the case of a CLO, the originating bank commonly continues to service the underlying loan portfolio and retains the equity tranche.

This is done for the following reasons:

FIGURE **15.2** *Conventional CDO Structure*
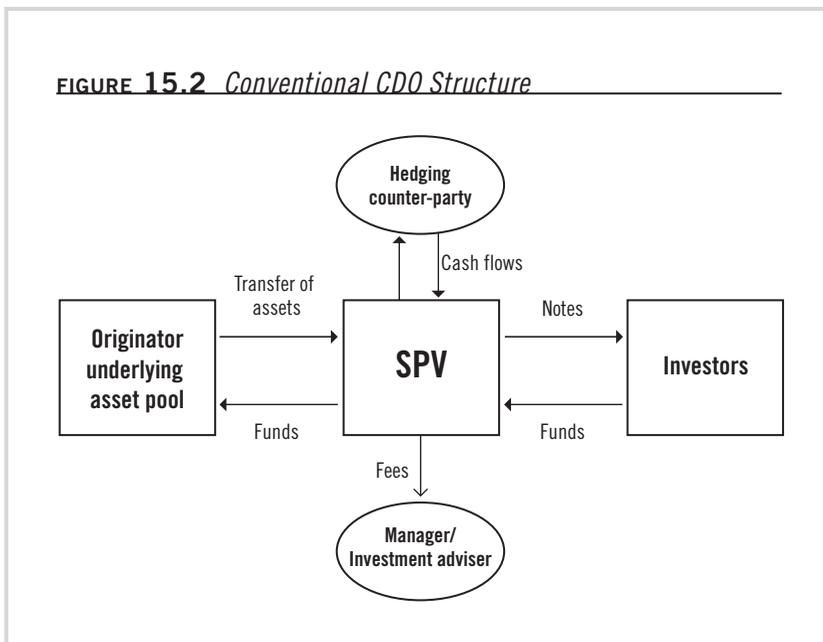
❑ The bank has detailed information on the loans that enables it to manage effectively the risk it retains.
❑ Having a financial interest in the performance of the loan portfolio, the bank remains motivated to service it.
❑ The return required by a potential purchaser of the equity tranche may be too high.

Structuring a conventional CDO may give rise to significant issues. For instance, the transfer of assets into the SPV may have adverse tax, legal, and regulatory impacts, depending on the jurisdiction in which the transfer of assets takes place and the details of the legislation pertaining to that jurisdiction. Another issue is reinvestment risk: in a conventional CDO, the originator receives cash, and if the main objective of the transaction is to transfer credit risk or to acquire protection for credit risk, this cash must be reinvested in other assets.

The multitranche structure, with its prioritization of cash flow payments to investors, provides the CDO with a credit enhancement. To enhance the credit of the senior notes, the originating bank may also use other mechanisms, such as credit insurance on the underlying portfolio, known as a credit wrap, and reserve accounts that absorb a loss before the equity tranche.

Synthetic CDO Structures

As with conventional CDOs, the main motivation for issuing synthetic CDOs is the desire to hedge or transfer credit risk, in order to achieve regulatory-capital relief or to obtain credit protection on an underlying asset pool. In a synthetic CDO structure, such as that illustrated in **FIGURE 15.3**, the reference pool of assets remains on the originator's balance sheet. In place of an asset transfer, the originator enters into a credit-default swap with the SPV that covers any losses in the underlying asset pool. This transfers the credit risk of the asset portfolio to the SPV. If its regulatory authority recognizes the credit-risk offset, this allows the bank to release regulatory and economic capital. In return for the credit protection, the originator bank pays the SPV a premium, typically in the form of a regular fee.

Notes issued in synthetic structures are organized by tranche. With the proceeds from the notes it issues to investors, the SPV purchases high-quality (AAA) liquid securities—for example, U.S. Treasuries, bank asset-backed paper such as credit card ABS, and German bonds, such as Pfandbriefe —to serve as collateral. This collateral will generate LIBOR-related interest and principal cash flows that the SPV passes on to the investors together with the swap premium, which creates an additional credit spread on the notes. The cash flows from the collateral may not match the payments due on the issued notes—for example, the bonds used as collateral may pay a fixed rate and the issued notes a floating one. To remedy this, the

FIGURE 15.3 *A Synthetic CDO Structure*
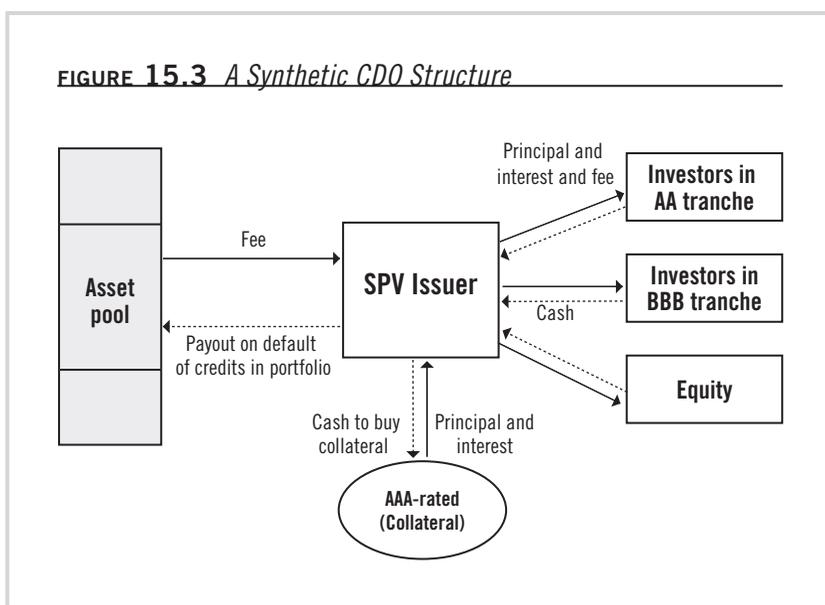
SPV usually enters into an interest rate swap. The swap counterparty may also sell the SPV other derivative instruments, such as interest rate caps, to manage possible cash flow risk. Such risk-exposure management requires careful attention, since the SPV's risk profile can have a significant impact on the credit risk of the notes issued to investors. In an unleveraged transaction, the size of the issue is equivalent to the credit protection the SPV offers on the reference pool of assets. For example, if the credit default swap is on a nominal of $300,000, the nominal value of the notes issued will be $300,000.

The notes issued to investors are linked to the credit risk of the portfolio through the credit-default swap—which usually has the same term to maturity as the notes—and to the credit derivative counterparty. The notes are therefore *credit linked.*

The payout from the credit-default swap is triggered by a credit event. The precise definition of credit event is important, since it may affect the notes' returns. It usually includes bankruptcy and failure to pay off the underlying credit. In a failure to pay, a grace period may be specified, so that default is not triggered if the payment is delayed for technical reasons, such as information technology issues. The International Swaps and Derivatives Association's definitions for a credit derivative transaction refer to *restructuring,* as a credit event. Its inclusion in the credit definitions of default swaps in synthetic CDO transactions depends on whether it would affect the regulatory-capital-relief treatment of the underlying asset pool. Restructuring refers to the process when the loan liabilities of a borrower are restructured, in terms set by its lenders, in times of financial difficulty.

If a credit event occurs, the SPV usually pays out a cash amount equal to the par value of the underlying assets covered by the credit default swap, less their post-default price. Less commonly, the SPV physically settles the credit-default swap by purchasing the defaulted assets at par value. The credit loss is then passed on to the investors according to the priority of the tranches they hold.

Motivation Behind CDO Issuance

Different types of institutions have different motives for originating CDO transactions. The main ones are to optimize regulatory capital, to obtain funding, to engage in arbitrage, or, on occasion, a combination of all three.

Balance Sheet–Driven Transactions

In a balance sheet–driven CDO the originating bank is trying to obtain off–balance-sheet treatment for existing assets to which bank capital has been allocated. This enables the originator to manage capital constraints and to improve its return on capital.

The originators in these situations are mainly commercial banks. The underlying asset pool may include commercial loans, both secured and unsecured, guarantees, and revolving credits. Although there is usually no intention to trade these assets, the pool may be subject to substitutions or replenishments over the life of the structure. The originating bank usually acts as investment adviser with respect to such changes, to maintain the underlying asset pool's quality and protect the note holders. In this regard, the rating agencies often require that the average credit quality of the asset pool be maintained.

Investor-Driven Arbitrage Transactions

In a CDO whose purpose is arbitrage, the underlying asset pool includes not only instruments that generate investment income but also some that provide the opportunity to generate value from active trading. The latter may be existing positions or may be acquired expressly for the CDO. The potential for trading profits depends on the quality and expertise of the CDO manager, who is usually the investment adviser. The aim is to profit from the spread between the investment- and subinvestment-grade markets. CDOs allow lower-rated debt to be repackaged as higher-rated notes, exploiting the fact that spreads between the two grades of debt are often greater than justified by the credit difference. The originator thus earns more on the risky debt than it pays to securitize it and to enhance its credit rating.

The profitability of an arbitrage-driven CDO depends on such factors as the following:

- ❏ the return required by the holders of the issued tranches
- ❏ the return of the underlying asset pool
- ❏ the expenses of managing the SPV

If the underlying portfolio performs well and its loss profile is more attractive than projected, because of better-than-expected default and recovery rates, the return to the equity holder after payments to the senior and junior tranche will be higher than expected. If, however, the underlying portfolio performs poorly and default and recovery rates are worse than projected, perhaps because of adverse economic conditions, the tranche returns will be lower than expected. Poor investment management will also have an adverse impact on the return to investors.

Fund managers create arbitrage CDOs from pools of high-yield bonds since this enables them to increase the size of their assets under management with comparatively small levels of equity. Their objective is to set up the CDO so that the return generated by the underlying pool of high-yield bonds is sufficient to pay off investors and provide them with a profit on the equity tranche on top of their management fee.

Analysis and Evaluation

A number of factors are important to consider when analyzing, evaluating, or rating a CDO. The basic ones are discussed in this section.

Portfolio Characteristics

The *credit quality of the underlying asset pool* is critical, since it determines the structure's credit rating. It is common to allocate an average rating to the initial reference pool. A constraint in structuring the transaction may be that any permitted changes to the pool should not lower the average rating. The portfolio's credit quality and its possible variability are used to project the pool's default frequency and loss rates. In some cases, the bank's own system for determining the credit risk of borrowers is a key part of the rating process. Particularly for unrated assets, investors should determine the internal rating system's accuracy by mapping it onto the rating agency's system.

The *diversity of the reference pool* plays a part in determining its credit risk. Diversity is measured in terms of the portfolio's concentration by industry group, obligor, and sovereign country. Broadly speaking, the greater the diversification, the lower the credit risk. A portfolio is assigned a diversity score based on its weighted average credit score. Each incremental credit exposure in the underlying asset pool gets a marginal score determined by the makeup of the entire credit portfolio. If the portfolio is concentrated in one category—say, an industry group—the marginal score attributed to the marginal credit is reduced to reflect the lack of diversity. As a result, an asset pool with a wide range of credit exposures has a higher diversity score. Change in the underlying asset pool may be constrained by a requirement that a minimum diversity score be maintained for the life of the CDO.

Cash Flow Analysis and Stress Testing

The cash flow profile of a CDO depends on the following issues:

❑ The spread between the interest earned on the collateral and the coupon paid on the securities issued

❑ The frequency of defaults in the underlying asset pool and their severity, in terms of recovery rate, plus the impact of losses on investors' principal

❑ The principal-repayment profile and expected amortization of the underlying loans

❑ The contingent payments under any credit-default swap used to transfer credit risk from the originator to another party, such as the SPV or an OECD bank. Under Basel I rules, if an investor takes out credit protection on a loan, and the protection is provided by a OECD bank, his capital charge changes from 100 percent to 20 percent.

❑ The contingent cash flows from any credit wrap or credit insurance on the underlying asset pool

❑ The cash flows receivable from or payable to a hedge counterparty under swap agreements or derivative contracts

❑ The premium received from the credit default swap counterparty

❑ Fees and expenses

The cash flows are tested to see how they are affected in both normal and stressed scenarios. The types of stress scenarios tested depend on the underlying asset pool.

Originator's Credit Quality

The impact of the originator's credit quality on the rating of the notes issued depends on the CDO structure. In a conventional structure, where the underlying assets are transferred from the originator to the SPV, the credit quality of the CDO notes is "delinked" from that of the originator, depending solely on the portfolio performance and the credit enhancement. In a synthetic structure, in contrast, the underlying asset pool remains on the originator's balance sheet. Investors, therefore, may be exposed to both the originating bank's credit quality and the portfolio performance. The rating of such a credit-linked CDO is capped by that of the originator, whose fiscal soundness determines the reliability of the interest and principal payments.

The senior tranches of a synthetic CDO, however, may be delinked from the bank's rating by using AAA-rated collateral and default swaps, as described above. The final rating is influenced by the credit rating of the default-swap provider and the extent to which the cash flows to investors are exposed to the risk of default by the asset pool.

Operational Aspects

In market value transactions, the portfolio manager's ability is key, since the performance of the underlying portfolio is critical to the structure's success. The originator's procedure for reviewing credit approvals of the borrowers of the underlying loans and monitoring the loans is another factor to consider. The better the credit-assessment and monitoring proce-

dure, the more comfortable investors can be with the integrity and quality of the underlying asset portfolio.

Review of Credit-Enhancement Mechanisms

Credit enhancements include reserve accounts, subordinated tranches, credit wraps, and liquidity facilities. Investors should consider the impact of the particular enhancements a structure uses. This will usually be observed by subjecting the CDO to stress scenarios that are designed to determine the effect on the cash flows.

Subordination. Each tranche's rights to and priority in receiving interest and principal payments are set out in an issue's offering circular, which provides a detailed description of the notes and their legal structure. In allocating cash flows, typically, fees and expenses are subtracted from the cash flows, then the most senior tranches are serviced, followed by the junior tranches, and finally the equity tranche. This method of cash flow is sometimes referred to as a *cash flow waterfall.*

Credit wrap. As explained above, the originating bank may buy credit insurance on the debt instruments of the underlying portfolio, to improve its credit quality.

Reserve accounts. The banks may also set aside cash reserves from the note proceeds in accounts, usually managed by the servicing agent or a specialized cash manager, which provide first-loss protection to investors by absorbing losses before the equity tranche.

Liquidity facility. A facility is an arrangement to provide a borrower with credit support. In the case of a CDO, this arrangement involves the originating bank ensuring that, should the underlying asset pool experience a temporary cash shortfall, the notes' interest and principal payments will still be made.

Legal Structure of the Transaction

A typical CDO structure is described in several legal agreements that formalize the roles played by the various counterparties to the deal. In addition to the offering circular, which presents the transaction details to investors, these documents include the following:

- ❑ The trustee agreements, which set out the responsibilities for administration and maintenance of books and records
- ❑ The sale agreement or credit-default swap agreement used to transfer credit risk
- ❑ The hedging agreements, such as interest rate or cross-currency swaps and other derivative contracts
- ❑ Guarantees or insurance, such as credit wraps on the underlying asset pool

Before the deal is closed, the SPV's incorporation documents are also reviewed, to ensure that it is bankruptcy remote and established in a tax-neutral jurisdiction.

Expected Loss

Rating a CDO involves a detailed analysis of its structure, including the elements discussed above. It also encompasses a quantitative assessment. This is often based on the expected loss (*EL*) to note holders. The expected loss calculated for each tranche is mapped to a table of expected losses and their corresponding rating to assign the tranche an appropriate rating. The tranche's credit rating is a key determinant in its pricing and marketability. Expected loss is calculated using equation (15.1).

$$EL = \sum_x p_x \left(L_x \right) \tag{15.1}$$

where

L_x = the loss on the notes under scenario x

p_x = the probability of the scenario x occurring

Note holders' expected losses are determined by considering the impact on their cash flows of the credit losses—losses from loan defaults—occurring in various scenarios, taking into account how such losses are allocated to the issue's tranches. The cash flows to the note holders depend on whether a default has occurred and the size of the resulting loss. The severity of the loss equals the par value of the note less the recovery rate. The probability of default may be inferred from the rating of the underlying credit exposures. Expected losses are calculated using Monte Carlo techniques, which simulate thousands of scenarios and cash flows and so require sophisticated computational models.

This page is intentionally blank

Selected Market Trading Considerations

Part Three presents the author's insights into trading, based on his experiences working as a gilt-edged market maker and sterling-bond proprietary trader. The topics covered include implied spot and market zero-coupon yields, yield-curve spread trading, and butterfly spreads.

This page is intentionally blank

The Yield Curve, Bond Yield, and Spot Rates

This chapter examines a number of issues relevant to participants in the fixed-income markets. The analysis presented is based on government-bond trading and is confined to generic bonds that are default-free, with no consideration given to factors that apply to corporate bonds, asset- and mortgage-backed bonds, convertibles, or other nonvanilla securities, or to issues such as credit risk and prepayment risk. Nevertheless, the principles adduced are pertinent to all relative-value fixed-income analysis.

Practical Uses of Redemption Yield and Duration

The drawbacks of duration and gross redemption yield (henceforth referred to simply as *yield*) for bond analysis are well documented. That different bonds, even vanilla government securities, can have their yields analyzed in a number of ways suggests, moreover, that acceptable return measures are lacking. When assessing the opportunities available in a market, investors often use that market's yield convention. The resulting multiplicity of yield-calculation methods—illustrated in **FIGURE 16.1**—makes bond comparisons problematic. Duration is another measure that can be defined in more than one way, again making comparison between different bonds difficult. This section discusses ways to mitigate the problems inherent in using yield and duration and how using analyses based on these methods should proceed.

FIGURE 16.1 *Bloomberg YA Screen Showing Alternative Yield Calculations for a U.S. Treasury*

```
                                                           Govt  YA

        DISCOUNT/YIELD  ANALYSIS      CUSIP    9127950E4
 TREASURY BILL     B 06/24/04        .92   / .91    ( 0.94 /93) BGN  @ 9:00
   DISCOUNT  O.91000  SETTLEMENT  DATE  3/26/2004
 MATURITY    6/24/2004    90 DAYS TO MATURITY      ISSUE:12/26/2003
 PRICE              99.7725000  CASHFLOW  ANALYSIS
 YIELD                    TAXED   FOR     1000M FACE AMOUNT
 CALCULATIONS            28.00%  PRINCIPAL[RND(Y/N)N]    997725.00
 U.S. TREASURY CONVENTION  0.927  0.668  REDEMPTION           1000000.00
 U.S. GOVT BOND EQUIVALENT 0.927  0.668  PROFIT                   2275.00
 SIMPLE INTEREST (ACT/ 360) 0.912 0.657  REPO RATE               0.990
 MEDIUM TERM CD  (ACT/ 360) 0.912 0.657  OVERNIGHT REPO EQUITY   0.912
                                         COST OF CARRY (PTS)     -0.078
                                         NET P&L                 -6.48
 SENSITIVITY  ANALYSIS
 DURATION(YEARS)              0.246  FURTHER  ANALYSIS
 ADJ./MOD DURATION            0.245  1 <GO> COST OF CARRY ANALYSIS
 RISK                         0.245  2 <GO> DISCOUNT EQUIVALENT TABLE
 CONVEXITY                    0.000  3 <GO> DISCOUNT/YIELD TABLE

     A   0.01   CHANGE IN
   DISCOUNT/ PRICE/ YIELD  D      U.S. Treasury convention yield
 RESULTS IN A -0.00250 CHANGE IN PRICE    based on a 366-day year
      AND A  0.01021 CHANGE IN YIELD
```

The Concept of Yield

Ideally, an instrument's yield indicates what return an investor can achieve by holding it. Such an ideal measure would be a function of the value of the initial investment, the holding period, and the value of the matured investment. It would also take into account any reinvestment of the income received during the holding period—that is, the effect of compounding. A yield measure having these properties may be defined as follows for a simple instrument such as a Treasury bill.

Consider a T-bill with a term of m days and a price of P. Equation (16.1) may be rearranged to compute the bill's true yield, rm.

$$P = \frac{100}{\left(1 + \frac{1}{2} rm\right)^n}$$

(16.1)

where

n = the number of interest periods from value date until maturity

The U.S. Treasury interest basis is semiannual, and the market uses an actual/actual day count. So, the value of n for the 90-day T-bill whose yield analysis as of March 25, 2004, is shown in figure 16.1 would be 90/183, where 183 represents the number of days in half a year, given that 2004, as a leap year, had 366 days. The bill was priced

at 99.7725, so its true yield is given by equation (16.2).

$$99.7725 = \frac{100}{\left(1 + \frac{1}{2} rm\right)^{90/183}} \tag{16.2}$$

$$rm = 0.927 \text{ percent}$$

The yields quoted on T-bills often differ from their true yields. This is because their yield calculations often assume simple rather than compound interest. Nevertheless, true yield is important for its application to longer-dated coupon bonds.

Yield has been defined in previous chapters as the discount rate that equates the sum of the present values of all a bond's cash flows to its observed market price. A vanilla bond, such as a U.S. Treasury, has m future cash flows—the coupon payments—each having a value C, equal to one-half the coupon rate applied to the face value. C_m is the principal payment. The sum total of the bond's discounted cash flows is given by equation (16.3).

$$PV = \frac{C_1}{\left(1 + \frac{1}{2} r\right)^{n_1}} + \frac{C_2}{\left(1 + \frac{1}{2} r\right)^{n_2}} + \ldots + \frac{C_m}{\left(1 + \frac{1}{2} r\right)^{n_n}} \tag{16.3}$$

where

C_i = the ith bond payment
C_m = the coupon payment plus the principal repayment
m_i = the number of days from the value date to maturity
n_i = the number of interest periods from the value date until C_i = m_i/
182 or 183
r = the discount rate

From this it is easy to derive the definition of true yield, which is the discount rate that equates a bond's current market price to the present value of its cash flows. The bond's market price is its *dirty* price—that is, the price including accrued interest. This is represented in (16.4), where *PV* is replaced by *P*, representing the clean price plus *AI*, representing accrued interest.

$$P + AI = \frac{C_1}{\left(1 + \frac{1}{2} rm\right)^{n_1}} + \frac{C_2}{\left(1 + \frac{1}{2} rm\right)^{n_2}} + \ldots + \frac{C_m}{\left(1 + \frac{1}{2} rm\right)^{n_i}} \tag{16.4}$$

Yield Comparisons in the Market

U.S. Treasury price quotes are in *ticks,* or thirty-seconds of a price point. A half tick is denoted by a plus sign. On May 10, 1994, the 10.25 percent Treasury bond maturing July 21, 1995, was quoted at 104-28+— in other words, an investor would pay $104.28625 for every $100 in face value. It pays coupons on January 21 and July 21. On May 11, 1994, the settlement date, it will have accrued 109 days of interest, for a total of 10.25 × 109/365 × 0.5, or 1.53048 for every $100 of face value. The dirty price of the bond on this date is thus 104-28+ plus 1.53048, or 106.421105.

The remaining bond cash flows are $5.125, on both July 21, 1994, and January 21, 1995, and $105.25, on July 21, 1995. January 21, 1995, however, is a Saturday, so the cash flow will not actually be received until Monday, January 23. The number of days between the value date, May 11, 1994, and the receipt of each cash flow is

July 21, 1994	71 days
January 23, 1995	230 days
July 21, 1995	436 days

The interest periods between each cash flow date and the value date number are

(71/183) or 0.387978
(230/183) or 1.256830
(436/183) or 2.382513

Plugging the derived values for price, accrued interest, cash flows, and interest periods into (16.4) gives

$$106.421105 = \frac{5.125}{\left(1 + \frac{1}{2}rm\right)^{0.387978}} + \frac{5.125}{\left(1 + \frac{1}{2}rm\right)^{1.256830}} + \frac{105.125}{\left(1 + \frac{1}{2}rm\right)^{2.38251}}$$

which can be rearranged to solve for rm = 0.073894, or 7.3894 percent.

The conventional yield—the one usually quoted—is almost invariably different from the true yield. This is because the conventional calculation derives the number of interest periods between the value date and the cash flows based on exact half-year intervals between payments, ignoring the delays that occur when the payment dates fall on nonbusiness days.

Measuring a Bond's True Return

The true yield measure derived in the previous section is not as straightforward as the one given earlier for the T-bill. Because a T-bill has only a single cash flow, its maturity value is known, so its return is easily calculated as its increase in value from start to maturity. Investors know that money put into a 90-day T-bill with a yield of 5 percent will have grown by 5 percent, compounded semiannually, at the end of three months. No such certainty is possible with coupon-bearing bonds. Consider: although the investors in the 90-day T-bill are assured of a 5 percent yield after ninety days, they don't know what their investment will be worth after, say, sixty days or at what yield they will be able to reinvest their money when the bill matures. Such uncertainties don't effect the return of the short-term bill, but they have a critical impact on the return of coupon bonds.

It would certainly help investors if they could analyze bonds as though they had single cash flows. Investors often buy bonds against liabilities that they must discharge on known future dates. It would be a comfort if they could be sure the bonds' returns would meet their liability requirements. Put very simply, this is the concept of *immunization.*

The difficulty in calculating a bond's return is that its future value is not known with certainty, because it depends on the rates at which the interim cash flows can be reinvested, and these rates cannot be predicted. A number of approaches have been proposed that get around this. These are described in the following paragraphs, assuming simple interest rate environments.

The simplest approach assumes, somewhat unrealistically, that the yield curve is flat and moves only in parallel shifts, up or down. It considers a bond to be a package of zero-coupon securities whose values are discounted and added together to give its theoretical price. The advantage of this approach is that each cash flow is discounted at the interest rate for the relevant term, rather than at a single "internal rate of return," as in the conventional approach. Given the flat yield curve, however, this approach reduces to (16.3). An example of its application is on the following page.

A bond's return is influenced by changes in the yield curve that occur after its purchase. Say the yield curve moves in a parallel shift to a new level, rm_2. In that case, the expected future value of the bond changes. Assuming s interest periods from the value date to a specified "horizon date," the new value of the bond on that horizon date is given by equation (16.5).

$$P(rm_2, s) = \frac{C_{m-s}}{\left(1 + \frac{1}{2}rm_2\right)^{n_1-s}} + \frac{C_{m-(s-1)}}{\left(1 + \frac{1}{2}rm_2\right)^{n_2-s}} + \ldots + \frac{C_m}{\left(1 + \frac{1}{2}rm_2\right)^{n_m-s}} \quad (16.5)$$

EXAMPLE: *Conventional Bond Pricing*

Given a value date of December 8, 2000, value a hypothetical bond paying a 5 percent semiannual coupon and maturing December 8, 2002.

On the value date, the bond has precisely four interest periods to maturity and no accrued interest. Its cash flows are 2.50, 2.50, 2.50, and 102.50. Assuming that the yield curve on December 7 is flat at 5 percent, its price is calculated as follows:

$$P + AI = \frac{C_1}{\left(1 + \frac{1}{2}rm\right)^{n_1}} + \frac{C_2}{\left(1 + \frac{1}{2}rm\right)^{n_2}} + \frac{C_3}{\left(1 + \frac{1}{2}rm\right)^{n_3}} + \frac{C_4}{\left(1 + \frac{1}{2}rm\right)^{n_4}}$$

$$= \frac{2.50}{(1 + 0.025)^1} + \frac{2.50}{(1 + 0.025)^2} + \frac{2.50}{(1 + 0.025)^3} + \frac{102.50}{(1 + 0.025)^4}$$

$$= 2.4390 + 2.3879 + 2.3215 + 92.8600$$

$$= 100.00$$

So, at a uniform—because of the flat curve—discount rate of 5 percent, the price of the bond is par.

The equation expresses the fact that the C_{m-i} cash flow contributes $\frac{C_i}{\left(1 + \frac{1}{2}rm_2\right)^{n_i - s}}$ to the bond's value on this future horizon date. If this cash flow is received ahead of the horizon date, $n_i - s$ is a negative exponent, meaning that C_i is *compounded* (rather than discounted) forward to the horizon date at the rate rm_2. If the cash flow is received after the horizon date, $n_i - s$ is positive, and C_i is *discounted* back to the horizon date at the same rate.

If s is small, a majority of the bond's cash flows take place after the horizon date, and a shift up in the yield curve—that is, $rm_2 > rm$—produces a lower future value, $P(rm_2, s)$, because most of the flows must be discounted at the higher rate. If s is large, so that most of the cash flows take place before the horizon date, an upward shift increases the value of these cash flows, because they can be reinvested at a higher rate of interest. When s is sufficiently large, this reinvestment gain matches, and may

even exceed, the loss suffered through revaluation of the discounting at the higher rate, resulting in a higher future value, $P(rm_2, s)$. The reverse occurs if $rm_2 < rm$. These changes in future value represent the *reinvestment risk* borne by the bondholder.

Between the short- and long-term horizon dates is one at which the net effect of the change in reinvestment rate on the bond's future value is close to zero. At this date, the bond behaves like a single-cash-flow or zero-coupon security, and its future value can be predicted with greater certainty, no matter what the yield curve does after its purchase. Defining this date as s_H interest periods after the purchase date and P_H as the value of the bond at that point, it can be shown that the bond's rate of return up to this horizon date is the value for rm_H that solves equation (16.6).

$$P + AI = \frac{P_H}{\left(1 + \frac{1}{2} rm_H\right)^{s_H}}$$

(16.6)

The left side of equation (16.6) is the bond's market price broken into clean price plus accrued interest, as it was in (16.4). In fact, it can be shown that rm_H is identical to the initial yield in (16.4), rm. The value of s_H that results in a stable future bond value is the Macaulay duration. At this point, assuming the existence of only one, parallel yield shift, a change in yield will not impact the future value of the bond. The bond's cash flows are immunized, and the instrument can be used to match a liability existing on that date.

Because of the analysis's assumed restrictions, however, investors applying it must continually adjust their portfolios if they wish to remain immunized. Fabozzi (1996) contains a very accessible discussion of the key issues involved in dynamically managing a portfolio. A number of other considerations also limit the use of duration in portfolio management. For instance, as Blake (1990) 5.8.1 points out, most Treasury bonds have durations of less than twelve years. This makes portfolio immunization difficult when liabilities are very long dated.

Zero-coupon bonds don't pose these problems, because their durations are identical to their terms to maturity. This potentially increases their attractiveness as investments. A five-year zero-coupon bond has a duration of five years when purchased; after two years, its duration is three years, no matter what interest rates have done. A long-dated zero-coupon bond can thus be safely used to match a long-dated liability.

Implied Spot Rates and Market Zero-Coupon Yields

The yield analysis described above considers coupon bonds as packages of zeros. How does one compare the yields of zero-coupon and coupon bonds? A two-year zero is clearly the point of comparison for a coupon bond whose duration is two years. What about very long-dated zero-coupon bonds, though, for which no equivalent coupon Treasury is usually available? The solution lies in the technique of stripping coupon Treasuries, which allows implied zero-coupon rates to be calculated, which can be compared with actual strip-market yields.

This section describes the relationships among spot interest rates and the actual market yields on zero-coupon and coupon bonds. It explains how an implied spot-rate curve can be derived from the redemption yields and prices observed on coupon bonds, and discusses how this curve may be used to compare bond yields. Note that, in contrast with the common practice, *spot rates* here refer only to rates derived from coupon-bond prices and are distinguished from *zero-coupon rates,* which denote rates actually observed on zero-coupon bonds trading in the market.

Spot Yields and Coupon-Bond Prices

As noted in chapter 2, a Treasury bond can be seen as a bundle of individual zero-coupon securities, each maturing on one of the bond's cash flow payment dates. In this view, the Treasury's price is the sum of the present values of all the constituent zero-coupon bond yields. Assume that the spot rates for the relevant maturities— $r_1, r_2, r_3, ... r_N$ —can be observed. If a bond pays a semiannual coupon computed at an annual rate of C from period 1 to period N, its present value can be derived using equation (16.7).

$$(16.7)$$

$$P = \frac{\frac{C_1}{2}}{\left(1 + \frac{1}{2}r_1\right)} + \frac{\frac{C_2}{2}}{\left(1 + \frac{1}{2}r_2\right)^2} + \frac{\frac{C_3}{2}}{\left(1 + \frac{1}{2}r_3\right)^3} + ... + \frac{\frac{C_{N-1}}{2}}{\left(1 + \frac{1}{2}r_{N-1}\right)^{N-1}} + \frac{\frac{C_N}{2} + 100}{\left(1 + \frac{1}{2}r_N\right)^N}$$

Equation (16.7) differs from the conventional redemption yield formula in that every cash flow is discounted, not by a single rate, but by the zero-coupon rate corresponding to the maturity period of the cash flow. To apply this equation, the zero-coupon-rate term structure must be known. These rates, however, are not always readily observable. Treasury prices, on the other hand, are and can be used to derive implied spot interest rates. (Although in the market the terms are used interchangeably, from this point on, *zero coupon* will be used only of observable rates and

FIGURE **16.2** *Ten Hypothetical Treasuries*

MATURITY DATE	YEARS TO MATURITY	COUPON (%)	YIELD TO MATURITY	PRICE
1-Sep-99	0.5	5.0	6.00	99.5146
1-Mar-00	1.0	10.0	6.30	103.5322
1-Sep-00	1.5	7.0	6.40	100.8453
1-Mar-01	2.0	6.5	6.70	99.6314
1-Sep-01	2.5	8.0	6.90	102.4868
1-Mar-02	3.0	10.5	7.30	108.4838
1-Sep-02	3.5	9.0	7.60	104.2327
1-Mar-03	4.0	7.3	7.80	98.1408
1-Sep-03	4.5	7.5	7.95	98.3251
1-Mar-04	5.0	8.0	8.00	100.0000

spot only of derived ones.) To see how the derivation works, consider the ten hypothetical U.S. Treasuries whose maturities, prices, and yields are shown in **FIGURE 16.2**. Assume that the yield curve is positive and that the securities' settlement date—March 1, 1999—is a coupon date, so none of them has accrued interest.

The first bond matures in precisely six months and thus has no intermediate cash flow before redemption. It can therefore be treated as a zero-coupon bond, and its yield of 6 percent taken as the 6-month spot rate. Using this, the 1-year spot rate can be derived from the price of a 1-year coupon Treasury. The principle of no-arbitrage pricing requires that the price of a 1-year Treasury strip equal the sum of the present value of the coupon Treasury's two cash flows:

September 1, 1999 $5
March 1, 2000 $5 + $100 = $105

The combined present values of these cash flows is given by equation (16.8).

$$PV_{Mar00} = \frac{5}{(1+r_1)} + \frac{105}{(1+r_2)^2} \qquad (16.8)$$

where
r_1 = one-half of the 6-month theoretical spot rate
r_2 = one-half of the 1-year theoretical spot rate

Plugging in the 6-month spot rate of 6 percent, divided by two, and the 1-year Treasury's observed market price, 103.5322, gives equation (16.9), which can be solved for r_2 as shown.

$$103.5322 = \frac{5}{(1.03)} + \frac{105}{(1+r_2)^2} \qquad (16.9)$$

$$103.5322 = 4.85437 + \frac{105}{(1+r_2)^2}$$

$$98.67783 = 105 / (1 + r_2)^2$$

$$(1 + r_2)^2 = 105 / 98.67783$$

$$(1 + r_2)^2 = 1.064069$$

$$(1 + r_2) = \sqrt{1.064069}$$

$$r_2 = 0.03154$$

The theoretical 1-year spot rate is twice r_2, or 0.06308, for an annualized bond-equivalent yield of 6.308 percent. This figure can now be used to calculate the theoretical 1.5-year spot rate. The cash flows for the 7 percent 1.5-year coupon Treasury are
    September 1, 1999 $3.50
    March 1, 2000 $3.50
    September 1, 2000 $103.50

The present value of this cash flow stream is

$$PV_{Sep00} = \frac{3.50}{(1+r_1)} + \frac{3.50}{(1+r_2)^2} + \frac{103.50}{(1+r_3)^3}$$

where
r_3 = one-half the 1.5-year theoretical spot rate

Plugging in one-half the 6-month and 1-year spot rates—0.03 and 0.03154, respectively—and the price of the 7 percent 1.5-year coupon Treasury gives equation (16.10), which can be solved for r_3.

$$100.8453 = \frac{3.5}{(1.03)} + \frac{3.5}{(1.03154)^2} + \frac{103.5}{(1+r_3)^3} \tag{16.10}$$

$$r_3 = 0.032035$$

This value doubled—6.407 percent—is the 1.5-year theoretical spot rate.

Repeating this process for all the Treasuries in figure 16.2 results in the implied spot rates shown in **FIGURE 16.3**.

The general relationship used to derive an implied spot rate for the Nth 6-month period, given earlier as 16.7, is repeated here, as (16.11), without the C subscripts. This can be rewritten as (16.12).

$$P_n = \frac{C/2}{\left(1+\frac{1}{2}r\right)} + \frac{C/2}{\left(1+\frac{1}{2}r_2\right)^2} + \frac{C/2}{\left(1+\frac{1}{2}r_3\right)^3} + \ \ldots\ + \frac{C/2+100}{\left(1+\frac{1}{2}r_N\right)^N} \tag{16.11}$$

$$P_N = \frac{C}{2}\sum_{t=1}^{N-1}\frac{1}{\left(1+r_t\right)^t} + \frac{C/2+100}{\left(1+r_N\right)^N} \tag{16.12}$$

FIGURE 16.3 *Bootstrapping Zero-Coupon Yields*

MATURITY DATE	YEARS TO MATURITY	YIELD TO MATURITY (%)	THEORETICAL SPOT RATE (%)
1-Sep-99	0.5	6.00	6.000
1-Mar-00	1.0	6.30	6.308
1-Sep-00	1.5	6.40	6.407
1-Mar-01	2.0	6.70	6.720
1-Sep-01	2.5	6.90	6.936
1-Mar-02	3.0	7.30	7.394
1-Sep-02	3.5	7.60	7.712
1-Mar-03	4.0	7.80	7.908
1-Sep-03	4.5	7.95	8.069
1-Mar-04	5.0	8.00	8.147

where

r_t = the already calculated theoretical spot rates $r_1, r_2, \ldots, r_{N-1}$

Equation (16.12) can be rearranged as (16.13), to solve for r_N.

$$r_N = \left[\frac{\dfrac{C}{2} + 100}{P_N - \dfrac{C}{2} \displaystyle\sum_{t=1}^{N-1} \dfrac{1}{\left(1 + r_t\right)^t}} \right]^{\frac{1}{N}} - 1 \qquad (16.13)$$

Implied Spot Yields and Zero-Coupon Bond Yields

Spot yields cannot be directly observed in the market. They can, however, be computed from the observed prices of zero-coupon bonds, or strips, if a liquid market exists in these securities. An implied spot yield curve can also, as the previous section showed, be derived from coupon bonds' prices and redemption yields. This section explores how the implied and actual strip yields relate to each other.

A spot rate may be thought of as the rate of return at purchase of a single-cash-flow security held to maturity. Alternatively, as the previous sections have shown, it can be viewed as the rate payable on a coupon bond whose yield is analyzed as a complex average of the spot yields of the individual zero-coupon cash flows into which the bond may be decomposed. Zero-coupon securities, in contrast, are actual market instruments that have been created, as explained in Choudhry (1999), by stripping individual cash flows from coupon bonds and trading the resulting cash flows separately. Those consisting of principal-redemption or residual cash flows are termed *principal strips*; the other cash flows are used to form *coupon* strips. The yields on these strips reflect market supply and demand. Zero-coupon yields and spot yields, although identical in theory, are thus different in practice.

Because it is affected by current demand, the yield of a particular zero-coupon bond at any time may differ from the equivalent-maturity spot yield. When investors value an individual zero-coupon bond less highly as a stripped security than as part of a coupon bond's theoretical package of zero-coupon cash flows, the strip's yield will be above the spot rate for the same maturity. The opposite happens when investors prefer to hold the zero-coupon security.

Despite their supply-and-demand-induced divergence from zero-coupon rates, implied spot rates are important because they enable inves-

FIGURE **16.4** *Gilt Market Gross Redemption True Yields and Implied Spot Yields on March 2, 1999*

TERM	GILT	TRUE YIELD	SPOT YIELDS	COUPON STRIP YIELD	PRINCIPAL STRIP YIELD
0.5	6% 10/8/1999	5.033	5.033	5.068	
1.5*	8% 7/12/2000	5.027	4.977	4.973	5.036
2.5	7% 6/11/2001	4.963	4.927	4.949	
3 *	7% 7/6/2002	4.878	4.836	4.91	4.867
4	8% 10/6/2003	4.845	4.789	4.824	
4.5*	6.50% 7/12/2003	4.735	4.687	4.786	4.691
5	6.75% 26/11/2004	4.709	4.658	4.798	
6 *	8.50% 7/12/2005	4.78	4.728	4.774	4.74
7 *	7.50% 7/12/2006	4.8	4.77	4.781	4.775
8 *	7.25% 7/12/2007	4.759	4.721	4.753	4.723
9	9% 13/10/2008	4.72	4.644	4.735	
10 *	5.75% 7/12/2009	4.604	4.542	4.723	4.518
11	6.25% 25/11/2010	4.695	4.665	4.711	
12	9% 12/7/2011	4.751	4.721	4.713	
13	9% 6/8/2012	4.787	4.727	4.713	
14	8% 27/9/2013	4.763	4.75	4.71	
16 *	8% 7/12/2015	4.711	4.659	4.705	4.670
18	8.75% 25/8/2017	4.729	4.695	4.682	
22 *	8% 7/6/2021	4.679	4.609	4.635	4.612
29 *	6% 7/12/2028	4.537	4.376	4.402	4.365

* Indicates strippable gilts

Source: Bloomberg

tors to assess relative value for both zero-coupon and coupon bonds. Consider **FIGURES 16.4** and **16.5**, a table and graph of the implied spot yields derived from U.K. government gilt prices on March 2, 1999, together with the coupon- and principal-strip yields for the same date. (The graph

FIGURE **16.5** *Spot and Strip Yields on March 2, 1999*
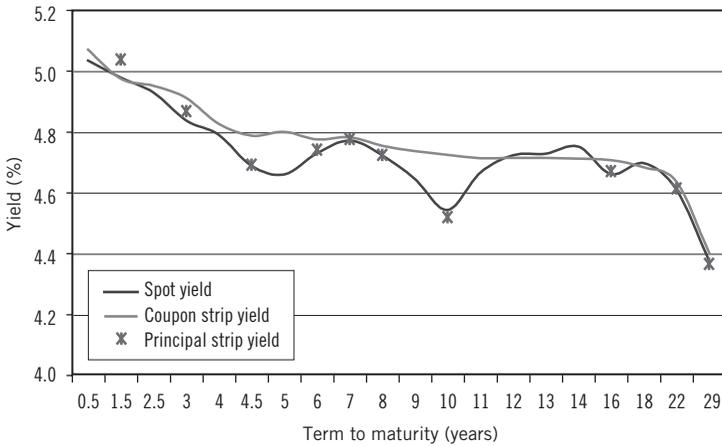
FIGURE **16.6** *Gilt Bond, Implied Spot and Coupon Strip Yields, March 2, 1999*
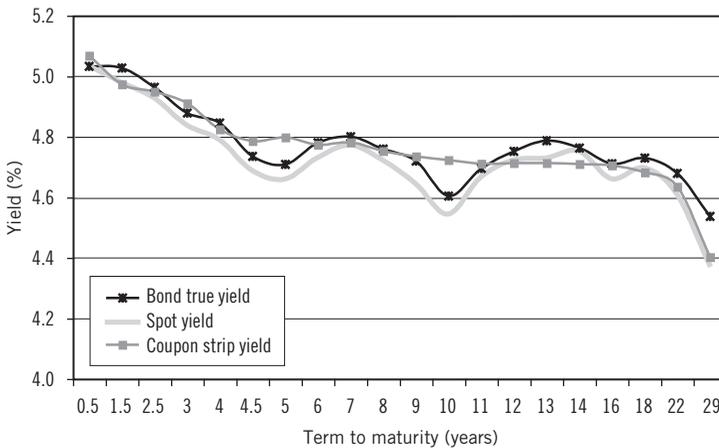
does not show a curve corresponding to the principal-strip yields because this would have necessitated making problematic conclusions about behavior in the long periods between the actual observed yields.) **FIGURE 16.6** illustrates the stripped yields against the true yield.

Two observations are worth making. First, the zero-coupon bonds are shown to be trading cheap relative to the spot curve throughout the term structure. This indicates that investors at the time were not prepared to hold strips unless they could earn a spread above their theoretical yields. This probably reflected the inverted yield curve during the period, which meant that strips would be expensive relative to coupon bonds of the same maturity. It should be noted, however, that strips were expensive relative to the spot curve from the 11- to the 15-year point on the curve. Second, principal strips trade at lower yields than coupon strips of the same maturity, reflecting the fact that investors prefer holding the former.

Determining Strip Values

The three most common ways to calculate a strip's value are to use

- ❏ the bond curve
- ❏ equivalent duration
- ❏ theoretical zero-coupon curve construction, also known as bootstrapping

The first method equates a strip's value with its spread to a bond having the same maturity. The main drawback of this rough-and-ready approach is that it compares two instruments with different risk profiles. This is particularly true for longer maturities. The second method, which aligns strip and coupon-bond yields on the basis of modified duration, is more accurate. The most common approach, however, is the third. This requires constructing a theoretical zero-coupon curve in the manner described above in connection with the relationship between coupon and zero-coupon yields.

When the bond yield curve is flat, the spot curve is too. When the yield curve is inverted, the theoretical zero-coupon curve must lie below it. This is because the rates discounting coupon bonds' earlier cash flows are higher than the rate discounting their final payments at redemption. In addition, the spread between zero-coupon and bond yields should decrease with maturity.

When the yield curve is positive, the theoretical zero-coupon curve lies above the coupon curve. Moreover, the steeper the coupon curve, the steeper the zero-coupon curve.

One argument against bootstrapping is that the theoretical zero-coupon yields obtained are too sensitive for real-world trading. This is because the spot curve derivation requires a coupon yield for every year, even if, for example, the yield curve is constructed only from 1-year, 10-year, and 30-year yields. The 30-year implied spot yield could be substantially higher

or lower depending on whether these maturity points are connected by a smooth curve or a straight line. This is particularly critical when there are few bonds between major yield points on the term structure—for example, eight liquid Treasuries between the 10- and 30-year maturities, with only two of them between the 20- and 30-year points. In such cases, filling in missing values using linear interpolation is inaccurate. To get around this problem, bond analysts use spline or other curve-smoothing techniques.

Strips Market Anomalies

Treasury analysts have observed some long-standing anomalies in the Treasury-strip market. These include the following:

Principal strips trade at a premium to coupon strips. Investors find principal strips more attractive because of their greater liquidity and, in some markets, for regulatory and tax reasons. This holds true even, at times, when their outstanding nominal amounts are lower than those of coupon strips.

The final principal strip trades expensive relative to their theoretical values. The shape of the strip yield curve might be expected to be similar to that of the coupon curve, which normally slopes gently upward. However, because investors, as noted above, prefer principal to coupon strips, final principal strips have greater weight than coupon strips.

The strips with the longest maturities are the most expensive. As in all well-developed strip markets, Treasury strips having the longest duration and the greatest convexity trade expensive relative to their theoretical values. Conversely, those with intermediate maturities tend to trade cheap to the curve. This is evident from a comparison of the Treasury strip and coupon curves.

Intermediate maturity coupons are often cheap relative to the curve. Because of client demand for longer maturities, market makers often find themselves with large quantities of intermediate-maturity coupon strips. In the Treasury strip market, 3- to 8-year coupon strips have traded cheap to the curve for this reason.

In contrast to the situation in most strip markets, Treasury strips with very short maturities do not trade expensive relative to the curve. When the yield curve is positive, short strips are often in demand because they enable investors to match liabilities without reinvestment risk and at a higher yield than they could get on coupon bonds of the same maturity. The Treasury strip yield curve, on the other hand, has been inverted from before there was a market. In other government strip markets, such as France's, however, short maturities of up to three years are often well bid.

Strips Trading Strategy

Market makers who strip Treasuries earn their profits through arbitrage that exploits the mispricing of the coupon bond. To preclude arbitrage opportunities, the bid price of an issued Treasury must be lower than the offer price of a synthetic one—that is, one reconstituted from a bundle of coupon and principal strips—and the Treasury's offer must be higher than the synthetic's bid. Otherwise, a risk-free profit can be obtained by selling one instrument while simultaneously buying the other and pocketing the difference.

The potential profit from stripping a Treasury coupon depends on current market Treasury yields and the implied spot yield curve. Consider a hypothetical 5-year, 8 percent Treasury trading at par—and therefore offering a yield to maturity of 8 percent—in the yield curve environment shown in figure 16.2. A market maker buys the Treasury and strips it with the intention of selling the resulting zero-coupon bonds at the yields indicated in figure 16.2.

FIGURE 16.7 on the following page shows the present values of the Treasury's cash flows, each discounted using the relevant market interest rate, and the present values of the strip cash flows, each discounted using the observed market yield corresponding to its maturity. A comparison of the two sets reveals an opportunity for arbitrage profit.

The fourth column shows how much the market maker paid for each of the cash flows by buying the entire package of them—that is, by buying the bond at a yield of 8 percent. The $4 coupon payment due in three years, for instance, cost $3.1616, based on the 8 percent (4 percent semiannual) yield. But if the assumptions embodied in the table are correct, investors are willing to accept a lower yield, of 7.30 percent (3.65 percent semiannual), for this maturity and pay $3.2258 for the three-year strip corresponding to the coupon payment. On this one coupon payment, the market maker thus realizes a profit of $0.0645, the difference between $3.2258 and $3.1613. The total profit from selling all the strips is $0.4913 per $100 nominal of the original Treasury.

What if, instead of the observed yields to maturity, investors required the theoretical spot yields from figure 16.3? **FIGURE 16.8** shows that, in this case, the total proceeds from the sale of the zero-coupon Treasuries would be approximately $100, representing no profit and thus rendering the stripping process uneconomical. These two scenarios demonstrate that profit opportunities exist where strip yields deviate from theoretical ones.

In practice, strip yields do differ from theoretical yields, indicating that there are (often very small) differences between derived and actual

FIGURE **16.7** *Zero-Coupon Bonds Stripped from Hypothetical 5-Year, 8 Percent Treasury, Valued at Market Yields*

MATURITY DATE	YEARS TO MATURITY	CASH FLOW	PRESENT VALUE AT 8%	YIELD TO MATURITY (%)	PRESENT VALUE AT YIELD TO MATURITY
1-Sep-99	0.5	4	3.8462	6.00	3.8835
1-Mar-00	1.0	4	3.6982	6.30	3.7594
1-Sep-00	1.5	4	3.5560	6.40	3.6393
1-Mar-01	2.0	4	3.4192	6.70	3.5060
1-Sep-01	2.5	4	3.2877	6.90	3.3760
1-Mar-02	3.0	4	3.1613	7.30	3.2258
1-Sep-02	3.5	4	3.0397	7.60	3.0809
1-Mar-03	4.0	4	2.9228	7.80	2.9453
1-Sep-03	4.5	4	2.8103	7.95	2.8164
1-Mar-04	5.0	104	70.2587	8.00	70.2587
			100.0000		100.4913

prices. Do these differences present arbitrage opportunities? Because of the efficiency and transparency of developed-country bond markets, the answer is usually no. The process of coupon stripping prevents Treasuries from trading at prices that are *materially* different from their theoretical prices based on the derived spot yield curve. When discrepancies arise, arbitrage activity causes them to disappear very quickly. In a liquid market such as that for U.S. Treasuries, the laws of supply and demand eliminate obvious arbitrage opportunities. Nevertheless, occasional opportunities do arise to exploit differences between actual market prices of strips and the theoretical prices implied by the benchmark coupon Treasury yield curve.

FIGURE **16.8** *Zero-Coupon Bonds Stripped from Hypothetical 5-Year, 8 Percent Treasury, Valued at Spot Yields*

MATURITY DATE	YEARS TO MATURITY	CASH FLOW	PRESENT VALUE AT 8%	THEORETICAL SPOT RATE (%)	PRESENT VALUE AT SPOT RATE
1–Sep–99	0.5	4	3.8462	6.000	3.8835
1–Mar–00	1.0	4	3.6982	6.308	3.7591
1–Sep–00	1.5	4	3.5560	6.407	3.6390
1–Mar–01	2.0	4	3.4192	6.720	3.5047
1–Sep–01	2.5	4	3.2877	6.936	3.3731
1–Mar–02	3.0	4	3.1613	7.394	3.2171
1–Sep–02	3.5	4	3.0397	7.712	3.0693
1–Mar–03	4.0	4	2.9228	7.908	2.9331
1–Sep–03	4.5	4	2.8103	8.069	2.8020
1–Mar–04	5.0	104	70.2587	8.147	69.7641
			100.0000		~100.0000

Case Study: Treasury Strip Yields and Cash Flow Analysis

This section illustrates the process of yield and cash flow analysis through the Bloomberg screens for the 4 percent Treasury maturing on February 15, 2014[1]—the 10-year benchmark during 2004—and its principal and coupon strips maturing on February 15, 2014.

FIGURE **16.9** on the following page shows the cash flows received by a holder of $1 million nominal of the Treasury, together with the corresponding spot rates. On the trade date—March 25, 2004, for settlement on March 26—the bond was priced at 99-19+, for a yield of 4.0479 percent, and its convexity was 0.770.

FIGURE **16.10** shows the cash flow for the Treasury's principal strip. Its yield is 4.0751 percent, corresponding to a price of $67.10027 per $100 nominal, which represents a spread above the gross redemption yield of the coupon Treasury. This relationship is expected, given a positive yield

FIGURE 16.9 *Cash Flow Analysis of the 4 Percent Treasury Due 2014*

CASHFLOW ANALYZER

US TREASURY N/B   T 4 02/15/14   98-28+ / 98-30+ (4.15 /14)

PRC/YLD:	99-19+	/ 4.0479	OR	SPOT,DUR,ADJ.DUR: S	SETTLE: 3/26/04
PRES VAL:		1000.489354	OR		
SPREAD:		31.166			

DISPLAY SPOT RATES BY   Workout: 2/15/14

FMC/IYC: 1 # 25 US TREASURY ACTIVES   Daytype: ACT/ACT   Redemp: 100.0000%

Side: Mid B/A/M   Pricing Date: 3/25/04   Frequency: 2   Face: 1000

IYC

CURVE # 25 12/29/4
2yr 1.50
3yr 1.86
5yr 2.68
10yr 3.28
30yr 8.59
10yr 4.18
30yPP 4FBB0
ACT/ACTe S
DAYTP FREQ
ACT/360e –
ACT/ACTe S

PAY DATE	CASH FLOW	SPOT	PAY DATE	CASH FLOW	SPOT
8/15/04	20.00	1.1212	8/15/10	20.00	3.0688
2/15/05	20.00	1.1212	2/15/11	20.00	3.1899
8/15/05	20.00	1.1181	8/15/11	20.00	3.3101
2/15/06	20.00	1.1181	2/15/12	20.00	3.4314
8/15/06	20.00	1.4419	8/15/12	20.00	3.5512
2/15/07	20.00	1.8775	2/15/13	20.00	3.6727
8/15/07	20.00	2.0811	8/15/13	20.00	3.7920
2/15/08	20.00	2.2872	2/15/14	1020.00	3.9136
8/15/08	20.00	2.4907			
2/15/09	20.00	2.6971			
8/15/09	20.00	2.8289			
2/15/10	20.00	2.9499			

TOT DURATION: 8.158
TOT CASHFLOW: 1400.00   ACC INT: 4.40 Prev Coup Date: 2/17/04

FIGURE 16.10 *Cash Flow Analysis of the 4 Percent Treasury 2014 Principal Strip, at a Yield of 4.075 Percent*

CASHFLOW ANALYZER

STRIP PRINC   SP 02/15/14   68-06 / 68-10 (4.23 /21)

PRC/YLD:	67.100357	/ 4.0751	OR	SPOT,DUR,ADJ.DUR: S	SETTLE: 3/26/04
PRES VAL:		671.003574	OR		
SPREAD:		16.023			

DISPLAY SPOT RATES BY   Workout: 2/15/14

FMC/IYC: 1 # 25 US TREASURY ACTIVES   Daytype: ACT/ACT   Redemp: 100.0000%

Side: Mid B/A/M   Pricing Date: 3/25/04   Frequency: 2   Face: 1000

IYC

CURVE # 25 12/29/4
2yr 1.50
3yr 1.86
5yr 2.68
10yr 3.28
30yr 8.59
10yr 4.18
30yPP 4FBB0
ACT/ACTe S
DAYTP FREQ
ACT/360e –
ACT/ACTe S

PAY DATE	CASH FLOW	SPOT	PAY DATE	CASH FLOW	SPOT
8/15/04	0.00	1.1212	8/15/10	0.00	3.0688
2/15/05	0.00	1.1212	2/15/11	0.00	3.1899
8/15/05	0.00	1.1181	8/15/11	0.00	3.3101
2/15/06	0.00	1.1181	2/15/12	0.00	3.4314
8/15/06	0.00	1.4419	8/15/12	0.00	3.5512
2/15/07	0.00	1.8775	2/15/13	0.00	3.6727
8/15/07	0.00	2.0811	8/15/13	0.00	3.7920
2/15/08	0.00	2.2872	2/15/14	1000.00	3.9136
8/15/08	0.00	2.4907			
2/15/09	0.00	2.6971			
8/15/09	0.00	2.8289			
2/15/10	0.00	2.9499			

TOT DURATION: 9.890
TOT CASHFLOW: 1000.00   ACC INT: 0.00 Prev Coup Date: 2/17/04

FIGURE 16.11 *Cash Flow Analysis of the February 2014 Coupon Strip, at a Yield of 4.089 Percent*

```
                                               Govt   PFC1
              CASHFLOW  ANALYZER                   page 1 of 1
STRIPS            $ 02/15/14    67-10+ / 67-14+  ( 4.04 /02) BGN  @10:00
PRC/YLD:      67.010/ 4.0090   OR   SPOT,DUR,ADJ.DUR: $    SETTLE: 3/26/04
PRES VAL:      670100.000000   OR
SPREAD:           -17.608            DISPLAY SPOT RATES BY  Workout: 2/15/14
FMC/IYC: 1 # 25 US TREASURY ACTIVES      Daytype: ACT/ACT  Redemp: 100.0000%
Side: Bid B/A/M  Pricing Date: 3/25/04   Frequency: 2   Face:    1000000
IYC
CURVE ♦ 25 | PAY DATE | CASH FLOW | SPOT  | PAY DATE | CASH FLOW   | SPOT
  3/25/ 4  | 8/15/04  |    0.00   | 0.9799| 8/15/10  |      0.00   | 3.1436
3mo  0.92  | 2/15/05  |    0.00   | 1.0343| 2/15/11  |      0.00   | 3.3042
6mo  0.98  | 8/15/05  |    0.00   | 1.0701| 8/15/11  |      0.00   | 3.4632
2yr  1.50  | 2/15/06  |    0.00   | 1.1092| 2/15/12  |      0.00   | 3.6241
3yr  2.17  | 8/15/06  |    0.00   | 1.5747| 8/15/12  |      0.00   | 3.7830
5yr  2.65  | 2/15/07  |    0.00   | 2.1888| 2/15/13  |      0.00   | 3.9440
10yr 4.05  | 8/15/07  |    0.00   | 2.3101| 8/15/13  |      0.00   | 4.1024
30yr 4.66  | 2/15/08  |    0.00   | 2.4319| 2/15/14  | 1000000.00  | 4.2638
           | 8/15/08  |    0.00   | 2.5522|
DAYTP FREQ | 2/15/09  |    0.00   | 2.6741|
ACT/360o   | 8/15/09  |    0.00   | 2.8257|
ACT/ACTo S | 2/15/10  |    0.00   | 2.9859|
TOT DURATION:      9.890
TOT CASHFLOW:   1000000.00    ACC INT:      0.00  Prev Coup Date: 2/15/04
```

curve. There is only one cash flow: the redemption payment, which equals $1 million for a holding of $1 million nominal. The principal strip's convexity is higher than the coupon bond's, as is its duration, again as expected.

Finally, **FIGURE 16.11** shows the cash flow analysis for the coupon strip maturing on February 15, 2014, and trading at a yield of 4.089 percent, for a price of 67.01 per $100 nominal. This illustrates an anomaly noted earlier: although the law of one price states that all strips maturing on the same date should cost the same—after all, why would investors require a different yield for a payment of interest than for one of principal?—principal strips in fact trade at lower yields than coupon strips, because they are more liquid and so more sought after.

Chapter Notes

1. In fact, this date is a Saturday, so the actual redemption proceeds will be paid on Monday, February 17, 2014.

This page is intentionally blank

Approaches to Trading

The term *trading* covers a wide range of activities. Market makers, who quote two-way prices to market participants, may also be asked to provide customer service and build retail and institutional volume. Alternatively, their sole responsibility may be to run their books at a profit and maximize return on capital.

Trading approaches are determined in part by the nature of the market involved. In a highly transparent and liquid market, such as that for U.S. Treasuries, price spreads are fairly narrow, so opportunities to profit from the mispricing of individual securities, although not nonexistent, are rare. Participants in such markets may engage in *relative value* trades, or *spread trading,* which involves relationships such as the yield spread between individual securities or the future shape of the yield curve. On the derivatives exchanges, a large volume of the trading is for hedging purposes, but speculation is also prominent. Bond and interest rate traders often use futures or options contracts to bet on the direction of particular markets. This is frequently the case with market makers who have little customer business—whether because they are newcomers, for historical reasons, or because they don't want the risk related to servicing high-quality customers. They speculate to relieve the tedium, often with unfortunate results.

Speculative trading is based on the views of the trader, the desk, or the department head. The view may be that of the firm's economics or research department—based on macro- and microeconomic factors affecting not just the specific market but the economy as a whole—or it may be that of the individual trader, resulting from *fundamental* and *technical*

analysis. Because credit spreads are key to the performance of corporate bonds, those running corporate-debt desks focus on the fundamentals driving those spreads, concentrating on individual sectors as well as on the corporations themselves and their wider environment. Technical analysis, or *charting,* is based on the belief that the patterns displayed by continuous-time series of asset prices repeat themselves. By detecting current patterns, therefore, traders should be able to predict reasonably well how prices will behave in the future. Many traders combine fundamental and technical analysis, although chartists often say that for technical analysis to work effectively, it must be the only method used.

Futures Trading

Because of the liquidity of their market and the ease and relative cheapness of trading them, futures and other derivatives are often preferred to cash instruments, for both speculation and hedging. The essential considerations in futures trading are the volatility of the associated commodity or financial instrument and the leverage deriving from the fact that the margin required to establish a futures position is a very small percentage of the contract's notional value.

Uncovered trades, made without owning the underlying asset, are speculative bets on the direction of the market. Traders who believe short-term U.S. interest rates are going to fall, for example, might buy Eurodollar contracts—representing the level, at contract expiration, of the interest rate on a 3-month deposit of $1 million in commercial banks located outside the United States—on the Chicago Mercantile Exchange (CME). The contracts' tick value—the price change associated with a movement of 0.005 of a percentage point in the interest rate—is $12.50. Say a trader buys one lot at 98.84, representing a future interest rate of 1.16 (100 – 98.84 percent) percent, and sells it at the end of the day for 98.85, representing a rate of 1.15 percent. That's a rise of 0.01, or two ticks, resulting in a profit of $25, from which brokerage fees are subtracted.

Traders can also bet on their interest rate views using a cash-market product or a forward rate agreement (FRA)—a contract specifying the rate to be received or paid starting at specified future date. Transactions are easier and cheaper, however, on the futures exchange, because of the low cost of dealing there and the liquidity of the market and narrow price spreads.

More common than directional bets are trades on the spread between the rates of two different contracts. Consider **FIGURES 17.1** and **17.2**, both of which relate to the prices for the CME Eurodollar contract on March 24, 2004 (contracts exist for every month in the year).

FIGURE 17.1 *CME 90-Day Eurodollar Synthetic Forward Rates*

FIGURE 17.2 *CME 90-Day Eurodollar Futures Strip Analysis*

Futures exchanges use the letters H, M, U, and Z to refer to the contract months March, June, September, and December. The June 2004 contract, for example, is denoted by "M4." Forward rates can be calculated for any term, starting on any date. Figure 17.1 shows the futures

prices on that day and the interest rate that each price implies. The *stub* is the interest rate from the current date to the expiry of the first, or *front month,* contract—in this case, the June 2004 contract. Figure 17.2 lists the forward rates for six months, one year, and so on from the *spot* date. It is possible to trade a strip of contracts replicating any term out to the maximum maturity of the contract. This may be done to hedge or to speculate. Figure 17.2 shows that a spread exists between the cash and futures curves. It is possible to take positions on cash against futures, but it is easier to trade only on the futures exchange.

Short-term money market rates often behave independently of the yield curve as a whole. Traders in this market watch for cash market trends—more frequent borrowing at a certain point in the curve, for example, or market intelligence suggesting that one point of the curve will rise or fall relative to others. One way to exploit these trends is with a *basis spread* trade, running a position in a cash instrument such as a CD against a futures contract. The best way, though, is with a spread trade, shorting one contract against a long position in another. Say you believe that in June 2004 3-month interest rates will be lower than those implied by the current futures price, shown in Figure 17.1, but that in September 2004 they will be higher. You can exploit this view by buying the M4 contract and shorting an equal weight of the U4—say one hundred lots of each. In doing so, you are betting not on the market direction but on the spread between two contracts. If rates move as you expect, you realize a profit. Because it carries no directional risk, spread trading requires less margin than open-position trading. Figure 17.2 presents similar trade possibilities, depending on your view of forward interest rates. It is also possible to arbitrage between contracts on different exchanges.

In the example above, you are *shorting the spread,* believing that it will narrow. Taking the opposite positions—short the near contract and long the far one—is *buying the spread.* This is done when the trader believes the spread will widen. Note that the difference between the two contracts' prices is not limitless: a futures contract's theoretical price provides an upper limit to the size of the spread or the basis, which, moreover, cannot exceed the cost of carry—that is, the net cost of buying the cash security today and delivering it into the futures market at contract expiry. The same principle applies to short-dated interest rate contracts, where cost of carry equals the difference between the interest cost of borrowing funds to buy the security and the income generated by holding it until delivery. The two associated costs for a Eurodollar spread trade are the notional borrowing and lending rates for buying one contract and selling another.

Traders who believe the cost of carry will decrease can sell the spread to exploit this view. Those with longer time horizons might trade the spread between the short-term interest rate contract and the long bond future. Such transactions are usually done by proprietary traders, because it is unlikely that one desk would be trading both 3-month and 10- or 20-year interest rates. A common example is a yield curve trade. Say you believe the U.S. dollar yield curve will steepen or flatten between the 3-month and the 10-year terms. You can buy or sell the spread using the CME 90-day Eurodollar contract and the U.S. Treasury 10-year note contract, traded on CBOT. Because one 90-day interest rate contract is not equivalent to one bond contract, however, the trade must be duration-weighted to be first-order risk neutral. The note contract represents $100,000 of a notional Treasury, although its tick value is $15.625; the Eurodollar contract represents a $1 million time deposit. Equation (17.1) calculates the hedge ratio, with $1,000 being the value of a 1 percent change in the value of each contract.

$$h = \frac{(100 \times tick) \times P_b^f \times D}{(100 \times tick) \times P_{short\ ir}^f} \tag{17.1}$$

where
  $tick$ = the tick value of the contract
  D = the duration of the bond represented by the long bond contract
  P_b^f = the price of the bond futures contract
  $P_{short\ ir}^f$ = the price of the short-term deposit contract

The notional maturity of a long-bond contract is always given as a range: for the contract on the 10-year note, for example, it is six to ten years. The duration used to calculate the hedge ratio would be that of the cheapest-to-deliver bond.

A *butterfly spread* involves two spreads between three contracts, the middle of which is perceived to be mispriced relative to the other two. The underlying concept is that the market will correct this mispricing, changing the middle contract's relation to the outer contracts. Traders put on a butterfly when they are not sure which contract or contracts will move to effect this adjustment.

Consider figure 17.1 again. The prices of the front three contracts are 98.84, 98.72, and 98.51. Traders may feel that the September contract, at a spread of +12.5 basis points to the June contract and –21 basis points to the December one, is undervalued. The question is, will this be corrected by a fall in the June and December prices or by a rise in the September one? The traders don't need to choose. If they believe that the spread be-

tween the June and September contracts will widen and the one between the September and December contracts will narrow, they can put on a butterfly spread by buying the first and selling the second. This is also known as *selling the butterfly spread.*

Yield Curves and Relative Value

Bond market participants take a keen interest in both the cash and the zero-coupon (spot) yield curves. In markets where an active zero-coupon bond market exists, the spreads between implied and actual zero-coupon yields also receive much attention.

Determinants of Government Bond Yields

Market makers in government bonds consider various factors in deciding how to run their books. Customer business apart, decisions to purchase or sell securities will depend on their views about the following:

❑ whether short- and long-term interest rates are headed up or down
❑ which maturity point along the entire term structure offers the best value
❑ which of the issues having that maturity offers the best value

These three factors are related but are affected differently by market-moving events. A report on the projected size of the government's budget deficit, for example, will not have much effect on two-year bond yields, but if the projections are unexpected, they could adversely affect long-bond yields. The type of effect—negative or positive—depends on whether the projections were higher or lower than anticipated.

For a first-level analysis, many market practitioners look no further than the traditional yield curve like the one shown in **FIGURE 17.3**. Investors with no particular views on the future shape of the curve or level of interest rates might adopt a neutral strategy, holding bonds with durations matching their investment horizons. If the curve is positive and they believe interest rates are likely to remain stable for a time, they might buy bonds with longer durations, thus picking up additional yield but increasing their interest rate risk.

Once investors have determined which part of the yield curve to invest in or switch into, they must select specific issues. To make an informed choice, they use relative value analysis.

Relative value analysis focuses on bond issues located in certain sectors, or "local" parts, of the curve. Because a bond's yield is a function not just of its duration—after all, two issues with near-identical

FIGURE 17.3 *U.S. Treasury Yield Curve, March 25, 2004*
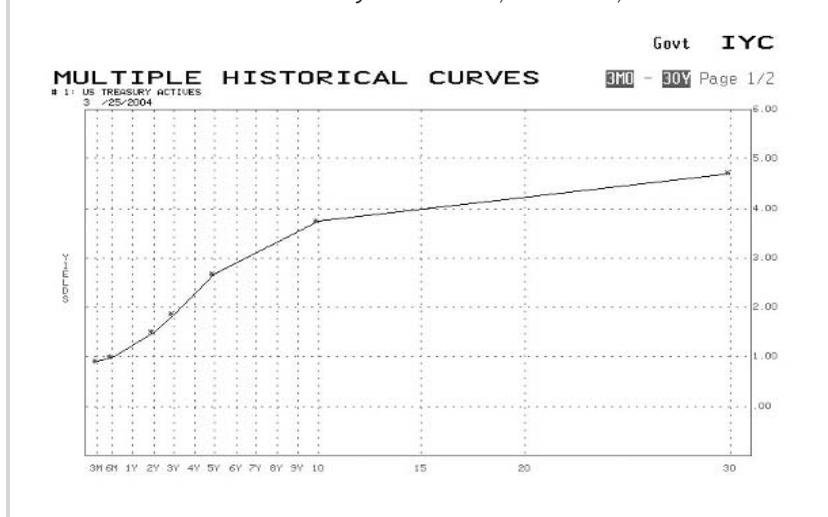
duration can have different yields—the analysis assesses other factors as well. These include liquidity, the interplay of supply and demand, and coupon rate, all of which affect yield. **FIGURE 17.4** illustrates the impact of coupon rate.

The figure shows that, when the curve is inverted, investors can pick up yield while shortening duration. This might seem an anomalous situation, but in fact, liquidity issues aside, the market generally disfavors bonds with high coupons, so these usually trade cheap to the curve.

As with any commodity, supply and demand also play important roles in determining bond prices, and therefore their yields. A shortage of issues at a particular point in the curve—the result, perhaps, of an effort to reduce public-sector debt—depresses yields for that maturity. On the other hand, when interest rates decline—ahead of or during a recession, say—and new bonds are issued with increasingly lower coupons, the stock of "outdated" high-coupon bonds increases and can end up trading at a higher yield.

Demand is a function primarily of investors' views of a country's economic prospects. It can also be affected, however, by government actions, such as the debt-buyback program instituted during the last years of the Clinton administration in response to the budget surplus. These buybacks reduced the supply of bonds in the market, artificially depressing yields because demand could not be met, especially for 30-year bonds, which the

FIGURE **17.4** *Yields of Bonds with Similar Durations but Different Coupons, Given an Inverted Curve*
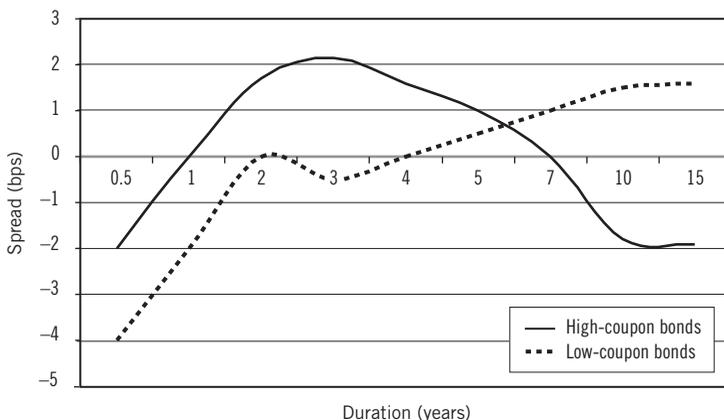
Treasury announced it was discontinuing. With the return to large budget deficits during the Bush administration, however, issuance has resumed, and the impact of low supply has abated.

Liquidity differences often produce yield differences among bonds with similar durations. Institutional investors prefer to hold the benchmark bond—the current 2-, 5-, 10-, or 30-year issue—which both increases its liquidity and depresses its yield. The converse is also true: because more-liquid bonds are easier to convert into cash if necessary, demand is higher for them, and their yields are thus lower. The effect of liquidity on yield can be observed by comparing the market price of a six-month bond with its theoretical value, derived by discounting its cash flow at the current 6-month T-bill rate. The market price—which is equal to the present value of its cash flow discounted at its yield—is lower than the theoretical value, reflecting the fact that the T-bill yield is lower than the bond yield, even though the two securities' cash flows fall on the same day. The reason is liquidity: the T-bill is more readily realizable into cash at any time.

A bond's coupon and liquidity, as well as its duration, thus help determine the yield at which it trades. Accordingly, analyses of relative value among bonds consider these factors in conjunction with others.

Characterizing the Complete Term Structure

As many readers undoubtedly have gathered, the traditional yield curve illustrated in Figure 17.4 is inadequate for analyzing the market. This is because it highlights only the curve's general shape, which is not a sufficient basis on which to make trading decisions. This section describes a technique for gaining a more complete, and useful, picture.

FIGURE 17.5 graphs the March 2004 bond par yield curve against the T-bill yield curve for the same date. The two curves in **FIGURE 17.6** represent the *low-coupon* and *high-coupon* yield spreads—that is, the yield differences between coupon bonds trading at par and, in the first case, bonds with coupons 100 basis points below the par yield and, in the second case, those with coupons 100 basis points above the par yield. By watching and comparing the curves illustrated in the two figures, investors can see the impact of coupons on the shape of the par yield curve and on yields at different maturity points.

Identifying Relative Value in Government Bonds

This section discusses the factors that must be assessed in analyzing the relative values of government bonds. Since these securities involve no credit risk (unless they are emerging-market debt), credit spreads are not among the considerations. The zero-coupon yield curve provides the framework for all the analyses explored.

FIGURE **17.5** *T-Bill and Par Yield Curves, March 2004*
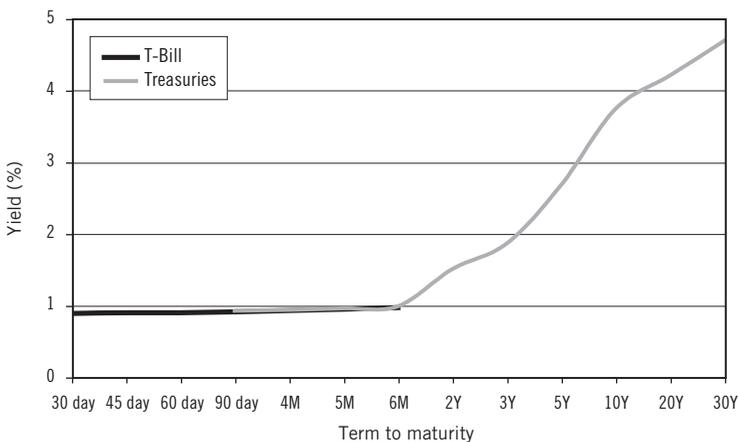
Source: Bloomberg

Treasury curve includes interpolated yields.

FIGURE **17.6** *Par and High- and Low-Coupon Yield Curves, March 2004*
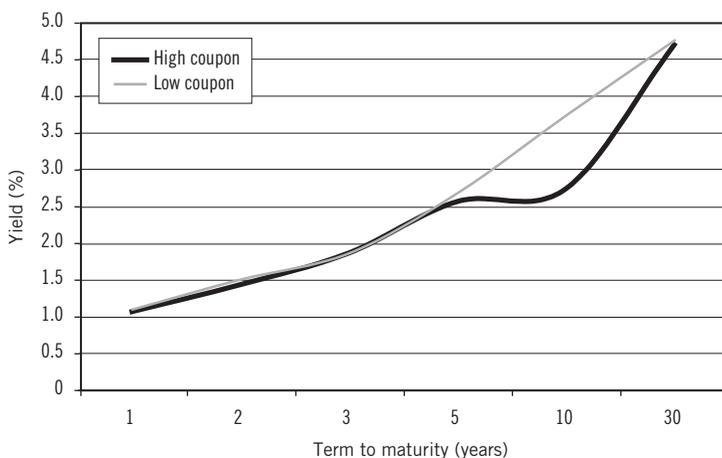
Source: Bloomberg

The objective of much bond analysis is to determine the relative values of individual securities and thus identify which should be purchased and which sold. Such decisions are, at the broadest level, a function of whether one thinks interest rates are going to rise or fall. Analysis in this sense identifies securities' absolute value. More locally based analysis focuses on specific sectors of the yield curve and is aimed at determining whether these will flatten or steepen, or whether bonds with similar durations are trading at large enough spreads to warrant switching from one to another. This type of analysis identifies relative value.

On rare occasions, assessing relative value is fairly straightforward. If the 3-year yield were 5.75 percent, the 2-year 5.70 percent, and 4-year 6.15 percent, for example, 3-year bonds would appear to be overpriced. Such situations, however, do not occur often in real life. More realistically, the 3-year bond's value would have to be assessed relative to much shorter- or longer-dated instruments. Comparing a short-dated bond with other short-term securities is considerably different from comparing, say, the 2-year bond to the 30-year. Although, in a graph, the smooth curve linking 1-year to 5-year yields appears to be repeated from the 5-year out to the 30-year, in reality the very short-dated sector of the yield curve often behaves independently of the long end.

One method of identifying relative value is to quantify the effect of coupon rates on bond yields. The relationship between the two is expressed in equation (17.2).

$$rm = rm_P + c \times \max\left(C_{PD} - rm_P, 0\right) + d \times \min\left(C_{PD} - rm_P, 0\right) \quad (17.2)$$

where
    rm = the yield of the bond being analyzed
    rm_P = the yield of a par bond of specified duration
    C_{PD} = the coupon of an arbitrary bond whose duration is similar to the par bond's
    c = a coefficient representing the effect of a high coupon on a bond's yield
    d = a coefficient representing the effect of a low coupon on a bond's yield

When the par bond yield is lower than the coupon of the bond having a similar duration—that is, $C_{PD} > rm_P$, (17.2) reduces to (17.3).

$$rm = rm_P + c \times \left(C_{PD} - rm_P\right) \quad\quad (17.3)$$

Equation 17.3 expresses the yield spread between a high-coupon bond and a par bond as a linear function of the spread between the first bond's coupon and the par bond's yield and coupon. In reality, this relationship may not be purely linear. The yield spread between the two bonds, for instance, may widen more slowly when the gap between the coupons is very large. Equation 17.3 thus approximates the effect of a high coupon on yield more accurately for bonds trading close to par.

The same analysis can be applied to bonds with coupons lower than the coupon of the par bond having the same duration.

A bond may be valued relative to comparable securities or against the par or zero-coupon yield curve. The first method is more appropriate in certain situations. It is suitable, for instance, when a low-coupon bond is trading rich to the curve but fair compared with other low-coupon bonds. This may indicate that the overpricing is a property not of the individual bond but of all low-coupon bonds.

The comparative value analysis can be extended from the local structure of the yield curve to groups of similar bonds. This is an important part of the analysis, because it is particularly informative to know the cheapness or costliness of a single issue relative to the whole yield curve.

Traders may use the technique described above to identify excess positive or negative yield spreads for all the bonds in the term structure, resulting in a list like that in **FIGURE 17.7**. From these, they might select two or more bonds, some of which are cheap and others expensive relative to the curve, then switch between them or put on a spread trade.

FIGURE **17.7** *Yields and Excess Yield Spreads for Five Treasuries and Less-Liquid Issues, March 24, 2004*

Source: Bloomberg

The benchmark securities in the figure are all expensive relative to the par curve, and the less-liquid bonds are cheap. The $2^5/8$ percent 2009 appears cheap, but the $6^1/2$ percent 2010, which has a shorter duration, offers a higher yield. This curious anomaly disappeared a few days later, so, if not taken advantage of immediately, the profit opportunity was lost.

Hedging Bond Positions

A hedge is a position in a cash or off-balance-sheet instrument that removes the market risk exposure of another position. For example, a long position in 10-year bonds can be hedged with a short position in 20-year issues or with futures contracts. The concept is straightforward. Implementing it effectively, however, requires a precise calculation of the amount of the hedge needed, and that can be complex.

Simple Hedging Approaches

Say an investor wishes to hedge a position in one of the bonds listed in figure 17.7 with a counterbalancing position in another of these issues. As explained in chapter 2, it is possible to calculate the size of the position required for a duration-weighted hedge using the ratio of the two bonds' basis point values, or BPVs—that is, the change in their prices corresponding to a 1–basis point change in yields. This approach is very common in

the market. It is based, however, on two assumptions that hinder its effectiveness: first, that the two bonds' yields have comparable volatility and, second, that changes in the yields of the two bonds are highly correlated. This implies that the changes in yields are highly positively correlated. Correlation refers to changes in direction, not absolute values. Highly correlated means that if Bond A increases from 10 to 12, and Bond B is at 79.25, its price will increase too. In situations where one or both of these assumptions fails to hold, the hedge is compromised.

The assumption of comparable yield volatility becomes increasingly unrealistic the more the bonds differ in terms of market risk and behavior. Say the position to be hedged is a $1 million holding of the 5-year issue in figure 17.7 and the hedging instrument is the 5-year bond. A duration-weighted hedge would consist of a short position in the 5-year. Even if the two bonds' yields are perfectly correlated, they might still change by different amounts if the bonds have different yield volatilities. Say the 2-year is twice as volatile as the 5-year. That means the 5-year yield moves only half as far as the 2-year in the same situation. For instance, an event causing the latter to rise 5 basis points would effect a mere 2.5-basis-point increase in the former. So a hedge calculated according to the two bonds' BPV and assuming an equal change in yield for both bonds would be incorrect. Specifically, the short position in the 5-year bond would effectively hedge only half the risk exposure of the 2-year position.

The assumption of perfectly correlated yield changes is similarly unrealistic and so causes similar misweightings. Although bond yields across the whole term structure are positively correlated most of the time, this is not always the case. Returning to the example, assume that the 2-year and 5-year bonds possess identical yield volatilities but that changes in their yields are uncorrelated. This means that a 1-basis-point fall or rise in the 2-year yield implies nothing about change in the 5-year yield. That, in turn, means that the 5-year bonds cannot be used to hedge 2-year bonds, at least not with any certainty.

Hedge Analysis

From the preceding discussion, it is clear that at least two factors beyond BPV determine the effectiveness of a bond hedge: the bonds' yield volatilities and the extent to which changes in their yields are correlated.

FIGURE 17.8 shows the standard deviations—that is, volatilities—and correlations of weekly yield changes for a set of Treasuries during the nine months to March 2004. Note that, contrary to the assumptions inherent in the BPV hedge calculation, volatilities are far from uniform, and yield changes are imperfectly correlated. The standard deviation of weekly yield changes is highest for the short-dated paper and declines throughout

FIGURE 17.8 *Yield Volatilities and Correlations, Selected Bonds, Nine Months to March 2004*

	SEGMENT					
	2-YEAR	**3-YEAR**	**5-YEAR**	**10-YEAR**	**20-YEAR**	**30-YEAR**
VOLATILITY (BP)	18.7	19.5	20.2	20.0	20.6	21.2
CORRELATION						
2-year	1.000	0.973	0.949	0.919	0.887	0.879
3-year	0.973	1.000	0.961	0.935	0.901	0.889
5-year	0.949	0.961	1.000	0.968	0.951	0.945
10-year	0.919	0.935	0.968	1.000	0.981	0.983
20-year	0.887	0.901	0.951	0.981	1.000	0.987
30-year	0.879	0.889	0.945	0.983	0.987	1.000

the period for longer-dated paper. Correlations, as might be expected, are highest among bonds in the same maturity sectors and decline as they move farther apart along the yield curve; for example, two-year bond yields are more positively correlated with five-year yields than with 30-year ones.

Hedges can be made more accurate by adjusting their weightings according to the standard relationship for correlations and the effect of correlation. Consider two bonds with nominal values M_1 and M_2. If the bonds' yields change by Δr_1 and Δr_2, the net change in the position's value is given by equation (17.4)

$$\Delta PV = M_1 BPV_1 \Delta r_1 + M_2 BPV_2 \Delta r_2 \qquad (17.4)$$

The change in net value of a two-bond position is a function of the two securities' nominal values, their volatilities, and the correlation between their yield changes. The standard deviation of such a position may therefore be expressed by equation (17.5).

$$\sigma_{pos} = \sqrt{M_1^2 BPV_1^2 \sigma_1^2 + M_2^2 BPV_2^2 \sigma_2^2 + 2M_1 M_2 BPV_1 BPV_2 \sigma_1 \sigma_2 \rho} \quad (17.5)$$

where

ρ = the correlation between the yield volatilities of bonds 1 and 2

Equation (17.5) can be rearranged as shown in (17.6) to solve for the optimum hedge value for any bond.

$$M_2 = -\frac{\rho BPV_1 \sigma_1}{BPV_2 \sigma_2} M_1 \qquad (17.6)$$

where

M_2 = the nominal value of the bond used to hedge nominal value M_1 of the first bond

The lower the correlation between the two bonds' yields—and, thus, the more independent changes in one are of changes in the other—the smaller the optimal hedge position. If the two bonds' yields exhibit identical volatility and change in lockstep—a correlation of 1—equation (17.6) reduces to equation (17.7), the traditional hedge calculation, based solely on BPV.

$$M_2 = \frac{BPV_1}{BPV_2} M_1 \qquad (17.7)$$

Summary of the Derivation of the Optimum-Hedge Equation

According to equation (17.5), the variance of a net change in the value of a two-bond portfolio is given by equation (17.8), which can be rewritten as (17.9), using the partial derivative of the variance σ^2 with respect to the nominal value of the second bond.

$$\sigma_{pos}^2 = M_1^2 BPV_1^2 \sigma_1^2 + M_2^2 BPV_2^2 \sigma_2^2 + 2M_1 M_2 BPV_1 BPV_2 \sigma_1 \sigma_2 \rho \qquad (17.8)$$

$$\frac{\partial \sigma^2}{\partial^2 M_2} = 2M_2 BPV_2^2 \sigma_2^2 + 2M_1 BPV_1 BPV_2 \sigma_1 \sigma_2 \rho \qquad (17.9)$$

Setting equation (17.5) to zero and solving for M_2 gives equation (17.10), which is the hedge quantity for the second bond.

$$M_2 = -\frac{\rho BPV_1 \sigma_1}{BPV_2 \sigma_2} M_1 \qquad (17.10)$$

This page is intentionally blank

The Black-Scholes Model in Microsoft Excel

The figure on the following page shows the spreadsheet formulas required to build the Black-Scholes model in Microsoft Excel. The Analysis Tool-Pak add-in must be available, otherwise some of the function references may not work. Setting up the cells in the way shown enables the fair value of a vanilla call or put option to be calculated. The latter calculation employs the put-call parity theorem.

Price of underlying	100
Volatility	0.0691
Option maturity	3 months
Strike price	99.5
Risk-free rate	5%

Microsoft Excel Calculation of Vanilla Option Price

CELL	C	D	
8	**Underlying price, *S***	100	
9	Volatility %	0.0691	
10	Option maturity years	0.25	
11	**Strike price, *X***	99.50	
12	Risk-free interest rate %	0.05	
13			
14			
15			**CELL FORMULAE:**
16	ln (*S*/*X*)	0.0050125418	=LN (D8/D11)
17	Adjusted return	0.0000456012500	=((D12-D9)^2/ 2)*D10
18	Time adjusted volatility	0.1314343943	=(D9*D10)^0.5
19	d_2	0.0384841662	=(D16+D17)/D18
20	*N*(d2)	0.5153492331	=NORMSDIST(D19)
21			
22	d_1	0.1699185605	=D19+D18
23	*N*(d1)	0.5674629098	=NORMSDIST(D22)
24	e^{-rt}	0.9875778005	=EXP(-D10*D12)
25			
26	**CALL**	6.1060184985	=D8*D23-D11*D20*D24
27	**PUT**	4.3700096476 *	=D26-D8+D11*D24

* By put-call parity, $P = C - S + Xe^{-rt}$

REFERENCES

Chapter 1—The Bond Instrument

Fabozzi, F. 1989. *Bond Markets, Analysis and Strategies.* New York: Prentice Hall, chap. 2.

———. 1993. *Bond Markets, Analysis and Strategies,* 2nd ed. New York: Prentice Hall.

———. 1997. *Fixed Income Mathematics,* 3rd ed. Princeton: McGraw-Hill, 190–192.

——— ed. 2002. *The Handbook of Fixed Income Securities,* 6th ed. Princeton: McGraw-Hill.

———. 1998. *Valuation of Fixed Income Securities and Derivatives,* 3rd ed. New Hope, PA: FJF Associates.

———. 1998. *Treasury Securities and Derivatives.* New Hope, PA: FJF Associates.

Higson, C. 1995. *Business Finance.* Oxford: Butterworth.

Questa, G. 1999. *Fixed Income Analysis for the Global Financial Market.* Chichester: John Wiley & Sons.

Stigum, M., and F. Robinson. 1996. *Money Market and Bond Calculations.* New York: Irwin.

Sundaresan, S. 1997. *Fixed Income Markets and Their Derivatives.* Cincinnati: South-Western.

Van Deventer, D., and K. Imai. 1997. *Financial Risk Analytics.* New York: Irwin, 9–11.

Weston, J. F., and T. E. Copeland. 1986. *Managerial Finance.* Florence, KY: Dryden.

Chapter 2—Bond Instruments and Interest Rate Risk

Butler, C. 1998. *Mastering Value-at-Risk.* London: FT Prentice Hall.

Burghardt, G. 1994. *The Treasury Bond Basis.* New York: McGraw-Hill.

Garbade, K. 1996. *Fixed Income Analytics.* Boston: MIT Press, chap., 3, 4, 12.

Macaulay, F. 1999. *The Movements of Interest Rates, Bond Yields and Stock Prices in the United States Since 1856.* London: RISK Classics Library.

Chapter 3—Bond Pricing and Spot and Forward Rates

Adams, K., and D. Van Deventer. 1994. Fitting Yield Curves and Forward Rate Curves with Maximum Smoothness. *Journal of Fixed Income* 4, 52–62.

Avellaneda, M., and P. Laurence. 2000. *Quantitative Modeling of Derivative Securities.* Milan: Chapman & Hall/CRC, chap. 10–12.

Baxter, M., and A. Rennie. 1996. *Financial Calculus.* Cambridge: Cambridge University Press, chap. 5.

Cambell, J., A. Lo, and A. MacKinlay. 1997. *The Econometrics of Financial Markets.* Princeton: Princeton University Press, chap. 10–11.

Cox, J., J. Ingersoll, and S. Ross. 1985. An Inter-Temporal General Equilibrium Model of Asset Prices. *Econometrica* 53.

Deacon, M., and A. Derry. 1994. Estimating the Term Structure of Interest Rates. *Bank of England Working Paper Series* No 24, July.

Duffie, D. 1992. *Dynamic Asset Pricing Theory.* Princeton: Princeton University Press.

Fabozzi, F. 1993. *Bond Markets, Analysis and Strategies,* 2nd ed. New York: Prentice Hall, chap. 5.

———. 1997. *Fixed Income Mathematics.* New York: McGraw-Hill.

Gujarati, D. 1999. *Basic Econometrics,* 3rd ed. New York: Irwin.

Hicks, J. 1946. *Value and Capital,* 2nd ed. Oxford: Oxford University Press.

Ingersoll, J. Jr. 1987. *Theory of Financial Decision Making.* Lanham, MD: Rowman & Littlefield, chap. 18.

James, J., and N. Webber. 2000. *Interest Rate Modeling.* Chichester: John Wiley & Sons.

Jarrow, R. 1996. *Modeling Fixed Income Securities and Interest Rate Options.* Princeton: McGraw-Hill, chap. 3.

Kitter, G. 1999. *Investment Mathematics for Finance and Treasury Professionals.* Chichester: John Wiley & Sons, chap. 3, 5.

McCulloch, J. 1971. Measuring the Term Structure of Interest Rates. *Journal of Business* 44, 19–31.

———. 1975. The Tax-Adjusted Yield Curve. *Journal of Finance* 30, 811–830.

Neftci, S. 2000. *An Introduction to the Mathematics of Financial Derivatives,* 2nd ed. Oxford: Academic Press.

Nelson, C. and A. Siegel. 1987. Parsimonious Modeling of Yield Curves. *Journal of Business* 60(4), 473–489.

Ross, Sheldon M. 1999. *An Introduction to Mathematical Finance.* Cambridge: Cambridge University Press.

Schaefer, S. 1981. Measuring a Tax-Specific Term Structure of Interest Rates in the

Market for British Government Securities. *Economic Journal* 91, 415–438.

Shiller, R. 1990. The Term Structure of Interest Rates. In Friedman, B., and F. Hanh, eds. *Handbook of Monetary Economics,* Volume 1. Amsterdam: North Holland, chap. 13.

Suits, D., A. Mason, and L. Chan. 1978. Spline Functions Fitted by Standard Regression Methods. *Review of Economics and Statistics* 60, 132–139.

Sundaresan, S. 1997. *Fixed Income Markets and Their Derivatives.* Cincinnati: South-Western.

Van Deventer, D., and K. Imai. 1997. *Financial Risk Analytics.* New York: Irwin.

Van Horne, J. 1995. *Financial Management and Policy,* 10th ed. New York: Prentice Hall.

Waggoner, D. 1997. Spline Methods for Extracting Interest Rate Curves From Coupon Bond Prices. Working Paper, *Federal Reserve Bank of Atlanta,* 97–10.

Chapter 4—Interest Rate Modeling

Baxter, M., and A. Rennie. 1996. *Financial Calculus.* Cambridge: Cambridge University Press, 57–62.

Black, F. 1995. Interest Rates as Options. *Journal of Finance,* December, 1371–1376.

Black, F., E. Derman, and W. Toy. 1990. A One-Factor Model of Interest Rates and Its Application to Treasury Bond Options. *Financial Analysts Journal,* Spring, 33–39.

Brennan, M. and E. Schwartz. 1979. A Continuous Time Approach to the Pricing of Bonds. *Journal of Banking and Finance,* July, 135–155.

Campbell, J. 1986. A Defense of Traditional Hypotheses About the Term Structure of Interest Rates. *Journal of Finance* 41, 183–193.

Chen, R-R, and L. Scott. 1992. Pricing Interest Rate Options in a Two-Factor Cox-Ingersoll-Ross Model of the Term Structure. *Review of Financial Studies,* Winter, 613–636.

Cox, J., J. Ingersoll, and S. Ross. 1985. A Theory of the Term Structure of Interest Rates. *Econometrica* 53, March, 385–407.

Fong, H. G., and O. Vasicek. 1991. Fixed Income Volatility Management. *Journal of Portfolio Management,* Summer, 41–46.

Gibbons, M., and K. Ramaswamy. 1994. The Term Structure of Interest Rates: Empirical Evidence. *Review of Financial Studies.*

Haug, E. 1998. *The Complete Guide to Option Pricing Formulas.* Princeton: McGraw-Hill, chap. 4.

Heath, D., R. Jarrow, and A. Morton. 1992. Bond Pricing and the Term Structure of Interest Rates: A New Methodology for Contingent Claims Valuation. *Econometrica* 60, January, 77–105.

Ho, T., and S-B. Lee. 1986. Term Structure Movements and Pricing Interest Rate

Contingent Claims. *Journal of Finance,* December, 1011–1029.

Hull, J. 1997. *Options, Futures and Other Derivatives,* 3rd ed. New York: Prentice Hall, 220–222.

Hull, J., and A. White. 1990. Pricing Interest Rate Derivative Securities. *Review of Financial Studies,* 573–592.

Ingersoll, J., Jr. 1987. *Theory of Financial Decision Making.* Lanham, MD: Rowman & Littlefield.

James, J., Webber, N. 2000. *Interest Rate Modeling.* Chichester: John Wiley & Sons, chap. 5–15.

Jarrow, R. 1996. *Modeling Fixed Income Securities and Interest Rate Options.* Princeton: McGraw-Hill.

Merton, R. 1993. *Continuous Time Finance.* Oxford: Blackwell, chap. 11.

Rebonato, R. 1996. *Interest Rate Option Models.* Chichester: John Wiley & Sons.

Sundaresan, S. 1997. *Fixed Income Markets and Their Derivatives.* Cincinnati: South-Western.

Van Deventer, D., and K. Imai. 1997. *Financial Risk Analytics.* New York: Irwin, chap. 5.

Vasicek, O. 1977. An Equilibrium Characterization of the Term Structure. *Journal of Financial Economics,* 5, 177–188.

Chapter 5—Fitting the Yield Curve

Anderson, N., F. Breedon, M. Deacon, A. Derry, and M. Murphy. 1996. *Estimating and Interpreting the Yield Curve.* Chichester: John Wiley & Sons.

Carleton, W., and I. Cooper. 1976. Estimation and Uses of the Term Structure of Interest Rates. *Journal of Finance,* September, 1067–1083.

Choudhry, M., D. Joannas, R. Pereira, and R. Pienaar. 2001. *Capital Market Instruments: Valuation and Analysis.* London: FT Prentice Hall.

Deacon, M., and A. Derry. 1994. Estimating the Term Structure of Interest Rates. *Bank of England Working Paper* No 24, July.

Gujarati, D. 1995. *Basic Econometrics,* 3rd ed. New York: McGraw-Hill.

James, J., and N. Webber. 2000. *Interest Rate Modeling.* Chichester: John Wiley & Sons.

Jordan, J., and S. Mansi. 2000. How Well Do Constant-Maturity Treasuries Approximate the On-the-Run Term Structure? *Journal of Fixed Income* 10:2, September 35–45.

McCulloch, J. H. 1971. Measuring the Term Structure of Interest Rates. *Journal of Business* January, 19–31.

Questa, G. 1999. *Fixed Income Analysis for the Global Financial Market.* Chichester: John Wiley & Sons.

Steeley, J. M. 1991. Estimating the Gilt-Edged Term Structure: Basis Splines and Confidence Intervals. *Journal of Business Finance and Accounting* 18, 513–530.

Suits, D., A. Mason, and L. Chan. 1978. Spline Functions Fitted by Standard Regression Methods. *Review of Economics and Statistics* 60, 132–139.

Svensson, L. 1994. *Estimating and Interpreting Forward Interest Rates: Sweden, 1992–1994.* Cambridge, MA: National Bureau of Economic Research.

Van Deventer, D., and K. Imai. 1997. *Financial Risk Analytics.* New York: Irwin.

Vasicek, O., and H. G. Fong. 1982. Term Structure Modeling Using Exponential Splines. *Journal of Finance* 37(2), May, 339–348.

Chapter 6—Forwards and Futures Valuation

For descriptions and analyses of bond futures contracts, basis, implied repo, and the cheapest-to-deliver bond, see Burghardt et al. (1994). Plona (1996) is a readable treatment of the European government bond basis.

Burghardt, G. 1994. *The Treasury Bond Basis.* New York: Irwin.

French, K. 1983. A Comparison of Futures and Forwards Prices. *Journal of Financial Economics* 12, November, 311–342.

Galitz, L. 1995. *Financial Engineering.* London: FT Pitman, chap. 3, 4, 6–8.

Grinblatt, M., and N. Jegadeesh. 2000. Futures vs. Forward Prices: Implications for Swap Pricing and Derivatives Valuation. In Jegadeesh, N., and B. Tuckman, eds. *Advanced Fixed-Income Valuation Tools.* New York: John Wiley & Sons.

Hull, J. 2000. *Options, Futures and Other Derivatives,* 4th ed. New York: Prentice Hall.

Jarrow, R., and G. Oldfield. 1981. Forward Contracts and Futures Contracts. *Journal of Financial Economics* 9, December, 373–382.

Kolb, R. 2000. Futures, Options and Swaps, 3rd ed. Oxford: Blackwell.

Park, H., and A. Chen. 1985. Differences Between Futures and Forward Prices: A Further Investigation of Marking to Market Effects. *Journal of Futures Markets* 5, February, 77–88.

Plona, C. 1996. *The European Bond Basis.* New York: McGraw-Hill.

Rubinstein, M. 1999. *Rubinstein on Derivatives.* London: RISK Books.

Chapter 7—Swaps

Bicksler, J., and A. Chen. 1986. An Economic Analysis of Interest Rate Swaps. *Journal of Finance* 41: 3. 645–655.

Brotherton-Ratcliffe, R., and B. Iben. 1993. Yield Curve Applications of Swap Products. In Schwartz, R., and C. Smith, eds. *Advanced Strategies in Financial Risk Management.* New York: New York Institute of Finance.

Das, S. 1994. *Swaps and Financial Derivatives,* 2nd ed. London: IFR Publishing.

Decovny, S. 1998. *Swaps,* 2nd ed. London: FT Prentice Hall.

Dunbar, N. 2000. Swaps Volumes See Euro Wane. *Risk* September.

Eales, B. 1995. *Financial Risk Management.* Princeton: McGraw Hill, chap. 3.

Fabozzi, F., ed. 1998. *Perspectives on Interest Rate Risk Management for Money Managers and Traders.* New Hope, PA: FJF Associates.

Gup, B., and R. Brooks. 1993. *Interest Rate Risk Management.* New York: Irwin.

Henna, P. 1991. *Interest-Rate Risk Management Using Futures and Swaps.* Chicago: Probus.

International Swaps and Derivatives Association. 1991. Code of Standard Working, Assumptions, and Provisions for Swaps. New York.

Jarrow, R., and S. Turnbull. 2000. *Derivative Securities,* 2nd ed. Cincinnati: South-Western.

Khan, M. 2000. Online Platforms Battle for Business. *Risk* September.

Kolb, R. 2000. *Futures, Options and Swaps,* 3rd ed. Oxford: Blackwell.

Li, A., and V. R. Raghavan. 1996. Libor-in-Arrears Swaps, *Journal of Derivatives* 3, Spring, 44–48.

Lindsay, R. 2000. High Wire Act. *Risk* August.

Marshall, J., and K. Kapner. 1990. *Understanding Swap Finance.* Cincinnati: South-Western.

Turnbull, S. 1987. Swaps: A Zero Sum Game. *Financial Management* 16, Spring, 15–21.

Chapter 8—Options

For a detailed discussion of the mathematical basis of the Black-Scholes model, readers should refer to the original account in Black and Scholes (1973); other good accounts are in Ingersoll (1987), Neftci (1996), and Nielsen (1999).

Black, F., and M. Scholes. 1973. The Pricing of Options and Corporate Liabilities. *Journal of Political Economy* 81, May–June, 637–659.

Black, F., E. Derman, and W. Toy. 1996. A One-Factor Model of Interest Rates and Its Application to Treasury Bond Options. In Hughston, L., ed. *Vasicek and Beyond.* London: Risk Publications.

Bookstaber, R. 1982. *Option Pricing and Strategies in Investing.* Reading, MA: Addison-Wesley.

Boyle, P. 1986. Option Valuation Using a Three Jump Process. *International Options Journal* 3, 7–12.

Brace, A., D. Gatarek, and M. Musiela. 1996. The Market Model of Interest Rate Dynamics. In Hughston, L., ed. *Vasicek and Beyond.* London: Risk Publications.

Briys, E., Bellalah, M., Mai, H-M; De Varenne, F. 1998. *Options, Futures and Exotic Derivatives.* Chichester: John Wiley and Sons.

Choudhry, M. 2001. *The Bond and Money Markets: Strategy, Trading, Analysis.* Oxford: Butterworth-Heinemann, chap. 42–49.

Cox, D., and H. Miller. 1965. *The Theory of Stochastic Processes.* Oxford: Chapman & Hall.

Cox, J., S. Ross. 1976. The Valuation of Options for Alternative Stochastic Pro-

cesses. *Journal of Financial Economics* 3, 145–166.

Cox, J., S. Ross, and M. Rubinstein. 1979. Option Pricing: A Simplified Approach. *Journal of Financial Economics* 7, October, 229–264.

Debreu, G. 1954. Representation of a Preference Ordering by a Numerical Function. In Thrall, R., C. Coombs, and R. Davis, eds. *Decision Processes.* Chichester: John Wiley & Sons.

Fama, E. 1965. The Behaviour of Stock Prices. *Journal of Business* 38, January, 34–105.

Harrison, J., and D. Kreps. 1979. Martingales and Arbitrage in Multi-Period Securities Markets. *Journal of Economic Theory* 20, 381–408.

Harrison, J., and S. Pliska. 1981. Martingales and Stochastic Integrals in the Theory of Continuous Trading. *Stochastic Processes and Their Applications* 11, 216–260.

Haug, E. G. 1998. *The Complete Guide to Option Pricing Formulas.* Princeton: McGraw-Hill.

Heston, S. 1993. A Closed Form Solution for Options With Stochastic Volatility With Application to Bond and Currency Options. *Review of Financial Studies* 6, 327–343.

Heston, S. 1993. Invisible Parameters in Option Prices. *Journal of Finance* 48 (3), 933–947.

Hull, J., and A. White. 1987. The Pricing of Options on Assets With Stochastic Volatilities. *Journal of Finance* 42, 281–300.

Hull, J., and A. White. 1988. An Analysis of the Bias Caused by a Stochastic Volatility in Option Pricing. *Advances in Futures and Options Research* 3, 29–61.

Ito, K. 1951. On Stochastic Differential Equations. *American Mathematical Society* 4, 1–51.

Klemkosky, R., and B. Resnick. 1979. Put-Call Parity and Market Efficiency. *Journal of Financial Economics* 34, December, 1141–1155.

Merton, R. 1973. Theory of Rational Option Pricing. *Bell Journal of Economics and Management Science* 4, Spring, 141–183.

———. 1976. Option Pricing When Underlying Stock Returns Are Discontinuous. *Journal of Financial Economics* 3, Jan–Mar, 125–144.

Neftci, S. 1996. *An Introduction to the Mathematics of Financial Derivatives.* Oxford: Academic Press.

Nielsen, L. T. 1999. *Pricing and Hedging of Derivative Securities.* Oxford: Oxford University Press.

Rendleman, R., and B. Bartter. 1979. Two State Option Pricing. *Journal of Finance* 34, 1092–1110.

Rubinstein, M. 1985. Nonparametric Tests of Alternative Option Pricing Models. *Journal of Financial Economics* 40, 455–480.

Scott, L. 1987. Option Pricing When the Variance Changes Randomly. *Journal of Financial and Quantitative Analysis* 22, 419–438.

Stein, E., and J. Stein. 1991. Stock Price Distributions With Stochastic Volatility: An Analytic Approach. *Review of Financial Studies* 4, 113–135.

Whaley, R. 1981. On the Valuation of American Call Options on Stocks with Known Dividends. *Journal of Financial Economics* 9, 207–211.

Chapter 9—Measuring Option Risk

Hull, J. 2000. *Options, Futures and Other Derivatives,* 4th ed. Princeton: Prentice Hall.

Kolb, R. 2000. *Futures, Options and Swaps,* 3rd ed. Oxford: Blackwell.

Galitz, L. 1995. *Financial Engineering,* rev. ed. London: FT Pitman.

Marshall, J., and V. Bansal. 1992. *Financial Engineering.* New York: New York Institute of Finance.

Tucker, A. 1991. *Financial Futures, Options and Swaps.* Eagan, MN: West Publishing.

Chapter 10—Credit Derivatives

Bessis, J. 1998. *Risk Management in Banking.* Chichester: John Wiley & Sons, 17–18.

Crosbie, P. 1997. Modeling Default Risk. *Credit Derivatives: Key Issues.* London: British Bankers Association.

Das, S. 1995. Credit Risk Derivatives. *Journal of Derivatives,* Spring, 7–23.

Das, S. 1997. *Credit Derivatives: Products, Applications and Pricing.* Singapore: John Wiley & Sons.

Francis, J., J. Frost, and J. G. Whittaker. 1999. *The Handbook of Credit Derivatives.* Princeton: McGraw-Hill.

Das, S., and P. Tufano. 1996. Pricing Credit Sensitive Debt When Interest Rates, Credit Ratings and Credit Spreads Are Stochastic. *Journal of Financial Engineering* 5, 161–198.

Gup, B., and R. Brooks. 1993. *Interest Rate Risk Management.* New York: Irwin.

Jarrow, R., and S. Turnbull. 1996. Credit Risk. In Alexander, C., ed. *Handbook of Risk Management and Analysis.* New York: John Wiley & Sons.

Kasapi, A. 1999. *Mastering Credit Derivatives.* London: FT Prentice Hall, chap. 4.

Longstaff, F., and E. Schwartz. 1995. Valuing Credit Derivatives. *Journal of Fixed Income* June, 6–14.

Lucas, D. 1995. Default Correlation and Credit Analysis. *Journal of Fixed Income* June 32–41.

Pierides, Y. 1997. The Pricing of Credit Risk Derivatives. *Journal of Economic Dynamics and Control* 5, 1579–1611.

Whittaker, G., and J. Frost. 1999. An Introduction to Credit Derivatives. *Journal of Lending and Credit Risk Management* May, 15–25.

Whittaker, G., and S. Kumar. 1996. Credit Derivatives: A Primer. In Dattatreya, R., ed. *Handbook of Fixed Income Derivatives.* Chicago: Probus.

Chapter 11—The Analysis of Bonds with Embedded Options

Bodie, Z., and R. Taggart. 1978. Future Investment Opportunities and the Value of the Call Provision on a Bond. *Journal of Finance* 33, 1187–2000.

Choudhry, M. 2001. *The Bond and Money Markets.* London: Butterworth-Heinemman.

Fabozzi, F. J. 1997. *Fixed Income Mathematics: Analytical and Statistical Techniques,* 3rd ed. Princeton: McGraw-Hill, chap. 16.

Kalotay, A., G. O. Williams, and F. J. Fabozzi. 1993. A Model for the Valuation of Bonds and Embedded Options. *Financial Analysts Journal,* May–June, 35–46.

Kish, R., M. Livingstone. 1992. The Determinants of the Call Feature on Corporate Bonds. *Journal of Banking and Finance* 16, 687–703.

Livingstone, M. 1993. *Money and Capital Markets,* 2nd ed. New York: NYIOF.

Mitchell, K. 1991. The Call, Sinking Fund, and Term-to-Maturity Features of Corporate Bonds: An Empirical Investigation. *Journal of Financial and Quantitative Analysis* 26, June, 201–222.

Narayanan, M. P., and S. P. Lim. 1989. On the Call Provision on Corporate Zero-Coupon Bonds. *Journal of Financial and Quantitative Analysis* 24, March, 91–103.

Questa, G. 1999. *Fixed Income Analysis for the Global Financial Market.* Chichester: John Wiley & Sons, chap. 8.

Tuckman, B. 1996. *Fixed Income Securities.* New York: John Wiley & Sons, chap. 17.

Van Horne, J. C. 1986. *Financial Management and Policy.* New York: Prentice Hall.

Windas, T. 1996. *An Introduction to Option-Adjusted Spread Analysis.* Princeton: Bloomberg Press.

Chapter 12—Inflation-Indexed Bonds

In addition to the references below, Bloomberg users can type ILB <go> 99 <go> *to see illustrations of yield and other calculations on U.S. TII securities.*

Anderson, N., F. Breedon, M. Deacon, A. Derry, and J. Murphy. 1996. *Estimating and Interpreting the Yield Curve.* Chichester: John Wiley & Sons.

Arak, M., and L. Kreicher. 1985. The Real Rate of Interest: Inferences From the New U.K. Indexed Gilts. *International Economic Review* (26) 2, 399–408.

Bootle, R. 1991. *Index-Linked Gilts: A Practical Investment Guide,* 2nd ed. Cambridge: Woodhead-Faulkner.

Brown, R., and S. Schaefer. 1994. The Term Structure of Real Interest Rates and the Cox, Ingersoll and Ross Model. *Journal of Financial Economics* 35:1, 3–42.

Brynjolfsson, J., and F. Fabozzi. 1999. *Handbook of Inflation Indexed Bonds.* New Hope, PA: FJF Associates.

Deacon, M., and A. Derry. 1994. Deriving Estimates of Inflation Expectations From the Prices of U.K. Government Bonds. *Bank of England Working Paper* no. 23, July.

———. 1998. *Inflation Indexed Securities.* London: Prentice Hall.

Fisher, T. 1930. *The Theory of Interest.* London: Pickering & Chatto Ltd.

Foresi, S., A. Penati, and G. Pennacchi. 1997. Reducing the Cost of Government Debt: The Role of Index-Linked Bonds. In de Cecco, M., L. Pecchi, and G. Piga, eds. *Public Debt: Index-Linked Bonds in Theory and Practice.* Cheltenham: Edward Elgar.

James, J., and N. Webber. 2000. *Interest Rate Modeling.* Chichester: John Wiley & Sons.

Waggoner, D. 1997. Spline Methods for Extracting Interest Rate Curves From Coupon Bond Prices. *Federal Reserve Bank of Atlanta Working Paper,* 97–10.

Wojnilower, A. 1997. *Inflation–Indexed Bonds: Promising the Moon.* New York: Clipper Group.

Chapter 14—Securitization and Mortgage-Backed Securities

Ames, L. 1997. Mortgage Backed Securities Analysis, in Fabozzi, F., *Handbook of Fixed Income Securities,* 4th ed. New York: McGraw Hill.

Anderson, G., J. Barber, and C. Chang. 1993. Prepayment Risk and the Duration of Default-Free Mortgage-Backed Securities. *Journal of Financial Research* 16, 1–9.

Arora, A., D. Heike, and R. Mattu. 2000. Risk and Return in the Mortgage Market: Review and Outlook. *Journal of Fixed Income,* June, 5–18.

Bear Stearns. 1994. *Asset-Backed Securities Special Report,* 5 December.

Bhattacharya, A., and F. Fabozzi, eds. 1996. *Asset-Backed Securities.* New Hope, PA: FJF Associates.

Fabozzi, F. J. 1997. *Fixed Income Mathematics: Analytical and Statistical Techniques,* 3rd ed. New York: McGraw-Hill.

———. 2000. *The Handbook of Mortgage-Backed Securities.* New Hope, PA: FJF Associates.

———, ed. 1998. *The Handbook of Structured Financial Products.* New Hope, PA: FJF Associates.

———, ed. 1987. *Mortgage-Backed Securities: New Strategies, Applications and Research.* London: Probus Publishing.

Fabozzi, F., and D. Jacob. 1996. *The Handbook of Commercial Mortgage-Backed Securities.* New Hope, PA: FJF Associates.

Fabozzi, F., C. Ramsey, and F. Ramirez. 1994. *Collateralized Mortgage Obligations.* New Hope, PA: FJF Associates.

Hayre, L., and C. Mohebbi. 1989. Mortgage Pass-Through Securities. In Fabozzi, F., ed. *Advances and Innovations in the Bond and Mortgage Markets.* Probus Publishing, 259–304.

Hayre, L., S. Chaudhary, and R. Young. 2000. Anatomy of Prepayments. *Journal of Fixed Income* June, 19–49.

Morris, D. 1990. *Asset Securitization: Principles and Practices.* New York: Executive Enterprise.

Schwartz, E., W. Torous. 1992. Prepayment and Valuation of Mortgage Pass-Through Securities. *Journal of Business* 15: 2, 221–240.

Sundaresan, S. 1997. *Fixed Income Markets and Their Derivatives.* Cincinnati: South-Western, chap. 9.

Tuckman, B. 1996. *Fixed Income Securities.* New York: John Wiley & Sons, chap. 18.

Waldman, M., and S. Modzelewski. 1985. A Framework for Evaluating Treasury-Based Adjustable Rate Mortgages. In Fabozzi, F., ed. *The Handbook of Mortgage-Backed Securities.* New York: Probus Publishing.

Chapter 16—The Yield Curve, Bond Yield, and Spot Rates

Blake, D. 1990. *Financial Market Analysis.* New York: McGraw-Hill.

Choudhry, M. 1999. *Introduction to the Gilt Strips Market.* London: Securities Institute (Services) Limited.

———. 2001. *The Repo Handbook.* Oxford: Butterworth-Heinemann.

Fabozzi, F. 1996. *Bond Portfolio Management.* New Hope, PA: FJF Associates, chap. 10–14.

Chapter 17—Approaches to Trading

Neftci (2000) is an excellent introduction to bond pricing and its relationship to spot and forward rates. Campbell, Lo, and MacKinley (1997) is a very readable book, worth purchasing for Chapter 10 alone, which is an excellent and readable study of term structure, providing proofs of some of the results discussed in this chapter.

Brennan, M., and E. Schwartz. 1979. A Continuous Time Approach to the Pricing of Bonds. *Journal of Banking and Finance* 3,134 *ff.*

———. 1980. Conditional Predictions of Bond Prices and Returns. *Journal of Finance* 35, 405 *ff.*

Choudhry, M. 2001. *Bond Market Securities.* London: FT Prentice Hall.

Campbell, J., A. Lo, and A. MacKinley. 1997. *The Econometrics of Financial Markets.* Princeton: Princeton University Press.

Fisher, L., and M. Leibowitz. 1983. Effects of Alternative Anticipations of Yield Curve Behaviour on the Composition of Immunized Portfolios and on Their Target Returns. In Kaufmann, G., eds. *Innovations in Bond Portfolio Management.* Greenwich, CT: JAI Press.

Neftci, S. 2000. A*n Introduction to the Mathematics of Financial Derivatives,* 2nd ed. Oxford: Academic Press.

"Out of Debt," *The Economist* February 12, 2000.

This page is intentionally blank

INDEX

ABOUT BLOOMBERG

Bloomberg L.P., founded in 1981, is a global information services, news, and media company. Headquartered in New York, the company has sales and news operations worldwide.

Bloomberg, serving customers on six continents, holds a unique position within the financial services industry by providing an unparalleled range of features in a single package known as the BLOOMBERG PROFESSIONAL® service. By addressing the demand for investment performance and efficiency through an exceptional combination of information, analytic, electronic trading, and Straight Through Processing tools, Bloomberg has built a worldwide customer base of corporations, issuers, financial intermediaries, and institutional investors.

BLOOMBERG NEWS®, founded in 1990, provides stories and columns on business, general news, politics, and sports to leading newspapers and magazines throughout the world. BLOOMBERG TELEVISION®, a 24-hour business and financial news network, is produced and distributed globally in seven different languages. BLOOMBERG RADIO^SM is an international radio network anchored by flagship station BLOOMBERG® 1130 (WBBR-AM) in New York.

In addition to the BLOOMBERG PRESS® line of books, Bloomberg publishes *BLOOMBERG MARKETS*® and *BLOOMBERG WEALTH MANAGER*® magazines. To learn more about Bloomberg, call a sales representative at:

London:	+44-20-7330-7500
New York:	+1-212-318-2000
Tokyo:	+81-3-3201-8900

FOR IN-DEPTH MARKET INFORMATION and news, visit the Bloomberg website at **www.bloomberg.com,** which draws from the news and power of the BLOOMBERG PROFESSIONAL® service and Bloomberg's host of media products to provide high-quality news and information in multiple languages on stocks, bonds, currencies, and commodities.

ABOUT THE AUTHOR

Moorad Choudhry is Head of Treasury at KBC Financial Products UK Limited. His responsibilities cover asset-liability management for the group's activities worldwide, including money-market funding, managing the asset-backed conduit program, and corporate bond repo trading.

He was previously a vice president in structured finance services with JPMorgan Chase Bank in London. Prior to that he worked as a gilt-edged market maker and treasury trader with ABN Amro Hoare Govett Sterling Bonds Limited, and as a sterling proprietary trader within the Treasury division at Hambros Bank Limited.

Dr. Choudhry is a visiting professor at the Department of Economics, London Metropolitan University, and a senior fellow at the Centre for Mathematical Trading and Finance, Cass Business School. He was educated at Claremont Fan Court School in Surrey, the University of Westminster, the University of Reading, Henley Management College, and Birkbeck, University of London.

His previous publications include *The Bond and Money Markets* (Butterworth Heinemann, 2001), *Capital Market Instruments* (FT Prentice Hall, 2001), *Analysing and Interpreting the Yield Curve* (John Wiley, 2004), *Structured Credit Products* (John Wiley, 2004), and numerous articles in academic and trade journals.

His research can be viewed at www.YieldCurve.com.