

III.1

Bonds and Swaps

III.1.1 INTRODUCTION

A financial *security* is a tradable legal claim on a firm's assets or income that is traded in an organized market, such as an exchange or a broker's market. There are two main classes of securities: primitive securities and derivative securities. A *primitive security* is a financial claim that has its own intrinsic price. In other words, the price of a primitive security is not a function of the prices of other primitive securities. A *derivative security* is a financial claim with a pay-off that is a function of the prices of one or more primitive securities.

This chapter focuses on *interest rate sensitive securities* that are traded in the debt markets, and on bonds and swaps in particular. We are not concerned here with the very short term debt markets, or *money markets* which trade in numerous interest rate sensitive instruments with maturities typically up to 1 year.¹ Our focus is on the market risk analysis of bonds and swaps, and at the time of issue most swaps have maturities of 2 years or more.

Virtually all bonds are primitive securities that are listed on exchanges but are traded by brokers in over-the-counter (OTC) markets. The exception is *private placements* which are like transferable loans. Since there is no secondary market, private placements are usually accounted for in the *banking book*, whereas most other interest rate sensitive securities are *marked to market* in the *trading book*. Forward rate agreements and swaps are derivative securities that are also traded OTC but they are not listed on an exchange.

By examining the relationships between bonds, forward rate agreements and swaps we explain how to value them, how to analyse their market risks and how to hedge these risks. Developed debt markets simultaneously trade numerous securities of the same maturity. For instance, for the same expiry date we may have trades in fixed and floating coupon bonds, forward rate agreements and swaps. Within each credit rating we use these instruments to derive a unique term structure for *market interest rates* called the *zero coupon yield curve*. We show how to construct such a curve from the prices of all liquid interest rate sensitive securities.

Bond futures, bond options and swaptions are covered in Chapters III.2 and III.3. However, in this chapter we do consider bonds with embedded options, also called *convertible bonds*. Convertible bonds are *hybrid securities* because they can be converted into the common stock of the company. These claims on the firm's assets share the security of bonds at the same time as enjoying exposure to gains in the stock price. The bond component makes them less risky than pure stock, but their value depends on other variables in addition to the stock price.

The outline of this chapter is as follows. In Section III.1.2 we introduce fundamental concepts for the analysis of bonds and associated interest rate sensitive securities. Here

¹ Money market instruments include Treasury bills and other discount bonds, interbank loans, sale and repurchase agreements, commercial paper and certificates of deposit. Exceptionally some of these instruments can go up to 270 days maturity.

we explain the difference between *discrete* and *continuous compounding* of interest and show how to translate between the two conventions. Then we introduce the terminology used for common types of market interest rates, describe the distinction between *spot* and *forward* interest rates and show how to translate between the spot curve and the forward curve.

Section III.1.3 begins with a brief catalogue of the different types of bonds that are commonly traded. We can distinguish bonds by the type of issuer, type of coupon (i.e. fixed or floating) and the bond maturity (i.e. the time until the claim expires) and we make the distinction between *fixed coupon bonds* and *floating rate notes*. Section III.1.4 examines the relationship between the price of a fixed coupon bond and its yield, and introduces the *zero coupon yield curve* for a given credit rating. We also examine the characteristics of the zero coupon spot yield curve and the *term structure of forward interest rates*.

Section III.1.5 examines the traditional measures of market risk on a single bond and on a bond portfolio. We introduce bond *duration* and *convexity* as the first and second order bond price sensitivities to changes in yield. These sensitivities allow us to apply Taylor expansion to approximate the change in bond price when its yield changes. Then we show how to approximate the change in value of a bond *portfolio* when the zero coupon curve shifts, using the *value duration* and the *value convexity* of the portfolio. Finally, we consider how to *immunize* a bond portfolio against movements in the yield curve.

Section III.1.6 focuses on bonds with semi-annual or quarterly coupons and floating rate notes. Section III.1.7 introduces *forward rate agreements* and *interest rate swaps* and explains the relationship between these and floating rate notes. We demonstrate that the market risk of a swap derives mainly from the fixed leg, which can be analysed as a bond with coupon equal to the swap rate. Examples explain how to fix the rate on a standard fixed-for-floating swap, and how the cash flows on a cross-currency basis swap are calculated. Several other types of swaps are also defined.

Section III.1.8 examines a bond portfolio's sensitivities to market interest rates, introducing the *present value of a basis point* (PV01) as the fundamental measure of sensitivity to changes in market interest rates.² We are careful to distinguish between the PV01 and the *dollar duration* of a bond portfolio: although the two measures of interest rate risk are usually very close in value, they are conceptually different.

Section III.1.9 describes how we *bootstrap* zero coupon rates from money market rates and prices of coupon bonds of different maturities. Then we explain how splines and parametric functions are used to fit the zero coupon yield curve. We present a case study that compares the application of different types of yield curve fitting models to the UK LIBOR curve and discuss the advantages and limitations of each model. Section III.1.10 explains the special features of *convertible bonds* and surveys the literature on convertible bond valuation models. Section III.1.11 summarizes and concludes.

III.1.2 INTEREST RATES

The future value of an investment depends upon how the interest is calculated. *Simple interest* is paid only on the principal amount invested, but when an investment pays *compound interest* it pays interest on both the principal and previous interest payments. There are two methods

² One basis point equals 0.01%.

of calculating interest, based on *discrete compounding* and *continuous compounding*. Discrete compounding means that interest payments are periodically accrued to the account, such as every 6 months or every month. Continuous compounding is a theoretical construct that assumes interest payments are continuously accrued, although this is impossible in practice. Both simple and compound interest calculations are possible with discrete compounding but with continuous compounding only compound interest rates apply.

Interest rates are divided into spot rates and forward rates. A *spot rate* of maturity T is an interest rate that applies from now, time 0, until time T . A *forward rate* is an interest rate starting at some time t in the future and applying until some time T , with $T > t > 0$. With forward interest rates we have to distinguish between the *term* (i.e. the time until the forward rate applies) and the *tenor* (i.e. the period over which it applies). The aim of this section is to introduce the reader to discretely compounded and continuously compounded spot and forward interest rates and establish the connections between them.

Since they are calculated differently, continuously compounded interest rates are different from discretely compounded interest rates and it is common to use different notation for these. In this text we shall use lower-case r for continuously compounded spot rates and capital R for discretely compounded spot rates. We use lower-case f for continuously compounded forward rates and upper-case F for discretely compounded forward rates. When we want to make the maturity of the rate explicit we use a subscript, so for instance R_T denotes the discretely compounded spot rate of maturity T and f_{nm} denotes the continuously compounded forward rate with *term* n and *tenor* m (i.e. starting at time n and ending at time $n + m$). For instance, the forward rate $f_{t,T-t}$ starts at time t and ends at time T and has tenor $T - t$.

The debt market convention is to quote rates in discretely compounded terms, with compounding on an annual or semi-annual basis. Since each market has its own *day count convention* the analysis of a portfolio of debt market instruments can be full of tedious technical details arising from different market conventions. For this reason banks convert market interest rates into continuously compounded interest rates because they greatly simplify the analysis of price and risk of debt market instruments.

III.1.2.1 Continuously Compounded Spot and Forward Rates

The *principal* N is the nominal amount invested. It is measured in terms of the local currency, e.g. dollars. The *maturity* T is the number of years of the investment. Note that T is measured in years so, for instance, $T = 0.5$ represents 6 months; in general T can be any finite, positive real number. Let r_T denote the continuously compounded T -maturity interest rate. Note that this is always quoted as an *annual* rate, for any T . Finally, let V denote the continuously compounded value of the investment at maturity.

It follows from Section I.1.4.5 and the properties of the exponential function (see Section I.1.2.4) that

$$V = N \exp(r_T T). \quad (\text{III.1.1})$$

Conversely, the *present value* of an amount V paid at some future time T is

$$N = V \exp(-r_T T). \quad (\text{III.1.2})$$

Formally, $\exp(r_T T)$ is called the *continuous compounding factor* and $\exp(-r_T T)$ is called the *continuous discount factor* for maturity T .

We now show how continuously compounded spot and forward rates are related. Let $r_1 (\equiv f_{0,1})$ denote the continuously compounded spot interest rate that applies for one period,

from time 0 until time 1. Let us say for simplicity that one period is 1 year, so the spot 1-year rate applies from now and over the next year. Then the 1-year forward interest rate $f_{1,1}$ is the 1-year rate that will apply 1 year from now. Denote by $r_2 (\equiv f_{0,2})$ the spot interest rate, quoted in annual terms, which applies for two periods (i.e. over the next 2 years in this example). The value of the investment should be the same whether we invest for 2 years at the current 2-year spot rate, or for 1 year at the 1-year spot rate and then roll over the investment at the 1-year spot rate prevailing 1 year from now. But the fair value of the 1-year spot rate prevailing 1 year from now is the 1-year forward interest rate. Hence, the compounding factors must satisfy $\exp(2r_2) = \exp(r_1) \exp(f_{1,1})$. In other words,

$$r_2 = \frac{r_1 + f_{1,1}}{2}. \quad (\text{III.1.3})$$

This argument has a natural extension to k -period rates. The general relationship between continuously compounded spot and forward compounding factors is that the k -period spot rate is the arithmetic average of the one-period spot rate and $k - 1$ one-period forward interest rates:

$$r_k \equiv f_{0,k} = \frac{f_{0,1} + f_{1,1} + \dots + f_{k-1,1}}{k}. \quad (\text{III.1.4})$$

III.1.2.2 Discretely Compounded Spot Rates

The discretely compounded analogue of equations (III.1.1) and (III.1.2) depends on whether simple or compound interest is used. Again let N denote the principal (i.e. the amount invested) and let T denote the number of years of the investment. But now denote by R_T the discretely compounded T -maturity interest rate, again quoted in annual terms. Under simple compounding of interest the future value of a principal N invested now over a period of T years is

$$V = N(1 + R_T T). \quad (\text{III.1.5})$$

However, under compound interest,

$$V = N(1 + R_T)^T. \quad (\text{III.1.6})$$

Similarly, with simple interest the present value of an amount V paid at some future time T is

$$N = V(1 + R_T T)^{-1}. \quad (\text{III.1.7})$$

but with compound interest it is

$$N = V(1 + R_T)^{-T}. \quad (\text{III.1.8})$$

Simple interest is usually only applied to a single payment over a fixed period of less than 1 year. For a period T of less than 1 year the *discretely compounded discount factor* is

$$\delta_T = (1 + R_T T)^{-1}, \quad (\text{III.1.9})$$

and when T is an integral number of years it is

$$\delta_T = (1 + R_T)^{-T}. \quad (\text{III.1.10})$$

When T is greater than 1 year but not an integral number of years the discretely compounded discount factor for maturity T is

$$\delta_T = (1 + R_T)^{-[T]} (1 + (T - [T]) R_T)^{-1}, \quad (\text{III.1.11})$$

where $[T]$ denotes the integer part of T .

Remarks on Notation

1. In this chapter and throughout this book we shall be discounting a sequence of cash flows at a set of future dates. Usually at least the first date occurs within 1 year, and frequently several dates are less than 1 year. The cash flow payments almost never occur in exactly integer years. Hence, the formula (III.1.11) should be used for the discount factor on these payment dates. However, to specify this will make the formulae look more complex than they really are. Hence, we shall assume that cash flows occur annually, or semi-annually, to avoid this burdensome notation. This is not without loss of generality, but all the concepts that we focus on can be illustrated under this assumption.
2. We remark that some authors use the notation B_T instead of δ_T for the discretely compounded discount factor of maturity T , because this also is the price of a pure *discount bond* of maturity T and redemption value 1. We shall see in the next section that a discount bond, which is also called a *zero coupon bond* or a *bullet bond*, is one of the basic building blocks for the market risk analysis of cash flows.
3. When market interest rates are quoted, the rate is specified along with type of rate (normally annual or semi-annual) and the frequency of payments. In money markets discretely compounded interest rates are quoted in annual terms with 365 days. In bond markets they are quoted in either annual or semi-annual terms with either 360 or 365 days per year. To avoid too many technical details in this text, when we consider discretely compounded interest rates we shall assume that all interest rates are quoted in annual terms and with 365 days per year, unless otherwise stated.

The frequency of payments can be annual, semi-annual, quarterly or even monthly. If the annual rate quoted is denoted by R and the interest payments are made n times each year, then

$$\text{Annual compounding factor} = \left(1 + \left(\frac{R}{n}\right)\right)^n. \quad (\text{III.1.12})$$

For instance, $(1 + \frac{1}{2}R)^2$ is the *annual* compounding factor when interest payments are semi-annual and R is the 1 year interest rate. In general, if a principal amount N is invested at a discretely compounded annual interest rate R , which has n compounding periods per year, then its value after m compounding periods is

$$V = N \left(1 + \frac{R}{n}\right)^m. \quad (\text{III.1.13})$$

It is worth noting that:

$$\lim_{n \rightarrow \infty} \left(1 + \frac{R}{n}\right)^n = \exp(R), \quad (\text{III.1.14})$$

and this is why the continuously compounded interest rate takes an exponential form.

EXAMPLE III.1.1: CONTINUOUS VERSUS DISCRETE COMPOUNDING

Find the value of \$500 in 3.5 years' time if it earns interest of 4% per annum and interest is compounded semi-annually. How does this compare with the continuously compounded value?

SOLUTION

$$V = 500 \times \left(1 + \frac{0.04}{2}\right)^7 = \$574.34,$$

but under continuous compounding the value will be greater:

$$V = 500 \times \exp(0.04 \times 3.5) = \$575.14.$$

III.1.2.3 Translation between Discrete Rates and Continuous Rates

The discrete compounding discount factor (III.1.11) is difficult to work with in practice since different markets often have different day count conventions; many practitioners therefore convert all rates to continuously compounded rates and do their analysis with these instead. There is a straightforward translation between discrete and continuously compounded interest rates.

Let R_T and r_T denote the discretely and continuously compounded rates of maturity T . R_T and r_T are equivalent if they provide the same return. Suppose there are n compounding periods per annum on the discretely compounded rate. Then equating returns gives

$$\exp(r_T T) = \left(1 + \frac{R_T}{n}\right)^{nT},$$

i.e. $\exp(r_T) = (1 + R_T/n)^n$, and taking logarithms of both sides gives the continuously compounded rate that is equivalent to R_T as

$$r_T = n \ln \left(1 + \frac{R_T}{n}\right). \quad (\text{III.1.15})$$

For instance, the continuously compounded equivalent of a semi-annual rate of 5% is $2 \ln(1.025) = 4.9385\%$.

Similarly, the discretely compounded rate that is equivalent to r_T is

$$R_T = n \left(\exp\left(\frac{r_T}{n}\right) - 1 \right). \quad (\text{III.1.16})$$

III.1.2.4 Spot and Forward Rates with Discrete Compounding

Recall that we use F_{nm} to denote an m -period discretely compounded forward interest rate starting n periods ahead, that is, n is the term and m is the tenor of the rate. For instance, if the period is measured in months as in the example below, then $F_{6,9}$ denotes the 9-month forward rate 6-months ahead, i.e. the interest rate that applies between 6 months from now and 15 months from now.

The argument leading to the relationship (III.1.3) between continuously compounded forward and spot rates also extends to discretely compounded forward and spot rates. Measuring periods in 1 year, we have

$$(1 + R_2)^2 = (1 + R_1)(1 + F_{1,1}) \quad (\text{III.1.17})$$

and in general

$$(1 + R_k)^k = (1 + F_{0,1})(1 + F_{1,1}) \dots (1 + F_{k-1,1}). \quad (\text{III.1.18})$$

The relationship (III.1.18) allows forward rates of various terms and tenors to be calculated from the spot rate curve, as the following examples show.

EXAMPLE III.1.2: CALCULATING FORWARD RATES (1)

Find the 1-year and 2-year forward rates, given that the discretely compounded annual spot interest rates are 5%, 6% and 6.5% for maturities 1 year, 2 year and 3 years, respectively.

SOLUTION

$$1\text{-year forward rate, } F_{1,1} = \frac{1.06^2}{1.05} - 1 = 7.0\%$$

$$2\text{-year forward rate, } F_{2,1} = \frac{1.065^3}{1.06^2} - 1 = 7.5\%$$

A variety of forward rates can be computed from any given spot rate curve, as shown by the following example.

EXAMPLE III.1.3: CALCULATING FORWARD RATES (2)

Given the spot rates in Table III.1.1, calculate a set of discretely compounded 3-month forward rates for 3 months, 6 months and 9 months ahead. Also calculate $F_{6,6}$, the 6-month forward rate that applies 6 months ahead.

SOLUTION

Table III.1.1 Discretely compounded spot and forward rates

Maturity (months)	Spot rates	Forward rates	
3	4.5%	$F_{0,3}$	4.5%
6	4.3%	$F_{3,3}$	4.05%
9	4.2%	$F_{6,3}$	3.92%
12	4.0%	$F_{9,3}$	3.30%

The 3-month forward rates are shown in Table III.1.1. Note that in the table, the forward rate matures at the time specified by the row and starts 3 months earlier. The forward rate from 0 months to 3 months is the same as the spot rate, and this is entered in the first row. Denote the spot and forward rates as above. Since the rates are quoted in annual terms, we have:

$$\left(1 + \frac{F_{0,3}}{4}\right) \left(1 + \frac{F_{3,3}}{4}\right) = \left(1 + \frac{F_{0,6}}{2}\right).$$

Thus the 3 month forward rate at 3 months (i.e. effective from 3 months to 6 months) is given by:

$$\begin{aligned} F_{3,3} &= 4 \times \left[\left(1 + \frac{F_{0,6}}{2}\right) \left(1 + \frac{F_{0,3}}{4}\right)^{-1} - 1 \right] \\ &= 4 \times \left[\left(1 + \frac{0.043}{2}\right) \left(1 + \frac{0.045}{4}\right)^{-1} - 1 \right] = 0.0405 = 4.05\%. \end{aligned}$$

This is the forward rate in the second row. Similarly, the 3-month forward rate at 6 months is given by the relationship

$$\left(1 + \frac{F_{0,6}}{2}\right) \left(1 + \frac{F_{6,3}}{4}\right) = \left(1 + \frac{3F_{0,9}}{4}\right).$$

Substituting in for the spot rates $F_{0,6}$ and $F_{0,9}$ and solving for the forward rate gives $F_{6,3} = 3.92\%$. Similarly,

$$\left(1 + \frac{3F_{0,9}}{4}\right) \left(1 + \frac{F_{9,3}}{4}\right) = (1 + F_{0,12}),$$

and this gives $F_{9,3} = 3.30\%$. And finally, $F_{6,6} = 3.62\%$ is found by solving

$$\left(1 + \frac{F_{0,6}}{2}\right) \left(1 + \frac{F_{6,6}}{2}\right) = (1 + F_{0,12}).$$

III.1.2.5 LIBOR

LIBOR stands for ‘London Inter Bank Offered Rate’. It is the interest rate charged for short term lending between major banks. It is almost always greater than the *base rate*, which is the rate that banks obtain when they deposit their reserves with the central bank. LIBOR is published each day (between 11.00 and 11.30) by the British Bankers’ Association (BBA). The BBA takes a survey of the lending rates set by a sample of major banks in *all* major currencies. Thus the LIBOR rates refer to funds not just in sterling but also in US dollars and all the major currencies. The terms for the lending rates are usually overnight, 1 week, 2 weeks, 1 month, and then monthly up to 1 year. The shorter rates up to 3 or 6 months are very precise at the time they are measured, but the actual rates at which banks will lend to one another continue to vary throughout the day.

London is the main financial centre in the world and so LIBOR has become the standard *reference rate* for discounting short term future cash flows to present value terms. It is used to find the fair (or theoretical) value of forwards and futures, forward rate agreements and swaps, bonds, options, loans and mortgages in all major currencies except the euro. For the euro the usual reference rates are the EURIBOR rates compiled by the European Banking Federation. Also emerging economies such as India are now developing their own reference rates, such as MIBOR in Mumbai.

III.1.3 CATEGORIZATION OF BONDS

The *principal* or *face value* of the bond is the amount to be repaid to the bond holder at maturity (if the bond is issued at a discount or premium, the principal is not the amount invested). The *coupon* on a bond determines the periodic payment to the bond holder by the issuer until the bond expires. The amount paid is equal to the coupon multiplied by the face value. When we price a bond we assume that the face value is 100.³ If the price is 100 we say the bond is priced *at par*; bonds are *below par* if their price is below 100 and *above par* if their price is above 100. On the *expiry date* of the bond, the bond holder *redeems* the bond with the issuer. The *redemption value* of the bond is often but not always the same as the face value.⁴ The *issuer* of a bond raises funds by selling bonds in the *primary market*.⁵ This is called *debt financing*. Between issue and expiry, bonds may be traded in the *secondary market* where trading is usually OTC.⁶

The bond market is a very old market that has evolved separately in many different countries. Each country has different conventions, and even within a country there are numerous variations on bond specifications. In the following we provide only a very general summary of the different types of straight bonds that are commonly traded, categorizing

³ Depending on the currency unit, bonds usually have a minimum amount invested much greater than 100. For instance, the minimum amount one can buy of Fannie Mae bonds is \$1,000. Nevertheless their market prices are always quoted as a percentage of the face value, i.e. based on the face value of 100.

⁴ For instance, there could be a bonus payment at maturity, in addition to payment of the face value.

⁵ Usually bonds are underwritten by banks and placed with their investors.

⁶ But some government bonds are exchange traded.

them by issuer and the type of coupon paid. Readers seeking more specific details are recommended to consult the excellent books by Fabozzi (2005) or Choudhry (2005, 2006).

III.1.3.1 Categorization by Issuer

Bonds are divided into different categories according to the priority of their claim on the issuer's assets, if the issuer becomes insolvent. For instance, *senior* bond holders will be paid before *junior* bond holders and, if the bond is issued by a company that has also sold shares on its equity, all bond holders are paid before the equity holders. This gives bonds a security not shared by stocks and so investing in bonds is less risky than investing in equities. This is one of the reasons why bond prices are typically less volatile than share prices.

Almost all bonds have a *credit rating* which corresponds to the perceived probability that the issuer will default on its debt repayments. The best credit rating, corresponding to an extremely small probability of default, is labelled AAA by Standard & Poor's and Aaa by Moody's. Only a few bonds, such as G10 government bonds and some corporate bonds, have this credit rating. Standard & Poor's credit ratings range from AAA to C, and the C rating corresponds to a 1-year default probability of 20% or more. Any bond rated below *investment grade* (BBB) is called a *junk bond*.

Table III.1.2 sets out the common types of bond issuers ordered by the size of the markets in these bonds, stating the terminology used when we refer to these bonds and listing some specific examples of each type of bond.

Table III.1.2 Examples of bonds

Issuer	Examples					
	Issuer	Currency	Maturity	Coupon	Yield ^a	Rating
Financial institutions (banks, insurance companies and financial companies)	HBOS	GBP	Apr. 2008	5.50%	6.23%	AA
	Goldman Sachs	USD	Nov 2014	4.25%	6.01%	AA–
	Deutsche Fin.	EUR	Jul. 2009	1.75%	4.33%	AA
	KFW Int. Fin.	JPY	Mar 2010	6.38%	1.00%	AAA
Governments (called sovereign, Treasury or government bonds)	Australia	AUS	Feb. 2017	6.00%	5.96%	AAA
	Greece	EUR	Jul 2017	4.30%	4.58%	A
	Brazil	GBP	Mar 2015	7.88%	5.96%	BB+
Government agencies (agency bonds ^b)	Fannie Mae	USD	Jun 2017	5.375%	5.53%	AAA
	Freddie Mac	USD	Nov 2017	5.125%	5.14%	AAA
Towns/regions (municipal bonds)	Dutchess County	USD	Sep 2050	5.35%	3.53%	AA
	NYSU Dev. Corp.	USD	Jan 2025	5.00%	3.13%	AA–
Corporations (corporate bonds)	General Motors	USD	Nov 2031	8.00%	7.46%	BB+
	Tokyo Motor	JPY	June 2008	0.75%	0.90%	AAA
	Boots	GBP	May 2009	5.50%	8.65%	BBB

^aYields quoted as of 25 August 2007, reopening auctions.

^bThe main suppliers of agency bonds are the US housing associations that issue bonds backed by mortgages. See www.fanniemae.com and www.freddie.com for further details.

Bonds can be further categorized as follows:

- *Domestic bonds* are issued in the country of domicile of the issuer. In other words, they are issued by a domestic borrower in their local market and they are denominated in the local currency.

- *Foreign bonds* are issued by a foreign borrower in the local currency. For instance, a US issuer can issue a sterling denominated bond in the UK market (this is called a *Yankee*).
- *Eurobonds* are *bearer bonds* that are issued on the international market in any currency.⁷ This is a broker's market with a self-regulatory body called the *International Capital Markets Association* (ICMA).⁸

III.1.3.2 Categorization by Coupon and Maturity

The coupon payments on a bond are usually made at regular intervals. Often they are made once every year, in which case we say the bond has *annual coupons*, or every 6 months for a *semi-annual coupon* bond. If the coupons are semi-annual then half of the coupon payment is made every 6 months.⁹

The majority of bonds pay a *fixed coupon*: for instance, a 5% annual 2015 bond pays the bond holder 5% of the face value every year until 2015. However, some bonds, called *floating rate notes* or *floaters* for short, have a variable coupon. The periodic payment is based on a reference rate such as the LIBOR rate prevailing at the time that the coupon is paid.¹⁰ For instance, a semi-annual floater pays 6-month LIBOR plus spread, and a quarterly floater pays 3-month LIBOR plus spread.

Bonds are further divided into *short term bonds* (2–5 years), *medium term bonds* (5–15 years) and *long term bonds* (over 15 years) where the numbers in parentheses here refer to the time between issue and expiry.¹¹ Interest rate sensitive securities that have less than 2 years to expiry at issue are generally regarded as money market instruments. Most money market instruments such as *US Treasury bills* pay no coupons. They are just priced at a discount to their par value of 100. If they are held until expiry the profit is guaranteed. For instance, if a US Treasury bill is purchased at 95 and held until expiry the return is $5/95 = 5.26\%$.¹² Even though very few bonds pay no coupon it is convenient to work with hypothetical zero coupon bonds of different maturities and different credit ratings.¹³ In particular, for reasons that will become evident in Section III.1.4.3, the yields on zero coupon bonds are used as *market interest rates* for different credit ratings.

III.1.4 CHARACTERISTICS OF BONDS AND INTEREST RATES

In the following we begin by using market interest rates to price a bond. We do not question how these interest rates have been derived. Then we explore the relationship between the price of a bond and the yield on a bond and show that two bonds with the same maturity but different coupons have different yields. However, all zero coupon bonds of the same maturity have the same yield. This *zero coupon yield* is the *market interest rate* of that maturity. But

⁷ Bearer bonds are like banknotes – they belong to the person who holds the bond in his hand. But the owners of other types of bonds must be registered.

⁸ The ICMA Centre also hosts the MTS bond database. See <http://www.mtsgroup.org/newcontent/timeseries/> and http://www.icmacentre.ac.uk/research_and_consultancy_services/mts_time_series for further details.

⁹ Some bonds also have quarterly payments, but these are usually floating rate notes.

¹⁰ LIBOR is usually greater than the *base rate*, which is the remuneration rate when banks make deposits with the central bank. See Section III.1.2.5 for further details.

¹¹ This varies according to the country: for example, the longest Italian bonds are around 10 years to expiry.

¹² Of course, the same applies to any fixed income security: if it is held to expiry the return is locked in. This does not, however, mean that fixed income securities have no interest rate risk! See Section III.1.8.3 for further details.

¹³ We also call zero coupon bonds *discount bonds* for obvious reasons.

we are in a typical ‘chicken and egg’ situation. We could use the market interest rates to price bonds, and then derive the market interest rates from bond prices. So which come first, the bond prices or the market interest rates?

The prices of bonds in the secondary market are determined by supply and demand. They are *not* priced using a formula. Just because we start this section with a formula that relates the fair price of a bond to market interest rates, this does *not* imply that such a formula is actually used by brokers to price bonds. Of course, brokers might use this formula if the market is not very liquid, but in general they set their prices by supply and demand.¹⁴

In the following we shall characterize a bond of maturity T_n by a sequence of non-negative cash flows, starting with a sequence of regularly spaced coupon payments $C_{T_1}, \dots, C_{T_{n-1}}$ and finishing with a cash flow C_{T_n} which consists of the redemption value plus the coupon paid at time T_n . We shall price the bond assuming that the coupon payments and redemption value are both based on a face value of 100. In other words, a coupon of 5% paid at time t means that $C_t = 5$ and the redemption value is 100. However, coupon payments are, in reality, based on the principal amount invested and the redemption value could be different from the principal, for instance if bonus payments are included.

III.1.4.1 Present Value, Price and Yield

Denote the sequence of cash flows on the bond by $\{C_{T_1}, \dots, C_{T_n}\}$. Since we assume these are based on the face value of 100 the cash flows on a fixed annual coupon bond of maturity T_n years paying a coupon of c percent are:

$$C_{T_i} = \begin{cases} 100c, & \text{for } i = 1, \dots, n-1, \\ 100(1+c), & \text{for } i = n. \end{cases} \quad (\text{III.1.19})$$

Suppose we have a set of discretely compounded spot market interest rates, $\{R_{T_1}, \dots, R_{T_n}\}$, where R_{T_i} is the rate of maturity T_i for $i = 1, \dots, n$. That is, the maturity of each spot rate matches exactly the maturity of each cash flow. Then we can calculate the present value of a cash payment at time T_i as the payment discounted by the discretely compounded T_i -maturity interest rate.

The *present value* of a bond is the sum of the discounted future cash flows. If based on a face value of 100 it is

$$PV = \sum_{i=1}^n C_{T_i} (1 + R_{T_i})^{-T_i}. \quad (\text{III.1.20})$$

More generally the present value is

$$PV = \frac{N}{100} \sum_{i=1}^n C_{T_i} (1 + R_{T_i})^{-T_i}, \quad (\text{III.1.21})$$

where N is the principal amount invested.

The difference between the present value (III.1.20) and the *market price*, P of a bond is that the present value is a fair or theoretical value, whereas the market price of the bond is set by supply and demand. The market price of a bond is always quoted relative to a face value of 100.

The market interest rates on the right-hand side of (III.1.20) will not normally be the same. There is likely to be a *term structure* of interest rates, i.e. where short rates are not the same as long rates (see Figure III.1.6 below, for instance). Nevertheless, provided that all the cash

¹⁴ A similar comment applies to any liquid market. For instance, the prices of standard European calls and puts are *not* set by the Black–Scholes–Merton formula. We use this formula, yes, but not to price options unless they are illiquid. And even then, we usually adjust the formula (see Section III.3.6.7).

flows are positive, we can find a *single* discount rate that, when applied to discount every cash payment, gives the market price of the bond. In other words, we can find y such that

$$P = \sum_{i=1}^n (1+y)^{-T_i} C_{T_i}. \quad (\text{III.1.22})$$

This y is called the discretely compounded *yield*, or *yield to redemption*, on the bond. It is a single number that represents the *internal rate of return* on the bond.¹⁵ In other words, it is the return that an investor would get if he held the bond until expiry.

Similarly, given a set of continuously compounded interest rates $\{r_{T_1}, \dots, r_{T_n}\}$, the continuously compounded yield is found by solving

$$\sum_{i=1}^n \exp(-r_{T_i} T_i) C_{T_i} = \sum_{i=1}^n \exp(-y T_i) C_{T_i} \quad (\text{III.1.23})$$

assuming the bond is fairly priced.

EXAMPLE III.1.4: CALCULATING THE PRESENT VALUE AND YIELD OF FIXED COUPON BONDS

Consider two bonds that pay coupons annually. Bond 1 has coupon 5% and maturity 3 years and bond 2 has coupon 10% and maturity 5 years. Suppose the market interest rates are as shown in Table III.1.3. Find the present value and, assuming this present value is also the market price, find the yield on each bond.

Table III.1.3 Market interest rates

Maturity (years)	1	2	3	4	5
Interest rate	4.0%	4.25%	4.5%	4.25%	4.20%

SOLUTION Table III.1.4 sets out the cash flow for each bond, which is then discounted by the relevant interest rate to give the present values in column 4 (for bond 1) and column 6 (for bond 2). The sum of these discounted cash flows gives the price of the bond, i.e. 101.42 for bond 1 and 125.59 for bond 2. Bond prices reflect the coupon rate. For instance, if bond 2 had a coupon of only 5% then its price would be only 103.5, as readers can verify using the spreadsheet for this example.

Table III.1.4 Bond prices

Maturity	Interest rate	Cash flow	Present value	Cash flow	Present value
1	4.0%	5	4.81	10	9.62
2	4.25%	5	4.60	10	9.20
3	4.5%	105	92.01	10	8.76
4	4.25%	0	0	10	8.47
5	4.20%	0	0	110	89.55
Price bond 1			101.42	Price bond 2	125.59

¹⁵ Loans normally have all positive cash flows, i.e. the interest payments on the loan, but interest rate swaps do not. Hence, we cannot define the yield on a swap, but we can define the yield on a loan.

Although bond 2 has a much higher price than bond 1, this does not mean that investing in bond 2 gives a lower, or higher, return than investing in bond 1. The rate of return on a bond is given by its yield, and this is mainly determined by the interest rate at the bond maturity because the redemption payment has the most effect on the yield.

Table III.1.5 Bond yields

Maturity	Yield	Cash flow	Present value	Yield	Cash flow	Present value
1	4.48%	5	4.79	4.22%	10	9.60
2	4.48%	5	4.58	4.22%	10	9.21
3	4.48%	105	92.05	4.22%	10	8.83
4	4.48%	0	0	4.22%	10	8.48
5	4.48%	0	0	4.22%	110	89.48
Price bond 1			101.42	Price bond 2 125.59		

Table III.1.5 takes the bond price as given in Table III.1.4 and uses an iterative method to back out the value of the yield so that the sum of the cash payments discounted by this single yield is equal to the bond price.¹⁶ We find that the yield on bond 1 is 4.48% and the yield on bond 2 is 4.22%. This reinforces our comment above, that the coupon rate has little effect on the bond yield. For instance, if bond 2 had a coupon of only 5% then its yield would still be close to 4.2%, because this is the 5-year spot rate. In fact, it would be 4.21%, as readers can verify using the spreadsheet.

III.1.4.2 Relationship between Price and Yield

The price of bond 1 in Example III.1.4 is 101.42 and its yield is 4.48%. What would the price be if its yield changed by 1%? If the yield increased to 5.48% the price would decrease by 2.73 to 98.69; and if the yield decreased to 3.48% its price would increase by 2.83 to 104.25. The price increases when the yield decreases, and conversely, so the price and yield have a negative relationship. And since the price does not change by the same absolute amount when the bond yield increases and decreases, the price and yield have a *non-linear* negative relationship. The *price–yield curve* is the set of all points (y, P) determined by the price–yield relationship (III.1.22). In other words, it is a convex decreasing curve which gives the price of a fixed coupon bond as a function of its yield.

EXAMPLE III.1.5: THE EFFECT OF COUPON AND MATURITY ON THE PRICE–YIELD CURVE

Plot the price–yield curve for the following annual coupon bonds:

- (a) coupon 5%, maturity 3 years;
- (b) coupon 10%, maturity 3 years;
- (c) coupon 10%, maturity 5 years.

SOLUTION Figure III.1.1 illustrates the price–yield curve, i.e. the relationship (III.1.22) between the price and the yield, for each bond. It verifies the following:¹⁷

¹⁶ In the Excel spreadsheet for this example we have done this using the RATE function in Excel. However the bond yield could also be found using the Excel Solver.

¹⁷ These properties of the price–yield relationship may also be proved algebraically from (III.1.22).

- *An increase in coupon increases the price at each yield.* For instance, fixing the yield at 6.5%, the 3 year 5% coupon bond has price 96.03 but the 3 year 10% coupon bond has price 109.27.
- *When the yield equals the coupon the bond price is 100.* And we say it is priced at *par*. For instance, if the yield is 10% then both of the 10% coupon bonds have price 100.
- *An increase in maturity increases the steepness and convexity of the price–yield relationship.* And, comparing the curves for the 10% coupon bonds (b) and (c), the price at a given yield increases for yields below the coupon and decreases for yields greater than the coupon.




Figure III.1.1 Price–yield curves

III.1.4.3 Yield Curves

Now we consider the ‘chicken and egg’ situation that we referred to at the beginning of this section. We can replicate fair bond market prices using market interest rates, but the market interest rate of some given maturity is derived from a bond yield of the same maturity, which in turn is derived from the bond price. But two bonds with the same maturity but different coupons have different yields. So which bond yield should we use to derive the market interest rate? The next example illustrates how to break this circular construction.

EXAMPLE III.1.6: COMPARISON OF YIELD CURVES FOR DIFFERENT BONDS

Consider three sets of bonds. Each set contains 20 bonds with the same coupon but different maturities, and the maturities are 1, 2, . . . , 20 years. The coupons are 0% for the first set, 5% for the second set and 10% for the third set. The market interest rates are given in Table III.1.6.

- Draw a graph of the fair prices of each set of bonds, against the maturity.
- Draw a graph of the yields of each set of bonds, against the maturity.

Indicate the market interest rates on each graph.

Table III.1.6 Some market interest rates

Maturity (years)	1	2	3	4	5	6	7	8	9	10
Market interest rate	6.00%	6.50%	7.00%	7.20%	7.20%	6.60%	6.00%	5.60%	5.40%	5.25%
Maturity (years)	11	12	13	14	15	16	17	18	19	20
Market interest rate	5.25%	5.20%	5.20%	5.00%	5.10%	5.00%	4.90%	4.90%	4.80%	4.80%

SOLUTION The spreadsheet for this example calculates the prices using (III.1.20) and then, given each bond price it uses the Excel RATE function to back out the yield from the formula (III.1.22). The resulting graphs (a) and (b) are shown in Figures III.1.2 and III.1.3. The market interest rates from Table III.1.6 are added to each graph and in Figure III.1.2 we use the right-hand scale for these.




Figure III.1.2 Bond price versus maturity

Figure III.1.2 shows that the bonds with coupon above the market interest rate are priced above par, and those with coupon below the market interest rate are priced below par. We already know this from the preceding subsection. But Figure III.1.2 gives us some additional information: the price of a bond increases with maturity if the coupon is above the market interest rates and decreases with maturity if the coupon is below the market interest rates.

A *yield curve* is a graph of bond yield versus maturity, keeping the coupon fixed. Bonds with different coupons have different yield curves, and this is illustrated by Figure III.1.3. Here we depict three *different* yield curves, one for each set of bonds, i.e. the bonds with coupons 0%, 5% and 10% and varying maturities. We also show the market interest rates, but these are identical to the *zero coupon yield curve*.

It is this set of zero coupon yields that we use as the basic set of interest rates for pricing interest rate sensitive instruments. Whenever we refer to ‘the’ yield curve, as in *the BBB yield curve*, we mean the zero coupon yield curve. The zero coupon yield curve




Figure III.1.3 Bond yield versus maturity

is a term structure, i.e. a set of yields of different maturities on zero coupon bonds of the same credit rating category (AAA, BBB, etc). For each credit rating we have a different yield curve (i.e. a different set of market interest rates). The other two yield curves shown in Figure III.1.3, which are derived from market prices of coupon bearing bonds, are called *par yield curves*. There is one par yield curve for each coupon and for each credit rating.

Now let us return to the ‘chicken and egg’ problem. If we could observe the market price of a zero coupon bond, then we could use this market price to derive the market interest rate of maturity equal to the bond maturity. The price of a zero coupon bond of maturity T is just 100 times the T period discount factor:

$$\delta_T = (1 + y_T)^{-T}, \quad (\text{III.1.24})$$

where y_T is the discretely compounded zero coupon yield with maturity T years. So if we observe the market price P_T of a zero coupon bond with maturity T we can back out the market interest rate as y_T from the formula

$$P_T = 100 (1 + y_T)^{-T}. \quad (\text{III.1.25})$$

That is, we simply set:

$$y_T = \left(\frac{100}{P_T} \right)^{1/T} - 1. \quad (\text{III.1.26})$$

But we cannot observe the price of a zero coupon bond in the market for every T . Although there are many discount (i.e. zero coupon) instruments that are traded in the money markets, these rarely have maturity beyond 1 year, or 18 months at the most. Hardly any zero coupon bonds are traded with maturities longer than this. Hence, to obtain a complete set of market interest rates covering all maturities we must use a model for fitting a zero coupon yield curve using the prices of *coupon* bearing instruments. These *yield curve fitting models* are described in Section III.1.9.

III.1.4.4 Behaviour of Market Interest Rates

We examine the time series properties of a single interest rate, and for this we choose the US Treasury 3-month spot rate. Figure III.1.4 shows a daily time series of the 3-month US Treasury spot rate from January 1961 to December 2006.¹⁸ After reaching a high of over 14% in the early 1980s, US dollar interest rates declined sharply. In 2003 the 3-month rate reached an historic low, dipping below 2% for the first time since the 1950s. After this time rates rose steadily to above 4% by mid 2006.




Figure III.1.4 US Treasury 3-month spot rate, 1961–2006

In continuous time we model the short rate as a *mean-reverting* stochastic process. But if we use a very long term average rate then the rate of mean reversion is extremely slow, as is evident just by looking at the time series in Figure III.1.4. To confirm this view, we estimate a general mean-reverting stochastic process of the form¹⁹

$$dr(t) = \varphi(\theta - r(t))dt + \sigma r(t)^{\gamma} dW, \quad (\text{III.1.27})$$

where r is the US Treasury 3-month rate. The parameter φ is called the *rate of mean reversion* and the parameter θ is the long term average value of r . Of course what is ‘long term’ depends on the data used to estimate the parameters. If we use the entire sample of interest rates shown in Figure III.1.4 the long term value will be close to 6.53%, since this is the average rate over the entire sample. And the speed of reversion to this long term average, if indeed there *is* a mean reversion to this average, will be very slow indeed. There might be a 40-year cycle in Figure III.1.4, so we are going to test for this below. Interest rates increased from 4% in 1961 to almost 16% in 1981, then returned to 4% in 2001. Estimating the simple one-factor model (III.1.27) will yield parameters that reflect this single, long cycle if it exists.

¹⁸ Available from <http://www.federalreserve.gov/releases/h15/data.htm>.

¹⁹ See Sections I.3.7 and II.5.3.7 for more details about mean-reverting stochastic processes and see Section III.3.8.4 for more details about their application to short term interest rate models.

The estimated coefficients and standard errors, based on the entire sample between 1961 and 2006, are shown in Table III.1.7.²⁰ Using the parametric form

$$dr(t) = (\alpha + \beta r(t)) dt + \sigma r(t)^\gamma dW(t), \quad (\text{III.1.28})$$

we have:

- long term interest rate, $\theta = -\alpha\beta^{-1}$;
- mean reversion rate, $\varphi = -\beta$;
- characteristic time to mean-revert, $\varphi^{-1} = -\beta^{-1}$.

Table III.1.7 Estimates and standard errors of one-factor interest rate model

	Coefficient	Est. s.e.
α	0.00262	0.0024
β	-0.00039	0.0006
γ	1.06579	1.8799
σ	0.00864	0.2141
Long term interest rate	6.76%	
Mean reversion rate	0.00039	
Characteristic time to mean-revert	10.30 years	
Volatility	13.66%	

Note that the estimate of γ is very close to 1, with a large standard error, so we may as well set this to 1 and assume that interest rates follow a *geometric* process rather than an arithmetic process. The interest rate volatility is assumed constant, and it is estimated as 13.66%. As expected, the long term interest rate is 6.76%, which is close to the sample average. And it takes a very long time to mean-revert. The characteristic time to mean-revert is $(0.00039)^{-1} = 2754$ days. But note that the estimated standard error (s.e.) on β is 0.0006 so we cannot reject the null hypothesis that $\beta = 0$, i.e. that the series is integrated.²¹ In short, our estimation of the general model (III.1.27) using the entire sample leads to the conclusion that the 3-month US Treasury rate follows a standard geometric Brownian motion and there is no mean reversion in these interest rates!

We can try to detect other cycles in interest rates at different frequencies using a two-factor model of the form

$$\begin{aligned} dr(t) &= \varphi_1(\theta(t) - r(t))dt + \sigma_1 r(t)^{\gamma_1} dW_1(t), \\ d\theta(t) &= \varphi_2(\omega - \theta(t))dt + \sigma_2 \theta(t)^{\gamma_2} dW_2(t). \end{aligned} \quad (\text{III.1.29})$$

In this model the level to which interest rates mean-revert is not assumed to be constant. Instead it is a mean-reverting process itself. This way we can try to detect a shorter cycle in interest rates if it exists. Again we estimate the models in a form that allows us to test for the presence of mean reversion. That is, we could estimate

$$\begin{aligned} dr(t) &= (\alpha(t) + \beta_1 r(t))dt + \sigma_1 r(t)^{\gamma_1} dW_1(t), \\ d\alpha(t) &= (\lambda + \beta_2 \alpha(t))dt + \sigma_2 \alpha(t)^{\gamma_2} dW_2(t), \end{aligned} \quad (\text{III.1.30})$$

²⁰ Many thanks to my PhD student Stamatis Leontsinis for providing these estimates, based on a generalized methods of moments algorithm.

²¹ A careful discussion on this point is given in Section II.5.3.7.

and set

- mean reversion level of interest rate, $\theta(t) = -\alpha(t)\beta_1^{-1}$;
- characteristic time to mean-revert, $\varphi_1^{-1} = -\beta_1^{-1}$;
- long term level of interest rate, $\omega = -\lambda\beta_2^{-1}$;
- characteristic time to mean-revert, $\varphi_2^{-1} = -\beta_2^{-1}$.

III.1.4.5 Characteristics of Spot and Forward Term Structures

We also examine the properties of a cross-section of interest rates of different maturities (i.e. a *term structure* of interest rates) at a specific point in time, and for this we choose the term structures of UK spot rates and the associated term structure of UK forward rates. Data are provided by the Bank of England.²²

Figure III.1.5 shows how the UK spot rate curve evolved between 4 January 2000 and 31 December 2007. On 4 January 2000 the 1-year rate was 6.32% and the 25-year rate was 4.28%. During the whole of 2000 and until mid 2001 the short rates were higher than the long rates. A downward sloping yield curve indicates that the market expects interest rates to fall. If the economy is not growing well and inflation is under control, the government may bring down base rates to stimulate economic growth. From the end of 2001 until mid 2004 the (more usual) upward sloping term structure was evident, after which the term structure was relatively flat until the beginning of 2006. Notice the downward sloping term structure persisted throughout 2006, so the market was expecting interest rates to fall during all this time. In the event it turned out that they were not cut until August 2007, when the UK followed the US interest rate cuts that were necessitated by the sub-prime mortgage crisis.




Figure III.1.5 UK spot rate curve, January 2000 to December 2007

²² The Bank provides daily data on the UK government and LIBOR yield curves for maturities between 6 months and 25 years, based on zero coupon gilts, and a separate short curve to use for risk free interest rates up to 5 years' maturity. The short curve is based on sterling interbank rates (LIBOR), yields on instruments linked to LIBOR, short sterling futures, forward rate agreements and LIBOR-based interest rate swaps. More information and data are available from <http://www.bankofengland.co.uk/statistics/ylcurve/index.htm>

Taking a vertical section through Figure III.1.5 on any fixed date gives the term structure of spot interest rates on that day. As is typical of interest rate term structures, the UK spot rates fluctuate between a downward sloping term structure, where short rates are above long rates (e.g. during 2000–2001 and 2006–2007) and an upward sloping term structure, where short rates are below long rates (e.g. during 2002–2003). By definition, the forward rates will be greater than spot rates when the spot curve is upward sloping, and less than spot rates when the spot curve is downward sloping. To illustrate this, Figure III.1.6 compares the term structure of UK spot rates with that of 6-month forward rates on (a) 2 May 2000 and (b) 2 May 2003.




Figure III.1.6 UK government spot and 6-month forward rates on (a) 2 May 2000 and (b) 2 May 2003

We conclude this section by referring to two chapters in *Market Risk Analysis*, Volume II, with detailed analyses of the statistical properties of term structures of interest rates. Specifically:

- the characteristics of different covariance and correlation matrices of US Treasury interest rates are analysed in Section II.3.5;
- the principal component analysis of the covariance and correlation matrices of UK spot and forward rates is discussed in Section II.2.3. In each case we analyse both the short curve of 60 different rates with monthly maturities up to 5 years, and the entire curve with 50 different maturities from 6 months to 25 years.

The CD-ROM for Volume II contains all the data and results in Excel.

III.1.5 DURATION AND CONVEXITY

This section describes two traditional measures of a fixed coupon bond's price sensitivity to movements in its yield. The first order price sensitivity is called the *duration*, and we introduce the two most common duration measures: the *Macauley duration* and the *modified duration*. The second order price sensitivity to changes in the yield is called the *convexity* of the bond. We explain how duration and convexity are used in first and second order Taylor approximations to the change in bond price when its yield changes, and how to obtain similar approximations to the profit and loss (P&L) of a bond portfolio when the yield curve shifts parallel.

The Macaulay duration is the maturity weighted average of the present values of the cash flows. It is measured in years and a zero coupon bond has Macaulay duration equal to its maturity. Under continuous compounding we show that the Macaulay duration is also a first order approximation to the percentage change in bond price per unit change in yield. However, under discrete compounding the first order approximation to the percentage change in bond price per unit change in yield is the modified duration. This is the Macaulay duration divided by the discrete compounding factor for the yield.

Duration increases with the maturity of the bond, and for two bonds of the same maturity the duration is higher for the bond with lower coupon, or for the bond with less frequent coupon payments. Since the price–yield relationship is non-linear, the duration also depends on the yield. The duration is related to the slope of the price–yield curve. In general, the lower the yield, the higher the duration. The bond convexity is the sensitivity of duration to changes in the bond yield. It measures the curvature of the price–yield curve. We conclude this section by explaining how modified duration and convexity are used to immunize a bond portfolio against interest rate changes.

III.1.5.1 Macaulay Duration

The Macaulay duration of a bond – or indeed of any sequence of positive cash flows – is the maturity weighted average of the present values of its cash flows,

$$D_M = \frac{\sum_{i=1}^n T_i P_{T_i}}{P}, \quad (\text{III.1.31})$$

where P is the price of the bond and

$$P_{T_i} = (1 + y)^{-T_i} C_{T_i}. \quad (\text{III.1.32})$$

Here y is the yield and C_{T_i} are the coupon payments for $i = 1, \dots, n-1$ and C_{T_n} is the redemption value which consists of the face value plus the coupon payment, all based on a face value of 100.²³

Thus a zero coupon bond always has Macaulay duration equal to its maturity, by definition. In general the Macaulay duration provides a single number, measured in years, that represents an average time over which income is received. It can be shown (see Example III.1.8 below) that if the yield curve shifts down the Macaulay duration represents a ‘break-even’ point in time where the income lost through reinvestment of the coupons is just offset by the gain in the bond’s value.

EXAMPLE III.1.7: MACAULAY DURATION

What is the Macaulay duration of a 4-year bond paying a coupon of 6% annually when the interest rates at 1, 2, 3 and 4 years are as shown in Table III.1.8? What would the Macaulay duration be if the coupon rate were 5%?

SOLUTION The answer to the first part is also shown in Table III.1.8. The price of this bond is 103.62 and its yield is calculated using the Excel RATE function as usual, giving 4.98%. The Macaulay duration is therefore

$$D_M = 381.27/103.62 = 3.68 \text{ years.}$$

²³ We could calculate a time weighted average present value based on discounting by market interest rates, but it is easier to use the yield. For a bond we know the market price, and so also the yield, thus we can calculate D_M without having to construct the zero curve.

Although this is not shown in the table, using the spreadsheet for this example the reader can change the coupon rate to 5% and read off the Macaulay duration. When the coupon drops from 6% to 5% the Macaulay duration increases to $372.57/100.06 = 3.72$ years.

Table III.1.8 Macaulay duration of a simple bond

Year	Spot rate	Cash flow	PV (Spot)	Yield	PV (Yield)	PV (Yield) $\times$ Year
1	4.5%	6	5.74	4.98%	5.72	5.72
2	4.75%	6	5.47	4.98%	5.44	10.89
3	4.85%	6	5.21	4.98%	5.19	15.56
4	5%	106	87.21	4.98%	87.28	349.10
Sum			103.62		103.62	381.27

This example has demonstrated that the duration and the coupon are inversely related. Thus if two bonds have the same maturity, the one with the higher coupon will have the smaller duration.

EXAMPLE III.1.8: MACAULAY DURATION AS A RISK MEASURE

Consider a 10-year bond with a 10% coupon and with the market interest rates as in Table III.1.9. Find the fair value of the bond at maturity, and at h years from now for $h = 9, 8, 7, \dots, 0$ (the value at $h = 0$ being the present value of the bond). Now suppose that the yield curve shifts up by 2%. Re-evaluate the future value of the bond at the same points in time as before, and draw a graph of the bond values versus time under (a) the original yield curve in Table III.1.9, and (b) when the yield curve shifts up by 2%. At which point do the two value lines intersect? Does this point change if the shift in yield curve is different from +2%?

Table III.1.9 A zero coupon yield curve

Maturity	1	2	3	4	5	6	7	8	9	10
Spot rate	4.00%	4.25%	4.50%	4.25%	4.20%	4.15%	4.10%	4.00%	4.00%	4.00%

SOLUTION In the spreadsheet for this example we calculate the bond's maturity value and then discount this value using the relevant forward rate. This gives the value of the bond h years from now. We do this both for the original interest rates shown in Table III.1.9 and then after shifting that yield curve up by 2%. The resulting graph is shown in Figure III.1.7. The reader may change the yield curve shift in the spreadsheet from 2% to any different value and will note that, whilst the grey line changes, the point of intersection remains constant at approximately 7.36 years. Then we show, using the Excel Solver, that the Macaulay duration of the bond is 7.36 years. That is, the point in time where the two lines intersect is always equal to the Macaulay duration. If interest rates increase the value of the bond (including all reinvested coupons) decreases, but the reinvestment income from the coupon stream




Figure III.1.7 Future value of a bond under two different yield curves

increases.²⁴ The Macaulay duration represents the future point in time where the loss from the fall in the bond price is just offset by the income gained from an increased interest on the coupon payments. This is why Macaulay duration has been used as a measure of sensitivity to market interest rates.

III.1.5.2 Modified Duration

The modified duration is the first order approximation to the *percentage price change per unit change in yield*, multiplied by -1 so that the modified duration is positive (since the price falls when the yield increases). For an *annual* bond the modified duration is the Macaulay duration divided by $1 + y$.

It can be shown that, under *continuous* compounding, the Macaulay duration is given by²⁵

$$D_M = \text{Macaulay duration} = -\frac{1}{P} \frac{dP}{dy}.$$

But this is only the case under continuous compounding. Under discrete compounding, we have

$$D = \text{Modified duration} = -\frac{1}{P} \frac{dP}{dy} = \frac{D_M}{1 + y}. \quad (\text{III.1.33})$$

Since discrete compounding is applied in the markets it is the modified duration, not the Macaulay duration, that is used as the first order approximation to the percentage change in bond price per unit change in yield.

²⁴ This analysis is based on the assumption that all realised spot rates (i.e. the spot rates at which coupons are reinvested) are equal to the current forward rates.

²⁵ With continuous compounding, $P = \sum_{T_i} P_{T_i} = \sum_{T_i} e^{-yT_i} C_{T_i}$ so $\frac{dP}{dy} = -\sum_{T_i} T_i P_{T_i}$ and $-\frac{1}{P} \frac{dP}{dy} = \frac{1}{P} \sum_{T_i} T_i P_{T_i} = D_M$.

III.1.5.3 Convexity

The duration gives a first order approximation to the percentage change in the bond price per unit change in yield. Thus it is related to the slope of the price–yield curve. Yet we know that the price and yield have a non-linear relationship. Hence an additional measure of interest rate sensitivity is often employed to capture the curvature of the price–yield relationship. We define

$$\text{Convexity} = \frac{1}{P} \frac{d^2 P}{dy^2}. \quad (\text{III.1.34})$$

For instance, we can use (III.1.34) to calculate the convexity for an *annual* bond of maturity T years, with T being an integer and so coupon payments are in exactly 1, 2, $\dots$, T years. Then

$$\frac{d^2 P}{dy^2} = \frac{2C_1}{(1+y)^3} + \frac{6C_2}{(1+y)^4} + \frac{12C_3}{(1+y)^5} + \dots + \frac{T(T+1)C_T}{(1+y)^{T+2}}, \quad (\text{III.1.35})$$

which follows on differentiating (III.1.22) twice with respect to y .

A *zero coupon* bond of maturity T has convexity

$$\text{Zero coupon bond convexity} = T(T+1)(1+y)^{-2}. \quad (\text{III.1.36})$$

This can be proved by differentiating $P = (1+y)^{-T}$ twice and then dividing by P . For coupon bonds the convexity increases as the coupon payments become more spread out over time.

Convexity captures the second order effect of a change in yield on the price of a bond. The convexity of a bond is its second order price sensitivity to a change in the yield, i.e. convexity measures the sensitivity of modified duration to an absolute change in the yield. See Example III.1.9 below for a numerical illustration.

III.1.5.4 Duration and Convexity of a Bond Portfolio

Modified duration measures the percentage change in bond price per unit change in yield. Since these cannot be added, how do we measure duration when we have several bonds in a portfolio?

The *value duration* of a bond is the *absolute* change in bond price per unit change in yield:

$$\text{Value duration} = D^\$ = -\frac{dP}{dy}. \quad (\text{III.1.37})$$

Similarly, the *value convexity* is defined in terms of absolute changes as

$$\text{Value convexity} = C^\$ = \frac{d^2 P}{dy^2}. \quad (\text{III.1.38})$$

The value duration and convexity derive their names from the fact that they are measured in the same currency unit as the bond.²⁶ The value duration is the modified duration multiplied by the bond price and the value convexity is the convexity multiplied by the bond price. Note that, like ordinary duration and convexity, they can be defined for *any* sequence of positive cash flows, so they can also be calculated for fixed income securities other than bonds.

It is useful to work with value duration and value convexity *because they are additive*. Suppose a portfolio has positions on k different fixed income securities and that the value

²⁶ Note that they are also commonly called the *dollar duration* and *dollar convexity*. But not all bonds are measured in dollars, so we prefer the terms value duration and value convexity.

durations for these positions are $D_1^{\$}, \dots, D_k^{\$}$ and their value convexities are $C_1^{\$}, \dots, C_k^{\$}$. Then the portfolio has value duration $D_p^{\$}$ and value convexity $C_p^{\$}$ given by the sum of the value durations for each position:

$$D_p^{\$} = D_1^{\$} + \dots + D_k^{\$} \quad (\text{III.1.39})$$

and

$$C_p^{\$} = C_1^{\$} + \dots + C_k^{\$}. \quad (\text{III.1.40})$$

For instance, consider a long position of \$2 million with value duration \$400 and value convexity \$15 million, and a short position of \$1.5 million with value duration \$500 and value convexity \$25 million. Then the value duration of the total position is -\$100 and the value convexity is -\$10 million.

Since each bond in the portfolio has a different yield, the interpretation of the value duration and value convexity is tricky. For these to represent yield sensitivities we must assume that the yields on all the bonds in the portfolio move by the same amount. In other words, value duration and value convexity are the first and second order price sensitivities of the portfolio, measured in value terms, to an identical shift in the yields on all bonds in the portfolio. This is explained in more detail in the next subsection.

III.1.5.5 Duration–Convexity Approximations to Bond Price Change

Consider a second order Taylor approximation of the bond price with respect to its yield:²⁷

$$P(y) \approx P(y_0) + \left. \frac{dP}{dy} \right|_{y=y_0} (y - y_0) + \frac{1}{2} \left. \frac{d^2P}{dy^2} \right|_{y=y_0} (y - y_0)^2.$$

Write $(y - y_0) = \Delta y$ and $P(y) - P(y_0) = \Delta P$, so ΔP is the change in bond price when the yield changes by Δy . Then the second order Taylor approximation may be written

$$\Delta P \approx \frac{dP}{dy} \Delta y + \frac{1}{2} \frac{d^2P}{dy^2} (\Delta y)^2. \quad (\text{III.1.41})$$

So using the definitions of value duration and convexity above, we have

$$\Delta P \approx -D^{\$} \Delta y + \frac{1}{2} C^{\$} (\Delta y)^2. \quad (\text{III.1.42})$$

Now suppose that we hold a bond portfolio and that we assume that all the yields on the bonds in the portfolio change by the same amount Δy . Then we can add up the approximations (III.1.42) over all the bonds in the portfolio to obtain

$$\Delta P \approx -D_p^{\$} \Delta y + \frac{1}{2} C_p^{\$} (\Delta y)^2, \quad (\text{III.1.43})$$

where now ΔP denotes the change in portfolio value, i.e. the portfolio P&L, when all yields change by Δy .

Bond traders' limits are usually based on value duration and value convexity and traders often use (III.1.43) to approximate the effect of a change in *market interest rates* on their portfolio P&L. That is, they assume that *all the bond yields change by the same amount when the zero coupon yield curve shifts parallel*. They must assume this, otherwise (III.1.43) would

²⁷ See Section I.1.6 for an introduction to Taylor approximation.

not hold, but of course when the zero curve shifts parallel different bond yields change by different amounts.

Another duration–convexity approximation, this time for a single bond, follows on dividing (III.1.41) by P , giving

$$\frac{\Delta P}{P} \approx \frac{1}{P} \frac{dP}{dy} \Delta y + \frac{1}{2P} \frac{d^2 P}{dy^2} (\Delta y)^2.$$

When we compare the above with the definitions (III.1.33) and (III.1.34) we see that we have obtained an approximation to the percentage change in bond price for a given change in yield. That is:

$$\frac{\Delta P}{P} \approx -\text{Modified Duration} \times \Delta y + \frac{1}{2} \times \text{Convexity} \times (\Delta y)^2. \quad (\text{III.1.44})$$

Notice that the signs on modified duration and convexity are opposite. Hence, for two cash flows with the same modified duration, the one with the higher convexity has a value that is less sensitive to adverse interest rate movements. That is, its value decreases less if interest rates increase and increases more if interest rates fall. Hence, a large and positive convexity is a desirable quality.

EXAMPLE III.1.9: DURATION–CONVEXITY APPROXIMATION

Find the modified duration and convexity of the two bonds in Example III.1.4. Hence, for each bond, use the duration–convexity approximation to estimate the change in bond price when the yield increases by 1%. Compare this estimate to the full valuation of each bond’s price based on the new yield.

SOLUTION We already know the price and the yield of the bonds from Example III.1.4. Now, using the formulae derived above, first we obtain the Macaulay duration (as the maturity weighted average of the present value of the cash flows) and then we divide this by $1 + y$ to obtain the modified duration. For the convexity we use (III.1.35) and (III.1.34). Finally, the approximate change in price of each bond is obtained on applying (III.1.44) with $\Delta y = +1\%$. The actual percentage price change is calculated using the exact relationship (III.1.22) for the original yield and the yield $+1\%$. The solution is summarized in Table III.1.10.

Table III.1.10 Duration-convexity approximation

	Bond 1	Bond 2
Macaulay duration	2.86	4.27
Modified duration	2.74	4.09
Convexity	10.31	22.24
<i>Duration–convexity approximation</i>		
Yield change	1%	1%
Percentage price change	−2.686%	−3.982%
Actual percentage price change	−2.687%	−3.985%

III.1.5.6 Immunizing Bond Portfolios

Bond portfolio managers often have a view about interest rates. For instance, they might believe that short rates will increase. Portfolio managers aim to balance their portfolios so

that they can profit from the movements they expect, and will make no losses from other common types of movements in interest rates. One of the most common movements is an approximately parallel shift in the yield curve, i.e. all interest rates change by approximately the same absolute amount.

Thus a portfolio manager who aims to ‘immunize’ his portfolio from such movements will hold a portfolio that has both zero duration and zero convexity. On the other hand, if he wants to take advantage of a positive convexity he will structure his portfolio so that the convexity is large and positive and the duration is zero.

EXAMPLE III.1.10: IMMUNIZING BOND PORTFOLIOS

Consider again the two bonds in Examples III.1.4 and III.1.9. Suppose the principal amount invested in bond 1 is \$1.5 million and the principal of bond 2 is \$1 million.

- Calculate the value duration and value convexity of each bond.²⁸
- Find the value duration and value convexity of a portfolio that is long bond 1 and short bond 2.
- Immunize this portfolio using bonds 3 and 4 with characteristics in Table III.1.11.

Table III.1.11 Two bonds²⁹

	Bond 3	Bond 4
Principal (\$)	1,000	10,000
Present value (\$)	1,200	10,780
Value duration (\$)	5	2
Value convexity (\$)	20,000	100,000

SOLUTION The solution to parts (a) and (b) is given in the spreadsheet and summarized in Table III.1.12.

Table III.1.12 Value duration and value convexity

	Bond 1	Bond 2	Portfolio
Value duration (\$)	416.41	514.00	−97.59
Value convexity (\$)	15,687,184	27,938,173	−12,250,989

For part (c) we therefore seek (w_1, w_2) such that

$$\text{zero duration} \Rightarrow 5w_1 + 2w_2 - 97.59 = 0.$$

$$\text{zero convexity} \Rightarrow 20,000 w_1 + 100,000 w_2 - 12,250,989 = 0.$$

²⁸ Note that we have computed value durations here as the actual price change for a 1 *basis point* change in yield, rather than with respect to a 1% change in yield. This is so that the value duration and the PV01 have similar orders of magnitude, as demonstrated by Table III.1.17.

²⁹ Recall that the present value is the bond price times the principal amount invested divided by 100. Hence, bond 3 has price 120 and bond 4 has price 107.8.

The solution is

$$\begin{pmatrix} w_1 \\ w_2 \end{pmatrix} = \begin{pmatrix} 5 & 2 \\ 20,000 & 100,000 \end{pmatrix}^{-1} \begin{pmatrix} 97.59 \\ 12,250,989 \end{pmatrix} = \begin{pmatrix} -32.05 \\ 128.92 \end{pmatrix},$$

which means

- selling $32.05 \times \$1200 = \$38,459$ of bond 3. Since its price is 120 this is equivalent to selling the principal amount of \$32,049 in bond 3, and
- buying $128.92 \times \$10,780 = \$1,389,755$ of bond 4. Since its price is 107.8 this is equivalent to buying the principal amount of \$1,289,198 in bond 4.

The resulting portfolio, whose value is approximately invariant to parallel shifts in the yield curve, is summarized in Table III.1.13.

Table III.1.13 An immunized bond portfolio

Portfolio	Bond 1	Bond 2	Bond 3	Bond 4
Principal (\$)	1,500,000	−1,000,000	−32,049	+1,289,198

We remark that a disadvantage of using the traditional sensitivities for hedging the risk of a bond portfolio is that duration and convexity hedging is limited to *parallel* movements of the zero coupon yield curve. More advanced bond portfolio immunization techniques based on *principal component analysis* are introduced and demonstrated in Section II.2.4.4. The advantage of these is that we can:

- hedge the portfolio against the *most common* movements in the yield curve, based on historical data;
- quantify how much of the historically observed variations in the yield curve we are hedging.

Hence, these techniques provide better hedges than standard duration–convexity hedges.

III.1.6 BONDS WITH SEMI-ANNUAL AND FLOATING COUPONS

For simplicity all the numerical examples in Sections III.1.4 and III.1.5 were based on fixed coupon annual bonds. When a bond does not have an annual coupon, the coupon is usually paid every 6 months, and it is called a *semi-annual bond*. For instance, a semi-annual bond with coupon 6% pays 3% of the principal amount invested every 6 months. The first part of this section extends the general formulae for the price, yield, duration and convexity of a bond to semi-annual bonds, and then we generalize the formulae to a fixed coupon *quarterly bond* (i.e. coupons are paid every 3 months).

Fixed coupon quarterly bonds are rare, but *floating rate notes*, which are bonds that have variable coupons, are usually semi-annual or quarterly. The characteristics of floating rate notes, or *floaters* as they are commonly termed, are different from the characteristics of fixed coupon bonds. For instance, if the coupon on a semi-annual floater, which is set six months in advance, is just the 6-month LIBOR rate at the time the coupon is set then the duration of the floater is very short. In fact a semi-annual floater has a duration of no more than 6

months, as we shall demonstrate in this section. Thus by adding long or short positions in floaters, bond traders can change the duration of their bond portfolios, for instance to stay within their trading limits. In addition to straight floaters we have reverse floaters, capped floaters, floored floaters and many other floaters that can have complex option-like features.

III.1.6.1 Semi-Annual and Quarterly Coupons

When dealing with semi-annual bonds the only thing to keep in mind is that the definitions are stated in terms of the *semi-annual yield*, which is one-half the value of the annual yield. In terms of the annual yield y and annual coupon c percent, the price–yield relationship (III.1.22) for a bond with maturity $T/2$ years becomes:

$$P = \sum_{t=1}^T \left(1 + \frac{y}{2}\right)^{-t} C_t, \quad (\text{III.1.45})$$

where

$$C_t = \begin{cases} 100 \times c/2, & \text{for } t = 1, \dots, T-1, \\ 100(1 + c/2), & \text{for } t = T. \end{cases} \quad (\text{III.1.46})$$

Note that we now use t to index the cash flows. Hence $t = 1, \dots, T$, but t is *not* the time in *years* when the cash flows are paid. These times are $T_1 = 1/2, T_2 = 1, T_3 = 1\frac{1}{2}, \dots, T_n = \frac{1}{2}T$.

EXAMPLE III.1.11: YIELD ON SEMI-ANNUAL BOND

A 3-year semi-annual bond pays an 8% coupon and has a price of 95. Find its semi-annual yield and its annual equivalent yield.

SOLUTION We use the Excel RATE function to back out the semi-annual yield from the formula (III.1.45) with $P = 95$, $T = 6$ and semi-annual coupons of 4%. This gives the semi-annual yield $\frac{1}{2}y = 4.98\%$ so the annual equivalent yield y is 9.96%.

The duration and convexity are calculated using slight modifications of the formulae given in Section III.1.5. For instance, the Macaulay duration on a semi-annual bond of maturity $T/2$ is³⁰

$$D_M = \frac{\frac{1}{2} \sum_{t=1}^T t P_t}{P}, \quad (\text{III.1.47})$$

where P is the price of the bond and

$$P_t = \left(1 + \frac{y}{2}\right)^{-t} C_t, \quad (\text{III.1.48})$$

with y the annual equivalent yield on the bond.

The modified duration of a semi-annual bond is the Macaulay duration divided by $1 + \frac{1}{2}y$. That is, we divide the Macaulay duration by the compounding factor for the *semi-annual* yield and not the annual yield on the bond. This way, the modified duration is still defined as the percentage change in bond price for a unit change in the *annual equivalent* yield. Thus for a semi-annual bond (III.1.33) becomes

$$D = -\frac{1}{P} \frac{dP}{dy} = \frac{D_M}{1 + y/2}. \quad (\text{III.1.49})$$

³⁰ We use t to represent the payment number, i.e. $t = 1, 2, \dots, T$, but the times these payments are made are $t/2 = 1/2, 1, \dots, T/2$. Hence the factor of $\frac{1}{2}$ in the formula (III.1.47).

The convexity of a semi-annual bond is still defined as the second derivative of the bond price with respect to the annual equivalent yield, so (III.1.34) still applies. Thus, for a semi-annual bond,

$$\text{Convexity} = \frac{1}{P} \frac{d^2 P}{dy^2}, \quad (\text{III.1.50})$$

But when the coupons are paid semi-annually the formula (III.1.35) for the second derivative of the bond price with respect to a change in the annual equivalent yield y needs to be modified to

$$\begin{aligned} \frac{d^2 P}{dy^2} &= \frac{C_1/2}{(1+y/2)^3} + \frac{3C_2/2}{(1+y/2)^4} + \frac{3C_3}{(1+y/2)^5} + \dots + \frac{(T(T+1)/4)C_T}{(1+y/2)^{T+2}} \\ &= \sum_{t=1}^T \frac{(t(t+1)/4)C_t}{(1+y/2)^{t+2}}, \end{aligned} \quad (\text{III.1.51})$$

with the semi-annual cash flows defined by (III.1.46).

EXAMPLE III.1.12: DURATION AND CONVEXITY OF A SEMI-ANNUAL BOND

Find the Macaulay duration, modified duration and convexity of a 3-year semi-annual bond that pays an 8% coupon and has a price of 95.

SOLUTION This is the bond for which we found the semi-annual yield of 4.98% in the previous example. In (III.1.47) and (III.1.48) we have $P = 95$, $P_t = 4 \times (1.0498)^{-t}$ for $t = 1, 2, \dots, 5$, and $P_6 = 104 \times (1.0498)^{-6}$. The spreadsheet for this example sets out the calculation for the Macaulay duration, and we obtain $D_M = 2.718$. Then we find the modified duration using (III.1.49), so we divide D_M by 1.0498, and this gives the modified duration $D = 2.589$. Finally, the spreadsheet implements (III.1.50) and (III.1.51) with $T = 6$, giving the convexity 8.34.

For reference, the formulae for price, duration and convexity of a fixed coupon *quarterly* bond are summarized below:

$$\begin{aligned} P &= \sum_{t=1}^T PV_t \\ PV_t &= \left(1 + \frac{y}{4}\right)^{-t} C_t \\ C_t &= \begin{cases} 100 \times c/4. & \text{for } t = 1, \dots, T-1 \\ 100(1+c/4) & \text{for } t = T \end{cases} \\ D_M &= \frac{(\sum_{t=1}^T tP_t)/4}{P} \\ \frac{d^2 P}{dy^2} &= \sum_{t=1}^T \frac{(t(t+1)/16)C_t}{(1+y/4)^{t+2}} \\ D &= -\frac{1}{P} \frac{dP}{dy} = \frac{D_M}{1+y/4} \\ \text{Convexity} &= \frac{1}{P} \frac{d^2 P}{dy^2} \end{aligned} \quad (\text{III.1.52})$$

III.1.6.2 Floating Rate Notes

A floating rate note (FRN) or *floater* is a bond with a coupon that is linked to LIBOR or some other reference interest rate such as a treasury zero coupon rate. A typical floater pays coupons every 3 months or every 6 months, but coupon payments are sometimes also monthly or annual.³¹ When a coupon payment is made the next coupon is set by the current reference rate, and the period between the date on which the rate is set and the payment date is referred to as the *lock-out period*.

Usually the maturity of the reference rate is the time between coupon payments. So, for example, the next coupon payment on a semi-annual floater is equal to the 6-month reference rate prevailing at the last coupon payment date plus a spread; for a quarterly floater the next coupon payment is equal to the 3-month reference rate at the time of the last coupon plus a spread. The spread depends on the credit quality of the borrower. For instance, since LIBOR is rated AA any borrower with a rating less than AA will have a positive spread over LIBOR and this spread increases as the borrower's credit rating decreases. Also, to account for the possible deterioration in the credit quality of the borrower over time, the spread increases with maturity of the floater, so that a 10-year floater has a larger spread than a 2-year floater.

We now demonstrate how to price a floater. It is slightly easier to write down the mathematics by assuming the coupons are paid annually, so we shall make that assumption merely for simplicity; it does not alter the concepts in the argument. Suppose we are at time zero and the next coupon, which has been fixed at c on the last coupon payment date, is paid at time t , measured in years. Also suppose there are T further annual coupon payments until the floater expires. So the subsequent payments are at times $t + 1, t + 2, \dots, t + T$ and the final payment at time $t + T$ includes the principal amount invested, N .

Denote the 1-year reference rate at time t by R_t and for simplicity assume this is also the discount rate, i.e. the LIBOR rate. Denote the fixed spread over this rate by s . Thus the coupon payment at time $t + n$ is $100(R_{t+n-1} + s)$. Now the price of a floater may be written

$$P_{t+T}^s = (B_{t+T}^s - B_{t+T}^0) + 100(1 + c - s)(1 + tR_0)^{-1}, \quad (\text{III.1.53})$$

where P_m^s is the price of the floater with spread s and maturity m , B_m^s is the price of a bond with fixed coupon c and maturity m and R_0 is the t -period discount rate at time 0.

To prove (III.1.53) we decompose the coupon payments into two parts: the payments based on the reference rate and the payments based on the fixed spread s over this reference rate. First consider the fixed coupon payments of $100s$ on every date, including the last. What is the discounted value of the cash flows $100\{s, s, \dots, s, s\}$ at times $t, t + 1, \dots, t + T$? Let us temporarily add 100 to the final payment, recalling that the discounted value of a cash flow of 100 at time $t + T$ is the value of a zero coupon bond expiring at time $t + T$. Adding this zero bond, the cash flows become $100\{s, s, \dots, s, 1 + s\}$. These are the cash flows on a bond with maturity $t + T$ and fixed coupon s . Hence, discounting the cash flows $100\{s, s, \dots, s, s\}$ to time zero gives $B_{t+T}^s - B_{t+T}^0$, which is the first term in (III.1.53).

Now consider the floating part of the payments. Working backwards from the final payment, we discount to time t all the floating coupon payments at $t + 1, t + 2, \dots, t + T$, plus 100 at $t + T$, assuming the spread is zero. At time $t + T$ the payment of $100(1 + R_{t+T-1})$ is discounted to time $t + T - 1$ using the reference rate R_{t+T-1} , so the discounted value is 100.

³¹ However, Fed funds floaters, for instance, can reset coupons daily because the reference rate is an overnight rate, while T-bill floaters usually reset weekly following the T-bill auction.

Similarly, at time $t + T - 1$ the coupon payment added to the discounted value of 100 of the final payment is $100(1 + R_{t+T-2})$. This is discounted to time $t + T - 2$ using the reference rate R_{t+T-2} , so the discounted value is again 100. Continuing to discount the coupon payment plus the discounted values of future payments until we reach time t gives 100. Note that this is independent of the interest rates. Now we have a cash flow of $100(1 + c)$, where c is the coupon fixed at the previous payment date, at time t . Discounting this to time zero gives the second term in (III.1.53).

EXAMPLE III.1.13: PRICING A SIMPLE FLOATER

A 4-year FRN pays annual coupons of LIBOR plus 60 basis points. The 12-month LIBOR rate is 5% and the 2-, 3- and 4-year discount rates are 4.85%, 4.65% and 4.5%, respectively. Find its price.³²

SOLUTION To calculate the first term in (III.1.53) we need to find the price of a 4-year bond paying annual coupons of 60 basis points and the price of a 4-year zero bond. These are based on (III.1.20) and are found in the spreadsheet as 86 for the coupon bond and 83.866 for the zero bond. Hence, the fixed part of the floater, i.e. the first term in (III.1.53), is 2.144. Since we are pricing the floater on a coupon payment date, the floating part of the floater, i.e. the second term in (III.1.53), is 100. This is because the payment in 1 year is discounted using the coupon rate, c .³³ Hence the price of the floater is 102.144.

The yield on a floater is the fixed discount rate that gives the market price when all cash flows are discounted at the same rate. Since a floater with zero spread is equivalent to an instrument paying 100 at the next coupon date, in present value terms, we can calculate the yield y on an annual floater using the formula³⁴

$$P_{t+T}^s = 100 \left(\sum_{\tau=t}^{t+T-1} s(1+y)^{-\tau} - (1-s)(1+y)^{-t+T} + (1+c-s)(1+y)^{-t} \right). \quad (\text{III.1.54})$$

We can thus back out the yield from the price using a numerical algorithm such as the Excel Solver. The Macaulay and modified durations are then calculated using the usual formulae (III.1.31) and (III.1.33).

EXAMPLE III.1.14: YIELD AND DURATION OF A SIMPLE FLOATER

Calculate the yield and the Macaulay and modified durations of the annual floater of the previous example. How would these quantities change if the floater paid 100 basis points over LIBOR?

SOLUTION The Solver is set up in the spreadsheet to back out the yield from (III.1.54) based on a price of 102.114, found in the previous example. The result is a yield of 4.98%, and then the usual formulae for the durations give a Macaulay duration of 1.03 years and a modified duration of 0.981. When we change the spread to 1% in the spreadsheet and repeat the exercise, the yield changes considerably, to 4.98%. However, the modified duration changes very little. The Macaulay duration becomes 1.049 years but the modified duration is 0.999, only 0.028 greater than it was for the floater with a spread of 60 basis points.

³² The solution assumes the 60bp spread is not a reflection of a lower than AA credit rating. If it is, then discounting should be not at LIBOR but at LIBOR + 60bp and the price would simply be 100.

³³ But note that the floating part would not be 100 if the floater were priced *between* coupon dates, because the cash flow of $100(1+c)$ would be discounted at a rate different from c , as it is in (III.1.53).

³⁴ To avoid awkward notation here we assume t and T are measured in years, as in our examples.

This example illustrates the fact that the duration of a floater is approximately equal to the time to the next coupon payment, in the case 1 year because we considered an annual floater priced on a coupon payment date. But floaters are typically monthly, quarterly or semi-annual, and are of course priced between coupon payments. Their duration is *always* approximately equal to the time to the next coupon payment. Hence, floaters have very short duration and are therefore often used by traders to change the duration of their bond portfolios.

The pricing formula (III.1.53) has important implications. It means that the market risk of a principal amount N invested in a floater has two components:

- The market risk on a portfolio where we buy a notional N of a bond with coupon s and are short N on a zero bond. This part of the market risk is often very low because the spread s is typically quite small. For instance, an AA borrower on a 5-year floater might pay 20–50 basis points over LIBOR.
- The market risk on a single cash flow at the next coupon date, which is equal to principal N plus the coupon payment on N , i.e. $N(1 + c - s)$. As with any cash flow, the market risk is due to fluctuations in the discount rate, because the present value of $N(1 + c - s)$ will fluctuate as spot rates change.

Hence, a floater has virtually no market risk beyond the first coupon payment. Clearly floaters with longer reset periods are more vulnerable to market risks.

III.1.6.3 Other Floaters

The simple floaters considered in the previous section were easy enough to price because the coupon payments were based on a spread over a discount rate of the same maturity as the time between coupons. But many floaters are not linked to a market interest rate of the same maturity as the coupon period. For instance, a quarterly floater could have coupons linked to a 12-month LIBOR rate. The valuation of these floaters is complex since it depends on the volatilities and correlations between interest rates. For instance, to value a quarterly floater with coupons linked to a 12-month LIBOR rate, we would need to estimate the volatilities of the 3- and 12-month LIBOR rates and their correlation.

A *reverse floater* has a coupon equal to a fixed rate *minus* the reference rate. For instance, the coupon could be $10\% - \text{LIBOR}$ or $12.5\% - 2 \times \text{LIBOR}$. Of course, no coupon is paid if the reverse floater coupon formula gives a negative value and since we cannot have negative coupons it is rather complex to price reverse floaters. The constraint that the coupon is non-negative imposes an embedded option structure on their price. Other floaters with embedded options are *capped floaters*, which have coupons that cannot rise above the level of the cap, and *floored floaters*, which have coupons that cannot fall below the level of the floor. Hence, a reverse floater is a special type of floored floater with the floor at zero.

III.1.7 FORWARD RATE AGREEMENTS AND INTEREST RATE SWAPS

This section introduces two related interest rate products that are very actively traded in OTC markets: forward rate agreements (FRAs) and interest rate swaps. Since these fixed income instruments are not exchange traded, their trading needs to be accompanied by legal documentation. In 1992 the International Swaps and Derivatives Association (ISDA) issued

standard terms and conditions for the trading of FRAs and swaps. Since then these markets have grown considerably, and the ISDA Master Agreement was thoroughly revised in 2002. The Bank for International Settlements estimate that by December 2006 the total notional value outstanding on interest rate swaps and FRAs was about 250 trillion US dollars, with a mark-to-market value of \$4.2 trillion.³⁵ The vast majority of this trading is on interest rate swaps rather than FRAs.

III.1.7.1 Forward Rate Agreements

A *forward rate agreement* is an OTC agreement to buy or sell a forward interest rate. For instance, a company entering into an FRA may wish to buy or sell the 3-month forward 3-month LIBOR rate, $F_{3,3}$. Consider an example of a company that has a loan of \$100 million from bank A on which it pays a floating interest rate, with interest payments every 3 months. As each interest payment is made the next payment is fixed based on the current 3-month spot LIBOR rate plus a spread. For instance, the payment due in 6 months will be determined by the spot rate in 3 months' time.

If the company is concerned that interest rates could rise over the next 3 months it may prefer to *fix* the payment that is due in 6 months *now*. This can be accomplished by exchanging the floating rate, which is whatever the 3-month spot LIBOR rate will be in 3 months' time, for a fixed rate which is determined when the FRA is purchased. Let us suppose the company buys this FRA from bank B when $F_{3,3} = 5\%$.³⁶ With this FRA, in 3 months' time the company and the bank make an exchange of interest rates. The company pays 5% and receives R_3 , i.e. the 3-month spot LIBOR rate in 3 months' time. So in effect, the company only pays 5% interest on its loan (plus the spread). In 6 months' time, when the interest on the loan is due, there is a single settlement to the company of

$$\frac{1}{4} (R_3 - 5\%) \times \$100\text{m},$$

which could of course be either positive or negative, depending on R_3 . However, this payment is actually made in 3 months' time, i.e. as soon as R_3 becomes known. So, discounted to that date, the payment is

$$\left(1 + \frac{R_3}{4}\right)^{-1} \times \frac{1}{4} (R_3 - 5\%) \times \$100\text{m}. \quad (\text{III.1.55})$$

For instance, if the 3-month spot LIBOR rate R_3 turned out to be 5.6%, the company would gain on the FRA, because they pay only 5% when the floating payment would have been 5.6%. In this case the payment to the company 3 months from now would be:

$$1.014^{-1} \times \frac{1}{4} \times 0.006 \times \$100\text{m} = 1.014^{-1} \times \$150,000 = \$147,929.$$

Note that there is no exchange of principal. The only cash flows are the interest payments on some notional principal amount.

³⁵ See the Statistical Annex to the BIS *Quarterly Review*, available from www.BIS.org

³⁶ Being a market rate this is quoted in annual terms, and so the equivalent interest rate over 3 months is 1.25%. Hence, the company knows now that it will pay interest of 1.25% on its loan (plus the spread, which is not swapped in the FRA) in 6 months' time. Also bank A and bank B could be the same bank.

III.1.7.2 Interest Rate Swaps

Most of the interest rate swaps market is in short term interest rate swaps. Bid–ask spreads on the swap rate are just a few basis points and swap rates serve as reference for AA fixed rates. They are used to manage interest rate risks without inflating the balance sheet,³⁷ and without creating large credit risks. They are often combined with ordinary debt instruments to form *structured products* that provide a match between the specific needs of borrowers and objectives of investors.

In a standard or *vanilla* interest rate swap, one party pays LIBOR and receives a fixed rate, the swap rate, from the other party over an agreed period. When interest payments are matching in time only one net payment is made on each payment date. There is no exchange of principal. The swap rate is set so that the initial value of the swap is zero.

For example, a typical short-rate vanilla interest rate swap could be to

- receive a fixed rate F ,
- pay a floating rate of 3-month LIBOR
- on a notional principal of \$10 million
- for a duration of 1 year, with four quarterly exchanges of payments.

The fixed rate F at which the swap is struck is called the *swap rate*. When a swap is struck, the swap rate is found by setting the net present value of the cash flows to zero. At each payment date the floating rate for the next payment is fixed according to the current spot rates. For this reason the floating leg payment dates are often called the *reset* or *fixing dates*.

This is an example of a standard *fixed-for-floating* interest rate swap, which is the most common interest rate hedging instrument. We have shown in Section III.1.6.2 that floaters have very short duration and hence very little market risk. By contrast, long bonds with fixed coupons have significant market risks. So fixed-for-floating swaps allow investors to alter the market risk profile of their portfolio.

In the following we show that a fixed-for-floating interest rate swap can be regarded as a series of FRAs where net interest rate payments are made at regular intervals but without necessarily the same periodicity on the fixed and floating legs. Financial markets are full of idiosyncrasies that make calculations just that little bit more complex. In the case of swaps, the fixed and floating leg payments are usually based on different frequencies and day count conventions: the fixed leg is based on the bond market convention and the floating leg is based on the money market convention.³⁸ To keep track of two different day count conventions is necessary in practice, but not useful for illustrating the main concepts. We therefore assume in the next example that all payments are at quarterly intervals.

EXAMPLE III.1.15: VALUING A SWAP

Consider a \$10 million 1-year receive fixed pay floating swap with quarterly payments where the floating rate is 3-month LIBOR. Suppose the swap rate is 6% and that we are 1 month into the swap. The LIBOR rates are currently at 5%, 5.5%, 6% and 6.5% at the 2-, 5-, 8- and 11-month maturities, respectively. The first payment on the floating leg was fixed 1 month

³⁷ Swaps are traded OTC and are therefore *off-balance sheet* instruments, in the sense that the notional does not appear on the balance sheet. However, under the EU accountancy standards IAS 39, all derivatives are marked-to-market on the balance sheet.

³⁸ So the fixed leg may have semi-annual or yearly payments, and these are based on bond market day count conventions, whilst the floating leg is usually based on LIBOR (or some other standard short term reference rate) and may have more frequent payments which are based on the money market day count convention.

ago, based on the (then) current 3-month forward rate of 5.5%. Find the value of the swap now. If we were to fix the swap rate now, what would be its fair value?

SOLUTION Although the floating payments are determined by the spot LIBOR at the previous fixing date, for valuation purposes we can assume that realised spot rates are equal to current forward rates. Consider Table III.1.14. The current spot interest rates are as given in the column headed LIBOR, the interest payments corresponding to a swap rate of 6% are given in the column headed 'Receive 6%' and their value discounted to today is given in the next column. The present value of the fixed leg is the sum of these discounted values, \$0.57852 million.

Table III.1.14 The value of a vanilla swap

Months	LIBOR	Receive 6%	Discounted	Forward rates	Pay Floating	Discounted
2	5%	0.15	0.14876	$F_{-1,3}$ 5.50%	0.13750	0.13636
5	5.5%	0.15	0.14664	$F_{2,3}$ 5.79%	0.14463	0.14139
8	6%	0.15	0.14423	$F_{5,3}$ 6.68%	0.16701	0.16058
11	6.5%	0.15	0.14157	$F_{8,3}$ 7.53%	0.18830	0.17771
Present value			0.58120			0.61605

The floating leg payments are shown on the right-hand side of the table. The payment in 2 months' time was fixed last month, at 5.5%. For the other floating leg payments we use (III.1.18) to obtain the relevant forward rates from the current LIBOR curve. These are shown in the column headed 'Forward rates'. When all the floating payments are discounted (using the spot rates) and summed we obtain the present value of the floating leg as \$0.61605 million. Thus, based on a swap rate of 6%, the net present value of the swap is \$0.58120m – \$0.61605m, i.e. –\$34,851.

If we were to set the swap rate now, we would find the value of the swap rate F so that the present value of the swap was zero. The swap rate only affects the fixed leg payments, so using the Excel Solver (or Goal Seek) we find the fixed rate that makes the net present value of the fixed leg equal to \$0.61605 million. It turns out to be 6.36%.

III.1.7.3 Cash Flows on Vanilla Swaps

The above example considered an *on-the-run swap*, i.e. one with standard vertices as payment dates. These are often used for hedging swaps portfolios because they have a liquid market. But in a typical swaps portfolio the swaps have payment dates at various times in the future. For this reason the cash flows need to be mapped to a fixed set of vertices. We deal with the principles of cash flow mapping in Section III.5.3. But before we can map the cash flows on an interest rate swap we need to know how to represent these cash flows.

A *vanilla swap* is a swap where the floating payments at the end of a certain period depend only on the forward rate prevailing over that period. We now show that we can decompose the cash flows on a vanilla swap into two series: a series of cash flows on a coupon bearing bond and a single cash flow at the first payment date. Figure III.1.8 illustrates the cash flows on the swap considered in Example III.1.15 with one modification: we have supposed there is an exchange of the notional principal at maturity. This is marked in grey on the figure. Such an exchange of principal does not actually happen in a vanilla swap, but by assuming this hypothetical exchange it becomes easier to decompose the cash flows on the




Figure III.1.8 Cash flows on a vanilla swap

swap. Since there is both a positive and a negative cash flow of N when we exchange the principal, this does not affect the value of the swap.

The fixed leg payments are shown above the line as positive cash flows and the floating leg payments are shown below the line as negative cash flows. Note that the first floating payment in Figure III.1.8, denoted F_2 , is already fixed by the three month spot rate at the time of the last payment, denoted $R_{-1,3}$ in the figure. So we use a solid arrow to depict this, whereas the subsequent floating payments have yet to be determined, and we use dotted arrows to depict these.

The floating payments at 5, 8 and 11 months, plus the principal at 11 months, all discounted to the 2-month vertex, together have value \$10 million. In other words, at the next reset date the value of the remaining floating leg payments plus principal, discounted to this reset date, is equal to the principal. This follows using the same logic as we used when valuing FRNs in Section III.1.6.2.³⁹ This is not a special case: it is *always* true that the principal plus floating payments after the first payment date, discounted to the first payment date, together have value equal to the principal, in this case \$10 million.

The implication is that there is no market risk on the floating leg payments after the first payment date. It does not matter how interest rates change between now and the first payment date, the discounted value of the future payments plus principal will always be equal to the principal. Hence, the only market risk on the floating leg of a swap is the risk from a single cash flow at the first payment date, equal to the notional N plus the first interest payment on N , i.e. $N(1 + R/4)$ where R is the fixed interest rate applied to the first interest payment. The market risk on the floating leg is due only to fluctuations in the discount rate as we move towards the next reset date, because the present value of $N(1 + R/4)$ will fluctuate as spot rates change.

Finally, consider the fixed leg in Figure III.1.8, where we have fixed payments determined by the swap rate (6%) and a payment of \$10 million principal at 11 months. This is equivalent to a face value of \$10 million on a hypothetical bond with coupon rate 6% paid quarterly. We conclude that the market risk of the fixed leg of a swap can be analysed as the market risk of a bond with coupon equal to the swap rate.

³⁹ At 11 months the payment of \$10.18830 million discounted to the 8-month vertex must have value \$10 million, because it will be discounted using the forward rate $F_{8,3}$ which is exactly the same as the forward rate used to calculate the interest payment of \$0.18830 million. So we have \$10.16701 million to discount from 8 months to 5 months, i.e. at the rate $F_{5,3}$ which is the same as the rate used to calculate the interest in the first place, and so the value discounted to 5 months is \$10 million. Now we have \$10.14463 million to discount to 2 months, which has value \$10 million.

III.1.7.4 Cross-Currency Swaps

A *basis swap* is an interest rate swap where two floating rates are exchanged on each payment date. These floating rates may be of different maturities, or they may be denominated in different currencies. In the latter case, if we add exchanges of principal, in the opposite direction to interest rate flows initially and in the same direction at maturity, we have the standard or *vanilla* cross-currency swap, also called a *cross-currency basis swap*. It can be regarded as two FRNs in different currencies, as illustrated in the following example.

EXAMPLE III.1.16: A CROSS-CURRENCY BASIS SWAP

On 6 July 2004 bank A entered into a 3-year floating basis swap with bank B based on a notional principal of \$100 million. Bank A pays 6-month US dollar LIBOR (actual/360) and receives 6-month sterling LIBOR (actual/365). Both LIBOR rates are set 6 months in advance. For instance, the interest rates for the payments on 6 July 2005 are fixed at the spot 6-month LIBOR rates on 6 January 2005. It is now 7 July 2007. We can look back over the past 3 years, recording the LIBOR rates and the exchange rate on the payment dates, and they are as shown in Table III.1.15. Explain the mechanisms of this *cross-currency basis swap* and calculate the net payments on each payment date.⁴⁰

Table III.1.15 USD and GBP 6-month LIBOR rates and spot USD/GBP exchange rate

Date	\$LIBOR	£LIBOR	USD/GBP
06/07/04	1.68%	4.73%	1.8399
06/01/05	2.63%	4.57%	1.8754
06/07/05	3.41%	4.39%	1.7578
06/01/06	4.39%	4.36%	1.7718
06/07/06	5.31%	4.60%	1.8365
08/01/07	5.13%	5.22%	1.9347
06/07/07	5.01%	5.86%	2.0131

SOLUTION In effect, bank A is borrowing \$100 million from bank B to make an investment in sterling LIBOR. Bank A will pay the interest on the \$100 million to bank B based on USD LIBOR and will finance these payments (hoping to make a profit) from the income it receives from its investment in sterling LIBOR. At the end of the 3-year period bank A pays back \$100 million to bank B.

The payments from the perspective of bank A are set out in Table III.1.16, where a negative number means a payment and a positive number means money is received. Thus on 6 July 2004, bank A receives \$100 million and, based on the spot exchange rate of 1.8399, invests $\pounds(100/1.8399)$ million = $\pounds 54,350,780$ at sterling 6-month LIBOR. On 6 January 2005 it makes its first payment to bank B on this loan, of:

$$\left(100 \times 0.0168 \times \frac{184}{360}\right) \text{m} = \$858,667.$$

⁴⁰ Note that 6 January 2007 was a Saturday.

But the same day its investment in UK LIBOR pays to bank A:

$$£\left(54,350,780 \times 0.0473 \times \frac{184}{365}\right) = £1,295,961.$$

The spot exchange rate on 6 January 2005 was 1.8754. Hence the net sum received by bank A on that date, in US dollars, is $-858,667 + 1,295,961 \times 1.8754 = 1,571,778$. This is shown in the last column of Table III.1.16.

Table III.1.16 Payments on a cross-currency basis swap

Date	Days	\$ LIBOR	Pay (\$m)	£ LIBOR	Pay (£m)	GBP/USD	Net payment (\$)
06/07/04		1.68%	100	4.73%	-54.350780	1.8399	0
06/01/05	184	2.63%	-0.858667	4.57%	1.295961	1.8754	1,571,778
06/07/05	181	3.41%	-1.322306	4.39%	1.231708	1.7578	842,790
06/01/06	184	4.39%	-1.742889	4.36%	1.202805	1.7718	388,241
06/07/06	181	5.31%	-2.207194	4.60%	1.175109	1.8365	-49,108
08/01/07	186	5.13%	-2.743500	5.22%	1.274042	1.9347	-278,611
06/07/07	179		-102.550750		55.742130	2.0131	9,663,732

We continue to calculate the payments on the \$100 million loan and the receipts on the £54,350,780 investment using the 6-month LIBOR rates that are fixed 6 months before the payment dates, and then calculating the net dollar payment by bank A. In fact, bank A did very well on the deal, since it only had to make a net payment on 6 July 2006 and on 8 January 2007.

It is very common that the main cash flow in a cross-currency swap arises from the change in the exchange rate over the period of the swap. And it is indeed the case here. The US dollar weakened considerably over the 3-year period, so bank A had an advantage holding sterling when it came to repaying the loan on 6 July 2007. On that date, it received interest of

$$£\left(54,350,780 \times 0.0522 \times \frac{179}{365}\right) = £1,391,350$$

and it received its principal of £54,350,780. Based on the exchange rate of 2.0131, it received \$112,214,482. After paying back the loan of \$100 million and the last interest payment on that of:

$$$(100 \times 0.0513 \times \frac{179}{360})m = $2,550,750,$$

bank A receives $\$112,214,482 - \$102,550,750 = \$9,663,732$ when the swap terminates.

In a cross-currency basis swap exchanges of principal take place initially and at maturity. The swap is therefore equivalent to an exchange of two FRNs, one in each currency. For instance, in Example III.1.16 the cash flows corresponded to the positions

- short a \$100 million 3-year floater with 6-month \$LIBOR coupons and
- long a £54,350,780 three-year floater with 6-month £LIBOR coupons.

The value of the swap is therefore equal to the difference in value of these two floaters, as illustrated in the example above.

Cross-currency basis swaps allow borrowers and investors to seek markets where they have a comparative advantage, including foreign markets. For instance, multinational companies with production costs and sales revenues in various countries might use a cross-currency basis swap to convert these cash flows into different currencies. They do not have to be based on two FRNs: for instance, the cash flows could be based on fixed interest rates in either currency.⁴¹

III.1.7.5 Other Swaps

In some interest rate swaps the floating payments at the end of a certain period are not determined only by the forward rate prevailing over that period. For instance, the payments could be based on a rate that has a tenor different than the time between payments. It may also be that the reset dates are not equal to the payment dates. Since it is no longer true that discounting the floating payments plus principal to next payment date always gives the value of the principal, the valuation of these swaps is much more complex. The swap rate will depend on the volatilities and correlations of the forward rates with tenor equal to the time between payments, and the interest rates used to determine the floating payments.

Another popular type of swap is a *total return swap*, in which two parties exchange returns on one asset for returns on another, including the capital gains or losses at every payment date. Two common types of total return swaps are the contract for difference and equity swap:

- A *contract for difference* is a short term total return swap on a share versus an interest rate. Here there is a single terminal payment linked to a principal N where the price appreciation of an investment of N in the share, plus any dividend, is swapped for the (fixed) interest on N .
- An *equity swap* is a long term total return swap on a share versus an interest rate, but now there are regular settlements. For instance, at time $t + 1$ the equity payer pays the total return on N_t invested in a share to the rate payer, who in return pays a floating rate plus a spread on N_t where N_t is the market value of the equity position at the beginning of the time interval.

These swaps are used by institutions, such as hedge funds, that require short equity positions. Alternatively, a fund manager may simply wish to hedge shares that he cannot sell against a fall in price. This second application is illustrated in the following example.

EXAMPLE III.1.17: A SIMPLE TOTAL RETURN SWAP

A UK fund manager thinks the price of Tesco shares will drop over the next month. So he borrows funds to buy 1 million shares in Tesco for 1 month. He uses this holding to enter a contract for difference with a bank. Tesco is currently quoted at 330–331p per share, LIBOR rates are flat at 4.5%, and the bank quotes 1-month LIBID at +40/– 50 basis points.⁴² After a month, Tesco is quoted at 320–321p. How much did the fund manager make on the contract?

⁴¹ In which case that leg of the swap is valued like a normal fixed coupon bond. In this case we have a cross-currency swap without basis risk.

⁴² Interest payments in contracts for difference are at the *London Interbank Bid Rate* (LIBID) which is quoted at LIBOR plus a spread for the buyer and minus a spread for the seller. So the UK fund manager, the seller of the contract, will receive interest of 4.5% less 50 basis points, i.e. 4.0%, in this case.

SOLUTION The principal is based on 1 million shares at 330p, hence $N = £ 3.3$ million. Hence the interest received over the 1-month period is $£ 3.3 \text{ million} \times 0.04/12 = £ 11,000$. In return the fund manager pays the bank the total return on the shares, but since they fell in price, the bank must pay the fund manager this negative return. Selling 1million shares at 330p and buying back at 321p makes a profit of $£ 90,000$ for the fund manager. So far, his total profit on the contract is $£ 101,000$, but we still have to subtract the funding cost for buying 1 million shares for 1 month. Assuming borrowing at LIBOR, the funding cost of the initial position is $£ 3.3 \text{ million} \times 0.045/12 = £ 12,375$. So the net gain to the fund manager was $£ 88,625$.

Many other non-standard interest rate swaps are traded, for instance:

- *constant maturity swap* – between a LIBOR rate and a particular maturity swap rate;
- *constant maturity treasury swap* – between a LIBOR rate and a particular maturity treasury bond yield;
- *extendable/puttable swaps* – one party has the option to extend or terminate the swap;
- *forward starting swap* – interest payments are deferred to a future date;
- *inflation linked swap* – between LIBOR and an inflation index adjusted rate;
- *LIBOR in arrears swap* – LIBOR fixings are just before each interest payment;
- *variable principal swap* – principal changes during the life of the swap;
- *yield curve swap* – between rates for two maturities on the yield curve;
- *zero coupon swap* – between a fixed or floating rate and a fixed payment at maturity.

Readers are recommended to consult Flavell (2002) for further details.

III.1.8 PRESENT VALUE OF A BASIS POINT

There is no such thing as a yield on a swap. That is, it is not possible to back out a unique discount rate which, when applied to all the cash flows, gives the present value of the swap. A unique yield exists only when all the cash flows are positive, but the cash flows on a swap can be negative. And since we cannot define the yield, we cannot define duration or convexity for a swap. So what should we use instead as a measure of sensitivity to interest rates?

In this section we introduce an interest rate sensitivity called the present value of a basis point move, or ‘PV01’ or ‘PVBP’ for short. It is the change in the present value of a sequence of cash flows when the yield curve shifts down by one basis point, i.e. when all zero coupon rates are decreased by one basis point. So the PV01 of a bond portfolio is almost, but not quite, the same as its value duration. The importance of PV01 is that it can be used for *any* sequence of cash flows, positive or negative. After explaining how to calculate PV01, including useful approximations under discrete and continuous compounding of interest rates, we end the section with a discussion on the nature of interest rate risk.

III.1.8.1 PV01 and Value Duration

In the case of a single bond, the PV01 measures *the absolute change in the value of the bond for a fall of one basis point in market interest rates*. More generally, we can define the PV01 for *any* sequence of cash flows given an associated set of spot rates of maturities matching these cash flows. We use the vector notation

$$\mathbf{c} = (C_{T_1}, \dots, C_{T_n})' \text{ and } \mathbf{r} = (R_{T_1}, \dots, R_{T_n})'$$

for the cash flows and for the interest rates, respectively. Denote the present value of the cash flows $\mathbf{c}$ based on the discount rates $\mathbf{r}$ by $PV(\mathbf{c}, \mathbf{r})$. Let $\mathbf{r}^- = \mathbf{r} - 0.01\% \times \mathbf{1}$ where $\mathbf{1}$ is the vector with all elements equal to 1. That is, $\mathbf{r}^-$ denotes the interest rates when each rate is shifted down by one basis point. Now we define:

$$PV01(\mathbf{c}, \mathbf{r}) = PV(\mathbf{c}, \mathbf{r}^-) - PV(\mathbf{c}, \mathbf{r}). \quad (\text{III.1.56})$$

For a single bond, the PV01 is similar to the value duration, which is the absolute change in the present value of the bond per unit absolute change in the bond *yield*. For a single bond, the PV01 and the value duration will only differ in the fifth decimal place or so. But when a cash flow corresponds to an entire portfolio of bonds, there is a subtle difference between value duration and the PV01. The PV01 is the *exact* cash flow sensitivity to a parallel shift in the zero coupon yield curve, whereas value duration is the *approximate* cash flow sensitivity to a parallel shift in the zero curve. The two would only coincide if a shift in the zero curve caused the yields on all the bonds in the portfolio to change by the same amount. Clearly this is very unlikely, so for a bond portfolio PV01 is a more precise measure of interest rate sensitivity than its value duration.

Moreover, value duration is not a concept that can be extended to cover all interest rate sensitive instruments other than bonds.⁴³ By contrast, the PV01 is very easy to calculate for *any* portfolio. We merely have to represent each instrument as a series of cash flows, net these cash flows, and then we can find the PV01 using the simple method described below. In short, PV01 is a more accurate and more general measure of interest rate sensitivity than value duration.




Figure III.1.9 $\delta 01$ as a function of maturity

Consider a single cash flow at time T , where R_T is the associated discount rate. For simplicity we assume that T is an integral number of years, but the extension to non-integer values of T should be obvious from our definitions in Section III.1.2.2. We define $\delta 01$, the *basis point sensitivity of the discount factor*, as

$$\delta 01_T = [(1 + R_T - 0.01\%)^{-T} - (1 + R_T)^{-T}]. \quad (\text{III.1.57})$$

⁴³ In some cases it can be extended, for instance we can define a discrete approximation called the *effective duration* for a convertible bond.

This is the change in the discount factor when its associated interest rate falls by one basis point. Figure III.1.9 depicts $\delta 01_T$ as a function of T for different levels of r , in each curve assuming the term structure of interest rates is flat. This shows that typically, $\delta 01$ will be less than 10 basis points, that $\delta 01$ increases with maturity and that this increase is more pronounced for low levels of interest rates (e.g. 2.5%) than for high levels of interest rates (e.g. 10%).

Using this notation, we can express the PV01 of a cash flow C_T at time T as

$$PV01_T = C_T \times \delta 01_T,$$

and the total PV01 of all the cash flows is the sum of these,

$$PV01 = \sum_{i=1}^n C_{T_i} \times \delta 01_{T_i}. \quad (\text{III.1.58})$$

EXAMPLE III.1.18: CALCULATING THE PV01 OF A SIMPLE BOND

Calculate the PV01 for £1 million notional in a 6% 4-year annual bond when the zero coupon rates are 4.5%, 4.75%, 4.85% and 5% at maturities of 1, 2, 3 and 4 years, respectively. Compare this with the value duration of the bond.⁴⁴

SOLUTION This bond was examined in Example III.1.7, where we computed its price, yield and Macaulay duration. Table III.1.17 extends the calculations of Table III.1.8 to compute the PV01 and the value duration of the bond. For the PV01, the present value of each cash flow is calculated twice, first based on the zero coupon yield curve and then again based on the same curve shifted down one basis point. The PV01 is the difference between the two values, i.e. £36,312.75. This tells us that if *zero coupon rates* were to shift down by one basis

Table III.1.17 PV01 for a bond

Years to maturity	Cash flow (£m)	Interest rate	Present value	Interest rate	Present value	PV01
1	6	4.5%	5.7416	4.49%	5.7422	549.49
2	6	4.75%	5.4682	4.74%	5.4692	1,044.19
3	6	4.85%	5.2053	4.84%	5.2068	1,489.64
4	106	5%	87.2065	4.99%	87.2397	33,229.42
			103.6216		103.6579	36,312.75
Years to maturity	Cash flow (£m)	Bond yield	Present value	Bond yield	Present value	Value duration
1	6	4.98%	5.7154	4.97%	5.7160	544.49
2	6	4.98%	5.4443	4.97%	5.4454	1,037.37
3	6	4.98%	5.1861	4.97%	5.1876	1,482.32
4	106	4.98%	87.2757	4.97%	87.3090	33,262.41
			103.6216		103.6579	36,326.59

⁴⁴ Note that we have computed value durations here as the actual price change for a 1 *basis point* change in yield, rather than with respect to a 1% change in yield. This is so that the value duration and the PV01 have similar orders of magnitude, as demonstrated by Table III.1.17.

point the bond price would rise by 0.03631, i.e. from 103.6216 to 103.6579, and we would make a profit of £36,312.75 per million pounds invested in the bond.

We also calculate the value duration by shifting the bond's *yield* down by one basis point. In Example III.1.7 we found that the bond had yield 4.98% based on a price of 103.62. The present value is calculated again, now based on a yield of 4.97%, and the value duration is the difference between the original value and the new value, i.e. £36,326.59. This tells us that if the *bond yield* were to fall by one basis point then the price of the bond would rise by 0.03633, i.e. from 103.6216 to 103.6579, and we would make £36,326.59 per million pounds invested in the bond.

This verifies that, for a single bond, the PV01 and the value duration are almost exactly the same. However, for reasons made clear above, PV01 is the preferred interest rate sensitivity measure.

III.1.8.2 Approximations to PV01

In this subsection we derive a first order approximation to $\delta 01$ that allows one to approximate the PV01 very quickly. If T is an integral number of years, the discount factor is $\delta_T = (1 + R_T)^{-T}$ and so

$$\delta 01_T \approx -\frac{d\delta_T}{dR_T} \times 10^{-4} = T(1 + R_T)^{-(T+1)} \times 10^{-4}.$$

This shows that a useful approximation to the PV01 of a cash flow C_T at time T , when T is an integral number of years, is

$$PV01_T \approx T C_T (1 + R_T)^{-(T+1)} \times 10^{-4}. \quad (\text{III.1.59})$$

Similar first order approximations can be derived when T is not an integral number of years. For instance, when T is less than 1 year,

$$\delta_T = (1 + TR_T)^{-1}$$

and

$$\delta 01_T \approx T(1 + TR_T)^{-2} \times 10^{-4}. \quad (\text{III.1.60})$$

This is simple enough, but the expression becomes a little more complex when T is greater than 1 year but not an integer.

Because of the unwanted technical details associated with working with discretely compounded rates it is common practice to price fixed income instruments and futures with continuously compounded rates using the relationship between discretely and continuously compounded rates derived in Section III.1.2.3. An approximation that is equivalent to (III.1.59) but that applies for continuously compounded interest rates of *any* maturity is derived by differentiating the continuous discounting factor $\delta_T = \exp(-r_T T)$ with respect to the continuously compounded interest rate r_T of maturity T . Thus

$$\delta 01_T \approx -\frac{d\delta_T}{dr_T} \times 10^{-4} = T \exp(-r_T T) \times 10^{-4} = T \delta_T \times 10^{-4}. \quad (\text{III.1.61})$$

Hence, a simple approximation to the PV01 for a cash flow at *any* maturity, under continuous compounding, is

$$PV01_T \approx T C_T \exp(-r_T T) \times 10^{-4}. \quad (\text{III.1.62})$$

EXAMPLE III.1.19: CALCULATING PV01

Suppose a cash flow has been mapped to vertices at 1 and 2 years with €10 million mapped to the 1-year vertex and €5 million mapped to the 2-year vertex. Suppose the 1-year zero rate is 4% and the 2-year zero rate is 4.5%. Use the above approximation to calculate the PV01 at each vertex and hence find the total PV01 of the mapped cash flow.

SOLUTION From (III.1.59),

$$PV01_1 \approx 10^{-4} \times 10 \times 10^6 \times (1.04)^{-2} = 1000 \times (1.04)^{-2} = 924.56,$$

$$PV01_2 \approx 10^{-4} \times 2 \times 5 \times 10^6 \times (1.045)^{-2} = 1000 \times (1.045)^{-2} = 876.30,$$

so the total PV01 is approximately €1800.86. For comparison we also compute the exact solution, which is summarized in Table III.1.18. With the original interest rates the present value of the cash flow is €14.194 million, but if interests fall by one basis point the present value will become €14.1958 million. Hence, the exact PV01 of the cash flow is €1801.07.

Table III.1.18 PV01 for a cash flow

Maturity (years)	Cash flow (€m)	Interest rate	δ	PV (€m)	Interest rate	δ	PV (€m)	PV01(€)
1	10	4.0%	0.9615	9.6154	3.99%	0.9616	9.6163	924.65
2	5	4.5%	0.9157	4.5786	4.49%	0.9159	4.5795	876.42
Totals				14.1940				14.1958
								1801.07

In the above example the approximation error from using (III.1.59) instead of exact valuation is only 0.01%. These approximations are clearly very accurate and they will prove useful later when we show how to perform a PV01 invariant cash flow mapping, in Section III.5.3.

III.1.8.3 Understanding Interest Rate Risk

Consider a cash flow of C_T dollars at time T . We know the present value of this cash flow today: it is just C_T dollars discounted to today using the appropriate discount rate, i.e. the spot zero coupon rate of maturity T . We also know its present value when we receive the cash at time T , as this requires no discounting. But at any time t between now and maturity, we do not know the present value exactly. There is an uncertainty on this future discounted value that depends on (a) how much the discount rate changes and (b) how sensitive the value is to changes in the discount rate. How do we assess these?

- (a) The standard deviation of the discount rate represents the possible changes in the future discount rate, as seen from today. If σ denotes the volatility of this forward rate then, assuming changes in the forward rate are *independent and identically distributed*, we may use the square-root-of-time rule to find the t -period standard deviation of the forward rate as $\sigma_{t,T-t}\sqrt{t}$.⁴⁵

⁴⁵ The square-root-of-time rule is explained in Section II.3.1.

- (b) The best forecast of the appropriate discount rate at time t is the forward zero coupon rate $F_{t,T-t}$ starting at time t with tenor $T - t$. The expected discounted value at time t of the cash flow at time T is C_T dollars discounted by this forward rate. The sensitivity of this value at time t to movements in the forward rate is measured by the PV01 that is given (approximately) by (III.1.59) with maturity $T - t$, assuming T and t are integers, i.e.

$$\text{PV01}_{t,T} \approx 10^{-4} C_T (1 + F_{t,T-t})^{-(T-t+1)} (T - t).$$

The interest rate risk of a cash flow of C_T dollars at time T refers to the uncertainty about the discounted value of this cash flow at some future time t . It can be measured by the standard deviation of this future discounted value, which is the product of (a) and (b) above. It is easier to express this assuming continuous compounding, since then for *any* values of T and t :

$$\text{StDev}(\text{PV}_t) \approx 10^{-4} \sigma_{t,T-t} \sqrt{t} (T - t) C_T \exp(-f_{t,T-t}(T - t)). \quad (\text{III.1.63})$$

where $\sigma_{t,T-t}$ is the volatility of the continuously compounded forward rate, $f_{t,T-t}$.

EXAMPLE III.1.20: STANDARD DEVIATION OF FUTURE PV

The current zero-coupon forward rates, starting in m months and ending in 12 months, and their volatilities are as shown in Table III.1.19 for $m = 1, 2, \dots, 11$. Calculate the discounted value in m months' time of a cash flow of £5 million received in 1 year. Also calculate the standard deviations of these discounted values.

Table III.1.19 Forward rates and their volatilities (in basis points)

m (months)	1	2	3	4	5	6	7	8	9	10	11
Forward rate (%)	4.70	4.75	4.80	4.85	4.80	4.85	4.80	4.75	4.70	4.65	4.60
Volatility (bps)	15	18	20	25	30	35	35	35	40	40	40

SOLUTION The expected discounted values of the cash flow in m months time are calculated by discounting at the appropriate forward rate, given in the second row of Table III.1.19. For instance the expected discounted value of the cash flow in 3 months' time is $5 \times \exp(-0.048 \times 0.75) = £4.82$ million, since there are 9 months of discounting. The corresponding PV01 is calculated using the approximation (III.1.62) and this is shown in the third row of Table III.1.20. For instance:

$$\text{PV01 of expected value in 3 months} = 100 \times 0.75 \times 5 \exp(-0.048 \times 0.75) = £362.$$

Note that we multiply by 100 because the cash flow is measured in millions and the PV01 refers to a movement of one basis point. Now the standard deviation of the future discounted value of the cash flow is $\text{PV01} \times \sigma \sqrt{t}$ where σ denotes the volatility of the forward rate. For instance:


$$\text{Standard deviation of value in 3 months} = 362 \times 20 \times \sqrt{0.25} = £3617.$$

Equivalently, to calculate the standard deviation directly we could have used (III.1.63) with $t = 1, 2, \dots, 12$. Take care to multiply the result by 100 to obtain the correct units of measurement, as we did for the PV01. Either way, we obtain the standard deviations of the future discounted values that are shown in the last row of Table III.1.20.

Table III.1.20 Expectation and standard deviation of future PV

Time (months)	1	2	3	4	5	6	7	8	9	10	11
Expected value (£m)	4.79	4.81	4.82	4.84	4.86	4.88	4.90	4.92	4.94	4.96	4.98
PV01 (£)	439	400	362	323	284	244	204	164	124	83	42
StDev of value (£)	1901	2943	3617	4658	5492	6039	5459	4688	4280	3019	1509

In Figure III.1.10 we use $\pm 1.96 \times \text{StDev}$ to obtain approximate 95% confidence limits for the value of the cash flow in the future. Note that there is no uncertainty about the discounted values now, and when we receive the cash flow. But between now and the time that the cash flow is received (shown on the horizontal axis) there is an uncertainty that arises from possible changes in the discount rate that will apply in the future, when we calculate the cash flow's present value. This is interest rate risk.

**Figure III.1.10** Uncertainty about future values of a single cash flow

In general, the interest rate risk of a fixed income portfolio is the uncertainty about the discounted value of all the portfolio's future cash flows, where this value is measured at some future date – the *risk horizon*. The uncertainty arises from unpredictable movements in the discount rate between now and the risk horizon, i.e. we do not know exactly what the discounted value of the cash flows will be in the future, because we do not know what the discount rate will be at some future point in time.

Our perception of interest rate risk depends on whether the interest is fixed or floating and it also depends upon the accounting framework. Fixed income instruments are entered into the trading book at their mark-to-market value. Their interest rate risk depends on the evolution of discount rates, as explained above. But if entered into the banking book, where accrual accounting is applied, fixed income instruments are *not* marked to market. Cash flows are valued at the time of payment, so fixed income instruments have no interest rate risk in accrual accounting. On the other hand, floating rate instruments have virtually no

interest rate risk in mark-to-market accounting, as we have shown in Section III.1.6.2. But they do have interest rate risk in accrual accounting, because the cash flow on any future payment date is uncertain.

Hence, the perceptions of interest rate risk of fixed and floating instruments differ depending on the accounting framework. In accrual accounting floaters are risky but fixed rates are not, whilst the opposite is true under mark-to-market accounting. It is this difference that drives the swap market.

III.1.9 YIELD CURVE FITTING

In this section we review the techniques that can be used to obtain a yield curve from the market prices of fixed income instruments such as short term discount bonds, FRAs and swaps. Section III.1.9.1 discusses the suitability of each type of instrument for yield curve fitting, and thereafter we focus on the influence that the chosen fitting technique has on the statistical properties of spot and forward interest rates. The simplest way to obtain a yield curve is to use a *bootstrap*. Section III.1.9.2 demonstrates how securities prices can be used to bootstrap the yield curve using a simple empirical example. The problem with bootstrapping is that even a small amount of noise in these securities prices can result in large spikes in the forward curve, especially at longer maturities. Therefore if the yield curve is to be used to make inferences on the volatility and correlation structure of interest rates it is better not to derive the yield curve using the bootstrap technique.

Sections III.1.9.3 and III.1.9.4 describes the semi-parametric and parametric models for yield curve fitting that are used by the Bank of England,⁴⁶ the European Central Bank⁴⁷ and the US Federal Reserve.⁴⁸ These techniques are also applied to the LIBOR market, as discussed in the comprehensive and excellent book by James and Webber (2000). LIBOR rates for a number of currencies are provided on the BBA CD-ROM.⁴⁹ The case study in Section III.1.9.5 demonstrates how the volatilities and correlations of forward LIBOR rates depend very much on whether we obtain the yield curve using natural splines, cubic splines or a parametric yield curve model.

III.1.9.1 Calibration Instruments

The US government short curve is obtained from the prices of Treasury bills, FRAs, swaps and liquid coupon bearing bonds in the AA rating category. The UK government curves are derived from UK government bonds (also called *gilts*), gilt sale and repurchase transactions (*gilt repos*), interbank loans, short sterling futures, FRAs and swaps. A LIBOR yield curve is usually constructed using a combination of spot LIBOR rates, FRAs, futures and swap rates. One would typically use the 1-month to 12-month LIBOR rates to estimate the short end of the curve. But LIBOR rates are not available for maturities longer than 1 year so other instruments with the same credit risk and liquidity as LIBOR rates need to be used for

⁴⁶ See <http://www.bankofengland.co.uk/statistics/yieldcurve/index.htm>.

⁴⁷ See <http://www.ecb.int/stats/money/yc/html/index.en.html>.

⁴⁸ See <http://www.federalreserve.gov/releases/h15/update/>.

⁴⁹ See <http://www.bba.org.uk>.

building the longer end of the LIBOR curve. To fit LIBOR rates in the medium and long range we could use interest rate futures, FRAs and/or swaps.

There are many combinations of bonds, futures, FRAs and swaps that can be used to construct a yield curve. Of course, if the data were perfect and markets were arbitrage free, the choice of instruments should not matter. However, given the inevitable noise in market data, the choice of securities *does* have an impact on the shape of the yield curve. We now discuss the relative advantages and disadvantages of using futures, FRAs and swaps for yield curve fitting.

Interest rate futures and FRAs typically cover a period up to 2 years. Both kinds of contracts enable an investor to lock in a rate of return between two dates and therefore provide information about the corresponding forward LIBOR rates. The main difference between these contracts is that the settlement on an FRA occurs at the contract maturity, while futures positions are marked to market on a daily basis, resulting in a stream of cash flows between the parties along the whole life of the contract. Since these cash flows are a function of the prevailing level of the LIBOR rate, futures provide *biased* information about forward rates. The bias results from the fact that the short party experiences a cash outflow in the event of an interest rate increase and a cash inflow in the event of a fall in interest rates. Therefore the short party will systematically seek financing when interest rates rise and invest the cash flows from marking to market as interest rates fall. The opposite is true about the long party on the futures contract. Thus the equilibrium futures price will be biased upward, in order to compensate the short party. The importance of the bias depends on the length of the contract and the LIBOR rate volatility expectations. Strictly speaking, one would need to use an interest derivative pricing model to value a futures contract. From this perspective, the use of futures as an input to a yield curve fitting model should be avoided.

Unlike futures, FRAs could be seen as instruments providing unbiased information about the forward LIBOR rates. However, the liquidity of FRAs is typically lower than that of the LIBOR rates and futures. Therefore, FRA quotes may be stale and fail to reflect the changes in the yield curve. Consequently, neither futures nor FRAs are ideal instruments to use for LIBOR yield curve estimation.

The information about the long end of the yield curve can be obtained from interest rate swaps, since these are liquid instruments that extend up to 30 years. For the purposes of yield curve fitting we decompose a swap into an exchange of a fixed coupon bond for a floating bond paying prevailing LIBOR rates.⁵⁰ At initiation the value of a swap is 0 and the value of a floater is equal to the notional, which means that the price of the fixed rate bond paying the swap rate must be 100 at the initiation of the swap. In other words, the swap rate can be regarded as a coupon rate on a par-coupon bond of a corresponding maturity, as shown in Section III.1.7.3.

III.1.9.2 Bootstrapping

Bootstrapping is the term we apply to an iterative coupon stripping technique, which is illustrated in the following example.⁵¹

⁵⁰ As explained in Section III.1.7.

⁵¹ See Miron and Swannell (1991) for further details.

EXAMPLE III.1.21: COUPON STRIPPING

Calculate zero coupon rates from the market prices of four zero coupon bonds and two coupon bearing bonds shown in Table III.1.21.

Table III.1.21 Six bonds

Maturity	1 month	2 months	3 months	12 months	2 years	3 years
Coupon	0	0	0	0	6%	10%
Market price	99.5	99.1	98.7	95	101	112

SOLUTION Dividing the price of a zero bond by 100 gives the discount factor and then putting $y_n = n^{-1}(\delta_n^{-1} - 1)$ when $n \leq 1$, gives the zero coupon yields shown in Table III.1.22.

Table III.1.22 Bootstrapping zero coupon yields

Maturity	1 month	2 month	3 month	12 months	2 years	3 years
Coupon	0	0	0	0	6%	10%
Price	99.5	99.1	98.7	95	101	112
Discount factor	0.995	0.991	0.987	0.95	0.90	0.85
Zero coupon yield	6.03%	5.45%	5.27%	5.26%	5.46%	5.56%

This table also shows zero coupon yields for the 2-year and 3-year maturities. These are obtained via a bootstrapping procedure as follows:

1. The price, cash flows and discount factors on the 2-year bond are related as

$$101 = 6 \times \delta_1 + 106 \times \delta_2.$$

We know δ_1 from the 12-month zero bond, so we can obtain δ_2 .

2. Moving on to the 3-year bond, now we have

$$112 = 10 \times \delta_1 + 10 \times \delta_2 + 110 \times \delta_3$$

which, knowing δ_1 and δ_2 , gives δ_3 .

3. Once we know the discount factors we can obtain the zero coupon yields, using

$$y_n = \delta_n^{-1/n} - 1$$

We can generalize the method used in the above example to state a general algorithm for bootstrapping a yield curve as follows:

1. Take a set of T liquid instruments.
2. Map the cash flows to a set of T standard maturities (see Section III.5.3).
3. Denote the mapped cash flows on the k th instrument by $\mathbf{c}_k = (C_{k1}, \dots, C_{kT})'$ and let $\mathbf{C}$ be the $T \times T$ matrix of cash flows, with k th column $\mathbf{c}_k$.
4. Take the market price of each instrument and label the $T \times 1$ vector of market prices $\mathbf{p} = (P_1, P_2, \dots, P_T)'$.

5. Denote the discount factor at maturity i by δ_i and let $\boldsymbol{\delta} = (\delta_1, \delta_2, \dots, \delta_T)'$ be the $T \times 1$ vector of discount factor.
6. Then $\mathbf{p} = \mathbf{C}\boldsymbol{\delta}$ or equivalently, $\boldsymbol{\delta} = \mathbf{C}^{-1}\mathbf{p}$. Hence, the market prices and cash flows of the instruments together determine the discount factor.

This algorithm has two shortcomings:

- It is restricted to using the same number of instruments as maturities in the yield curve.
- We only obtain zero coupon yields at the same maturities as the instruments, so we cannot price instruments with other maturities.

The second problem, that we only obtain discount rates at maturities for which liquid instruments exist, can be solved by proceeding as follows:

7. Interpolate between the discount factors (e.g. using splines, as explained in the next subsection) to obtain discount rates for every maturity.
8. Obtain the zero coupon yield of any maturity n as described above.

But still the bootstrap has tendency to *overfit* the yield curve so any noise in the original data is translated into the bootstrapped forward rates. This often results in spikes in the forward curve, especially at longer maturities, which gives a misleading indication about the volatility and correlation structure of forward rates.

When fitting a yield curve we should take into account the market prices of *all* liquid instruments having the same credit rating as the curve, and typically there will be more than one instrument for each maturity. In practice, more instruments than maturities are used for yield curve construction. In that case the inverse cash flow matrix $\mathbf{C}^{-1}$ cannot be found exactly and instead of bootstrapping we must use a best fitting technique, such as minimizing the root mean square error (RMSE) between the prices based on the fitted curve and the market prices. We describe the two most popular yield curve fitting methods in the following subsections.

III.1.9.3 Splines

Spline-based yield curve techniques fit a curve to the data, and this curve is composed of many segments that can move almost independently between the fixed *knot points*, save for some constraints imposed to ensure that the overall curve is smooth. To limit the number of parameters used it is important to use a low order polynomial to model the curve between the knot points, hence the popularity of *cubic splines* which were first applied to yield curve fitting by McCulloch (1975). Cubic spline interpolation was introduced and applied to interpolate between interest rates of different maturities in Section I.5.3.3.⁵²

Cubic splines may also be applied to fit a yield curve by fitting the corresponding discount factors, or by fitting the market prices of all securities in the calibration set. The discount factors are linear in the parameters and the calibration can be done by a simple regression where the objective is to minimize the weighted squared differences between the model prices of the securities in the calibration set and their market prices. However, the most straightforward spline fitting method, natural cubic splines, are commonly found to be numerically unstable.⁵³

⁵² See also Lancaster and Salkauskas (1986).

⁵³ See Anderson and Sleath (1999) for further details.

An alternative is to use *basis splines* (B-splines), as in Steely (1991), where each segment of the yield curve between the knot points $\{x_1, \dots, x_m\}$ is built using a linear combination of some fundamental curves. To define these fundamental curves, which are called *basis functions*, we set:

$$B_{i,1}(\delta) = \begin{cases} 1, & \text{if } x_i \leq \delta < x_{i+1}, \\ 0, & \text{otherwise,} \end{cases}$$

for $i = 1, \dots, m-1$. Then the n th order basis functions, for $n = 2, 3, \dots, N$, are obtained by setting

$$B_{i,n}(\delta) = \left(\frac{\delta - x_i}{x_{i+n-1} - x_i} \right) B_{i,n-1}(\delta) + \left(\frac{x_{i+n} - \delta}{x_{i+n} - x_{i+1}} \right) B_{i+1,n-1}(\delta), \quad (\text{III.1.64})$$

where δ denotes the discount factor.

Clearly the order of the basis functions used to fit the curve between knot points can be smaller when there are many knot points than when there are only a few knot points. But as the number of knot point increases, the smoothness of the fitted yield curve will deteriorate. Hence, there is a trade-off between using enough knot points to ensure a good fit with low order basis functions and using so many knot points that the fitted curve is not smooth. However, it is possible to impose a roughness penalty on the optimization algorithm,⁵⁴ and our case study below demonstrates why this is essential.

III.1.9.4 Parametric Models

Introduced by Nelson and Siegel (1987) and later extended by Svensson (1994), parametric yield curve fitting models impose a functional form on the instantaneous forward rate curve that captures its typical ‘humped’ shape. The *Nelson and Siegel model* assumes that the instantaneous forward rate of maturity n is parameterized as

$$f(n; \beta_0, \beta_1, \beta_2, \tau) = \beta_0 + \left(\beta_1 + \beta_2 \left(-\frac{n}{\tau} \right) \right) \exp \left(-\frac{n}{\tau} \right). \quad (\text{III.1.65})$$

The curve (III.1.65) has only one ‘hump’, which reflects the typical shape of the interest rate term structure. However, sometimes the term structure develops *two* humps, and in that case the parameterization (III.1.65) is too restrictive. The *Svensson model* has two additional parameters, which allow for an additional hump, with

$$\begin{aligned} f(n; \beta_0, \beta_1, \beta_2, \beta_3, \tau_1, \tau_2) = & \beta_0 + \left(\beta_1 + \beta_2 \left(-\frac{n}{\tau_1} \right) \right) \exp \left(-\frac{n}{\tau_1} \right) \\ & + \beta_3 \left(-\frac{n}{\tau_2} \right) \exp \left(-\frac{n}{\tau_2} \right). \end{aligned} \quad (\text{III.1.66})$$

In both (III.1.65) and (III.1.66) the short rate is $\beta_0 + \beta_1$ and the curve is asymptotic to β_0 at the long end. The discount factors are non-linear in the parameters and the model calibration requires a non-linear least squares algorithm. In the case study below we have used the Levenberg–Marquardt algorithm, described in Section I.5.4.1, with the objective of minimizing the sum of the squared differences between the model and the market prices of the instrument calibration set.

⁵⁴ See Fisher et al. (1995), Waggoner (1997) and Anderson and Sleath (1999).

The Federal Reserve Board has made available to the public the entire US Treasury yield curve from 1961.⁵⁵ The spot curve is estimated daily using the Svensson model, and spot and forward rates, monthly up to 30 months, are provided. The Excel spreadsheet also gives the Svensson parameter estimates for each day so that users can infer rates of any maturities from the key rates at monthly maturities.

III.1.9.5 Case Study: Statistical Properties of Forward LIBOR Rates⁵⁶

In this study we compare the applications of B-splines and the Svensson model, for calibration of the UK LIBOR yield curve. Both yield curve fitting models were applied to 703 daily observations on UK LIBOR rates, FRAs and swap rates between October 1999 and June 2002⁵⁷. For each day all the available market instruments were used to build the discount curve and the associated spot curve. These were then used to obtain model LIBOR rates, FRAs and swap rates. The model rates were then compared to the market data to obtain the RMSE.

Figures III.1.11 and III.1.12 compare the yield curves calibrated using the two procedures. There are noticeable differences between the two curves, especially at the long end and especially when rates are volatile (for instance, during the last 6 months of 2001). At the short end, where LIBOR rates are observed in the market, the differences are usually tiny and of the order of 5–10 basis points, but at the very long end where the market instruments consist of only a few long term swaps, the differences can exceed 50 basis points at times, as shown in Figure III.1.13.




Figure III.1.11 UK LIBOR curve (Svensson model)

⁵⁵ See Gurkaynak et al. (2006) and <http://www.federalreserve.gov/pubs/feds/2006/200628/200628abs.html>.

⁵⁶ Many thanks to my PhD student Dmitri Lvov for providing these results.

⁵⁷ The following contracts were used: 1-month to 12-month LIBOR rates; 3-month FRAs starting in 2, 3, ..., 9 months; 6-month FRAs starting in 1, 2, ..., 6, 12 and 18 months; 9-month FRA starting in 1 month; 12-month FRAs starting in 2, 3, 6, 9 and 12 months; and swap contracts maturing in 2, 3, ..., 10 years.




Figure III.1.12 UK LIBOR curve (B-splines)




Figure III.1.13 Difference between Svensson rates and B-spline rates

Comparison of the Models' Goodness of Fit

The average RMSE over the entire sample period were very similar: 2.59 for the B-spline and 2.55 for the Svensson model. However, the time series plots of these RMSEs reveal some interesting characteristics. Figure III.1.14 compares the RMSE obtained from the two models for every day during the sample.⁵⁸ The lowest errors occurred during the year

⁵⁸ We tested several different B-spline and natural cubic spline procedures with different choices of knot points. When the knot points were the same, all splines on discount factors produced virtually identical RMSE on every day during the sample. So henceforth we present results for the B-spline on discount factors, and the optimal choice of knots had several at the medium to long maturity.

2000 and greater errors were experienced during the latter half of 2001, when interest rates were unusually volatile, especially using the Svensson model. In view of this time variation it is difficult to draw general conclusions about the relative performance of in-sample fit from the two yield curve fitting procedures. However, further information can be obtained from the historical volatilities and correlations of daily changes in log forward LIBOR rates, as these differ considerably depending on the choice of the technique used to fit the yield curve.




Figure III.1.14 RMSE errors from the model calibrations

The Influence of Yield Curve Model on Forward Rate Volatilities

In the LIBOR model forward rates are usually assumed to follow correlated *geometric* processes and so we estimate historical forward rate volatilities using log returns, not changes in forward rates.⁵⁹ Figure III.1.15 depicts the historical forward rate volatility term structures estimated using the entire sample based on the two yield curve models.

The implied volatilities of forward rates that are backed out from cap prices often display a hump shape, with maximum volatility usually lying somewhere between the 1- and 2-year maturities. So the empirical evidence on historical forward rates does not appear to support traders' views on interest rates at all. The short maturity forward rates tend to be *less* volatile than long maturity rates. However, to a large extent this is due to the presence of measurement errors, which are greater for longer maturity rates since one observes relatively few data points, and these measurement errors are compounded when spot rates are translated into discrete forward rates. The net result is an increasing volatility of forward rates with respect to maturity.

⁵⁹ See Section III.3.8.5 for further details about the LIBOR model.




Figure III.1.15 Forward rate volatility estimates based on B-splines and Svensson

The forward rates have different volatility characteristics depending on the yield curve fitting model used.

- The Svensson forward rate volatilities are influenced by the quality of fit of a two-hump parametric form. Further analysis that the reader may wish to apply using the data in the case study work book shows that the Svensson model produced highly volatile long rates at the end of 2001, when the goodness-of-fit statistics implied that the Svensson model was giving the worst in-sample fit.
- The B-spline forward rate volatilities are constrained to be lower at the knot points (in order to maximize the goodness of fit, several knots were placed between the 5-year and 8-year maturities). Although the McCulloch natural spline volatilities are not shown, they have similar characteristics to the B-spline volatilities, again being dominated by the choice of knot points.

The Influence of Yield Curve Fitting on Forward Rate Correlations

The correlation structure of the semi-annual forward rates obtained from the Svensson model is related to the functional form imposed on the instantaneous forward rates. The 'hump' shape of the Svensson forward rate curve induces correlations between short and long maturity forward rates that can be higher than correlations between medium and long maturity forward rates. Forward rate correlations do not necessarily decrease monotonically with the maturity spread, nevertheless the Svensson model imposes exactly this sort of smooth pattern on the forward rate correlations. The empirical correlations between Svensson semi-annual forward rates of different maturities are shown in Figure III.1.16 based on equal weighting over the entire sample period. The empirical correlations have a less structured pattern when the yield curve is fitted using a spline model, and the B-spline correlations are shown in Figures III.1.17. Here, correlations are distorted and very much depend on the knot points chosen.




Figure III.1.16 Forward rate correlation estimates (Svensson model). (See Plate 1)




Figure III.1.17 Forward rate correlation estimates (B-spline model). (See Plate 2)

Comparison with Bank of England Forward Rate Curves

Daily and monthly data on UK government and commercial liability yield curves for maturities between 6 months and 25 years, and the short curve for monthly maturities from 1 month to 5 years are available from the Bank of England.⁶⁰ The curves are fitted using a cubic B-spline algorithm with robustness penalty, as explained in Anderson and Sleath (1999).

⁶⁰ See <http://www.bankofengland.co.uk/statistics/yieldcurve/index.htm>

The robustness penalty has the effect of smoothing the curve between the knot points, thus reducing the possibility for ill-behaved volatilities and correlations such as those obtained above from the B-spline procedure. The short forward curve, which contains 60 monthly maturities up to 5 years, is very well behaved: the forward rate volatilities exhibit a humped shape and the forward rate correlations are very smooth. However the entire forward curve, containing semi-annual rates up to 25 years, still has quite a rough volatility and correlation structure, despite the smoothing in their algorithm.




Figure III.1.18 Bank of England forward curve – volatilities




Figure III.1.19 Bank of England forward curve – correlations. (See Plate 3)

For instance, Figures III.1.18 and III.1.19 show the term structure of forward rate volatilities and correlations respectively, based on daily log returns over the period January 2005 to September 2007. As before the volatilities and correlations are based on equally weighted averages of returns and cross products over the entire sample. The robustness penalty does provide some smoothing, but the irregular features in volatility and correlation that result from the use of a B-spline are still apparent.

Conclusions

Goodness-of-fit criteria are inconclusive as to the best fitting yield curve model because this varies over time. We have presented some empirical evidence indicating that B-splines may fit better during volatile periods, but otherwise the Svensson model provides a very close fit. The important learning point to take away from this study is that forward rate volatilities and correlations inherit their structure from the yield curve fitting model: the Svensson forward rate volatilities and correlations are the best behaved since their structure is determined by the exponential functional form of the model; by contrast, in spline methods, volatilities and correlations are distorted by the choice of knot points.

Whilst implied forward rate volatilities that are backed out from the market prices of caps often exhibit a humped shape, we have found no empirical support for hump-shaped volatility term structures in historical forward rates, except at the very short term where the rates are anyway observed in the market. The lack of liquid market instruments of medium to long maturities introduces significant measurement error in calibrated forward rates, so that historical volatilities increase with maturity whatever yield curve fitting model is applied.

Likewise, when historical forward rate correlations are used to calibrate the lognormal forward rate model, their correlation matrix is usually assumed to have a simple parameterization where correlation decreases monotonically with the maturity gap between the futures.⁶¹ Such a parameterization is more appropriate for forward rate correlations that are derived from the Svensson yield curves than those derived from B-spline fitting procedures, since the choice of knot points has a marked effect on forward rate correlations when the yield curve is fitted using spline functions.

III.1.10 CONVERTIBLE BONDS

Convertible bonds are hybrid derivative instruments that simultaneously offer the potential upside of equities and a limited downside, since the investor holds a bond. The pay-off to holding a convertible bond is asymmetric, in that the increase in the bond's value corresponding to an increase in the stock price is greater than the decrease in the bond's value corresponding to an equivalent decrease in the stock price. Another reason why convertible bonds offer attractive investment opportunities is their inherent diversification. A convertible bond offers exposure to equity volatility as well as to equities, interest rates and credit spreads. This diversification benefit is significant since volatility is negatively correlated with share prices and credit spreads, especially during periods of market crisis. Thus convertible

⁶¹ See Alexander (2003).

bonds are a growing asset class and at the time of writing the global market has an estimated notional size of over \$613 billion.⁶² Over half of this is in US bonds, over \$150 billion is in European bonds and \$62 billion is in Japanese bonds.

III.1.10.1 Characteristics of Convertible Bonds

The purchaser of a convertible bond has the option to convert it into shares in the same company, usually after a certain *lock-in period*.⁶³ The *conversion ratio* is

$$\text{Conversion ratio} = \frac{\text{Face value of bond}}{\text{Conversion price of share}}. \quad (\text{III.1.67})$$

For instance, if the face value is \$1000 and the conversion price of the share is \$50 then the conversion ratio is 20. That is, 20 shares are received for each bond redeemed. Often, but not always, the conversion ratio is fixed at the time the bond is issued.

The decision whether to convert depends on the conversion ratio, the current prices of the bond and the share and the credit risk of the issuer. Convertible bonds carry a significant credit risk for two reasons. First, they are often issued by growth firms, i.e. firms that may not have a high credit rating but whose share prices are thought likely to rise in the future. These firms can issue a convertible bond with a relatively low coupon, thus limiting their payments, but the conversion option still makes them attractive to investors. Secondly, convertible bonds are frequently classed as junior or subordinated debt.

The conversion feature provides a floor for the bond price, which varies with the stock price. For example, suppose a bond with face value \$1000 may be converted at any time to 20 shares, and that the stock is currently trading at a price of \$52.50. Then the value of the bond cannot be less than $\$52.50 \times 20 = \1050 so the price of the bond cannot fall below 105 at this time.

Most convertibles have a *put feature* whereby the bond holder has the right to return the bond to the issuer at a fixed *put price*, which is usually well below par. This provides a fixed floor for the price of the convertible, which does not vary over time. For instance, if the put price of the above convertible is 85, then the price can never fall below 85 (otherwise the bond holders would put the bond back to the issuer). Similarly, most convertibles are also *callable* by the issuer on certain dates, at which times the issuer has a right to buy back the convertible at a given *call price* which is usually well above par.⁶⁴ This price provides a fixed cap on the bond price. For instance, if the call price is 125 then the bond price can never exceed this, otherwise the issuer would call the bond. If a convertible is called the bond holder is usually given a *call notice period* of between 15 days and several months during which the bond holder can elect to convert at any time.

Call features are attractive to issuers for several reasons: they cap investors' profits from a rise in share price; they lessen the uncertainty about issuers' future liabilities; and they are a means to force conversion, typically when the firm can refinance at a cheaper rate. On the other hand, call features make the bond less attractive to prospective investors. For this reason the contract usually specifies that the bond is not callable for the first few years

⁶² See www.ubs.com/2/e/us_library/gis/white_papers/Benefits_of_Convertible_Bonds.pdf.

⁶³ For instance, a 10-year convertible bond may only be convertible after 5 years. And some bonds are only convertible on specific *conversion dates*.

⁶⁴ The effective call price is the clean call price, which is fixed in the bond's covenant, plus accrued interest.

after issue (the *hard call protection* period) and following this there may also be *soft call protection*, where the bond can only be called if the share price exceeds a certain threshold that is typically significantly greater than the call price. Another form of soft call protection is a *make-whole provision*.⁶⁵ Call protection is designed to increase the attraction of the bond to investors.

According to Grimwood and Hodges (2002), the most common type of US convertible bond contract in the MTS database at the ICMA Centre has a maturity of 15.0 years, pays a 6% semi-annual coupon and is hard-callable for the first time within 3 years. Of these contracts, 72% have a hard no-call period, and 53% have a put clause. There are many Japanese convertibles in the database, of which 88% have a hard no-call period, 91% have a soft no-call period, 23% have a put clause, 78% are cross-currency and 56% have a conversion ratio re-fix clause.

III.1.10.2 Survey of Pricing Models for Convertible Bonds

The option to convert into shares creates a premium for convertibles, so that the price of a convertible bond is never less than the equivalent straight bond price. But the price of a convertible bond is very difficult to calculate because it depends on so many factors other than interest rates and credit spreads. The price also depends on the conversion ratio, the probability of default, the recovery rate, the stock price behaviour in the event of default, and the dynamics of the stock price. Of all these highly uncertain quantities the stock price dynamics and default behaviour are perhaps the most important.⁶⁶

With so many uncertain factors affecting the price of a convertible bond, the pricing of these securities has been the subject of considerable academic research. Ingersoll (1977a) derived a closed-form pricing formula for callable convertible bonds but this is based on very simple assumptions about the evolution of the risk factors and can only be applied to bonds having very simple features. McConnell and Schwartz (1986) observed that when the stock price is modelled as a diffusion the discount rate must be adjusted, otherwise there is no possibility of default. This observation inspired several *blended discount* approaches to convertible bond valuation. Notably, Derman (1994) considered a stock price binomial tree where the discount rate at each time step is a weighted average of the risky discount rate and the risk free rate, with weight determined by the probability of conversion. In this framework the default event is not explicitly modelled; however, compensation for credit risk is included through a credit-adjusted discount rate.

One of the most important papers, by Tsiveriotis and Fernandes (1998), provided a rigorous treatment of Derman's ideas by splitting the bond value into equity and bond components, each discounted at different rates. This approach has since been extended to include interest rate and foreign exchange risk factors by Landskroner and Raviv (2003a, 2003b), who applied a blended discount model to price domestic and cross-currency inflation-linked convertible bonds.

All these papers assume that the stock price will not jump if the issuer announces that it is bankrupt, which is of course rather unrealistic. The most recent convertible bond valuation models developed by Davis and Lischka (1999), Takahashi et al. (2001), Ayache et al. (2003), Bermudez and Webber (2003) and Andersen and Buffum (2004) include a stock

⁶⁵ For instance, if a call is made between coupon payments then the accrued interest on the coupon will be paid to the bond holder.

⁶⁶ See Grimwood and Hodges (2002).

price jump on default. Most of these models incorporate *equity-linked hazard rates* that are driven by the stock price diffusion and are calibrated to the initial term structure of interest rates, often via the Hull and White (1990) model.

Some of the most interesting research on convertible bonds seeks to explain the issuer's call policy after the call protection period. During the call notice period the issuer effectively gives the investor a put on the share. Thus the optimal call price (i.e. the share price at which it becomes optimal for the issuer to call) should be such that the conversion price is just greater than the effective call price plus the premium on the put. However, it is quite possible that the stock price will decline during the call notice period and the investors will choose to receive the cash instead of converting. Since the issuer typically calls because it wants to give out stock, it will call only when the stock is very far above the effective call price. Thus issuers often wait until the conversion price is significantly higher than the optimal call price before issuing the call.

Many reasons have been proposed for this *delayed call* phenomenon. These include stock price uncertainty during the call notice period and the issuer's aversion to a sharp stock price decline (Ingersoll, 1977b; Butler, 2002; Asquith, 1995; Altintig and Butler, 2005; Grau et al., 2003); the preferential tax treatment of coupons over dividends, which acts as an incentive to keep the convertible bonds alive (Constantinides and Grundy, 1987; Campbell et al., 1991; Asquith and Mullins, 1991; Asquith, 1995); signalling effects, whereby convertible bonds calls convey adverse information to shareholders that management expects the share price to fall (Harris and Raviv, 1985; Mikkelsen, 1985); and issuers preferring to let sleeping investors lie (Dunn and Eades, 1984; Constantinides and Grundy, 1987).

Although the stock price dynamics are more important for the bond valuation than those of the interest rate or credit spread, relatively few papers incorporate realistic stock price dynamics into the pricing model. Most of the literature on convertible bond valuation assumes that equity volatility is constant. An exception is Yigitbaşıoğlu and Alexander (2006) who obtain arbitrage-free price bounds and hedge ratios for convertible bonds when there is uncertainty about the long term stock volatility. Their results also explain the delayed call feature of issuer's call policies.

III.1.11 SUMMARY AND CONCLUSIONS

This chapter introduces the theory and practice of pricing, hedging and trading standard interest rate sensitive instruments. Only a basic knowledge of finance is assumed, and although the treatment is concise several numerical examples and empirical studies have been provided, with accompanying interactive Excel spreadsheets on the CD-ROM.

We began by introducing the fundamental concepts of *discrete compounding* and *continuous compounding* of interest. International bond and money markets use different *day count conventions* for accruing interest under discrete compounding, and so banks usually convert all discretely compounded rates to equivalent continuously compounded rates, which are easier to deal with mathematically. *Spot* interest rates are those that apply from now, and *forward* interest rates apply from some time in the future. Given a curve of spot interest rates of different maturities, we can find the corresponding forward rate curve; conversely, given a forward rate curve and the one-period spot rate, we can find the spot rate of other maturities. The standard reference curve for spot and forward rates is the London Interbank Offered Rate or LIBOR curve.

Bonds may be categorized by maturity, coupon and type of issuer. Bonds are issued by financial institutions, governments, government agencies, municipalities and corporates. Almost all bonds have a credit rating which corresponds to the perceived probability that the issuer will default on its debt repayments. Bonds are further categorized by the priority of their claim on the assets in the event of the issuer's default, and according to the domicile of the issuer. The majority of bonds pay a *fixed coupon*, but some bonds, called *floating rate notes* or *floaters*, have a variable coupon based on a reference rate such as LIBOR. In addition to straight floaters we have reverse floaters, capped floaters, floored floaters and many other floaters that can have complex option-like features. Bonds are further divided into *short, medium and long term* bonds and those that have less than 2 years to expiry at issue are generally regarded as *money market* instruments. Most of these pay no coupons, but are priced at a discount to their *par value* of 100. Very few bonds pay no coupon but the yields on hypothetical zero coupon bonds are used as *market interest rates* for different credit ratings.

The prices of liquid bonds are set by supply and demand, and their fair prices correspond to the present value of all the cash flows on the bond when discounted using market interest rates. Given any set of positive cash flows we can define a unique yield as the constant discount rate that, when applied to these cash flows, gives the market price of the bond. The price of a bond has a convex, decreasing relationship with the yield, and when the yield is equal to the coupon the bond is priced at par.

The *Macaulay duration* is the time-weighted average of the present value of the cash flows on the bond and, under continuous compounding, the Macaulay duration is also the first order sensitivity of the bond price to a unit change in its yield. Under discrete compounding this sensitivity is called the *modified duration*. We can use a second order Taylor expansion, the *duration-convexity approximation*, to approximate the change in present value of the cash flows when the yield changes by a small amount. When expressed in nominal terms we call these sensitivities the *value duration* and *value convexity*. Value durations can be added to obtain a single duration measure for a bond portfolio, which represents the change in portfolio value when all the bond yields change by the *same* amount. Similarly, value convexity is also additive over all bonds in the portfolio.

Floating rate notes have very short duration. Thus by adding long or short positions in floaters, bond traders can easily adjust the duration of their bond portfolios. But the problem with duration and convexity hedging, where a bond portfolio is constructed having zero value duration and zero value convexity, is that it only hedges against *parallel* movements of the zero coupon yield curve. It does not hedge the portfolio against the most common movements in the yield curve, based on historical data. For this we need to base bond portfolio immunization on *principal component analysis*.

A *forward rate agreement* is an OTC agreement to buy or sell a forward interest rate. In a standard *interest rate swap* one party pays LIBOR and receives a fixed rate, the swap rate, from the other party over an agreed period. This may be regarded as a sequence of forward rate agreements. A *vanilla swap* is a swap where the floating payments at the end of a certain period depend only on the interest rate prevailing over that period. We can decompose the cash flows on a vanilla swap into two series: the fixed leg of a vanilla swap can be regarded as a bond with coupon equal to the swap rate and the floating leg can be regarded as a single cash flow at the first payment date.

A *basis swap* is an interest rate swap where two floating rates are exchanged on each payment date. These floating rates may be of different maturities, or they may be denominated

in different currencies. Unlike interest rates swaps, a *cross-currency basis swap* involves exchanges of principal. It can be regarded as two floating rate notes in different currencies. Another popular type of swap is a *total return swap*, in which two parties exchange returns on one asset for returns on another, including the capital gains or losses at every payment. Many other non-standard interest rate swaps are traded.

It is not possible to calculate the yield or duration of a swap. However, the *present value of a basis point* (PV01) applies to all types of cash flows. The PV01 is the sensitivity of the present value of the cash flow to a shift of one basis point in market interest rates. That is, the whole zero coupon yield curve shifts parallel by one basis point. The PV01 should be distinguished from the value duration, which measures the price sensitivity when the yields on all the bonds in the portfolio change by the same amount. The *interest rate risk* of a cash flow refers to the uncertainty about the discounted value of this cash flow at some future time t . It can be measured by the standard deviation of this future discounted value. This is the product of the standard deviation of the discount rate, in basis points, and the interest rate sensitivity of the expected discounted value of the cash flow at time t , which is approximately equal to its PV01. Hence, at time t the interest rate risk of a cash flow at time T is approximately equal to the PV01 of the cash flow times the standard deviation of the forward rate of term t and tenor $T-t$.

We have reviewed the *yield curve fitting* techniques that can be used to obtain a zero coupon yield curve from the market prices of fixed income instruments. We discussed the suitability of each type of instrument for yield curve fitting, and focused on the influence that the chosen fitting technique has on the statistical properties of spot and forward interest rates. A case study compared the application of *basis splines* with the *Svensson model* for the calibration of the UK LIBOR yield curve. We showed that forward rate volatilities and correlations inherit their structure from the yield curve fitting model: the Svensson forward rate volatilities and correlations are the best behaved since their structure is determined by the exponential functional form of the model; by contrast, in spline methods, volatilities and correlations are mainly determined by the choice of knot points, even after imposing a roughness penalty.

Convertible bonds are hybrid derivative instruments that simultaneously offer the potential upside of equities and a limited downside since the investor holds a bond. A convertible bond offers exposure to equity volatility as well as to equities, interest rates and credit spreads. Most convertible bonds have *call features* that protect the issuer, but call protection policies and *put features* increase their attraction for investors. The price of a convertible bond is very difficult to calculate because it depends on so many factors other than interest rates and credit spreads. Perhaps the most important risk factors for convertible bonds are the stock price dynamics and the stock price behaviour in the event of an issuer default. We concluded the chapter with a survey of the very considerable academic literature on pricing models for convertible bonds.