1
Definition and valuation of the underlying instruments

1.1 INTRODUCTION

The main task of this book is to show how existing option models can be understood, analysed and implemented in order to price and risk-manage 'exotic' interest-rate options. The term 'exotic' is, by itself, far from being unambiguous: while, on the one hand, the financial press tend to extend the adjective to any 'derivative' instrument (thereby including even futures contracts), some traders in the US$ derivatives market, on the other hand, might consider a 10-year American, step-up-coupon swaption with exit penalty a 'commoditised' plain-vanilla instrument. In the context of this book, 'exotic' will be taken to refer to any option whose value depends in an important way on the evolution of the yield curve as a whole. Very often, this is tantamount to using the term 'exotic' for any option whose value cannot be reduced to a closed-form expression, such as, for instance, Black's (1976) celebrated formula. Models, such as the Hull and White (see Chapter 13), which do afford closed-form solutions for certain types of options (typically calls or puts on discount bonds) are, of course, also treated, but the emphasis is laid on the applications that go beyond the important but limited cases for which exact formulae exist.

Before embarking on the treatment of these exotic interest-rate options it is important to have a clear understanding of how the underlying instruments are priced, and of the rather subtle interplay between the prices of the 'underlying' and of the 'derivative' instruments. This is the task undertaken in the present chapter, and in Chapter 5. The attempt has been made to make the treatment as self-contained as possible, and therefore Black's formula is derived in Section 5.2. Readers totally unfamiliar with the fundamentals of option pricing might, however, find it profitable to read at least some chapters from any of the many introductory texts (I cannot think of a better one than Hull's (1992) classic). Finally, the overall treatment would have been more elegant if a completely axiomatic approach had been adopted, by deriving theoretical results first, and applications thereafter. It was feared, however, that in so doing more could be lost in financial intuition than gained in elegance. Chapters 1 to 5 therefore make use of an 'intuitive' understanding of what arbitrage is in order to derive important results; arbitrage is 'properly' defined in Appendix B. Clearly, the derivations in Chapters 1 and 5 are less general than they could have been if presented in the context of the martingale approach. However, the reader would have missed the important issue of the very motivation for yield curve models, presented in Chapters 3 and 5. The simple proofs presented in this chapter should therefore be seen in this light, and profitably revisited before embarking on Part Three.

Broadly speaking, pricing an exotic interest-rate option can be looked at in two distinct ways. Complex interest-rate options can be regarded as tools which allow the investor to express 'sophisticated' views about the future evolution of interest rates. These views, as mentioned in the introduction, can go well beyond simple directional positions ('rates are going to go "up" or "down"'), for which plain-vanilla instruments such as futures and swaps are probably more suited, but can, for instance, express predictions about the steepening or flatting of a yield curve, independently of, or in conjunction with, a change in level; or about the precise timing of a certain yield curve move. In all these cases the investor will reap his rewards if his views were correct *and if these views were different from the market's*. Therefore, in this framework, pricing the market exactly (i.e. recovering the observed market prices within the bid/offer spread) is neither logically necessary, nor, to some extent, desirable. If a model could be trusted to give a fundamentally correct, albeit necessarily simplified, description of economic reality then discrepancies between model and market values would point to possible trading opportunities. Equilibrium models, which attempt to describe the economy as a whole, belong in this class.

At the opposite end of the spectrum there is a pure no-arbitrage approach: in this framework the user who needs to price an exotic interest-rate option will hedge his position using either plain-vanilla options (such as caps or European swaptions) or the underlying 'cash' instruments (bonds, forward rate agreements (FRAs), swaps or futures); the correctness of the prices of the latter will not be questioned, and the value of the exotic options will be regarded as the cost of the replicating hedging portfolio. Clearly, exact recovery of the actually traded market prices of the underlying instrument now becomes all-important. At the risk of oversimplifying the issue, one approach can be seen as an attempt to *explain* prices, the other as using these prices as exogenously given building blocks.

Whatever the approach might be, it will be necessary to recover the implied market values for the underlying instruments using the more sophisticated approaches described in later chapters. (In the context of no-arbitrage models, the procedure is normally referred to as the calibration or parametrisation of an interest rate model.) In turn, the market prices of futures, FRAs, swaps, or bond prices are all a function of the discount curve. A discussion is therefore undertaken in this chapter of the issues underlying the construction of a reliable yield curve, and the evaluation from the latter of the market prices of the instruments mentioned above. An excellent (and lengthy) treatment of these topics as far as LIBOR curves are concerned can be found in Miron and Swannell (1991). The next section deals with the conceptual steps involved in the creation of a discount curve, and highlights the methodological similarities and differences in constructing bond and LIBOR yield curves.

1.2 DEFINITION OF SPOT RATES, FORWARD RATES, SWAP RATES AND PAR COUPON RATES

There exist a number of equivalent descriptions of a given market yield curve. Ultimately, spot rates, yields, swap rates, par coupon curves, etc., are all shorthand notations for cash flows that will occur at some future time. These cash flows, in turn, can either be certain, or conditional upon pre-specified states of the world being attained in the future. A title to a single certain cash flow of magnitude N at a known time in the future T is called a **discount bond**, and the amount N is referred to as the **principal**, or **notional**, or **face value.** In certain cases, i.e. especially for short maturities T, these instruments are traded directly in the market. What is more common, however, is to have securities that promise to pay a stream of certain payments at times $\{t_i\}$, often referred to as **coupons**, and a (usually larger) payment on the last payment date (the **maturity date**). These latter instruments are normally called **coupon-bearing bonds**. From the definition given, it is clear that a coupon-bearing bond can be regarded as a collection of discount bonds with notionals equal to the coupon of the bond, each maturing on one of the different coupon-payment dates, plus an additional discount bond with notional equal to the last (maturity) payment of the coupon-bearing bond. The valuation of the latter can therefore be reduced to the valuation of a collection of pure discount bonds. In the rest of this book, the price at time t of a discount bond maturing at time T will be denoted by $P(t, T)$, and the price at time t of a coupon-bearing bond paying coupons at times $\{t_i\}$, and the principal amount at time T by $Bnd(t, \{t_i\}, T)$. Since the problems connected with the possibility of default are not touched upon in this book, and since, as seen, coupon-bearing bonds are reducible to a suitable bundle of discount bonds with different face values, from the conceptual point of view the collection of prices $P(t, T)$ for any $T \geq t$ fully describes the value to be associated to any collection of certain future cash flows. Since, as mentioned before, pure, risk-free discount bonds are rarely traded in the market, their prices have to be imputed from the prices of the coupon-bearing bonds. Despite these practical difficulties the conceptual advantages arising from assuming that the prices of a continuum of discount bonds are indeed directly available are such that it will always be assumed in the following chapters that this continuum of prices (otherwise known as the **discount function**) have already been obtained. How this can be accomplished is explained in the following sections.

A security that entitles the holder to a given cash flow if a particular, pre-specified, state of the world is attained at one or more future dates is called a **contingent claim** (a more precise definition is given in Appendix B). In some situations, despite the fact that the cash flows are uncertain at time t_0, they can be replicated exactly by entering suitable strategies which require only (positive or negative) holdings of discount bonds. If arbitrage (see Appendix B and Chapter 6) is to be avoided, knowledge of the discount function at time t_0 therefore completely determines the prices of this class of contingent claims. Notice that, since the payoffs of these simple contingent claims can be replicated using instruments that pay *certain* cash flows in the future, no statistical assumptions have to be made regarding the probability of occurrence of any future state of the world. In particular, the value of these instruments is independent of any volatility. Contingent claims that can be replicated in this fashion (e.g. FRAs, swaps) are treated in this chapter.

For more general contingent claims, no-arbitrage and knowledge of the discount function are not sufficient to determine their value, and one must also make suitable assumptions about the probability distributions of the random variable(s) that determine the future cash flows. (It will actually turn out that, because of no-arbitrage, assumptions only need to be made regarding the second, or higher, moments of these distributions.) Not surprisingly, therefore, the prices of these contingent claims will turn out to be given (Appendix A, Section A.7) by suitable expectations, taken with respect to the appropriate probability distributions, of future cash flows. Some of these expectations can be easily evaluated with closed-form expressions, using the market discount function and the variance of the random variable(s) which determine the payoffs. These simple, volatility-dependent, contingent claims which admit closed-form solutions (caps, floors and European swaptions) are also treated in this chapter. The rest of the book deals with contingent claims for which such simple solutions are not available (or, as will be argued in Chapter 5, are of limited use), and whose future cash flows depend on pre-specified future realisations of the discount function. Despite the fact that this definition is perfectly self-consistent, it is more common to express the conditions that trigger future payments in terms of rates. These are therefore defined as follows.

The time-t **continuously compounded discrete spot rate of maturity** T, $R_c(t, T)$, is defined by

$$P(t, T) \equiv \exp[-R_c(t, T)(T - t)] \tag{1.1}$$

$$R_c(t, T) \equiv -\frac{\ln[P(t, T)]}{T - t} \tag{1.1'}$$

The quantity R_c just defined is often referred to as the (continuously compounded) yield of the discount bond $P(t, T)$. Given the existence of a full zoology of yields (flat, current, gross, redemption, etc.), often imprecisely and sometimes inconsistently defined, the term 'yield' is by and large avoided in this book.

The time-t **continuously compounded discrete forward rate** spanning the period $[T \;\; T + \Delta t]$, $f(t, T, T + \Delta t)$, is defined by

$$\frac{P(t, T + \Delta t)}{P(t, T)} \equiv \exp[-f(t, T, T + \Delta t)\Delta t] \tag{1.2}$$

$$f(t, T, T + \Delta t) \equiv -\frac{\ln(P(t, T + \Delta t)) - \ln(P(t, T))}{\Delta t} \tag{1.2'}$$

The limits as $T \to t$, $\Delta t \to 0$ in Equations (1.1′) and (1.2′) define the instantaneous short rate, $r(t)$, and the instantaneous forward rate, $f(t, T)$, respectively.

$$r(t) = \lim_{t \to T} R(t, T) \tag{1.3}$$

$$\boldsymbol{f(t, T)} = \lim_{\Delta t \to 0} -\frac{\ln(P(t, T + \Delta t)) - \ln(P(t, T))}{\Delta t} = -\frac{\boldsymbol{\partial \ln(P(t, T))}}{\boldsymbol{\partial T}} \tag{1.3'}$$

i.e. **the instantaneous forward rate as seen from the yield curve at time t is equal to (minus) the logarithmic derivative of the time-t price of a discount bond of maturity T with respect to its maturity.**

Finally, from (1.3′) one can write

$$\int_t^T \mathrm{d}\ln P(t, s) = -\int_t^T f(t, s)\,\mathrm{d}s = \ln(P(t, T)) - \ln(P(t, t)) \tag{1.4}$$

But, since $P(t, t) = 1$,

$$-\int_t^T f(t, s)\,\mathrm{d}s = \ln P(t, T) \tag{1.5}$$

and, finally,

$$\boldsymbol{P(t, T) = \exp - \int_t^T f(t, s)\,\mathrm{d}s} \tag{1.6}$$

The (logarithm of the) price at time t of a discount bond maturing at time T is equal to the integral over maturities of the instantaneous forward rates as seen from the time-t yield curve.

In complete analogy with the definitions given above, in the case of simple (rather than continuous) compounding, the **simply compounded spot rate** is defined to be

$$R_s(t, T) = \frac{1/P(t, T) - 1}{T - t} \tag{1.7}$$

and the **simply compounded forward rate** spanning the period $[T_1 \;\; T_2]$, $F(t, T_1, T_2)$ is then defined as

$$F(t, T_1, T_2) = \frac{P(t, T_1)/P(t, T_2) - 1}{T_2 - T_1} \tag{1.8}$$

Table 1.1 Summary of Definition of Rates

Name	Symbol	Definition
Continuously-compounded discrete spot rate	$R_c(t, T)$	$-\ln P(t, T)/(T-t)$
Continuously-compounded discrete forward rate	$f(t, T, T+\Delta t)$	$-[\ln P(t, T+\Delta t) - \ln P(t, T)]/\Delta t$
Instantaneous forward rate	$f(t, T)$	$-\partial \ln P(t, T)/\partial T$
Instantaneous spot rate	$r(t)$	$f(t, t)$
Simply-compounded spot rate	$R(t, T)$	$[1/P(t, T) - 1]/(T-t)$
Simply-compounded forward rate	$F(t, T_1, T_2)$	$[P(t, T_1)/P(t, T_2) - 1]/(T_2 - T_1)$
τ-period-compounded spot rate	$R_\tau(0, t)$	$\nu[P(0, t)^{-1/\nu t} - 1]$

Between the limiting cases of simple and continuous compounding one can define the τ-period compounded spot rates, $R_\tau(0, t)$, implicitly given by

$$P(0, T) = 1/(1 + R_\tau(0, t)/\nu)^{\nu t},$$

where $\nu \equiv 1/\tau$.

Having clarified these alternative equivalent ways of describing the value of future cash flows, the next section will tackle the task of associating a value to a contingent claim whose payoffs can be replicated using strategies involving pure discount bonds. It will be assumed throughout that no market frictions are present.

Table 1.1 summarises the definitions of the various rates.

1.3 THE VALUATION OF PLAIN-VANILLA SWAPS AND FRAS

A **plain-vanilla interest-rate swap** is an agreement whereby two parties undertake to exchange, at known dates in the future, a fixed for a floating set of payments (often referred to as the fixed and floating legs of a swap). The **fixed leg** is made up by payments B_i:

$$B_i = N_i X \tau_i \quad (1.9)$$

where N_i is the notional principal of the swap outstanding at time t_i, τ_i, usually referred to as the frequency or the tenor of the swap, is the fraction of the year between the $(i-1)$th and the ith payment (therefore approximately equal to $\frac{1}{2}$ or $\frac{1}{4}$ for a semi-annual or quarterly swap)[1], and X is the fixed rate contracted at the outset to be paid by the fixed-rate payer at each payment time. For a plain-vanilla swap each fixed payment B_i occurs at the end of the accrual period, i.e. at time t_{i+1}. See Figure 1.1.

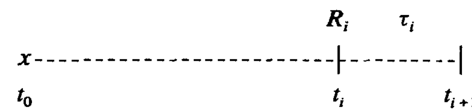

Figure 1.1 The timing of cash flows for a plain-vanilla swap: the realisation at time t_i (reset time) of the spot rate R_i spanning the period $[t_i\ t_{i+1}]$ determines a floating payment per unit principal at time t_{i+1} (payment time) of magnitude $R_i\tau_i$. The tenor τ_i is given by the number of days between t_i and t_{i+1} divided by 360 or 365, as dictated by the appropriate conventions. For a plain-vanilla swap, the fixed payment per unit principal $X\tau_i$ also occurs at time t_{i+1}

If we denote by $P(0, t)$ the price of a discount bond maturing at time t, the present value of each fixed payment B_i is given by:

$$PV(B_i) = N_i X \tau_i P(0, t_{i+1}) \quad (1.10)$$

As for the **floating leg**, each payment A_i, also occurring at time t_{i+1}, is given by

$$A_i = N_i R_i \tau_i \quad (1.11)$$

where R_i is a shorthand notation for the τ-period spot rate (i.e. the 3-month or 6-month LIBOR rate, for a quarterly or semi-annual swap, respectively) prevailing at time t_i, and covering the period from t_i to t_{i+1}: $R_i = R(t_i, t_i + \tau)$. Times t_i and $t_i + \tau$ are normally referred to as the **reset and payment times** for the ith period, respectively. Clearly, the realisations at times $\{t_i\}$ of these spot rates are not known at time 0, and therefore the net present value at a generic time t of each of these floating payments is given by

$$PV(A_i) = E[N_i R_i \tau_i P(t, t_{i+1})], \quad (1.12)$$

where the symbol $E[\cdot]$ denotes expectation. Needless to say, while, at time 0, the magnitudes of the fixed-leg payments are known, and a certain value can therefore be associated to them (Equation (1.10)), the realisations of the τ-period spot rates R_i at time $1, 2, \ldots, n$ are *not* known, and, therefore, for the moment we do not know what value to associate to expression (1.12).

Let us now consider the following strategy: let us purchase at time 0 (today) a discount bond maturing at time t_i, $P(0, t_i)$, and sell (go short of) a discount bond maturing at time t_{i+1}, $P(0, t_{i+1})$. At time t_i the resulting portfolio will have a value $V(t_i)$:

$$V(t_i) = P(t_i, t_i) - P(t_i, t_{i+1}) = 1 - P(t_i, t_{i+1}) \quad (1.13)$$

which, assuming simple compounding over the period $[t_i\ t_{i+1}]$, is equal to

$$V(t_i) = 1 - \frac{1}{1 + R_i\tau_i} = \frac{R_i\tau_i}{1 + R_i\tau_i} \quad (1.13')$$

By Equation (1.12) above, the payer of the floating leg will have to make a payment at time t_{i+1} of present value $V'(t_i)$ at time t_i equal to

$$V'(t_i) = \frac{R_i \tau_i}{1 + R_i \tau_i} \tag{1.14}$$

Therefore $V(t_i) = V'(t_i)$, i.e. the payoff arising from one reset of the floating leg can be perfectly and certainly met by entering the long/short bond strategy suggested before. **At time 0 the commitment to pay R_i in the floating leg and the strategy of holding a bond $P(0, t_i)$ and shorting a bond $P(0, t_{i+1})$ must therefore have the same value:**

$$\boldsymbol{P(0, t_i) - P(0, t_{i+1}) = R_i \tau_i P(0, t_{i+1})} \tag{1.15}$$

It follows that, to avoid arbitrage (precisely defined in Appendix B, but intuitively understandable at this stage as the capability of making certain money without any risk), one can value the floating leg of a swap by setting the unknown quantities R_i equal to the value

$$R_i = \frac{P(0, t_i)/P(0, t_{i+1}) - 1}{\tau_i} \tag{1.16}$$

(notice carefully that no statements about probability distributions or expectations have been used in the argument); but Equation (1.16) is simply the well-known definition of a simply compounded forward rate spanning the period $[t_i \; t_{i+1}]$, $F(0, t_i, t_{i+1})$ (see Equation (1.8)). Therefore, to avoid arbitrage **the a priori unknown cash flows in the floating leg must be set equal to the projected forward rates**. (The same result can be obtained in a more general way using the approach of Chapter 7, Section 7.5, where it is shown that, if one uses bonds of maturities t_{i+1} from the current term structure as numeraire, forward rates are driftless.) Notice carefully that the quantities $PV(A_i)$ in Equation (1.12) are stochastic variables, and, as such, possess a certain variance. It is only the present value of each floating reset *plus* the accompanying strategy of long/short bonds that has no variance, and is therefore amenable to a purely deterministic evaluation at time 0.

The **equilibrium swap rate** is then defined as the fixed rate X such that today's present value of the fixed and floating legs are the same:

$$\sum PV(B_i) = \sum N_i X \tau_i P(0, t_{i+1}) = \sum PV(A_i) = \sum N_i F_i \tau_i P(0, t_{i+1}) \tag{1.17}$$

The **equilibrium swap rate** is therefore equal to

$$X = \frac{\sum N_i F_i \tau_i P(0, t_{i+1})}{\sum N_i \tau_i P(0, t_{i+1})} \tag{1.18}$$

i.e. **it is a weighted average of the projected forward rates**. This can be seen more clearly by setting

$$w_i = \frac{N_i \tau_i P(0, t_{i+1})}{\sum N_i \tau_i P(0, t_{i+1})} \tag{1.19}$$

which allows one to rewrite Equation (1.18) as

$$X = \sum F_i w_i \tag{1.20}$$

This expression will be used later in the section.

By the way the equilibrium rate has been obtained it follows that entering an equilibrium swap (i.e. a swap struck at the equilibrium rate) today has zero cost, since, by definition, the two parties have undertaken to exchange legs of identical value. After an equilibrium swap has been entered, the swap itself will in general no longer have zero value, since interest rates will not, in general, have followed the implied forward curve. It is in fact easy to show that the only values of the joint realisations at time t_1 of the projected forward rates which preserve zero value for an equilibrium swap initiated at time 0 are the values for the same forward rates implied by the yield curve at time 0, i.e.

$$F(t_1, t_i, t_i + \tau_i) = F(0, t_i, t_i + \tau_i) \tag{1.21}$$

where the full notation $F(t, T, T + \tau)$ has been employed to indicate the forward rate from time T to time $T + \tau$, as seen from the yield curve at time t. Notice that, in Equation (1.21), the term $i = 1$ implies that

$$F(t_1, t_1, t_1 + \tau) = R_s(t_1, \tau_1) = F(0, t_1, t_1 + \tau) \tag{1.22}$$

since $F(t, t, T) = R_s(t, T)$; therefore, **in general** (i.e. barring fortuitous cancellations), **for the time-0 equilibrium swap to retain its zero value, the realisation at time t_i of a spot rate must equal the time-0 projected forward rate.**

For the payer of the fixed rate the present value of the swap at time t will be given by

$$NPV_{\text{swap}}(t) = -\sum N_i X \tau_i P(t, t_{i+1}) + \sum N_i F_i \tau_i P(t, t_{i+1}) \tag{1.23}$$

where the F_is are now the forward rates calculated from the discount curve at time t. As mentioned after Equation (1.16) above, notice carefully that the argument underpinning the derivation is of no-arbitrage nature, and no claim has been made about expectations of future rates.

At time t, the second term on the RHS in Equation (1.23) is equal, by definition, to the equilibrium swap rate prevailing at time t, X_t, times the present value of the fixed leg. Therefore Equation (1.23) can be rewritten as

$$NPV_{\text{swap}}(t) = (X_t - X_0) \sum N_i \tau_i P(t, t_{i+1}) = (X_t - X_0) \sum_i B_i \tag{1.23'}$$

which expresses the net present value of a swap at time t as the difference between the prevailing equilibrium swap rate, X_t, and the swap rate originally contracted, X_0, times the fixed leg of the swap.

Further insight into the equilibrium rate can be obtained by expanding the numerator in Equation (1.18) making use of relation (1.16). For a swap with n resets, and final maturity at time t_{n+1}, one can write

$$\sum_{i=1,n} N_i F_i \tau_i P(0, t_{i+1}) = \sum_{i=1,n} N_i \tau_i \left(\frac{P(0, t_i)/P(0, t_{i+1}) - 1}{\tau_i}\right) P(0, t_{i+1})$$
$$= \sum_{i=1,n} N_i [P(0, t_i) - P(0, t_{i+1})] \qquad (1.24)$$

which, for constant principals $N_i = 1$, can be written

$$\sum_{i=1,n} N_i F_i \tau_i P(0, t_{i+1}) = \sum_{i=1,n} P(0, t_i) - P(0, t_{i+1})$$
$$= P(0, t_1) - P(0, t_2) + P(0, t_2) - P(0, t_3) \ldots$$
$$+ P(0, t_n) - P(0, t_{n+1})$$
$$= P(0, t_1) - P(0, t_{n+1}) \qquad (1.24')$$

For a spot-starting plain-vanilla swap the present value of the floating leg of a unit-principal swap is therefore equal to

$$\sum \boldsymbol{PV}(\boldsymbol{A_i}) = \boldsymbol{P(0, 0) - P(0, t_{n+1}) = 1 - P(0, t_{n+1})} \qquad (1.24'')$$

Notice that the frequency of the plain-vanilla swap does not affect the value of the floating leg, which is therefore the same not only for different day count conventions (ACT/360, ACT/365, 30/360, etc.)[1], but also for a monthly or a yearly swap with same final maturity. **The equilibrium swap rate is therefore given by the ratio of two portfolios of discount bonds**:

$$X = \frac{\boldsymbol{1 - P(0, t_{n+1})}}{\sum \boldsymbol{\tau_i P(0, t_{i+1})}} \qquad (1.25)$$

or, in general, for a swap starting at a future time t, its value at time 0 is

$$X = \frac{\boldsymbol{P(0, t) - P(0, t_{n+1})}}{\sum \boldsymbol{\tau_i P(0, t_{i+1})}} \qquad (1.25')$$

and, therefore, **any yield curve model capable of pricing discount bonds exactly must recover the market swap rates correctly for any choice of the model volatility**. (Notice that, since Eq. 1.24″ holds for a plain-vanilla — i.e. constant-principal — swap, the notional amount has been cancelled out.)

Notice carefully that this would no longer be the case if a simple modification to the payment provisions were introduced, for instance if the rate determining

the payoff in Equation (1.11) were still the τ-period LIBOR reset at time t_i, but the corresponding payments were to occur at any time other than t_{i+1}. In this case the cancellation in Equation (1.24) between $P(0, t_{i+1})$ in the numerator and the denominator would no longer take place, and one would be left with a dependence on the volatility input of a given yield curve model (see Chapter 7, Sections 7.7 and 7.8 for a discussion of the LIBOR-in-arrears case).

Since a FRA is simply a one-period swap (or, conversely, since a swap is a series of FRAs), the expressions just derived also price this simpler instrument: the net present value of a forward rate agreement, whereby one party undertakes to pay the other at time t_{i+1} (per unit principal) the difference between the τ-period LIBOR resetting at time t_i and an agreed fixed rate, X, times the fraction of the year τ covered by the LIBOR spot rate, is simply given by

$$PV(FRA) = [P(0, t_i) - P(0, t_{i+1})] - XP(0, t_{i+1})\tau_i \qquad (1.26)$$

Let us now consider the issuer of a coupon-bearing bond, who, against receipt today of £1, undertakes to pay a fixed coupon X with frequency $1/\tau$, and to repay the principal at maturity t_{n+1}. The coupon liability incurred by the issuer clearly has the same present value as the fixed leg of a swap. As for the principal to be paid back (redeemed) at maturity and received upfront today, their combined present values, $PV(A')$, are given by

$$PV(A') = 1 - P(0, t_{n+1}) \qquad (1.27)$$

Comparing Equation (1.27) with expression (1.24″) one can therefore conclude that **the floating leg of a swap has exactly the same value and plays exactly the same role as the receipt of the proceeds from issuing the bond today, and the accompanying commitment to repay the principal at maturity**.

Whether one is dealing with plain-vanilla swaps or with bullet bonds (i.e. bonds paying regular coupons, without any call or put provisions), their pricing can therefore be completely reduced to a suitable manipulation of pure default-free discount bonds. More precisely, in view of the above it is clear that, apart from credit considerations, the present value of a bond, Bnd, paying n coupons X at regular intervals every τ years until a final maturity at time t_n is given by

$$PV(Bnd) = \sum_{i=1,n} X\tau_i P(0, t_i) + 100P(0, t_n) \qquad (1.28)$$

Using bond conventions, what has been referred to before as the equilibrium rate is called the **par coupon**, which can therefore be defined as that **particular coupon that prices the bond today exactly at par**. As pointed out before, by moving the last term on the RHS of Equation (1.28) to the left one immediately finds the bond equivalent of the present value of the floating leg of a swap. Just as an equilibrium swap struck with a rate X at time t_0 will in general have a non-zero value at a later time t (see Equation (1.23)), a par bond issued at time

t_0 will, in bond terminology, trade **at a premium** or **at a discount** at later time t, according to whether the RHS of Equation (1.24) above will add up to more or less than 100, respectively.

1.4 OBTAINING THE DISCOUNT FUNCTION FROM A SET OF SPANNING FORWARD OR SWAP RATES

In order to give a description of the discount function in terms of traded market quantities, let us first of all define spanning forward rates as a particular set of discrete forward rates such that the first coincides with the discrete spot rate, and the maturity of the ith coincides with the beginning (expiry) of the $(i+1)$th (see Figure 1.2).

Similarly, spanning swap rates can be defined as a set of swap rates such that the first covers the period from spot to a given maturity, and the ith spans the period from the ith reset of the spot swap rate to the same final maturity (see Figure 1.3). This set of swap rates is particularly important because it is made up of the swap rates underlying a Bermudan swaption (see Chapter 2, Section 2.3).

From the relationships obtained in the previous sections, using forward rates one can then immediately write for the discount function at the discrete reset times:

$$P(t_0, t_i) = \Pi_{i=1,n}[1/(1 + F(t_0, t_i, t_{i+1})\tau_i)] \qquad (1.29)$$

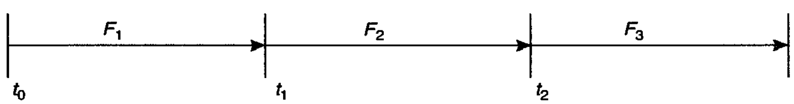

Figure 1.2 Three spanning forward rates covering the period from spot (t_0) to time t_3. Notice that the maturity of each forward rate coincides with the beginning (expiry) of the next

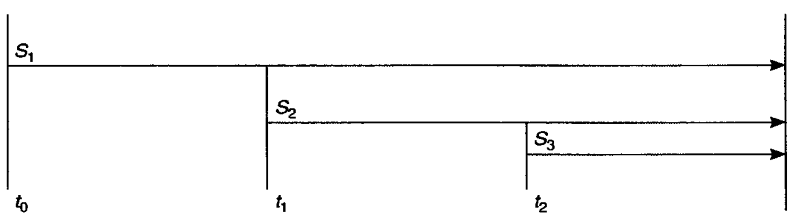

Figure 1.3 Three spanning swap rates covering the period from spot (t_0) to time t_3. In this example the first swap has three resets, the second two and the third just one. Notice that the maturity of all the swap rates is the same, and that the start of a generic swap other than the first coincides with the first cash-flow time from the previous swap

Also spanning swap rates, however, allow the recovery of the discount function sampled at the same discrete points and can, therefore, be regarded as equivalently fundamental building blocks. To see how this can be accomplished, let us write the equilibrium spanning swap rates for the example in Figure 1.3 as

$$S_3 = (P(t_0, t_2) - P(t_0, t_3))/(P(t_0, t_3)\tau) \qquad (1.30)$$

$$S_2 = (P(t_0, t_1) - P(t_0, t_3))/((P(t_0, t_3) + P(t_0, t_2))\tau) \qquad (1.31)$$

$$\begin{aligned} S_1 &= (P(t_0, t_0) - P(t_0, t_3))/((P(t_0, t_3) + P(t_0, t_2) + P(t_0, t_1))\tau) \\ &= (\;1 \quad - P(t_0, t_3))/((P(t_0, t_3) + P(t_0, t_2) + P(t_0, t_1))\tau) \end{aligned} \qquad (1.32)$$

where, in each equation, the numerator is the present value of the floating leg, and the denominator is the present value of the fixed leg for a unit coupon (the same tenor has been assumed for every period to lighten notation). Notice that Equation (1.32) does not introduce any additional unknowns, by virtue of the fact that $P(t_0, t_0) = 1$, and therefore one can solve for the three unknowns $P(t_0, t_1)$, $P(t_0, t_2)$ and $P(t_0, t_3)$. In general, n discrete discount factors can always be found, for any viable set of n spanning equilibrium swap rates, as the solution of a linear $[n \times n]$ problem.

This result is important, not simply because it shows that equilibrium swap rates are yet another set of equivalent quantities that can be used to describe the yield curve, but because it shows that the somewhat privileged role enjoyed by forward rates in term structure modelling might be more the result of historical accident than of some fundamental financial reason. A whole body of literature has in fact taken as the starting point a continuum of traded discount bonds, whose logarithmic derivative is directly linked to a set of instantaneous forward rates. These 'building blocks' have been regarded as the continuum-time counterpart of their discrete LIBOR equivalent, and, as such, have provided a natural proxy for the modelling of FRA-based instruments. For swap-based products, however (such as Bermudan swaptions, or options on Constant Maturity Swaps), swap rates constitute the most natural set of state variables, and the relationships presented above therefore afford the necessary link with the discount function. The issue of the 'best' choice of state variables (forward rates versus swap rates) will be addressed in detail in the chapter devoted to the BGM/J approach (see Chapter 18).

1.5 THE VALUATION OF CAPS, FLOORS AND EUROPEAN SWAPTIONS

It will be important for the following to link the variance of the equilibrium swap rates obtained above with the variances of and correlations among the underlying forward rates. To see this more precisely, it is necessary to define first of all a cap, and then a swap option (swaption in the following).

A **cap** is a collection of caplets. A **caplet**, in turn, is a contract which pays at time t_{i+1} the difference between the τ_i-period spot rate resetting at time t_i, R_i, and a strike price K multiplied by the year fraction τ_i between t_i and t_{i+1}, if this difference is positive, and zero otherwise:

$$\text{Caplet}(t_{i+1}) = \text{Max}[R_i - K, 0]\tau_i \tag{1.33}$$

The present value at time t_i of this payoff is given by

$$\text{Max}[R_i - K, 0]P(t, t_{i+1}) = \text{Max}[R_i - K, 0]\frac{1}{1 + R_i\tau} \tag{1.34}$$

where use has been made of definition (1.7).

No-arbitrage arguments, presented in detail in Appendix B, show that, for valuation purposes, the unknown future value of the rate resetting at time t_i must be set equal to today's implied forward (the analogy with the case of the FRA is in this respect complete). If, in addition, one accepts the log-normal assumption for the forward rates, one can very easily arrive at the Black model, in which a caplet of expiry t_i struck at K is seen as a call on the forward rate:

$$\text{Caplet} = [F(t_0, t_i, t_{i+1})N(h_1) - KN(h_2)]P(t_0, t_{i+1}) \tag{1.35}$$

where $N(\cdot)$ denotes the cumulative normal distribution,

$$h_{1,2} = \frac{\ln\left(\frac{F}{K}\right) \pm \frac{1}{2}\sigma^2(t_i - t_0)}{\sigma\sqrt{t_i - t_0}} \tag{1.35'}$$

σ is the percentage volatility of the forward rate, and the indices 1 and 2 correspond to the + and − signs, respectively. (See Chapter 5 for details of the derivation.) Notice that, despite the fact that the Black formula was originally derived for the particular case of an option on a futures commodity contract, in the context of this book any valuation formula for a call option obtained under the assumption of (i) log-normal distribution for the underlying variable, and (ii) absence of drift in the process for the percentage increment dx/x of the underlying variable will be referred to as a Black formula. The same terminology will therefore apply irrespective of whether the variable x is, for example, a forward rate or a forward bond price. Needless to say, stating that a certain forward rate is indeed driftless will have to be justified on the basis of no-arbitrage arguments (see Chapter 7).

Completely similar definitions hold for a floor(let), which can be seen as a put (rather than a call) on a forward rate.

Let us now consider a call expiring at time T and struck at $1/(1 + X\tau)$ on a discount bond maturing at time $T + \tau$. At option expiry its payoff will be given by

$$\text{Payoff} = \text{Max}\left[P(T, T + \tau) - \frac{1}{1 + X\tau}, 0\right]$$

$$= \text{Max}\left[\frac{1}{1 + R\tau} - \frac{1}{1 + X\tau}, 0\right] \tag{1.36}$$

where R now indicates the realisation at time T of the τ-period spot rate. Equation (1.36) can be rearranged to give

$$\text{Payoff} = \text{Max}\left[\frac{(X - R)\tau}{(1 + R\tau)(1 + X\tau)}, 0\right] \tag{1.37}$$

But, since the quantity $1/(1 + X\tau)$ is known at the outset, the payoff at time T of the call on the discount bond can be written as

$$\text{Payoff} = \frac{1}{1 + X\tau}\frac{\text{Max}[(X - R)\tau, 0]}{1 + R\tau} \tag{1.38}$$

i.e. it is identical, to within the proportionality factor $1/(1 + X\tau)$, to the payoff, paid at time $T + \tau$, from a put on the τ-period spot rate resetting at time T, or, in other terms, to a floorlet resetting at time T and paying at time $T + \tau$. Conversely, **a put (expiring and paying at time T) on the same discount bond is equivalent, to within the same proportionality factor, to a T-expiry caplet on the τ-period rate** (see Equation (1.34)). The equivalence just established between a caplet (floorlet) and a put (call) on a discount bond will be of great relevance in the context of the calibration of interest-rate models, whenever closed-form solutions are available for calls or puts on discount bonds. (See, in particular, Chapters 11, 13 and 14.)

Going back to Equation (1.35′), one can notice that, given the market practice of pricing caps using the Black formula, there is a one-to-one correspondence between the prices and the volatilities that enter the Black equation. It is therefore common in the market to express the value of a caplet in terms of the 'implied volatility'. One can therefore say that **the market expresses, via a complete set of cap prices, its views about the volatility of the underlying forward rates** and, at the same time, about how the imperfections of the Black model can be accounted for by adjusting the volatility input.

A European swaption is then defined as a contract that gives the holder the right at time t_i to enter a swap (i.e. to pay or receive the fixed rate over the life of the swap) of a given frequency, starting at time t_s and maturing at time t_m at a pre-known rate K:

$$\text{Max}[X - K, 0]B \quad \text{(payer's swaption)} \tag{1.39}$$

$$\text{Max}[K - X, 0]B \quad \text{(receiver's swaption)} \tag{1.39'}$$

with $B = \sum_{k=1,n} P(t_i, t_{i+k})\tau_k$. Swaptions are priced in the market assuming that forward swap rates, given by expression (1.25) of Section 1.3, are log-normally distributed, and using the Black model as applied to the forward swap rate X:

$$\text{Payer's swaption} = [X(t_0, t_i, t_{i+1})N(h_1) - KN(h_2)]B$$

(A justification of the theoretical soundness of this formula, i.e. of the reason why one is justified in assuming no drift for the swap rate, will be given in Chapter 7.) The relevant volatility for the Black formula as applied to swaptions is now the volatility of the forward swap rate X. It is shown in Appendix A that, despite the fact that B is in itself a stochastic quantity, as long as one can hedge one's position in the swaption using the forward bond B, the volatility of B does not enter the valuation formula. This is the exact counterpart of the statement that the volatility of the discount bond $P(t_0, t_{i+1})$ in Equation (1.35) does not enter the valuation formula for a caplet. It would be surprising if this were not the case, since a one-period swaption is a caplet. Notice, however, that since the swap rate is given by a linear combination of forward rates (Equation (1.20)), if one assumes the latter to be log-normally distributed — as implied by the Black formula as applied to caplets (Equation (1.35)) — then, strictly speaking, a swap rate cannot be log-normally distributed as well. In other terms, while both the cap and the swaption Black valuation formulae can be soundly justified *independently*, the simultaneous pricing of both caps and swaptions using the Black formula (as usually done in the market) is logically inconsistent.

The Black formulae for caps and swaptions reported above constitute the market standard for the valuation of these plain-vanilla instruments. Whatever the 'true' distributions might be, as long as the log-normal distributions are matched, as they are by the pricing procedure, to the first two moments, the impact of this inconsistency is quite small.

Having defined the payoffs of caps and swaptions, one can revisit Equation (1.20), which expresses a swap rate as a linear combination of forward rates. Let us denote the volatilities of forward rate F_i by σ_i, and let us make, for the moment, the assumption that these volatilities are constant over time. This important assumption, whose profound implications are discussed at length in Chapter 4, will be relaxed later on. If these forward rates F_i are correlated stochastic quantities characterised by a known covariance matrix (see Chapter 3), then, as long as one can regard the variability with interest rates of the coefficients $\{w\}$ to be much smaller than the variability of the forward rates $\{F\}$ (in general an excellent approximation), the variance of the swap rate is simply given by

$$\text{Var}[X] = \sum_i \sum_j w_i w_j \sigma_i \sigma_j \rho_{ij} \tag{1.40}$$

where ρ_{ij} is the correlation coefficient between forward i and forward j. Equation (1.40) shows that, if the covariance matrix of the forward rates is known, then the volatility of the swap rate can be immediately determined. Conversely, within the limits of the approximation stated above, from the variances of a complete series of swap rates one can, at least in principle, directly obtain the underlying covariance matrix. If the volatilities (standard deviations per unit time) of the swap rates are imputed from the traded prices of swaptions, then one can, at least in principle, obtain the market-implied covariance matrix. More precisely, let us consider two caplets, resetting at times t_1 and t_2, and spanning the periods τ_1 and τ_2. Their market prices will indicate (by inversion of Black's formula) the volatilities, σ_1 and σ_2, of the underlying forward rates, F_1 and F_2. See Figure 1.4. Let us then consider the market price of the option to enter the swap covering the period $\tau_1 + \tau_2$. This price will reveal the volatility of the swap rate (again by inverting Black's formula). Via Equation (1.28), and since σ_1 and σ_2 are known from the caplet prices, this will uniquely determine ρ_{12}.

Let us now extend the range of instruments to include a caplet resetting at time t_3. The new underlying forward rate F_3 will introduce correlation terms both with F_1 and with F_2, ρ_{13} and ρ_{23}. *Two* new swaptions are, however, brought into play by the additional forward rate, i.e. the option to enter the swap covering the period $\tau_1 + \tau_2 + \tau_3$, and the option to enter the swap spanning $\tau_2 + \tau_3$. Therefore, Equation (1.40), coupled with the assumption of constant swap rate volatilities, allows, at least in principle, the determination of the full covariance matrix by including the market volatilities of more and more swap rates. This apparently cumbersome procedure is of crucial importance in the calibration of two-factor models, and will be used, and referred to, extensively in following chapters.

In the present and previous sections (i) swap rates and bond prices have been defined; (ii) their valuation has been expressed in terms of pure discount bonds; (iii) the payoffs and the valuation of the plain-vanilla European options on FRAs and swaps have been defined; (iv) the link between the volatilities of FRAs and

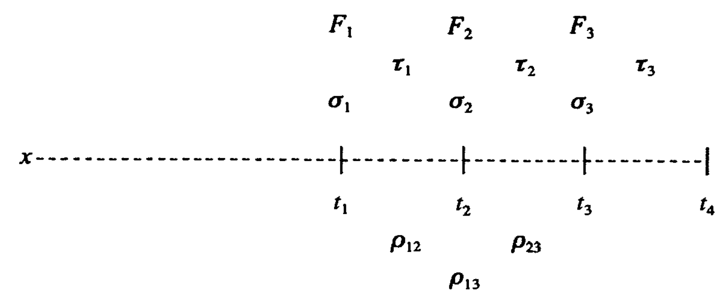

Figure 1.4 The market prices of caplets resetting at times t_1, t_2 and t_3 determine the volatilities σ_1, σ_2 and σ_3 of forward rates F_1, F_2 and F_3, respectively. The European swaption spanning the period $[t_1\ t_3]$ involves forward rates F_1 and F_2, and therefore brings in the correlation ρ_{12}. The European swaption spanning the period $[t_1\ t_4]$ involves forwards F_1, F_2 and F_3, and therefore brings in the additional correlations ρ_{13} and ρ_{23}. The European swaption spanning the period $[t_2\ t_4]$ involving forward rates F_2 and F_3 is only affected by the correlation ρ_{23}. Under the assumption of constant volatilities, the market prices of the three caplets and the three swaptions together therefore completely specify, at least in principle, σ_1, σ_2, σ_3 and ρ_{12}, ρ_{13}, ρ_{23}

swap rates has been highlighted. In order to compute the value of a plain-vanilla swap and of a bullet bond, however, it has been assumed that a market discount function, i.e. a continuum of prices $P(0, t)$, were known. These prices, however, are not directly available in the market, but have to be distilled from the prices of the instruments actually traded. The next section will therefore tackle the complementary problem of determining the prices $P(0, t)$ used in this section, or, as it is commonly said, of determining the market discount function. Strictly speaking, the **discount function** $P(0, t)$ is the function that gives the price today of a discount bond maturing at time t. The **term structure of interest rates** is the function associating to each maturity the value of the (suitably compounded) spot rate from time 0 to the desired maturity. The **yield curve** is the curve which displays the yields of coupon-bearing bonds as a function of maturity. Since, as is well known (see, e.g., Schaefer (1977) and Section 1.6), these yields depend in general on the coupon of the bond, the yields of par coupon bonds are normally used. Since any of these curves can be obtained from the other (see Section 1.6), the three terms will be used interchangeably.

Although conceptually similar, the creation of the bond and of the LIBOR discount curves are different in important ways. Bonds, in fact, are issued, either by corporate entities or by government issuers, at discrete time intervals with discrete maturities. At any point in time, therefore, one can obtain from the market the prices of bonds which had been priced at par on the day of their issuance, but currently trade at a discount or at a premium. In the swap market, on the other hand, every day one has access to the equilibrium swap rates for swaps of any of a range of benchmark maturities. It is as though par bonds of all these maturities were issued every day. Since, as will be shown in the next section, knowledge of the par coupon curve (i.e. of the curve describing the par coupon for any possible maturity) is tantamount to knowledge of the (discrete) discount function, it is clear that, for the LIBOR curve, obtaining the latter is little more than an exercise in skilful and careful, but conceptually straightforward, interpolation. Obtaining a reliable LIBOR discount curve is no trivial matter, but the difficulties mainly lie in accounting properly for different conventions in the deposit, futures and swap markets and in interpolating between benchmark maturities in an efficient and consistent way.

When it comes to the bond market, however, the comparative paucity of data changes the very nature of the problem, and forces upon the user the need to employ best-fit estimation techniques. As a result, while major trading houses agree on the LIBOR rates obtained from the discount factor typically to within a basis-point (i.e. a percentage of a percentage point) equivalent in their plain-vanilla rates, much bigger discrepancies arise in the bond curve estimation (to the point, for instance, that the *Financial Times* Gilt curve published on Mondays is always quoted with its source). The next sections therefore tackle the problem of showing how, at least conceptually, the two different term structures of interest rates can be obtained.

1.6 DETERMINATION OF THE DISCOUNT FUNCTION: THE CASE OF BONDS — LINEAR MODELS

Bonds are issued by government and corporate entities. Their market prices reflect not only the implied riskless discount function, but a series of important additional factors, such as the creditworthiness of the issuer, liquidity, tax regimes, institutional preferences or regulatory restrictions. All these are topics of great interest, and a thorough treatment would require a book in its own right. However, for the pricer of options in a given market (say, options on LIBOR instruments, or options on government bonds of similar coupons, or options on bonds issued by corporates of similar credit quality) the scope can be restricted by using homogeneous instruments for the estimation procedure. This is not a perfect, or even a wholly consistent, solution, but it is adequate for most practical purposes.

Even those investors most superficially acquainted with the bond market have certainly come in contact with the concept of gross redemption yield (GRY), i.e. the internal rate of return of a bond: if $Bnd(T)$ is the price of a bond of maturity T paying a coupon X every τ years, the GRY is defined as

$$Bnd(T) = \sum_{i=1,n} \frac{X\tau_i}{\left(1 + \frac{GRY}{\nu}\right)^{\nu i}} + \frac{100}{\left(1 + \frac{GRY}{\nu}\right)^{\nu n}} \tag{1.41}$$

where $\nu = 1/\tau$, and, to avoid unnecessary complications, it has been assumed that the valuation is made at the beginning of a coupon period (no accrued interest). The final nail in the coffin of the GRY was probably driven as early as 20 years ago by Schaefer (1977), who clearly showed the inadequacies of this bond statistic to convey any but the most imprecise information about the bond itself. Without repeating the arguments, it will suffice here to say that, by applying the concept of GRY to two bonds of different coupon and maturity, one **discounts payments occurring at the same point in time by different implied discount factors** (the terms $1/(1 + GRY/\nu)^{\nu i}$), since in general the GRY will be different for different bonds, **but one discounts payments from the same bond at different points in time by the same rate**. In reality one would wish to do exactly the opposite.

In order to provide a more useful description of the characteristics of a bond, one can employ the definition of the τ-period-compounded spot rates $R(0, t)$, given by

$$P(0, t) = \frac{1}{\left(1 + \frac{R(0, t)}{\nu}\right)^{\nu\tau}} \tag{1.42}$$

and thereby obtain a more satisfactory description of the yield curve, since these rates are no longer bond-specific, but can be used for discounting *any* cash flow occurring at a given point in time. While Equation (1.42) is simply a definition, an argument similar to the one presented in Section 1.2 (i.e. the discrete-time counterpart of Equation (1.6)) shows that the discount bond price $P(0, T)$ can be expressed in terms of forward rates as

$$P(0,T) = \frac{1}{\prod_{i=0,n-1} (1 + F(0, t_i, t_{i+1})\tau)} \tag{1.43}$$

where $T = n\tau$, and $F(0, t_i, t_i + \tau)$ is the forward rate obtainable from the term structure at time 0, spanning the period from time t_i to time $t_i + \tau$, i.e.

$$F(0, t_i, t_i + \tau) = \frac{P(0, t_i)/P(0, t_{i+1}) - 1}{\tau} \tag{1.44}$$

with $t_0 = 0$. A further equivalent description of the yield curve can be given by supplying the par-coupon curve. The par-coupon curve, it will be remembered from the previous section, is the bond equivalent of the equilibrium swap rate (Equation (1.25)), i.e. the ratio of the floating to the fixed leg. Avoiding again, for conceptual simplicity, the case of non-integer periods, and imposing, to lighten notation, $\tau = 1$, let us then consider the par coupon $X(1)$ paid by a bond maturing in one year's time; by rearranging Equation (1.24) one obtains

$$X(1) = \frac{P(0,0) - P(0,1)}{P(0,1)} = \frac{1 - P(0,1)}{P(0,1)} \tag{1.45}$$

whence $P(0, 1)$, i.e. the discount factor out to time t_1, can be immediately recovered; moving to a bond issued at par today and paying the par coupon $X(2)$, one can write the ratio of the floating to the fixed leg as

$$X(2) = \frac{1 - P(0,2)}{P(0,1) - P(0,2)} \tag{1.45'}$$

but, since $P(0, 1)$ is known from Equation (1.45), $P(0, 2)$ can be readily solved for. By this boot-strapping procedure the whole par-coupon curve can be obtained. Therefore, at least conceptually, **supplying the par-coupon curve is equivalent to supplying the discount function** (at least at the discrete points where the coupons are paid).

We have arrived at the conclusion that, while from the *GRY* the discount curve cannot be recovered, **either the set of spot rates $R(0, t)$, or the set of forward rates $F(0, t, t + \tau)$, or the par-coupon curve all give access via simple manipulations to the discount function**. Despite the fact that conceptually there is therefore no distinction between one description and the other, practically the differences are important, since the discount function can be estimated using linear methods, while any of the other equivalent quantities require non-linear procedures. This can be seen as follows.

Let us expand the discount function $P(0, t)$ on the basis of an arbitrarily chosen set of s basis functions $g_k(t)$:

$$P(0,t) = \sum_{k=1,s} \alpha_k g_k(t) \tag{1.46}$$

where α_k is the (as yet unknown) weight corresponding to the kth basis function. Therefore the price of a bond of maturity T_n, Bnd_n, paying (again for simplicity of notation) annual coupons can be written as

$$\begin{aligned} PV(Bnd_n) &= \sum_{i=1,n} XP(0, t_i) + 100P(0, t_n) \\ &= \sum_{i=1,n} X \sum_{k=1,s} \alpha_k g_k(t_i) + 100 \sum_{k=1,s} \alpha_k g_k(t_n) \\ &= \sum_{k=1,s} \alpha_k G_{n,k} \end{aligned} \tag{1.47}$$

with $G_{n,k} \equiv \sum_{i=1,n} X g_k(t_i) + 100 g_k(t_n)$, or, in vector notation,

$$\mathbf{Bnd} = \mathbf{GA} \tag{1.47'}$$

where **Bnd** is the (*nbonds*,1) vector containing the prices of the *nbonds* bonds available in the market, **A** is the $(s, 1)$ vector containing the coefficients of the s basis functions, and **G** is the (*nbonds*,*s*) matrix of coefficients $G_{n,k}$ of Equation (1.47). Equation (1.47') justifies the term 'linear' applied to this type of approach.

If $nbonds = s$, and providing that $\text{Det}[\mathbf{G}] \neq 0$, the system provides a unique and 'perfect' solution, where 'perfect' in this context simply means that all the bonds are exactly priced. More generally, and realistically, $nbonds \gg s$, and therefore

$$\mathbf{A} = (\mathbf{G}^T\mathbf{G})^{-1}\mathbf{G}^T\mathbf{Bnd} \tag{1.48}$$

will provide the least-square estimator of A, i.e. of the set of coefficients $\{\alpha\}$ such that the sum of the square of the differences between the market and the model prices, i.e. the prices calculated using the last line of Equation (1.47), is minimised. (The superscripts T and -1, as usual, indicate the transpose and the inverse of a matrix.) It is also very easy to incorporate constraints such as the requirement that the discount function should be equal to 1 at the origin, $P(0, 0) = 1$, either via the formal use of Lagrange multipliers, or by solving explicitly for one of the coefficients in Equation (1.47).

Since even low-power personal computers can effect the matrix inversion required by Equation (1.48) in a few seconds, the linear approach just described seems to afford a very efficient and rapid way of estimating the discount function. In practice, however, this strategy is fraught with practical difficulties, largely stemming from the very nature of the target function, i.e. the discount function $P(0, t)$. In order to be able to say something more precise, one must examine more carefully the basis functions used for the expansion of the discount function. Many different choices have been proposed (cubic splines, B-splines, Chebychev polynomials, exponential splines, Bernstein polynomials, etc.) — see, e.g., Anderson *et al.* (1996). In general, all these basis functions promise, and indeed deliver, a continuous and continuously differentiable discount function at every point. Mathematical continuity can, however, bear surprisingly little resemblance to an intuitive requirement of continuity, especially when it comes to derivatives. These in turn are particularly important since it has been shown in Section 1.2 that

$$f(0, t, t+\varepsilon) = -\frac{\ln P(0, t+\varepsilon)}{\ln P(0, t)}\frac{1}{\varepsilon} = -\frac{\ln P(0, t+\varepsilon) - \ln P(0, t)}{\varepsilon} \simeq -\frac{\partial \ln P(0, t)}{\partial t} \quad (1.49)$$

i.e. **forward rates are given by (minus) the logarithmic derivative with respect to maturity of the discount bond covering the period from time t to time $t + \varepsilon$.** Unfortunately, functions of very smooth appearance can have still continuous, but not so smooth, derivatives, as shown in Figure 1.5.

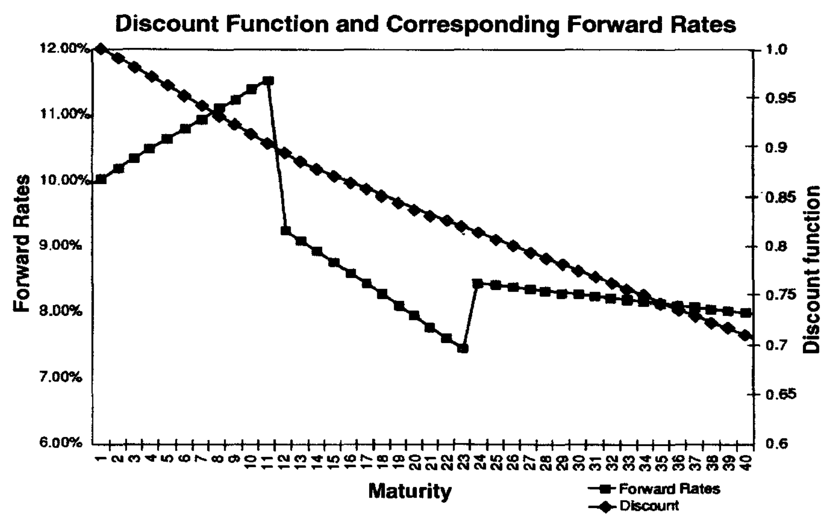

Figure 1.5 A smooth discount function, and the not-so-smooth forward rates that it implies

If this is the case, forward rates will also display a rapidly changing, and therefore intuitively implausible, behaviour. For the same reason, forward rates are prone to display a very pronounced dependence on such 'technical details' as the positioning of the knots for B-splines, the number of basis functions, or small variations in the vector of input bond prices. A thorough treatment of these important issues would require more than a chapter. The essential point that should have been conveyed, however, is that direct modelling of the discount function gives rise to a mathematically simple but practically very delicate exercise. The success of the approach will largely depend on the intrinsic suitability of the basis function chosen to describe the behaviour of the discount function. To use a not terribly original, but once again very apt, expression, obtaining reliable forward rates from fitted discount functions can be more of an art than a science.

1.7 DETERMINATION OF THE DISCOUNT FUNCTION: THE CASE OF BONDS — NON-LINEAR MODELS

The difficulties highlighted in the previous section have prompted some researchers to model the spot or forward rates themselves, rather than the discount function. More precisely, one can attempt to expand, say, the continuously compounded spot rates $R(t)$ using a chosen set of basis functions:

$$R(t) = \sum_{k=1,s} \alpha_k g_k(t) \quad (1.50)$$

giving rise to the present value of the bond price (neglecting again for simplicity non-integer periods):

$$\begin{aligned} PV(Bnd_n) &= \sum_{i=1,n} XP(0, t_i) + 100P(0, t_n) \\ &= \sum_{i=1,n} X \exp(-R(t_i)t_i) + 100\exp(-R(t_n)t_n) \\ &= \sum_{i=1,n} X \exp\left[\left(\sum_{k=1,s} \alpha_k g_k(t_i)\right) t_i\right] + 100\exp\left[\left(\sum_{k=1,s} \alpha_k g_k(t_n)\right) t_n\right] \end{aligned} \quad (1.51)$$

Once again the goal will be to change the coefficients $\{\alpha\}$ of expansion (1.50) until, for instance, the sum of the squares of the deviations between model and market bond prices will be minimised. Since, however, the coefficients now appear as arguments of an exponential function, the problem is no longer linear, and there exists no strategy guaranteed to yield the absolute minimum (i.e. the set of parameters that actually minimises the squares of the deviations). Non-linear minimisation algorithms are considerably more laborious to implement.

More importantly, those methods that ensure fast convergence are also the most prone to get stuck in local minima. (Conjugate gradients and simulated annealing are two apt and extreme examples of this feature.) To make matters worse, only the most inefficient methods (multi-dimensional simplex) permit the search for a minimum without knowledge of the derivatives of the target function with respect to the optimisation parameters. All the methods which ensure reasonably quick convergence do require (numerical or analytic) evaluation of the derivatives. These, in turn, can be quite time consuming to evaluate, and often algebraically very tedious. 'Smart' methods (such as conjugate gradients) promise stunningly good convergence properties: for a truly parabolic minimum the number of iterations needed should be exactly equal to two! In reality, no minimum is truly parabolic, and fast methods tend to get very close to their target very soon, and use an often infuriatingly large number of iterations to home in to the 'true' solution within the required degree of tolerance.

Slower methods 'ramble around' the coefficient space in a less purposeful manner in the first stages of the optimisation, but, because of this very weakness, are in general better able to explore competing minima.

In the cases of both linear and non-linear methods, the number of basis functions to be used to model either the discount factors or the (spot or forward) rates should be determined on the basis of an enlightened compromise between an oversimplistic description of the target function, and the dangerous tendency to chase every mispriced bond in the market.

For linear models (and, with some qualifications, for non-linear models as well), a useful indication of how many coefficients one 'should' use can be provided by the $\overline{R}^2$ criterion. Let us first define, for a quantity y (the target function, i.e. in this case the bond prices) 'explained' by k variables x_i, $(i = 1, k)$, the quantity R^2 as the ratio between the variance of y explained by the regressors and the total variance of y, i.e.

$$R^2 = \frac{\sum[y(mod) - y(avg)]^2}{\sum[y(obs) - \dot{y}(avg)]^2} \tag{1.52}$$

where $y(mod)$ indicates the model prediction for the bond price, $y(obs)$ the corresponding market price, and $y(avg)$ the average of the bond prices. Then $\overline{R}^2$ is given by

$$\overline{R}^2 = \frac{R^2 - k(n-1)}{(n-1)/(n-k-1)} \tag{1.53}$$

where k is the number of regressors (parameters), and n the number of observations (the number of bond prices). The imprecise but intuitively suggestive interpretation of the maximum of this statistical indicator is that it displays the 'optimal' number of parameters needed to describe a certain set of data.

Once again, the issue of non-linear minimisation cannot be dealt with in a satisfactory way in a limited space. More simply, the purpose of this section

is to highlight the type of issues faced by those users who need to model the bond discount function in order to calibrate the interest-rate models described later on. The different problems encountered when the underlying instrument is LIBOR-based are mentioned in the following section.

1.8 DETERMINATION OF THE DISCOUNT FUNCTION: THE CASE OF THE LIBOR CURVE

Every day interest-rate market screens display information about three different but interlinked markets: LIBOR borrowing rates, 3-month futures prices and equilibrium swap rates. (Equilibrium swap rates have been defined and treated in Sections 1.2 and 1.3.)

LIBOR (London Inter-Bank Offered Rates) are the rates at which high-credit financial institutions can **borrow** in the interbank market for a series of possible maturities, ranging from overnight to, usually, 12 months. (The lending rate is denoted by the acronym LIBID.) In most markets, it is conventional to quote these rates on a **simple compounding** basis.

As for interest-rate (short money) futures, at expiry they assume the value of 100 minus (the 3-month LIBOR at expiry multiplied by 100). Day by day, the buyer and seller of a futures contract make or receive a payment in the so-called margin account equal to the difference between the strike price X and the prevailing market value of the futures contract. On expiry date the cash flow occurring between the buyer and the writer of the futures contract is simply the last of the margin-account payments. In other words, all gains and losses are realised, day by day, as they occur. Positive balances in the margin account accrue interest. Even neglecting the effect of this reinvested cash, the equilibrium value of a futures contract (i.e. the value that gives zero value to the contract itself) is different from the equilibrium rate of the underlying FRA. In the latter case, in fact, the payment (notionally)[2] occurs three months after the expiry of the contract itself. Since, at inception of the contract, the future discount factor from expiry to three months hence is unknown, the relationship between the equilibrium FRA rate and the implied equilibrium futures rate cannot be achieved by discounting by a deterministic quantity (i.e. by a quantity known at contract inception). The correction term which accounts for this difference in payoff timing can easily be obtained using the formalism presented in Chapter 7, i.e. by making use of suitable numeraires, and depends on the volatility of the underlying forward rate and on the correlation between this forward rate and the discounting zero-coupon bond. This correction term is normally referred to as the 'futures/FRAs adjustment'. Very good approximations exist for this term (see, e.g., Doust (1995), Vaillant (1995)). For very long-expiry contracts (the US$ market has 3-month futures prices extending to 10 years!) the difference between (100 – futures price)/100 and the FRA rate can be as large as 30 basis points, for reasonable choices of the volatility and correlation inputs (see Vaillant

(1995)). The initial margin — i.e. the amount of money which both parties have to pay into the margin account at contract inception — is not accounted for by this correction term, but is of smaller import.

It is not the purpose of this chapter to go into the intricacies of exact futures pricing. It should be clear, however, that a great amount of care has to be exercised in order to splice together information available from different markets, and that there exists a spectrum of maturities over which more than one type of instrument is traded (LIBOR deposit rates overlap the first futures contracts, and the later-expiry futures contracts overlap the swap market). In theory, once proper account is taken of the exact payoffs, it should be immaterial which type of instruments should be used, since no-arbitrage should ensure the equivalence of the results, however obtained. In practice, not all instruments enjoy the same degree of liquidity; it is therefore common practice to make use of the most liquid assets over the different segments of the yield curve. For many markets this often means using deposit rates to reach the expiry of the first futures contract, and then using a market-dependent number of liquid futures contracts before switching to the market swap rates. Due to this overlap, careful interpolation techniques must be used in order to ensure a 'smooth transition' (and therefore, as discussed above, smooth forward rates). However delicate these techniques might be (see, e.g., Miron and Swannell (1991)), it is clear that the problem is, at least conceptually, much simpler than the estimation of the bond discount function: in the LIBOR market, as mentioned above, day by day there exists the equivalent of a whole series of par-coupon bonds. Not surprisingly, as mentioned above, the agreement amongst practitioners (or, at least, 'good' practitioners) about the LIBOR curves is far greater than the agreement about the government par-yield curve.

For the remainder of the book, the 'market' discount function will therefore be taken as uncontroversially available to market participants, despite the fact that, in the light of the last two sections, the assumption is actually more justifiable for the LIBOR curve. The first sections of this chapter have made use of this market discount function to price some elementary underlying instruments. In the chapters of Part Four the price of discount bonds will play a crucial role in the calibration of the various option models.

ENDNOTES

1. Different markets have different conventions (Actual/360, Actual/365, 30/365, etc.) to measure the fraction of the year. Details of the different conventions are clearly presented in Miron and Swannell (1991).
2. In practice, FRAs settle in cash at the start of the accrual period (time t_i) on the basis of the present value of the maturity (time t_{i+1}) payoff:

$$\text{Payoff}(t_i) = (R_i - K)\frac{1}{1 + R_i\tau}$$

2

Exotic interest-rate instruments: description and valuation issues

2.1 INTRODUCTION

In this chapter important classes of exotic interest-rate instruments are presented with a view to illustrating their financial rationale; to describing them as simply as possible in terms directly amenable to the treatment to be found in the subsequent chapters; and to highlighting the valuation issues that must be taken into account when applying in practice any of the models described in Part Four. No attempt will be made to indicate '*the* right way' to price any of these complex instruments, since it is the contention of the writer that such a question in isolation is essentially misguided: if there were *one* correct approach it would, after all, hardly make sense to study and apply several pricing models. *All* the models available at the moment are, in some important respect, 'wrong', and it is only by appreciating their shortcomings that intelligent use of them can be made. The emphasis of this book is therefore in furnishing the reader with the analytic tools necessary to strike an informed balance between such incommensurable criteria as realism of description of financial reality, robustness of parameter estimation, ease of calibration, etc. Rather than presenting a recipe book, the attempt has been made, in this and the following chapters, to illustrate what the essential features of a product or a model are, so as to achieve as effective a 'match' between the two as possible.

In this light, it has also been deemed pointless to present a catalogue or 'encyclopedia' of all the known, or even important, exotic interest-rate products. Rather, a choice has been made by selecting those representative types of product that share the same pricing issues of a whole class: for instance, LIBOR-in-arrears swaps have been described in some detail, despite the fact that, *per se*, they are not very important products from the financial point of view, because they introduce in the simplest manner the pricing issues shared by the 'hedge-payoff-mismatch'