1 Product Fundamentals

1.1 CHAPTER OVERVIEW

In this chapter we consider the features of a number of instruments that will be the focus of subsequent sections. The coverage is not intended to be comprehensive; the aim is to make sure that the reader is armed with sufficient terminology to be able to understand the more detailed concepts that will follow. Pricing and risk management will be the subject of Chapters 2 and 3, respectively.

This chapter starts with a discussion of the main "cash" (i.e., non-derivative) markets of fixed income, inflation and credit. The coverage then widens to incorporate the derivative building blocks, namely futures, forwards, swaps and options. Within this section the material occasionally leans towards the detail of specific products in certain asset classes that are considered key. However, the discussion relating to options is asset class neutral to keep the chapter size manageable.

Readers with a good knowledge of these subjects can skip this chapter but we would suggest a quick skim of the pages just in case a review is needed!

1.2 BOND FUNDAMENTALS

A key building block for the first part of the text will be bonds. A bond is an IOU that evidences the indebtedness of a borrower. Borrowers comprise mainly sovereign and corporate entities, although there have been issues made by individuals such as the pop star David Bowie.

1.2.1 Fixed income structures

Although bonds have many different forms we will initially focus on standard ("vanilla") structures. In return for borrowing a given sum of money, the issuer of the bond will pay a series of contractual interest payments to the owner of the instrument. When bonds were issued in physical form, the owner would detach a small coupon and present this to a bank appointed on behalf of the borrower as their eligibility to receive interest. As a result of this practice, interest payments on bonds have become termed coupons. At the maturity of the instrument the investor will be repaid the value stated on the face of the bond, but this may not be the sum that was originally paid to acquire the asset. This is because bonds are traded on a price basis, which is quoted as a percentage of the face value. Bonds are priced by present valuing all of the future cash flows, but this concept will be considered in Chapter 2. Suffice to say that with a limited amount of any bond in issue, the relative attractiveness of the fixed coupon will be the key determinant of how much an investor will pay to acquire the bond. If a bond has a fixed coupon of 5% but investors could earn a greater return on an equivalent investment (equivalent in terms of maturity and the risk of default), the bond will have to be priced at less than its face value in order to make the investment attractive. If it were priced at say 95.00 and the investor held the instrument to maturity, they would be repaid 100% of the face value and would enjoy a capital gain of just over 5% over the period. The opposite would be true for a bond that has a relatively attractive coupon. Through the interaction of demand and supply, investors will seek to possess the bond, which will drive up its price. If held to maturity the investor will incur a capital loss but will have earned an above-market interest rate. The market uses the concept of a yield, which captures any capital gain or loss in addition to the receipt of a particular coupon.

1.2.2 Floating-rate notes

Floating-rate notes (FRNs) are interest-bearing securities that pay a variable coupon on a regular basis (usually quarterly). The coupon is usually a spread to a given margin relative to an interest rate index such as LIBOR (London Interbank Offered Rate) or Euribor. For example, the instrument may pay 3-month USD LIBOR + 0.15% (15 basis points). The instrument is economically equivalent to a series of consecutive fixed-term bank deposits, where the interest rate is reset on a periodic basis. The fixed percentage margin over the specified interest rate index is referred to as the quoted margin. The quoted margin is a function of the issuer's default risk relative to the interbank rate to which the interest payments are referenced. The better the credit rating the lower the quoted margin and vice versa.

FRN issuance is driven by the desire of the issuer to match their assets and liabilities. For example, banks will tend to be big issuers of FRNs (which will represent a liability) as the assets that the bond proceeds are used to purchase will tend to pay a variable rate of interest (e.g., mortgages). This ensures that if interest rates change, interest costs and income will move in tandem. The concept of banks being able to borrow on a LIBOR basis will become key to much of the analysis that follows. This is because investment opportunities are often analysed based on the return they generate relative to LIBOR. FRN investors will include many different entities:

- Bank treasuries with excess cash who are looking to match floating-rate liabilities.
- Central Banks, retail investors and credit-conscious fund managers will buy sovereign-issued FRNs.
- Money market funds and corporates can earn an enhanced yield compared to alternatives such as cash and commercial paper.

1.2.3 Inflation

Definitions

Although most people would argue that they understand the concept of inflation, both authors have found that in reality a number of market participants often struggle when trying to verbalize a definition. Inflation represents rising prices, deflation falling prices and disinflation is where price increases slow down.

Within the inflation world a nominal frame of reference looks at investments in terms of cash paid without taking into account the loss of purchasing power. So if an item costs €1 today, with 2% inflation it will cost €1.02 by the end of the year. Alternatively we could say that at the end of the year, €1 will only buy 0.98 of the item. How would this relate to bonds? Consider a 1-year bond that pays a principal of €100 plus one interest payment of €5 at its maturity. The real value of this final cash flow will depend on what happens to prices over the period. If an investor expected inflation to be 3% then it will cost €103 in 1 year to buy something that costs €100 presently. However, the bond will pay a cash flow of €105 and so you expect to have €2 of extra purchasing power – a 1.94% increase in purchasing power.

The Fisher equation is used extensively by the market to express the relationship between the yields on nominal bonds and expected inflation. The equation expresses the relationships as:

$$(1+n) = (1+r)(1+f)(1+p)$$

where:

n = yield on nominal bond

r = real yield on inflation-linked bond

f = inflationary expectations

p = risk premium

However, the market has shortened the expression:

$$n = r + f + p$$

$$n = r + \text{bei}$$

where:

bei = breakeven inflation

In essence, the formula states that the yield on a nominal bond is made up of three components:

- A required real yield that investors demand over and above expectations of inflation.
- Inflationary expectations over a particular period of time ("breakeven inflation").
- A factor that captures the combination of a risk premium and a liquidity discount.
  - The risk premium is the compensation an investor earns for accepting undesirable inflation risk when holding nominal bonds. One interpretation is that it represents the risk premium demanded by nominal bond investors for unexpected inflation.
  - The liquidity discount represents the yield premium that investors demand to hold a less liquid inflation-linked bond.

However, the third component is generally considered to be difficult to disaggregate and so is generally ignored by the market.

The breakeven rate can be thought of as the average rate of inflation that will equate the returns on an inflation-linked bond and a comparator nominal bond issue of the same return. To illustrate how it should be interpreted, consider the following example. Suppose there are a nominal 5-year sovereign bond that is yielding 4.5% and an inflation-linked sovereign bond of the same maturity whose yield on a real basis is 1.5%. Using the principles of the Fisher equation this implies a breakeven inflation rate of 3.0%. An investor could use the value of breakeven inflation to assess which bond should be purchased:

- If the investor expects inflation to average less than 3.0% over the period, they should hold the nominal bond.

- If the investor expects inflation to average more than 3.0% over the period, they should hold the inflation-linked bond.
- If the investor expects inflation to average 3.0% over the period, they will be indifferent between the two assets.

Arguably the difficulty experienced by practitioners in trying to grasp the concept of inflation lies in defining the concept of a real yield. If one looked at the Fisher equation, a simple but somewhat unsatisfactory definition of real yields is simply the difference between nominal yields and inflation expectations. We present three other definitions:

- A real rate of interest reflects the amount earned or paid after taking into account the impact of inflation.
- It is the market clearing rate of return in excess of expected future inflation that ensures supply meets demand for a particular investment opportunity.
- The return for forgoing consumption today to consume more goods and services tomorrow.

Real yields should also:

- Reflect the growth in an economy's productivity.
- Represent the rate at which investments are rewarded. Investments compete for capital on the basis of the real yield they offer given their associated risk.

What can be even more confusing is when real rates of interest become negative, an example of which occurred in the US Treasury market in 2008. This happened when inflation expectations were higher than nominal interest rates. These negative real yields were attributable to:

- Slower economic growth prospects, which lowered rates of expected returns across investments.
- The US Federal Reserve was expected to cut interest rates such that inflation would be greater than nominal rates.
- A "flight to quality" by investors, which drove up the price of government securities, reducing their nominal returns.

So in general terms, negative real yields could occur if:

- An asset is not considered a productive use of capital.
- The asset is attractive but faces excess demand relative to its supply. As a result, its price rises and the nominal return falls.
- The existence of negative real yield can create an incentive to drive capital to other more potentially attractive investments.

Inflation-linked bonds

An inflation-linked bond is one whose value is linked to movements in a specific price index in order to maintain its purchasing power. An inflation index measures the way in which prices change. This is achieved by analysing and recording thousands of prices for a selection of goods and services on a monthly basis. Inflation figures for a particular month are then typically issued two to three weeks later. Some of the goods and services will carry a higher weighting, reflecting the fact that consumers will spend more money on some items than others. The basket and the constituent weightings are revised on an annual basis. The most common inflation index used is the consumer price index (CPI) for the respective country of issue, although each country will typically calculate and quote a number of indices. In the USA the "Treasury Inflation Protected Securities" (TIPS; also sometimes referred to as the Treasury Inflation Indexed Securities – TIIS) reference their return to the consumer price all urban non-seasonally adjusted inflation index. In Europe a common index is the Harmonised Index of Consumer Prices (HICP) for all items excluding tobacco, while the UK mainly uses the Retail Price Index (RPI).

1.3 REPURCHASE AGREEMENTS

One important aspect of the fixed income world relates to how the purchase of a bond will be financed. It would be fair to say that most banks will not have large piles of cash lying around idle and so will look to manage their cash efficiently. The implication is that the cash required to purchase an asset will need to be borrowed and the proceeds from any sale will be reinvested. The most popular technique used in the financing of fixed income transactions is the repurchase agreement or "repo". A repo involves the simultaneous sale and future repurchase of an asset. The seller of the asset buys it back at the same price at which it was sold. On the second leg of the transaction the seller pays the buyer interest on the implicit loan that has been created. This interest is termed the repo rate.

The main cash flows associated with a typical repurchase agreement are illustrated in Figure 1.1.

A reverse repo is the opposite of a repo. From this perspective the transaction is viewed as the purchase of an asset for cash, with an agreement to resell at some future date. The market distinguishes between two different types of repo in relation to the asset that is transferred. A specific repo involves a bond that is specified by the two counterparties, whereas a general collateral ("GC") transaction involves a bond that meets some pre-agreed criteria.

Economically, the repo can be viewed as a collateralized loan rather than a pair of securities trades. Legally, however, the transaction is a sale and repurchase, which will

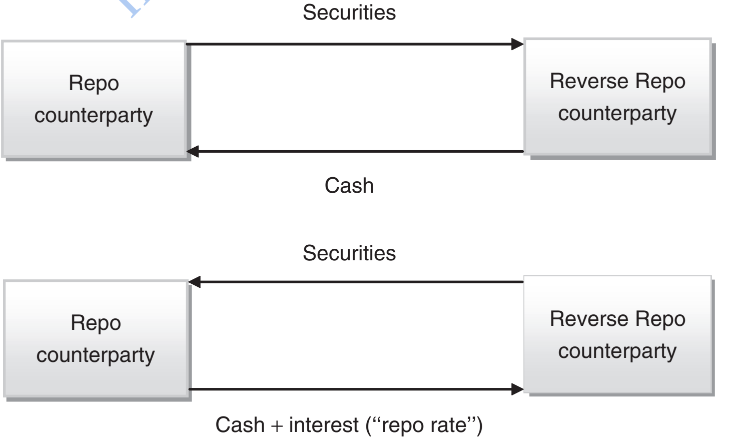

Figure 1.1 Repurchase agreements.

have important implications in the event of the default of one of the counterparties. If the securities had merely been pledged, then the default of the repo counterparty would result in the reverse repo counterparty becoming an unsecured creditor. However, the sale and repurchase structure means the reverse repo counterparty has the right of close out and set off – they get to keep the securities in lieu of the money lent. Similar principles would apply if the reverse repo counterparty were to fail.

Repos are quoted on a bid and offer basis. From a quoting institution's perspective a quote may be expressed as:

Bid	Offer
3.98%	3.92%
Buy securities	Sell securities
Earn interest	Pay interest

Although the convention of a high bid/low offer price may appear counter-intuitive, it allows for the market maker (i.e., the quoting institution) to make a profit through earning more interest than they would pay if they were able to execute offsetting trades simultaneously.

Appreciating that the transaction economically resembles a collateralized loan gives an insight into the popularity of the transaction. The interest that is payable on the second leg of the transaction will be lower than that of an unsecured borrowing and as a rule of thumb the rate that is agreed by the two counterparties is about 1/8th less than the LIBOR rate of the equivalent maturity.

It is important to appreciate that the legal title of the bonds is transferred to the reverse repo counterparty as part of the first leg of the transaction. This will allow them to sell on the bonds as part of an unrelated transaction if necessary. However, any economic benefit or risk is retained by the repo counterparty. This has a number of implications:

- If the bond issuer defaults over the period of the repo they will receive the security back but will still be forced to repay the price agreed in the first leg of the transaction.
- If the issuer of the bond being repo'd defaults, the repo counterparty will receive back the asset but will still be obliged to pay the original price agreed on the first leg of the transaction.
- If the bond pays a coupon during the period, this will have to be remitted back to the repo counterparty immediately.

Suppose that a bank is bullish on the prospects of the value of a particular bond and decides to use the repo mechanism to finance its purchase. The steps in the transaction are:

- Buy the bond for an agreed value and an agreed cash amount (an outright purchase).
- Sell the bond under repo and receive the market value with an agreement to repurchase the bond at a future date.
- The cash proceeds received from the first leg of the repo are used to settle the outright purchase.
- When the repo matures the bank retakes delivery of the bond and then sells it in the open market to any counterparty.
- The proceeds received from this sale are used to settle the outstanding principal and interest amount due under the repo.

It can be confusing as to why an investor would buy a bond outright and then sell it under repo to pay for it. However, the key to grasping the logic of this trade is to recall that all of the economic benefit of the transaction is retained by the repo seller (i.e., the outright buyer of the bond). So as long as the final sale generates sufficient cash to cover the initial purchase and the interest on the repo, the transaction will show a profit.

A similar procedure could be used if the market participant thought that a particular bond was going to fall in value:

- The target bond is purchased under a repo transaction.
- The bond is sold to a market participant in an outright sale.
- At the maturity of the repo the trader buys back the bond in the market to satisfy his commitment to redeliver the bond under the second leg of the repo.
- The proceeds of the repo (initial price plus interest received on the cash leg) are used to pay for the purchase of the bond.

As in the bullish scenario, as long as the cash received from selling the bond is greater than the cash paid to buy it, the transaction will be profitable.

Although this section is designed to give the reader an awareness of the key issues associated with a repurchase agreement, there is one particular aspect of the market that is worth highlighting. On occasion certain bonds will be in very high demand in the market and as a result the asset will "go on special" in the repo market. The excess demand for the bond may occur as a result of traders being very bearish in relation to a particular issue and there is significant demand to obtain the bond using the repo mechanism. Another example, which will be considered later, is that the bond futures contract may require a particular government bond to be delivered if it is held to its final maturity.

The impact of specialness in the repo market will result in repo rates going down. Intuitively, it would seem that the relative scarcity of an asset would cause rates to rise, but this is not the case. The participant who needs to take delivery of the asset will buy it under repo and deliver cash in return. Given the scarcity of the asset the cash that he has now lent out will only earn a very low rate of interest; this is the "cost" he must pay. Looked at from the repo seller's perspective, if they own the asset, they are able to profit from its scarcity by borrowing money at very low rates of interest. Depending on the level of demand for the asset, it is possible for the repo rate to turn negative; that is, the buyer of the bond in the repo transaction gets back less cash than they initially forwarded. This would occur if the penalty costs for failing to deliver are greater than the reduction in their repo proceeds.

1.4 CREDIT FUNDAMENTALS

One fundamental distinction made in the fixed income world is the importance of credit risk. This is defined as the risk that an entity will be unable to repay interest or principal due on monies that have been borrowed. The probability that an issuer will repay a particular debt is assessed by independent rating agencies, of which Standard and Poor's, Moody's and Fitch are examples. For example, Standard and Poor's defines a credit rating as an "independent opinion of the general creditworthiness of an obligor or an obligor's financial obligation based on relevant risk factors".[1] Each of the rating agencies applies different methods to assess this creditworthiness and express it using a mixture of letters and numbers. For example, Standard and Poor's express credit ratings for both short and long-term instruments. For long-term credit ratings, the AAA designation reflects the strongest credit quality while D reflects the lowest. It is also possible to add a degree of granularity to the credit ratings by adding a plus or minus sign to show the relative standing with the major rating categories from AA to CCC.[2] Obligations rated as BBB− or better are termed by the market as "investment grade", while ratings lower than this threshold are termed "high yield". This is an important distinction, as some investors may have restrictions on the nature of the assets in which they can invest.

From a market perspective an investor who buys a bond with a certain element of credit risk is rewarded in the form of an enhanced return. That is, they will earn a certain percentage amount over and above the so-called default-free return. This enhanced return is referred to as a credit spread. A generalized approach to estimating this spread can be stated in the following relationship:

$$\text{Credit spread} = \text{Probability of entity defaulting} \times \text{loss incurred in the event of default}$$

However, there are a number of different ways in which this credit spread is measured, and this will be addressed in Chapter 3.

1.5 DERIVATIVE FUNDAMENTALS

A derivative is defined as an instrument that derives its value from the price of an underlying asset. The three main building block instruments that comprise the derivative world are forwards/futures, swaps and options. So, taking crude oil as an example, the market trades crude oil futures and forwards, crude oil swaps and crude oil options. Derivatives can be traded on an organized exchange or directly between counterparties on an over-the-counter (OTC) basis.

1.5.1 Futures

A future is an exchange-traded contract that fixes a price on the trade date for delivery of an asset at some future time period. An interest rate future fixes an interbank rate for some future time period – say the 3-month rate in 3 months' time. A bond future fixes the price of a bond for delivery at some future time period. An example of a bond future referenced to German sovereign Bunds is given in Table 1.1.

Table 1.1 Contract specifications for Euro Bund future

Trading unit	€100,000 nominal value, notional Bund, 6% coupon
Delivery months	March, June, Sept, Dec
Delivery day	The 10th calendar day of the respective delivery month (at seller's choice)
Quotation	Per €100 nominal (in decimals to 2 places)
Minimum price movement	0.01 (1 tick = €10)
Last trading day	11.00 a.m., two trading days prior to delivery date

Source: Eurex.

Although the detail of the Bund future will be considered later, there are a number of general features that are worth highlighting:

- Futures are generally traded in fixed amounts (€100,000 in this case), although there are exceptions to this such as futures on equity indices. The monetary value of this type of future changes in line with the value of the index.
- The contract is linked to a specific underlying asset so that both counterparties know exactly what will be delivered.
- Upon expiry of the contract the underlying can be delivered according to an agreed schedule of dates (in the case of the Bund it expires on the 10th calendar day of March, June, September and December).
- The underlying asset may be physically delivered (e.g., Bund futures) or cash settled, where the nature of the underlying asset makes it operationally impractical (e.g., FTSE 100 equity index).
- The smallest price movement is predefined by the exchange and is referred to as a "tick". This tick movement will have an associated monetary value. In the case of the Bund, since the contract size is €100,000 and the tick is defined as 0.01%, the tick value is €10.

Another feature of exchanges is the requirement of both counterparties to post collateral. Termed "margin", this is generally seen in two forms. Initial margin is posted at the outset of the trade, while variation margin is the mechanism whereby profits and losses are transferred between entities on a daily basis. To facilitate the settlement of exchange-traded contracts, a central clearing house will act as the counterparty to both sides of the transaction. So once a transaction is executed between two entities, the clearing house will become the buyer to every seller and the seller to every buyer. This feature removes the counterparty credit risk that would result if a transaction were executed on an OTC basis. However, it is clear that this argument is somewhat flawed in that each original party to the trade has merely transferred its credit exposure to the clearing house. However, the clearing house is often very heavily capitalized in order to mitigate this potential default risk.

1.5.2 Forwards

A forward contract is economically equivalent to a futures contract in that it will involve the fixing of a price at the point of execution for delivery at some future date. An entity trading an OTC forward will not be faced with the constraints of contract standardization that are a feature of exchange-traded contracts. Forwards allow the user greater flexibility in specifying deal parameters such as transaction size and maturity dates. Although something of a generalization, the majority of forward contracts will be cash settled. So, a cash-settled bond forward would fix the price of the bond for future delivery but the final settlement would not require the exchange of the asset for cash. Instead, the seller of the contract will pay a cash sum equal to the current market value of the bond and the buyer will pay the fixed price originally agreed upon. A forward deal is a contractual commitment which cannot be terminated unless both parties to the deal agree mutually to end the transaction.

Forward rate agreements

A forward rate agreement (FRA) is an OTC transaction that fixes a single interest rate for a single period at an agreed date in the future. The start of the period the rate will be fixed for and its length are negotiated between the contract buyer and seller. So an FRA transaction that locks in the 3-month rate in 3 months' time is referred to as a 3/6 or 3s6s transaction. The first number indicates the effective date of the transaction, the final number the maturity and the difference between the two indicates the tenor of the interest rate that is being fixed. Interest rate tenors will typically reflect those most commonly traded in the cash markets and so will have a maximum maturity of 12 months. The effective and maturity dates for FRAs could extend as far as 5 years depending on the currency.

These instruments have never been adopted by the corporate community to hedge exposures and are arguably most often used by traders as a way of expressing a view on expected short-term interest rate movements. Schofield recalls a conversation with an FRA trader where the dealer pointed out that he created his quote based on where he thought the Central Bank rate would be at some future date (plus a few basis points to reflect the difference in credit risk).

These instruments are quoted on a bid and offer basis and so a hypothetical quotation could be:

3/6	3.11%–3.12%
6/9	3.15%–3.16%
9/12	3.25%–3.26%

The interpretation of the quotation from a market maker's perspective is:

Bid	Offer
Buy FRA	Sell FRA
Pay fixed rate	Receive fixed rate
Receive LIBOR	Pay LIBOR

From this quoting convention we can start to see that an FRA is a contract for difference, which involves an exchange of cash flows. On the trade date the parties to the deal agree a fixed contract rate for an agreed future period and then will make or receive compensation depending on the actual level at which LIBOR settles. As we will show in the next section, an FRA can be thought of as a single-period interest rate swap.

To illustrate the concept, consider the following example. Suppose that 3-month interbank rates are 3.00% and the market believes that Central Bank rates will increase over the next year. The trader sees the market quoting a 9/12 rate as 3.25%–3.26%. He believes that actual 3-month rates in 9 months' time will be lower than this and so decides to sell the FRA at the bid price of 3.25% (he is a market user not a market maker) on a notional of USD 10m. This will contract him to receive 3.25% and pay the prevailing LIBOR rate in 9 months' time. If market rates evolve as per his view (i.e., 3-month LIBOR is lower than 3.25%) he will end up being a net receiver of cash.

Let us say that 9 months later, the 3-month LIBOR rate fixes at 3.20%. The parties to the FRA agreement can calculate the settlement amount due. The market user who sold the contract expects to be a net receiver of 5 basis points per annum on a USD 10 million notional amount. However, there is something of a quirk in the settlement convention.

Normally, interest rate contracts will settle in arrears but in the FRA market the settlement takes place as soon as LIBOR fixes. This means that the recipient of the cash flow has use of the funds in advance of normal market practice and as a result the settlement amount is present valued. The discount rate is the same LIBOR rate used in the numerator of the equation. The contract settles according to the following formula:

$$\frac{\dfrac{(+/-\text{ Settlement rate } -/+\text{ FRA Contract rate})}{100} \times \text{Notional amount} \times \dfrac{\text{No. of days}}{\text{Day basis}}}{1 + \left(\dfrac{\text{Settlement rate}}{100} \times \dfrac{\text{No. of days}}{\text{Day basis}}\right)}$$

The +/− signs are used to indicate if the participant is a receiver (+) or payer (−) of a particular rate. This will quickly allow the participants to decide who will be the receiver or payer of the cash settlement. Since interest rates are quoted as a percentage per annum it is necessary to pro rate the settlement according to the tenor of the interest rate. The day basis will either be 360 or 365 depending on the currency of the transaction.

So in our example the settlement amount would be:

$$\frac{(-3.20\% + 3.25\%) \times 10{,}000{,}000 \times \dfrac{90}{360}}{1 + \left(3.20\% \times \dfrac{90}{360}\right)} = \frac{1250}{1.008} = 1240.08$$

The calculation assumes an exact 90-day quarter.

1.5.3 Swaps

Interest rate swaps

In its most basic form an interest rate swap consists of a periodic exchange of cash flows with one referenced to a fixed rate while the other is referenced to a floating rate of interest, such as a particular maturity of LIBOR (Figure 1.2). Swaps are traded on a notional amount basis, which is usually fixed. The notional amount of a swap is merely a reference value and does not represent an actual cash flow. It will simply determine the magnitude of any cash flow that is subsequently exchanged. Swaps are typically long term, with maturities that may extend out to 30 or 50 years. Although the deals have a long-term maturity, the exchange of cash flows will take place on a more frequent basis. The cash flows are calculated on a simple interest basis and are paid in arrears. Each market has adopted its own conventions as to the frequency of these payments. For example, in the USD market the convention is a semi-annual payment of fixed for a quarterly payment of LIBOR. Where the payment dates coincide, it is market convention for the cash flows to be netted. However, since the transaction is OTC, all of the terms and conditions are negotiable and so there are many different variations of the simple vanilla "fixed/float" structure.

Figure 1.2 Illustration of fixed vs. floating interest rate swap.

To illustrate the principles involved, let us assume that both fixed and floating are paid semi-annually, with the rates for the period being 5.00% and 4.50%, respectively. We will assume that the cash flows are denominated in GBP, the notional amount is £10 million and that in the 6-month period there are 182 days.

The fixed cash flows will therefore be:

$$£10{,}000{,}000 \times 5.00\% \times 182/365 = £249{,}315.07$$

The floating cash flows will be:

$$£10{,}000{,}000 \times 4.50\% \times 182/365 = £224{,}383.56$$

Since the two payments coincide, there will be a net payment of £24,931.51 in favour of the receiver of fixed.

Swaps are quoted on a bid–offer basis. So if we were to analyse a typical quote from the perspective of a market maker (i.e., the institution giving the quote), it may look as follows:

Bid	Offer
4.5050%	4.5450%
Pay fixed	Receive fixed
Receive LIBOR	Pay LIBOR

A market user (i.e., the institution requesting the quote) would interpret the same values in the opposite manner. The key learning point from this is that the quotation is given in terms of the fixed rate. Since many investment banks will be running "matched positions" (i.e., they will try and structure their portfolio such that they will have a mix of pay and receive positions that are profitable overall), the LIBOR cash flows are assumed to cancel out.

Some practitioners will also say that they are "long" or "short" the swap. This is not our favoured quoting convention but for the sake of completeness, we include a short explanation. At the bid price the market maker is said to be "long" the swap – it is a "buying" position. However, the convention assumes that the market maker is buying a stream of LIBOR cash flows, for which they will pay a single fixed rate. By the same logic the offer price represents a short or "selling" swap position in that the market maker is delivering a stream of LIBOR cash flows for which their compensation is a single fixed price.

Schofield recalls one swaps class where an experienced swaps dealer disagreed vehemently with these definitions, stating that the entire swap market-making community defined "long" and "short" in the opposite way. After a flurry of phone calls, which yielded no consensus, he realized that he had been arguing at cross purposes with the participant. It would seem that some market makers reasonably viewed the offer side of the market as similar to buying a bond. The buyer receives a fixed coupon and finances the purchase at a LIBOR cash flow. Hence at the offer side of the quote perhaps it would be more helpful to describe a market maker as being long the market (as opposed to the swap) and being short the market at the bid price. It is no shame to use the terms "pay and receive fixed" to describe one's intention, and will ensure no costly mistakes are made!

It is also worth mentioning one other aspect of swap quotes. The customer may well end up paying more or receiving less than the quoted interbank rates as many banks will apply a dealing margin. There are two main factors that will impact the margin, which comprise the liquidity and credit charge. The liquidity charge is the "cost" of hedging the interest rate risk of entering into a single swap. It may not be possible for the bank to immediately offset the risk of a new swap by entering into an equal and opposite position. Alternatively, if the swap transaction is executed on a large notional amount, it may be difficult to hedge the entire position with one transaction. The credit charge has become more important in recent years and has recently become termed the "credit value adjustment". In simple terms, if there is a greater risk of a client defaulting this will have to be reflected in the price that they pay. Most interbank and hedge fund clients will have cash collateral agreements and so the credit charge may be fairly small. But if the collateral they provide is of lower quality or less liquid there will still be a credit charge, although this will be less than that applied to unsecured accounts.

In the next chapter we will show that there is a linkage between swaps and financially settled forwards in that a swap can be thought of as a multi-period, cash-settled forward contract, executed at a uniform fixed rate.

Asset swaps

An asset swap is a combination investment package where an investor buys a fixed-rate bond and simultaneously enters into a "pay fixed" interest rate swap. Asset swaps can be structured in a number of ways, the details of which will be covered in Chapter 5. Here we will consider one variation, which is the "par in, par out" (or just "par–par") structure. Here the investor pays 100% of the face value of the bond (i.e., its par value) at the start of the transaction, holds the bond to maturity and then receives par from the issuer at maturity. If the market value of the asset is anything other than par at the point of purchase this will create an advantage to either the buyer or seller. If the bond is trading below par this means the investor will be disadvantaged as he will "overpay" for the bond. For bonds that are trading at a price greater than their face value this bestows a cash flow advantage on the investor (i.e., if the bond is trading at 120 the investor only pays 100). The second element of the asset swap structure is that the fixed rate on the swap is set equal to the coupon on the bond. Again, since it is unlikely that these two parameters will be exactly equal it will create an advantage for one of the parties.

Figure 1.3 shows the cash flows associated with an asset swap and gives a visual depiction of the rationale for entering into this type of transaction. Since the fixed coupon on the bond and the fixed rate on the swap are equal and opposite to each other (and have the same maturity), the cash flows have no net economic impact on the investor. As a result, the investor owns a structure that pays them LIBOR plus or minus a spread, which makes the structure economically equal to a floating-rate note.

In the previous section we had shown that the majority of interest rate swaps involved a LIBOR cash flow without any associated spread. However, in the asset swap package the spread to LIBOR acts as a balancing mechanism to ensure that any advantage or disadvantage incurred as a result of the investor paying par and entering into an off-market swap is returned over the life of the deal. In this way the entire package will then become an equitable exchange of cash flows.

An asset swap structure has more credit than interest rate exposure. Suppose that interest rates were to rise. The bond element of the structure would lose money but since the investor

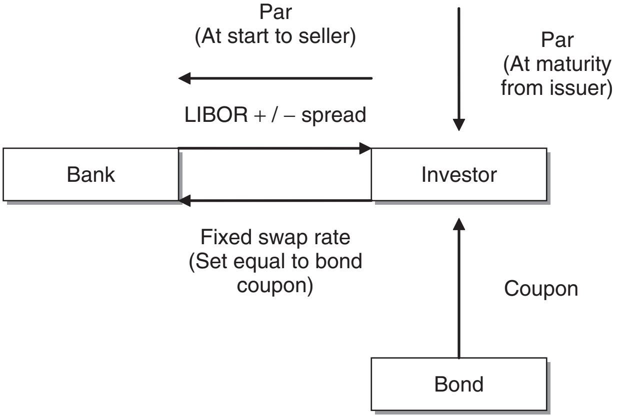

Figure 1.3 Asset swap package.

is paying fixed on a swap, this deal now becomes a more attractive transaction and will therefore increase in value. As a result, the two elements more or less cancel each other out. The same effect in the opposite direction would happen for a fall in interest rates. However, if interest rates remain unchanged but there is a perception that the issuer is more likely to default then the bond element will lose value with no offsetting profit on the swap. Overall there will be a net loss.

There are a number of reasons why an investor may wish to enter into an asset swap package:

- They may wish to reduce the market risk of holding a fixed-rate bond (bond market risk will be covered in Chapter 3).
- Since floating-rate notes offer an investor credit exposure rather than interest rate exposure, the investor may wish to take a view on how this component will evolve.
- The corporate entity to which the asset swap is linked does not have any floating-rate debt in issue and so the investor may have to create this synthetically using the asset swap.
- If the fixed-rate bond is trading cheap to its fair value then asset swapping the asset will create an attractively priced FRN.
- It is possible to buy a fixed-rate bond in, say, USD and asset swap it using a currency swap where the fixed rate is also USD but the LIBOR cash flows are denominated in, say, EUR.

Overnight index swap

An overnight index swap (OIS) is an interest rate swap where the floating leg of the swap is equal to the average of an agreed overnight index such as SONIA or EONIA (see Appendix 2.1 for more detail on these indices). As we will show in Chapter 2, the fixed leg of the swap is set at such a level that the transaction will be considered an equitable exchange of cash flows.

Figure 1.4 shows one simple application of the swap. A bank has agreed to take money on deposit on a fixed rate but can use the swap to transform the nature of their interest rate risk. They enter into an OIS where they receive a fixed rate and pay an overnight index rate. If we assume that the two fixed rates cancel out then the bank accepting the

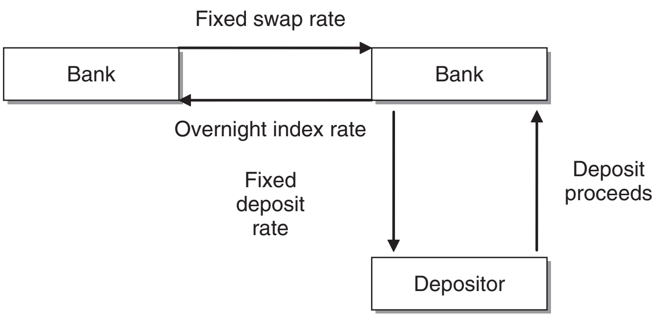

Figure 1.4 Overnight index swap.

deposit has acquired the funds and on a net basis is paying the overnight index rate. For deals with a maturity of less than one year, the fixed versus compounded floating payments are exchanged at maturity. For OIS transactions of greater than one year the payments are exchanged annually.

Credit default swaps

A credit default swap (CDS) is a bilateral contract that allows one entity to buy protection against the possibility that a particular reference entity (or basket of reference entities) will suffer a specific credit event. The buyer of this credit protection pays a fixed premium that is typically paid quarterly (although some markets pay semi-annually) to a protection seller. The protection seller will agree to "make whole" the protection buyer by agreeing to pay an amount of compensation if the agreed credit event occurs. The fee that the protection buyer agrees to pay is referred to as a premium or a spread. Upon the occurrence of the credit event the premium payments stop and the contract will terminate. The buyer of protection is considered to be short the credit risk as economically their position is equivalent to selling the credit risky asset. The protection seller is long the credit risk; similar to the buyer of a bond they are accepting the risk that a particular entity will suffer a credit event and are being paid a regular cash flow as compensation. By convention the protection seller is sometimes termed an investor as this position can be viewed as being economically equivalent to buying a bond.

Diagrammatically the CDS can be represented as in Figure 1.5.

In the credit fundamentals section of this chapter (Section 1.4) we derived an intuitive approach to pricing credit that suggested the magnitude of the spread was a function of the expected loss and the probability of default. In the USA and Europe, CDS trade with fixed coupons: 100 and 500 basis points for the USA; 25, 100, 500 and 1000 basis points for Europe. Higher-yielding names will trade with higher coupons, while investment grade names will trade with lower coupons. These coupons will not necessarily reflect the current market value of the spread and so an upfront cash adjustment will be necessary. Suppose an investment grade name is trading with a fixed coupon of 100 basis points but the market believes that the current value of the spread is 90 basis points. The buyer of protection will be required to pay the spread on a quarterly basis but is locking into a value that is 10 basis points higher than the market spread at the time the deal is executed. As a result he will receive an upfront cash adjustment of 10 basis points per annum, discounted to reflect the

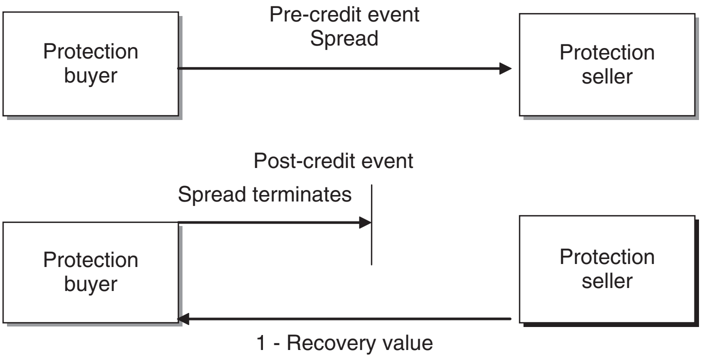

Figure 1.5 Credit default swaps.

time value of money and the probability that the company may default. This is covered in more detail in Chapter 3.

Two quoting methods have evolved for single-name CDS transactions. For investment grade names, the quote is given in terms of a "par spread". In effect this spread is a reflection of the market's current perception of the creditworthiness of the reference entity. Once the trade is executed, an upfront adjustment is paid or received and all subsequent cash flows are based on the agreed fixed coupon. For high-yield names the quotation is given in terms of "points upfront" but again all subsequent cash flows are executed on a fixed coupon basis. There is a standard market methodology to convert between par spreads and points upfront.

Originally, physical settlement was the norm upon the advent of a credit event and it required the protection seller to deliver the agreed notional amount of the transaction to the protection buyer. In return, the protection buyer delivered an asset issued by the agreed reference entity that conformed to the terms of the agreed contract. This could have been a defaulted asset or a non-defaulted asset that was considered pari passu. Physical settlement was preferred because banks typically ran a "matched position", i.e., they held an offsetting position with identical terms and conditions. As a result, if a credit event were to occur, physical settlement would allow the bank to pay and receive the notional amounts on the offsetting contracts and receive and deliver the agreed obligation. In this situation there is no price risk on the delivered obligation as long as the notional amounts on the transactions were equal. The market value of the delivered obligation was irrelevant as the notional amounts that are associated with the trades are for a fixed monetary value and the transfer of the obligation involves nothing other than a change of title of an asset.

Cash settlement was less popular as a bank holding a matched position would be subject to price risk. In the early days of the market, if a credit event were to occur, unless the institution were able to ensure that the settlement of both legs coincided there was no guarantee that the value of the delivered obligation on one side of the transaction would match that of the obligation received.

As the CDS market grew, inevitably the total notional amount of outstanding transactions outstripped the supply of outstanding obligations. This meant that if a credit event occurred, the protection buyers (who are not obligated to actually own obligations issued by the reference entity) often encountered difficulties in sourcing the required obligations. To avoid the introduction of price risk in the settlement of the CDS contract the market has now moved to an auction process where a panel of traders agree on a post-default value of the deliverable obligations. This single value is then used by all market participants, removing the potential price mismatch.

The CDS contract has a number of key characteristics.

Reference entity: It is very important for the two counterparties to agree the exact legal entity on which protection is being bought or sold.

Reference obligation: Stated in the deal confirmation is a debt obligation issued by the reference entity that identifies the exact nature of the credit risk being transferred. This is because different obligations within a capital structure will have different degrees of credit risk depending on the amount that would be recovered by the lender in the event of a default. Typically the reference obligation represents the senior unsecured portion of the reference entity's capital structure, but some transactions may be based on a different component (i.e., subordinated debt) in order to express a different view. Knowing the reference obligation will allow the contracts to be fairly priced (as the price must reflect the potential loss to the protection seller) and will also determine the nature of the asset that will be delivered if a credit event is activated.

Credit events: The market participants will agree a number of market standard events that will cause the contingent compensation to be paid by the protection seller. These are:

- bankruptcy
- failure to pay
- restructuring (of which there are different variations)
- obligation acceleration/default
- repudiation/moratorium (for sovereign reference obligations).

Not all contracts will include all of the default terms. Over time, each market has developed particular conventions as to which terms should be included. For example, in the USA the convention is to trade bankruptcy and failure to pay.

Obligations: Once the credit events have been specified the entities must agree the population of the issuer's obligations that could lead to a credit event being triggered. The market standard is typically "borrowed money" (although there are differences between markets). This is defined to include such things as bonds, loans and certificates of deposit.

Deliverable obligations: If a credit event is triggered and the deal is physically settled then the protection seller will deliver a cash sum equal to the agreed notional amount. In return the buyer will be required to deliver an acceptable asset issued by the reference entity. This need not be the reference obligation as the supply of this component may be limited. As a result, the market allows for some leeway in terms of the actual asset. Although a number of criteria apply, the most significant is that it should not be subordinated to the reference obligation as those assets possess a different degree of credit risk.

In recent times the credit derivatives market and principally the CDS component have experienced considerable growth. The spread of the CDS is viewed by market participants as representative of the "pure" credit risk of a particular reference entity. As such it is now used as a benchmark to assess and price credit risk in the bond markets. CDS instruments will trade with many different maturities even where there is no debt of an equivalent maturity. This has given rise to the concept of the credit curve, which is a representation of CDS spreads of different maturities.

1.5.4 Vanilla options

An option is a contract that gives the holder the right but not the obligation to buy or sell an asset at a pre-agreed price in the future. Essentially, it is a forward contract that allows the buyer to walk away if at maturity the market rates that prevail make execution of the deal unattractive. An option that gives the holder the right to buy an underlying asset is referred to as a call option, while the right to sell the asset is referred to as a put option. The price at which the counterparty agrees to deal is referred to as either the strike rate or the exercise price. When the buyer ("holder") of the option can use the option is a function of whether the contract is European, American or Bermudan ("semi-American") in style. A European-style option allows the holder to exercise the option only at expiry. An American-style option allows the holder to exercise the option at any time prior to its stated maturity. A Bermudan option allows the holder to exercise the option according to a pre-agreed schedule of dates.

The decision to exercise the option will depend on whether the option is in-, out-of- or at-the-money. An in-the-money (ITM) option is an option where the strike rate is more favourable than the underlying price. If this were the case it would be logical to exercise the option. An out-of-the-money (OTM) option is one where the strike rate is less favourable than the underlying price and so the option would not be exercised. An at-the-money (ATM) option is one where the underlying market price is equal to the strike price. In this case the holder would be indifferent as to whether they would exercise the option.

Assume an investor buys a 3-month European-style call option on an asset which is trading at a price of 100. Suppose that the strike rate for the option is also 100 and that the agreed premium is 5 units. If at maturity the underlying asset is trading at 90, the holder would not exercise the option to buy at the strike. The option would lapse as it is OTM and their losses would be equal to the premium paid. If the underlying asset was 110 at expiry, the holder would exercise the option. They would deliver a cash amount equal to the strike (100) and take delivery of the asset, which has a current value of 110. As a result they would have a profit of 5 units. This is calculated as the difference between the strike price of the option (100) and the expiry value of the underlying asset (110), less the premium paid (5). Note that if the underlying asset is between 100 and 105 at maturity the option would still be ITM as a result of our definition; however, the holder will not have broken even. The holder would still exercise the option as in this range of prices they will be seeking to minimize their losses and recoup some of the premium paid. An option seller is faced with a profit and loss profile opposite to that of the buyer. If the option is not exercised they will retain the premium but their losses will increase as the underlying price rises. However, the use of the terms ITM, OTM and ATM would remain the same and they are defined by convention from the buyer's perspective. So, if the underlying price is 90 at the expiry of the option, the option is OTM but is the preferred outcome from the seller's perspective.

The outcome for the buyer of a put with the same characteristics as the call would follow similar principles. However, a put option will be ITM if the strike price is greater than the underlying price, and OTM if the strike price is less than the underlying price.

From the profiles illustrated in Figure 1.6, it can be seen that for buyers of options, the maximum loss can never exceed the premium paid. The buyer's maximum loss represents the maximum profit that could be made by the seller of the option. The maximum profit for the option buyer appears to be unlimited, which is true for the call profile but not for

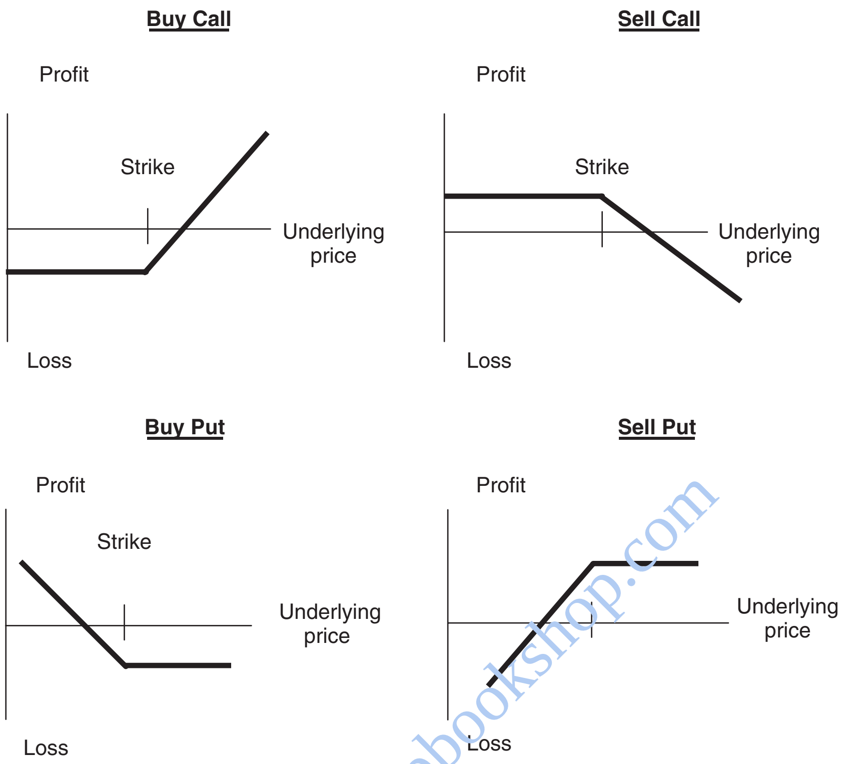

Figure 1.6 Profit and loss profiles for options at expiry.

the put. The maximum payoff on a put option is the difference between the strike rate and zero as the price of the underlying asset cannot go negative. Selling options can lead to unlimited losses for sellers of calls and significant losses for put writers, with the latter's losses "benefiting" from the restraint of a zero price boundary.

The payoff of calls and puts at expiry is often represented using the following expressions:

$$\text{Expiry payoff of call option} = \text{Max}(\text{underlying price} - \text{strike}, 0)$$

$$\text{Expiry payoff of put option} = \text{Max}(\text{strike} - \text{underlying price}, 0)$$

So for the holder of the call option the payoff will be the greater of either the underlying price less the agreed upon strike, or zero. For the holder of the put option they will exercise the option if the strike rate less the underlying price is greater than zero. These expressions do not measure the profit or loss of the position as they do not take into account the premium paid.

The premium is normally paid at the outset of the option and will be expressed in the same units as the underlying asset. So, if the option is referenced to an interest rate the premium will be paid in percentage points; crude oil options will be expressed in USD per barrel while options on equity indices will be quoted in index points.

Interest rate options

The two most common types of option within the interest rate world are cap/floor structures and swaptions. A cap structure is a strip of OTC interest rate call options on a series of forward rates, all traded with a single strike. The cap gives the buyer protection against an agreed index or reference rate such as LIBOR of a stated maturity rising above a pre-agreed strike rate. The term "cap" is the collective name for the component options, which are individually referred to as caplets. The premium payable on a cap structure is simply the sum of the individual caplet premia. It can be paid as a lump sum upfront or amortized over the life of the transaction.

Suppose that a company has taken in USD 100m for a period of 2 years on which 3-month USD LIBOR is payable. The Treasurer decides to insure himself against the possibility of an unfavourable rise in rates and so buys a 2-year cap referenced to 3-month LIBOR. Although the underlying transaction covers eight 3-month periods, there are by convention $n - 1$ options, i.e., seven caplets. This is because LIBOR interest rates are always set at the start of the period to which they apply but any associated payment is made at the end. As a result it would be impossible to buy protection against a known interest rate for the first period of the cap structure. The payoff on a cap for any single period is:

$$\text{Max}((\text{Reference rate} - \text{strike rate}) \times \text{notional} \times \text{days/day basis}, 0)$$

If we were to assume that the previously mentioned borrower agreed a strike rate of 5.00% and that in the second period, 3-month LIBOR fixed at 5.25%, then assuming a 92-day period and a 360-day year they would receive a sum equal to:

$$\text{USD } 100\text{m} \times 0.25\% \times 92/360 = \text{USD } 63{,}888.89$$

This could be used to offset the cost of the underlying borrowing which would now incur an interest charge of 5.25%. Similar to the underlying loan the payout on the cap would be known at the start of the quarter but payable in arrears. The net cash flows received under the cap combined with the payment on the underlying loan would result in a net interest cost of 5.00% – equal to the strike on the option. If the value of LIBOR had been equal to or less than the 5.00% strike rate then there would be no receipt under the terms of the option. The borrower would let the option lapse and enjoy the benefits of borrowing at a lower rate.

A floor gives the holder the protection again an agreed reference rate of interest falling below a pre-agreed strike. Again, a floor is a collective name for the component options which individually are referred to as floorlets. The mechanics of the floor are the same as the cap, but the payoff to a holder would only occur if the reference interest rate was less than the strike rate:

$$\text{Max}((\text{Strike rate} - \text{reference rate}) \times \text{notional} \times \text{days/day basis}, 0)$$

Swaptions

A swaption is an option that gives the holder the right but not the obligation to enter into an interest rate swap. A payer swaption allows the holder to pay fixed in a swap of a predefined maturity and so the buyer of a payer swaption would benefit if rates were to rise. A receiver

Table 1.2 The positions that result from the exercise of swaptions

	Receivers	Payers
Buy	Right to receive fixed on an interest rate swap	Right to pay fixed on an interest rate swap
Sell	Obligation to pay fixed on an interest rate swap	Obligation to receive fixed on an interest rate swap

swaption allows the holder to receive fixed in an interest rate swap and would benefit if rates were to fall. The different permutations are shown in Table 1.2.

If one were to draw the "hockey stick" at maturity payoffs for these options similar to those documented in Figure 1.6, the purchase of a payer swaption would resemble a long call, while the purchase of a receiver swaption would resemble a long put option. There is a slight difference in that the non-horizontal part of the profile will display some curvature as the instrument into which the option is exercised exhibits a non-linear profit and loss profile (the pricing of swaps is explained in greater detail in Chapter 2).

Although there are exceptions, the majority of swaptions are cash settled at expiry. So, rather than enter into an actual interest rate swap the buyer will receive the current market value of an interest rate swap with a fixed rate equal to the strike of the swaption. Swaptions are quoted in terms of the option maturity followed by the tenor of the swap. For example, an option to enter into a 5-year swap, 1 year in the future, would be written as 1y × 5y or "1 into 5".

1.5.5 Exotic options

The description of options has so far concentrated on those that are classified as being "vanilla", which means that their expiry profit and loss profiles at maturity conform to those shown in Figure 1.6. However, there is a larger family of options that are designated as being exotic. Somewhat unhelpfully, the only definition of an exotic option is one whose profit and loss expiry profile does not conform to the four vanilla building-block positions. The two most common exotic options are barriers and binaries.

Barrier options

A barrier option is an option that has an additional price performance feature, sometimes referred to as a trigger, which if hit by a movement in the spot price will result in an option position being either activated ("knocked in") or deactivated ("knocked out"). Although not universally adopted language, it is useful to classify barrier options in terms of the position of the barrier in relation to the spot price. A standard barrier option is one where the barrier is placed in the OTM region. A reverse barrier option represents an instance where the barrier is placed in the ITM region. Within either of these two categories it is possible to categorize the structure according to whether it is knocked in or knocked out, and then according to whether it is a call or a put. This results in 16 different permutations, which are shown in Table 1.3. It is possible to extend the number to 32 by distinguishing between buyers and sellers, but we will restrict our analysis to the 16 key positions. It is market convention to refer to these barrier options using terms such as "down and out". Take, for example, a standard knock out call option. Suppose an option has a strike rate of 100 and a

Table 1.3 Taxonomy of barrier option positions

	Standard knock ins	Standard knock outs	Reverse knock ins	Reverse knock outs
Calls	Down and in Barrier below spot	Down and out Barrier below spot	Up and in Barrier above spot	Up and out Barrier above spot
Puts	Up and in Barrier above spot	Up and out Barrier above spot	Down and in Barrier below spot	Down and out Barrier below spot

barrier at 90. Unless the spot rate trades at 90 prior to maturity, the option is for all intents and purposes a European-style call option. As soon as 90 trades, the option position is terminated and no further rights or obligations accrue. Hence the name; if the spot rate goes down, the option position gets knocked out.

Arguably the main motivation for executing a barrier option is one of cost. Although an extensive treatise on barrier pricing is beyond the scope of this book, since there is a possibility that the option will either survive or be terminated the price of a barrier option will always be cheaper than that of an equivalent European-style contract.

Binary options

A binary option is also sometimes referred to as a "digital", "all or nothing" or "bet" option. If exercised, the structure will pay out a fixed sum irrespective of how deeply the option is in-the-money. There are some variations on this basic description. An "at expiry" binary option is a European-style option that pays out a fixed amount only if the option is ITM at expiry. Since the payout at maturity is limited to a fixed amount, they will be cheaper than an equivalent non-binary option. A "one touch" option is an American-style digital option, which pays out a fixed amount at the point that the strike is hit, which could be any time prior to expiry. A "no touch" option is a digital option that pays out a fixed sum at expiry if the underlying price does not touch the strike.

Depending on their style, digital options may or may not be path dependent. Path dependent means that the magnitude of the option's payout is affected by the movement of the underlying price prior to expiry. One touch and no touch options are path dependent and will cost more than the European digital. As a rule of thumb, American-style digitals will cost twice as much as European digitals due to the increased probability of exercise.

Since the option has a fixed payout and a fixed premium cost there is no need to have a principal amount. The premium is often quoted as a percentage of a 100% payout. So if the trader wished to have a payout of USD 1m, the premium might be quoted as 10% (USD 100,000). Somewhat confusingly, the strike rate on a binary option may be referred to as a "barrier" and the words call and put may be replaced by "up" and "down". Reference is also made to "knock ins" and "knock outs" for touch and no touch structures, respectively.