1
An Introduction to Derivative Products

SYNOPSIS The purpose of this chapter is to outline the main features of derivatives and provide a description of the main ways in which they are priced and valued.

This chapter is divided into two distinct sections that cover:

- The key features of the derivative "building block" products
- The principles of how each of the products is priced and valued.

The coverage is not particularly mathematical in style, although numerical examples are included where it helps to illustrate the key principles.

In the first section the fundamental concepts of the *main derivative products* are considered. The products covered include:

- Futures
- Forwards
- Swaps
- Options (mostly "vanilla" with some "exotic" coverage)

In the second section the focus is on the *pricing of derivatives*. The approach considers that all of the building block markets are linked through mathematical relationships and describes how the price of one product can be derived from another.

One of the unique elements of pricing commodity derivatives is the existence of the *convenience yield*, which is explained in conjunction with the concepts of *contango* and *backwardation*.

Two extra themes are developed in the pricing section that are relevant to other parts of the book. The first is a discussion on *put–call parity*, which will help the reader to understand how some structures are created. This idea is then developed to outline the potential *sources of value* in risk management solutions.

The chapter concludes with a description of the main measures of *option risk management* – the Greeks.

When analysing derivatives it is convenient to classify them into three main building blocks:

- Forwards and futures
- Swaps
- Options.

However, within the option category it is possible to make a distinction between two subcategories, the so-called "plain vanilla" structures (that is, options that conform to a basic accepted profile) and those that are considered "exotic", such as binaries and barriers.

For ease of illustration we will use gold in the following examples.

1.1 FORWARDS AND FUTURES

A forward contract will fix the price today for delivery of an asset in the future. Gold sold for spot value will involve the exchange of cash for the metal in two days' time. However, if the transaction required the delivery in say 1 month's time it would be classified as a forward transaction. Forward contracts are negotiated bilaterally between the buyer and the seller and are often characterised as being "over the counter".

The forward transaction represents a contractual commitment; so if gold is bought forward at, say, USD 430.00 an ounce but the price of gold in the spot market is only USD 420.00 at the point of delivery, I cannot walk away from the forward contract and try to buy it in the underlying market. However, it is not impossible to terminate the contract early. This could be achieved by agreeing a "break" amount, which would reflect the current economic value of the contract.

A futures contract is traded on an organised exchange with the New York Mercantile Exchange being one example. Economically a future achieves the same result as a forward by offering price certainty for a period in the future. However, the key difference between the contracts is in how they are traded. The contracts are uniform in their trading size, which is set by the exchange. For example, the main features of the contract specification for the gold future are listed in Table 1.1.

There are some fundamental differences between commodity and financial products traded on an exchange basis. One of the key differences is that futures require collateral to be deposited when a trade is executed (known as initial margin). Although different exchanges will work in different ways, the remittance of profits and losses may take place on an ongoing basis (variation margin) rather than at the maturity of the contract. An example of this is detailed in the chapter on base metals.

Table 1.1 Gold futures contract specification

Trading unit	100 troy ounces
Price quotation	US dollars and cents per troy ounce
Trading hours	Open outcry from 8.20am until 1.30pm New York time (electronic trading is also available)
Trading months	Trading is conducted for delivery during the current calendar month; the next two calendar months; any February, April, August and October falling within a 23-month period; and any June and December falling within a 60-month period beginning with the current month.
Minimum price fluctuation	USD 0.10 (10c) per troy ounce (USD 10.00 per contract).
Last trading day	Trading terminates at the close of business on the third to last business day of the maturing delivery month.
Delivery period	The first delivery day is the first business day of the delivery month, the last delivery day is the last business day of the delivery month.
Margin requirements	Margins are required for open futures positions.

Source: NYMEX.

Settlement of financial futures is often for a single date specified by the exchange, such as the third Wednesday in March, June, September or December. For commodity futures settlement could be for any day within the ensuing three months (see "trading days" section in the above specification). By offering delivery on any day for the current and two successive months, this commodity future possess a feature of the forward market – the flexibility to settle for a variety of dates. Another difference is the concept of grade and quality specification. If one is delivering a currency, the underlying asset is homogeneous – a dollar is always a dollar. However, because metals have different shapes, grades and quality, there must be an element of standardisation to ensure that the buyer knows what he or she is receiving. Some of the criteria that NYMEX apply include:

- The seller must deliver 100 troy ounces (±5%) of refined gold.
- The gold must be of a fineness of no less than 0.995%.
- It must be cast either in one bar or in three 1-kilogram bars.
- The gold must bear a serial number and identifying stamp of a refiner approved and listed by the Exchange.

1.2 SWAPS

In a swap transaction two parties agree to exchange cashflows, the sizes of which are based on different price indices. Typically, this is represented as an agreed fixed rate against a variable or floating rate. Swaps are traded on an agreed notional amount, which is not exchanged but establishes the magnitude of the fixed and floating cashflows. Swap contracts are typically of longer-term maturity (i.e. greater than one year) but the exact terms of the contract will be open to negotiation. For example, in many base metal markets a swap transaction is often nothing more than a single period forward, which allows for the transaction to be cash settled, involving the payment of the agreed forward price against the spot price at expiry.

The exact form may vary between markets, with the following merely a sample of how they may be applied in a variety of different commodity markets.

- *Gold*: Pay fixed lease rate vs receive variable lease rate.
- *Base metals*: Pay fixed aluminium price vs receive average price of near dated aluminium future.
- *Oil*: Pay fixed West Texas Intermediate (WTI) price vs receive average price of near dated WTI future.

Swaps will usually start as spot and so become effective two days after they are traded. However, it is also possible for the swap to become effective at some time in the future – a forward starting swap. The frequency with which the cashflows are settled is open to negotiation but they could vary in tenor between 1 month and 12 months. Where the payments coincide there is a net settlement between the two parties. One of the features of commodity swaps that is not shared by financial swaps is the use of an average rate for the floating leg. This is because many of the underlying exposures that commodity swaps are designed to hedge will be based on some form of average price.

The motivation for entering into a swap will differ between counterparties. For a corporate entity one of their main concerns is risk transference. Consider a company that has to purchase a particular commodity at the market price at regular periods in the future. To offset the risk that the underlying price may rise, the company would receive a cashflow under the swap based on movements in the market price of the commodity and pay a fixed rate. If the counterparty to the transaction were an investment bank, the latter would now have the original exposure faced by the corporate. The investment bank would be receiving a fixed rate and paying a variable rate, leaving it exposed to a rise in the price of the underlying commodity. In turn the investment bank will attempt to mitigate this exposure by entering into some form of offsetting transaction. The simplest form of this offsetting deal would be an equal and opposite swap transaction. In order to ensure that the bank makes some money from this second transaction, the amount it receives from the corporate should offset the amount paid to the offsetting swap counterparty.

Swaps are typically traded on a bid–offer spread basis. From a market maker's perspective (that is the institution actually giving the quote) the trades are quoted as follows:

Bid	*Offer*
Pay fixed	Receive fixed
Receive floating	Pay floating
Buy	Sell
Long	Short

Although the terms "buy" and "sell" are often used in swap quotes the actual meaning is often confusing to anyone looking at the market for a first time. The convention in all swap markets is that the buyer is receiving a stream of variable cashflows for which the price is a single fixed rate. Selling a swap requires the delivery of a stream of floating cashflows for which the compensation is a single fixed rate.

1.3 OPTIONS

A forward contract offers price certainty to both counterparties. However, the buyer of a forward is locked into paying a fixed price for a particular commodity. This transaction will be valuable if the price of the commodity subsequently rises, but will be unprofitable in the event of a fall in price. An option contract offers the best of both worlds. It will offer the buyer of the contract protection if the price of the underlying moves against him but allows him to walk away from the deal if the underlying price moves in his favour.

This leads to the definition of an option as the right but not the obligation to either buy or sell an underlying commodity at some time in the future at a price agreed today. An option that allows the holder to buy the underlying asset is referred to as a *call*. Having the right to sell something is referred to as a *put*. Options may be either physically settled (that is the commodity is actually delivered/received) or cash settled. The price at which the underlying is traded if the option is exercised is referred to as the *strike* price. The strike price is negotiated between the option buyer and seller. Cash settlement involves the seller paying the buyer the difference between the strike and the spot price at the point of exercise. Cash settlement is advantageous to the buyer as it may be more convenient to either source or deliver the commodity separately; the option would simply offer price protection.

Options come in a variety of styles relating to when the holder can actually exercise his right. A European style option allows the holder to exercise the option only on the final maturity date. An American style option allows the holder to exercise the option at any time prior to final maturity. A Bermudan option allows the holder to exercise the option on a pre-agreed set of dates prior to maturity.

An option that is "in-the-money" (ITM) describes a situation where it would be more advantageous to trade at the strike price rather than the underlying market price. Take for example an American style option to buy gold at USD 400 an ounce when the current spot price is, say, USD 425; the option to buy at the strike is more attractive than the current market price. Where the option is "out-of-the-money" (OTM) the strike is less attractive than the market price. If the same American style option had a strike rate of USD 430 the higher strike makes the option less attractive than being able to buy the underlying at a price of USD 420. Finally an option where the strike is equal to the current market price is referred to as being "at-the money" (ATM).

Since options confer rights to the holder a premium is payable by the buyer. Typically this is paid up front but certain option structures are constructed to be zero premium or may involve deferment of the premium to a later date. Premiums on options are quoted in the same units as the underlying asset. So since physical gold is quoted in dollars per troy ounce, the premium will be quoted in the same manner.

Many of the derivatives strategies based on options that are discussed and illustrated within the text are based on the value of the option at maturity. These are presented in Figure 1.1.

In the top left-hand part of Figure 1.1, the purchase of a call option is illustrated. If, at expiry of the option, the market price is lower than the strike, the option is not exercised

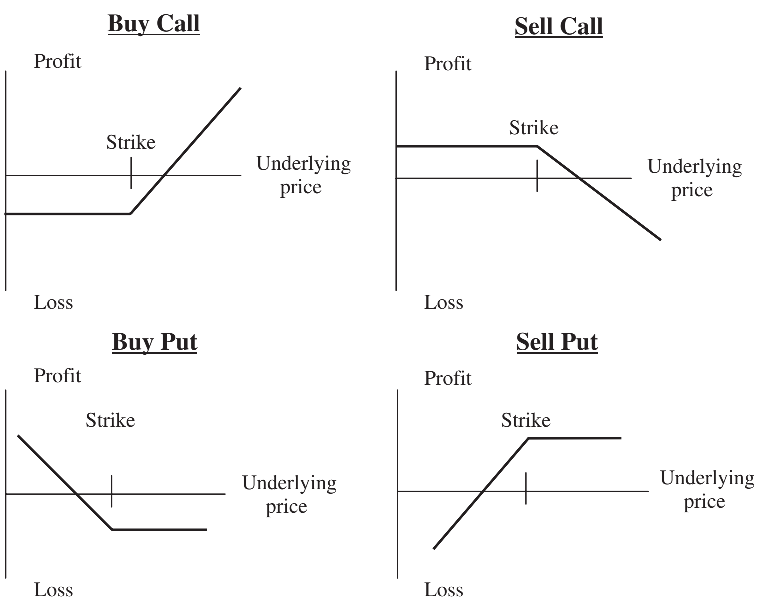

Figure 1.1 Profit and loss profiles for options at expiry

and the buyer loses the premium paid. If the underlying price is higher than the strike price the option is exercised and the buyer receives the underlying asset (or its cash equivalent), which is now worth more in the underlying market than the price paid (i.e. the strike price). This profit profile is shown to the right of the strike price. On the other side of the transaction there is the seller of a call option (top right-hand quadrant of Figure 1.1). The profit and loss profile of the seller must be the mirror of that of the buyer. So in the case of the call option the seller will keep the premium if the underlying price is less than the strike price but will face increasing losses as the underlying market price rises.

The purchase of a put option is illustrated in the bottom left-hand quadrant of Figure 1.1. Since this type of option allows the buyer to sell the underlying asset at a given strike price, this option will only be exercised if the underlying price falls. If the underlying price rises, the buyer loses the premium paid. Again the selling profile for the put is the mirror image of that faced by the buyer. That is, if the underlying price falls, the seller will be faced with increasing losses but will keep the premium if the market price rises.

Exotic options are a separate class of options where the profit and losses at exercise do not correspond to the plain vanilla American/European styles. Although there is a proliferation of different types of exotic options (many of which will be introduced in the main text), it is worth introducing two key building blocks, which feature prominently in derivative structures.

A binary option (sometimes referred to as a "digital") is very similar to a simple bet. The buyer pays a premium and agrees to receive a fixed return. Very often the strike rate on the digital is referred to (somewhat confusingly) as a "barrier". With a European style call option the holder will deliver the strike price to the seller and receive a fixed amount of gold. However, the value of the gold will depend on where the value of gold is trading in the spot market upon exercise. With a binary option the buyer will receive a fixed sum of money if the option is exercised irrespective of the final spot level.

The purchaser of a barrier option will: (1) start with a conventional "plain vanilla" option that could subsequently be cancelled prior to maturity (known as a "knock-out"), or (2) start with nothing and be granted a plain vanilla option prior to the maturity of the transaction (known as a "knock-in").

The cancellation or granting of the option will be conditional upon the spot level in the underlying market reaching a certain level, referred to either as a "barrier" or a "trigger".

The position of the barrier could either be placed in the out-of-the-money region or in the in-the-money region. This will be above or below the current spot price, as we will show below. The former are referred to as "standard" barriers with the later known as "reverse" barriers. This could result in what may initially seem like a bewildering array of possibilities, and Figure 1.2 summarises the concepts.

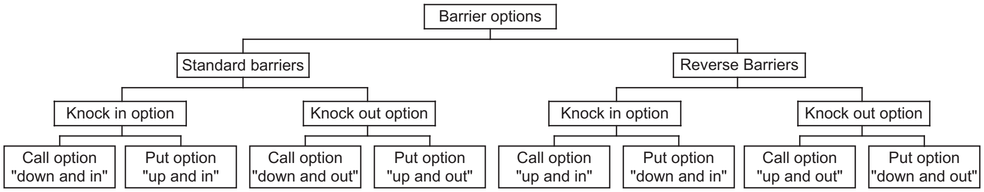

Figure 1.2 Summary of barrier options

To illustrate the concept further, let us return to the option example illustrated earlier and concentrate on analysing a call option. We will assume that the option is out-of-the-money and the current market conditions exist:

Spot	USD 425
Strike	USD 430
Maturity	3 months

The purchaser of a standard knock-in barrier option would be granted a European style option if spot hit a certain trigger. Since it is a standard barrier option, the trigger has to be placed in the out-of-the-money region so it would be set at, say, USD 420. Consequently, spot has to reach USD 420 or below before the option is activated ("knocked in"), hence the name "down and in". If the purchaser started with a standard barrier call option with the trigger at USD 420, it would be a "down and out". That is, if spot were to fall to USD 420 or lower, the option contract would be cancelled. A reverse knock-in call option would have the barrier placed in the money, say at USD 435. A purchaser of such an option would have a contract that would grant a call option with a strike of USD 430 if spot hits USD 435. The final example would be a reverse knock-out call option, with the trigger again set at USD 435. Here the purchaser starts the transaction with a regular call option which would be cancelled if spot reached USD 435 – an "up and out" contract.

Options arguably offer great flexibility to the end user. Depending on their motivation it could be argued that option usage could be categorised in four different ways. Firstly, options can be used to take a directional exposure to the underlying market. So, for example, if a user thought the underlying price of gold was to rise, he could buy physical gold, buy a future or buy a call option. Buying the gold requires the outlay of proceeds, which may need to be borrowed; buying a future reduces the initial outlay of the physical but will incur a loss if the future's price falls. Buying a call option involves some outlay in the form of premium but allows for full price participation above the strike and limited downside if the price falls. The second usage for options is an asset class in its own right. Options possess a unique feature in implied volatility and this can be isolated and traded in its own right. The focus of this type of strategy is how the option behaves prior to its maturity. The third motivation, which is particularly relevant to the corporate world, is as a hedging vehicle that allows a different profile than that of the forward. With a little imagination it is possible to structure solutions that will offer differing degrees of protection against the ability to profit from a favourable movement in the underlying price. The final motivation concerns options as a source of outperformance. For example, if an end user owns the underlying asset (e.g. central bank holdings of gold) he can use options to exceed some performance benchmark such as money market deposit rates. From a hedger's perspective, options could be used to outperform an ordinary forward rate.

1.4 DERIVATIVE PRICING

The purpose of this chapter is to arm the readers with sufficient knowledge to enable them to follow the main pricing issues referred to in the main text. It is not intended to be an exhaustive treatment of all aspects of derivative pricing. Readers interested in delving into more detail should refer to the bibliography for a list of suggested titles.

The principles of pricing commodity derivatives can often differ from financial products such as bonds or equities and these key differences will be highlighted as appropriate. Again, for the sake of simplicity, gold will be the main focus of the chapter.

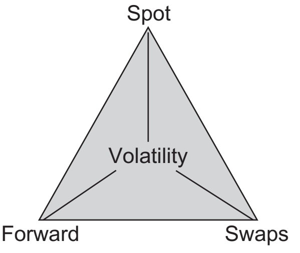

Figure 1.3 The relative value triangle

1.4.1 Relative Value

One of the key themes of trading all financial assets is the concept of relative value. This is defined as the optimum way to take exposure to a particular asset. If I wish to take exposure to the gold market, which instruments (or combination thereof) would give me the greatest return? This approach regards the spot, forward, swap and option markets as being interrelated through a series of mathematical relationships and therefore allows the trader to identify the particular market that offers an enhanced return/reduction in cost.

This approach can be encapsulated by the relative value triangle shown in Figure 1.3. A feature of the relative value triangle is not only the mathematical relationships that are implied between each of the instruments, but also the different trading strategies that exist by reading down each side of the triangle. For example, trading the spot market against the forward (or future) is referred to as a basis trade. The aim of the remaining part of this chapter is to illustrate the mathematical relationships between each of the components of the triangle.

1.5 THE SPOT–FORWARD RELATIONSHIP

1.5.1 Deriving forward prices: market in contango

Within the commodities world, there are two ways of describing the state of a forward market – contango or backwardation. Contango describes a situation where the price for forward delivery is higher than the price for spot delivery, while backwardation exists when the forward price is below the spot price. Although both of these states exist in the pricing of traditional financial products, the role of the underlying physical markets in commodities is much more important, particularly when demand exceeds supply.

We will use the example of a gold producer who approaches a bank asking for a price for delivery of gold in, say, 6 months. The price quoted by the bank is not a guess and neither is it a forecast of where it thinks the price of gold will be at the time of delivery. Rather the price quoted will be driven by the cost of hedging the bank's own exposure. This illustrates one of the key maxims of derivative pricing – the cost of the product is driven by the cost of the hedge.

If the bank does not hedge its price exposure then in 6 months' time it will take delivery of gold at the pre-agreed price and will then be holding an asset whose market value could be lower (or higher) than the price paid to the producer.

To avoid the risk of a fall in the gold price, the bank executes a series of transactions on the trade date to mitigate the risk. Since the bank is agreeing to receive a fixed amount of gold in the future, it sells the same amount in the spot market to another institution – say another investment bank. However, the bank has sold a quantity of metal now that it will not take delivery of until a future time. To fulfil this spot commitment it can borrow the gold until it receives the gold from the producer. The gold could be borrowed from a central bank that would receive interest at maturity. Having sold the gold spot and borrowed to cover the sale, the bank is now holding dollar proceeds. Since the bank would be seeking to manage its cash balances prudently, these dollars would now be invested until the producer delivers the gold.

As a result it is possible at the inception of the forward trade to identify all the associated cashflows, allowing the bank to quote a "fair value" or theoretical price that will ensure no loss at the point of delivery, irrespective of the prevailing price. In this example the maximum amount the bank will pay the producer cannot exceed:

- proceeds received from the spot sale *plus*
- the interest received from the dollar deposit *less*
- the interest paid to the lender of gold.

A simple example may help to illustrate the point. We will assume that the producer asks for a 6-month (182-day) forward price. For simplicity we will calculate the forward price for a single ounce. In the cash market gold is trading at USD 425.40 per ounce, so the dealer agrees to sells 1 ounce. In order to complete the spot delivery he borrows the same amount from the local central bank for 6 months at a lease rate of 0.11570% per annum. The dollars received from the spot sale are put on deposit for 6 months at a LIBOR to earn, say, 3.39% per annum. The interest cost of borrowing the metal is USD 0.2488 (spot × lease rate × 182/360) and that USD 7.29 is earned from the cash deposit (spot sale proceeds × 6-month LIBOR × 182/360). So the maximum amount he can afford to pay the producer is USD 432.4418. This is calculated as spot sale proceeds plus interest on LIBOR deposit minus the borrowing fee (USD 425.40 + USD7.29 − USD 0.2488). This fair value is a breakeven price for the trader and so may be adjusted to build in an element of profit.

The forward price is therefore the spot price plus the cost of carrying an underlying hedge. It is important to note that the shorter the time to maturity the smaller will be the differential between the spot and forward price since the hedge is carried for a shorter period. Indeed, if we were to recalculate the forward price applicable for a fixed date in the future on a regular basis, the differential would reduce every day (all other things being equal) until the final date when spot and forward become the same thing and the two prices will have converged.

The observed forward price of a commodity is kept in line with its fair value by the possibility of arbitrage. Take the previous example where the fair value of gold for 6-month delivery was USD 432.44. It is unlikely that the market price would be significantly different. Let us assume that a market price of USD 425.00 was observed. With the fair value of the instrument calculated at USD 432.44, the commodity would be described as being "cheap to fair value". In this situation an arbitrager could:

- buy the commodity forward, paying USD 425.00 upon delivery;
- short the underlying in the spot market to earn USD 425.40;
- invest the cash proceeds at LIBOR, earning 3.39% for 6 months to earn USD 7.29;
- borrow gold in the lease market in order to fulfil the short spot sale paying a 6-month lease rate, which equates to a cash amount of USD 0.2488;

- repay the gold borrowing upon receipt of the metal under the terms of the forward contract.

The arbitrager would end up with a net profit of USD 7.44, the difference between the theoretical value of the forward contract (USD 432.44) and the market value of the forward (USD 425.00).

1.5.2 Deriving forward prices: market in backwardation

Backwardation describes a situation where the prices for shorter-dated contracts are higher than those of longer-dated contracts. Forward pricing theory dictates that the market maker quotes a forward price such that all expenses are passed on to the customer as well as any income benefits he may have derived from carrying an underlying hedge. With gold storage on an unallocated basis is not consider to be a significant expense and so is traditionally not included in forward pricing considerations. However, with base metals warehousing and insurance costs included, this should give us the following relationship:

$$\text{Forward price} = \text{Spot price} + \text{LIBOR} + \text{Warehousing/Insurance costs}$$

This would also suggest that the fair value of a forward contract should always be greater than the spot value. However, many commodity markets move into backwardation (e.g. some base metals, crude oil), which suggests that the previous equation is incorrectly specified. Over time, to explain this apparent anomaly, the market has added an extra expression referred to as the "convenience yield". This is defined as the premium that a consumer is willing to pay to be able to consume the commodity now rather than at some time in the future, and the equation now reads:

$$\text{Forward price} = \text{Spot price} + \text{LIBOR} + \text{Warehousing/Insurance costs} - \text{Convenience yield}$$

The magnitude of the convenience yield will vary according to the physical balance of demand for the underlying commodity. If the commodity is in very short supply, its value will rise, moving towards zero in "normal" supply conditions. If, however, the minimum value of the convenience yield is zero, there is no maximum as its value is driven by the consumer's need to obtain the physical commodity immediately. For example, as it is difficult and expensive to store oil, the cost of closing a refinery for, say, 3 months would be very high. As a result, the consumer is willing to pay a premium for spot delivery.

The concept of convenience yield will seem particularly strange to someone new to the commodities market. Market practitioners tend to view it as a financial mathematician's tool to try to describe a market behaviour they never witness in traditional financial products such as bonds and equities and so cannot explain it; in reality no one in the real world uses the convenience yield. An article in *Risk* (November 2006) described it as the

> flow of services and benefits that accrues to an owner of a physical commodity, but not to an owner of a contract for future delivery of the commodity. This can come in the form of having a secure supply of raw materials and hence, eliminating the costs associated with stock outs.

Since, however, it is an intangible element that allows the equation to balance, it doesn't really explain backwardation to any degree of satisfaction.

In the case of a backwardated market the forward market price is lower than the spot price suggesting that the contract is mispriced or trading "cheap to fair value". If this is the case, a speculator should be able to buy the cheap forward contract, sell it for spot value and hold the combined position to maturity. This strategy, which is very common in the financial markets, would allow the arbitrager to earn the difference between the mispriced forward and its theoretical value. The reason this cannot happen, and why the market will remain in a prolonged state of backwardation, is that when selling the contract in the spot market, the participant will need to obtain the commodity to fulfil the selling commitment. Since the availability of the commodity in the spot market is very scarce, these supplies simply cannot be obtained. Hence, this apparent mispricing will persist for prolonged periods, as there is no mechanism to exploit the potential arbitrage.

If we were to plot the prices of the commodity for various times to delivery we would derive a forward curve, which in a backwardated market would be negatively sloped. That is, shorter-dated contracts would have a higher price than longer-dated contracts. Another way of explaining the slope of the curve is to consider the nature of the activities of the participants at certain maturities. Intuitively it is reasonable to suggest that the producer of a commodity would be more likely to sell his production forward, while a consumer is more likely to want to buy the commodity in the spot or near months. This combination of longer-dated forward sales and shorter-dated purchases combines to create a negatively sloped curve. This poses another question: Who are the shorter-dated sellers and who are the longer-dated buyers? This role is filled by entities that have no underlying economic commodity exposure but are willing to take views on the slope of the curve. For example, let us assume that the forward curve for a particular commodity is steeply inverted but a hedge fund believes that the slope between the 3-month and 12-month forward will gradually flatten. They could execute a trade that would involve the simultaneous sale of a short-dated forward (or future) and the purchase of the longer-dated contract. If the price differential between the two contracts does narrow as expected, a reversing transaction could be initiated to close out the original exposure at a profit.

1.6 THE SPOT–FORWARD–SWAP RELATIONSHIP

The price of any swap, irrespective of the underlying asset class, is the fixed element of the transaction. Since we are focusing on the gold market the fixed element would be a fixed lease rate with the opposing leg being a variable lease rate. (The mechanics of the transaction are explained in Chapter 3.)

To calculate the price of a gold swap the first starting point is to appreciate that all swaps should be considered an equitable exchange of cashflows on the day they are traded. That is, the present value of the expected payments must equal the present value of the expected receipts. To a reader new to the world of swaps, this seems a strange situation – a transaction that does not seem to have any initial profit. However, note that the fair price of the swap was described in terms of expected cashflows. Profits and losses will arise as actual payments are crystallised – these could be substantially different from those expected at the start of the transaction.

Since we are trying to solve for an unknown fixed rate, our analysis of swap prices starts with floating or variable cashflows. The aim on the floating side is to calculate the present value of the future cashflows, which are linked to a series of yet to be determined unknown lease rates. To solve this problem we can revert to the forward market and derive a series of lease rates that we expect to occur at different points in time in the future–forward lease rates.

At this point it is necessary to take two short diversions to see how forward rates are calculated. Let us assume that the 6-month lease rate is 0.11571% and the 12-month lease rate is 0.18589%. If a market participant were looking to deposit gold for 1 year, we have a choice between

- one 12-month deposit at 0.18589% or
- a 6-month deposit at 0.11571%, which would be rolled over at the end of 6 months, earning whatever the 6-month rate is at that time.

When posed with the question of which decision the lender should make, a common reaction is that the choice depends on the lender's view on what interest rates are likely to be for the final 6-month period. However, if a market participant is offered two choices that are identical in terms of maturity and credit risk, they must offer the same potential return. If the two strategies offered different returns, the investor would opt for that choice that offered the higher yield. As other market participants identify this, the excess returns will gradually diminish until the advantage disappears. It is therefore possible to calculate a lease rate for the 6- to 12-month period whose value will make the two choices equal to each other. This rate is referred to as the forward lease rate and can be calculated using the formula:

$$\text{Forward rate}_{b\times c} = \frac{(\text{Lease rate}_{a\times c} \times \text{Days}_{a\times c}) - (\text{Lease rate}_{a\times b} \times \text{Days}_{a\times b})}{\text{Days}_{b\times c} \times (1 + \text{Lease rate}_{a\times b} \times (\text{Days}_{a\times b}/\text{Day basis}))}$$

where:

$\text{Lease rate}_{a\times c}$ = the lease rate from spot to final maturity
$\text{Days}_{a\times c}$ = number of days from spot to final maturity
$\text{Lease rate}_{a\times b}$ = lease rate from spot to start of forward period
$\text{Days}_{a\times b}$ = number of days from spot to start of forward period
$\text{Days}_{b\times c}$ = days from start of forward period to final maturity
Day basis = 360 days.

Applying the 6- and 12-month lease rates given above, and assuming 182 days for the first 6-month period and 183 days for the second 6-month period, the value for the forward rate is given as

$$\text{Forward rate}_{6\times 12} = \frac{(0.18589\%_{0\times 12} \times 365_{0\times 12}) - (0.11571\%_{0\times 6} \times 182_{0\times 6})}{183_{6\times 12} \times [1 + 0.11571\%_{a\times b} \times (182_{0\times 6}/360)]} = 0.255537\%$$

The interpretation of the forward rate depends on its usage. In the market for trading short-term interest rates, the forward rate is used as a mechanism for trading expectations of future expected movements in central bank rates. In that sense one interpretation is that the forward rate is simply the market's current "best guess" over future cash rates. For example, if we applied that logic to the previous calculation, we could say that although current 6-month rates are 0.11571%, the market expects them to be 0.2555375% in 6 months' time. However, it does not mean that actual lease rates at that time will take that value; the actual value will only be known at the start of the period. Also, forward rates are notoriously bad predictors of actual rates but despite this they are still very popular for the purpose of trading future views.

The second interpretation is that of a breakeven rate. In the previous example, the forward rate is clearly a rate that equates the two investment alternatives. This brings us back to the initial question: Where does the lender place the gold? Since the forward rate can be thought of as a breakeven rate, the two choices would appear to be equal. The correct decision for the lender is driven by where he thought the actual lease rate would be at the start of the 6×12 period. If he thought the lease rate was going to be greater than the implied forward rate, he would invest in the two 6-month strategies. If he thought the lease rate was going to be less than the implied forward, he would execute the single 12-month deposit. Note that it is not necessarily an issue of whether rates will rise or fall, it is more a case of where actual rates will be in relation to the implied forward rate.

The other piece of information necessary to price a swap is a series of discount factors. Discount factors have a value between 0 and 1 that can be used to give a present value to a future cashflow. The discount factor can be applied to a future cashflow using the following simple relationship:

$$\text{Present value} = \text{Future value} \times \text{Discount factor}$$

The source of these discount factors are yields on zero coupon instruments that have the same degree of credit risk as the cashflow to which they will be applied. A zero coupon instrument pays no cashflow until maturity with the buyer's return usually in the form of a capital gain. However, interest-bearing deposits may also be in zero coupon form if they only have two cashflows – the initial and final movement of funds.

The reason these instruments are used to present value cashflows is that the investor's exact return can be calculated. This return is different from that offered by an interest-bearing instrument that pays a series of cashflows prior to maturity. Although a yield to maturity could be calculated for such an instrument, it would not be a true measure of the overall return as this measure requires the interim interest payments to be reinvested at the yield that prevailed at the start of the transaction. This is referred to as reinvestment risk.

The only problem with zero coupon yields is the lack of available market observations. As a result, the analyst is often forced to use mathematics to transform the yield on an interest-bearing instrument into a zero coupon equivalent. A comprehensive treatment of this subject can be found in either Galitz (1996) or Flavell (2002).

If we were, for example, pricing interbank interest rate swap cashflows, then LIBOR (London InterBank Offered Rate) interest rates would be appropriate. LIBOR discount factors can be derived from three principal sources:

- Short-term LIBOR deposits, which are zero coupon in style as the transactions only have two cashflows, one at the start and one upon maturity.
- Short-term interest rate futures, also zero coupon as they are priced off expectations of future LIBOR.
- Interest rate swaps, which require some mathematical manipulation since they suffer from the reinvestment issue noted earlier. Readers new to the pricing of swaps may be suspicious of pricing swaps from existing swaps, as they cannot reconcile the circularity. However, the fact is that a liquid market exists for the instruments with banks willing to quote for a variety of maturities, which therefore allows new swaps to be priced.

Table 1.2 Swap pricing inputs

Time period	LIBOR/swap rates	Cash lease rates	Lease rate discount factor	Forward lease rate
0.25	3.13%	0.09%	0.999775	
0.50	3.39%	0.12%	0.999400	0.1500%
0.75	3.60%	0.16%	0.998839	0.2249%
1.00	3.79%	0.19%	0.998104	0.2947%
1.25	3.89%	0.21%	0.997441	0.2658%
1.50	3.98%	0.22%	0.996705	0.2952%
1.75	4.07%	0.24%	0.995896	0.3252%
2.00	4.09%	0.25%	0.995012	0.3553%

Source: Barclays Capital; intermediate rates interpolated.

However, for gold lease swaps the situation is somewhat easier as it is possible to obtain lease rates as long as 10 years, which by convention are zero coupon in style. These lease rates can then be easily manipulated into discount factors.

To calculate the discount factors from zero coupon instruments of a maturity of up to one year, the necessary formula is:

$$\text{Discount factor}_t = \frac{1}{1 + (\text{Zero coupon rate}_t \times \text{Fraction of year})}$$

So, the calculation for the 9-month discount factor becomes:

$$\text{Discount factor}_{0.75} = \frac{1}{1 + (0.16\%\text{x } 0.75)} = 0.998801$$

(There is a small rounding difference between this result and the figure used in Table 1.2, which was derived using a spreadsheet.)

To calculate the swap rates beyond one year the formula is:

$$\text{Discount factor}_t = \frac{1}{(1 + \text{Zero coupon rate}_t)^n}$$

In this case n is defined as the number of years (or fraction thereof) from the effective date of the swap until the time of the cashflow. Therefore, for the 2-year discount factor the calculation is (again with a small rounding difference):

$$\text{Discount factor}_2 = \frac{1}{(1.0025)^2} = 0.995019$$

Table 1.2 provides an overview of a typical lease rate swap. The rates were observed in the market for value 11 April 2005. The terms of the swap are:

Underlying asset	Gold
Notional amount	50,000 ounces
Maturity	2 years
Settlement	Quarterly for both fixed and floating

Floating payments	Based on the 3-month lease rate at the start of each period
Fixed payments	Based on a single fixed rate of 0.2497%
Payment timing	Payments to be made in arrears
Base price	USD 425.40 (current spot rate)

In Table 1.2:

- The second column comprises LIBOR deposit rates up to 9 months and swap rates thereafter.
- The third column contains lease rates of different maturities observed from the market.
- The lease rate discount factor column represents discount factors of different maturities, which have been derived from the lease rates in the third column.
- The final column represents 3-month forward lease rates. For example, the rate of 0.15% represents the 3-month rate in 3 months' time. The lease rate discount factors in column 4 are used in the following mathematical relationship:

$$\text{Forward rate}_{a\times b} = \frac{\text{Short dated discount factor}}{\text{Long dated discount factor}} - 1 \times 4$$

$$\text{Forward rate}_{0.25\times 0.50} = \frac{0.999775}{0.999400} - 1 \times 4 = 0.15\%$$

Table 1.3 details the swap cashflows required to determine the fixed rate on the swap.

The floating cashflows in the fourth column of Table 1.3 are calculated as:

$$\text{Notional amount} \times \text{Spot value of gold} \times \text{Forward lease rate} \times \text{Fraction of a year}$$

The first floating cashflows payable at the end of the first quarter used the current 3-month lease rate of 0.09%. Hence, as a worked example, the floating payment due at the end of the first year would be:

$$50{,}000 \times \text{USD } 425.40 \times 0.2947\% \times 0.25 = \text{USD } 15{,}670 \text{ (small rounding difference)}$$

Table 1.3 Swap cashflows

	Fixed cashflows		Floating cashflows	
Time period	Gross	PV	Gross	PV
0.25	USD 13,275	USD 13,272	USD 4,786	USD 4,785
0.50	USD 13,275	USD 13,267	USD 7,974	USD 7,970
0.75	USD 13,275	USD 13,260	USD 11,957	USD 11,943
1.00	USD 13,275	USD 13,250	USD 15,668	USD 15,639
1.25	USD 13,275	USD 13,241	USD 14,079	USD 14,043
1.50	USD 13,275	USD 13,232	USD 15,671	USD 15,620
1.75	USD 13,275	USD 13,221	USD 17,263	USD 17,193
2.00	USD 13,275	USD 13,209	USD 18,855	USD 18,761
TOTALS		**USD 105,953**		**USD 105,953**
	NET PRESENT VALUE		0	

Note that the magnitude of the spot price is irrelevant and will not alter the fixed rate, which is eventually derived. Its main purpose is to convert the cashflows into a USD value. The third and fifth columns of Table 1.3 simply gives the present values of the cashflows by multiplying the gross cashflows by the discount factor of the appropriate maturity. Using the 1-year floating values in Table 1.3, this yields:

$$\text{USD } 15{,}668 \times 0.998104 = \text{USD } 15{,}638 \text{ (small rounding difference)}$$

The remaining floating cashflows are calculated in a similar fashion. The present values of the floating cashflows are then summed to give a value of USD 105,953. Since the swap has to be considered as an equitable exchange of cashflows at its inception – implying a net present value (NPV) of zero – we know that the present value of the fixed side must also be USD 105,953. The calculation of the present value of each fixed cashflow is:

$$\text{Notional amount} \times \text{Spot rate} \times \text{Fixed rate} \times \text{Fraction of year} \times \text{Discount factor}$$

This means that we have to solve for the unknown fixed rate by iteration. The single fixed rate that returns a present value of USD 105,953 for the fixed cashflows is found to be 0.2497%.

By looking at the fixed rate in relation to the forward lease rates in the final column of Table 1.2, it can be seen that the former is a weighted average of the latter. By looking at the fixed rate in this manner we receive an insight into the essence of swaps. Ignoring any underlying economic exposure, an entity paying a fixed rate must believe that over the life of the swap he will receive more from the floating cashflows. In other words the fixed rate payer must believe that actual lease rates will rise faster than those currently implied by the forward market. The opposite would be true of a fixed rate receiver; that is, he expects actual future lease rates to be below those currently implied by the forward market.

1.7 THE SPOT–FORWARD–OPTION RELATIONSHIP

Probably one of the most documented areas of finance is that of option pricing. Since the aim of this chapter is to give readers a basic understanding of where the value of a derivative instrument comes from, the analysis will avoid excessive discussions on the mathematics of options and concentrate on the intuition. Those interested in understanding the mathematics are recommended to refer to a variety of texts such as Galitz (1996), Tompkins (1994) and Natenberg (1994). A more recent and commodity specific text has been written by Geman (2005).

An option's premium is primarily dependent on:

- the expected payout at maturity
- the probability of the payout being made.

Although at an intuitive level these concepts are easy to understand, the mathematics behind the principles is often complex and discourages many readers. To determine the premium on an option, a variety of inputs are required, as is an appropriate model. Those inputs will include

Table 1.4 Option pricing parameters

Spot	USD 425.40
Strike	USD 425.40
Time to maturity	6 months
Six month LIBOR	3.39%
Six month lease rate	0.11571%
Six month forward price	USD 432.40
Implied volatility	15%

- The spot price
- The strike price
- Time to maturity
- The cost of carrying the underlying asset as a hedge (i.e. any income earned through holding the underlying asset less any expense incurred)
- The implied volatility of the underlying asset.

Table 1.4 shows the values that will be used for the following option pricing examples. We will assume that the option is European in style and, given the parameters set out in the table, a call option would be in-the-money since the strike price is less than the forward price. Using a Black–Scholes–Merton model, the premium is estimated at USD 21.50 per troy ounce (option premiums are quoted in the same format as the underlying asset).

An option premium can be decomposed into two elements: time value and intrinsic value. The intrinsic value can be thought of as the amount of profit the buyer would make if he were to exercise the option immediately. However, the term is defined in such a way that it does not take into account the initial premium paid. The intrinsic value for a call option can be expressed in the following manner:

$$\text{Intrinsic value of a call} = \text{MAX}(0, \text{Underlying price} - \text{Strike})$$

For put options it is expressed as

$$\text{Intrinsic value of a put} = \text{MAX}(0, \text{Strike} - \text{Underlying price})$$

In both expressions the underlying price is the forward price for European style options and spot for American style. In the original definition of intrinsic value prior to maturity, the difference between the underlying price and the strike price should be present valued since exercise of the European style option could only take place at expiry. However, when analysing options intuitively, traders would probably disregard this aspect.

Time value is the extra amount that the seller charges the buyer to cover him for the probability of future exercise – a sort of uncertainty charge. Time value is not paid or received at expiry of the option as, for buyers, it will fall to zero over the option's life. When an option is exercised the holder will only receive the intrinsic value. Time value is only paid or received upon the sale or purchase of an option prior to maturity. This explains why early exercise of an option is rarely optimal. If a purchased option is no longer required, it would usually be more efficient to neutralise the exposure by selling an opposite position in the market in order to receive the time value.

An understanding of the two option premium components is vital in order to understand the logic of interbank transactions. Tompkins (1994) shows clearly that movement in the underlying market will drive intrinsic value, while time value is influenced primarily by time and implied volatility.

1.8 PUT–CALL PARITY: A KEY RELATIONSHIP

Although pricing models provide one linkage between the underlying price, the forward rate and the option premium, the concept of put–call parity is an alternative representation. (Tompkins, 1994, provides a detailed analysis of the concept.)

Put–call parity is a concept that attempts to link options with their underlying assets such that arbitrage opportunities could be identified. The conditions of put–call parity will hold as long as the strike, maturity and amount are the same.

Although put–call parity varies according to the underlying asset, in its simplest form it can be represented by the expression:

$$C - P = F - E$$

where C = price of a call option
P = price of a put option
F = the forward price of the underlying asset
E = the strike rate for the option.

For the purposes of this case study the formula will be truncated to

$$C - P = F$$

where F is redefined as a position in the underlying, in this case a forward position. Each of the symbols will be expressed either with a "plus" or a "minus" to indicate a buying or a selling position, respectively. By way of the previous expression, one could therefore link the underlying market and the option by rewriting the equation as:

$$+C - P = +F$$

That is, buying a call and selling a put is equivalent to a long position in the underlying. Although this may seem a very dry concept, the relationship is frequently used in financial engineering to create innovative structures.

An option's "Greeks" (see section 1.10) are measures of market risk that have been developed over time to help market participants to manage the risk associated with the instruments. Each of the option model inputs has a related Greek whose value can be calculated numerically (see Tompkins, 1994, or Haug, 1998) or derived by perturbing the appropriate option model. The majority of the Greeks look at how a change in the relevant market input affects the value of the option premium – that is, they are "first order" effects; measures such as gamma are second order in nature.

1.9 SOURCES OF VALUE IN A HEDGE

One way in which a put call could be applied is in the construction of hedges. Let us take an example of a client wishing to buy a commodity forward (e.g. an automotive hedger).

If the hedger wanted to buy the commodity at a strike rate equal to the current forward rate, then the price of a call and a put will be the same. Not unreasonably, he may wish to achieve a strike rate E that is less than the current forward rate F, which would suggest that, for the relationship to hold, the price of the call will increase relative to the put. The challenge therefore becomes to obtain a more favourable hedging rate by cheapening the value of the call option.

This can be achieved using three possible strategies:

- Buy a call on a notional amount of N and sell a put with a notional amount of $2N$ (sometimes referred to as a ratio forward).
- Buy a reverse knock-out call option with the barrier set at a level that is unlikely to trade.
- Buy a short-dated call option and finance this by selling a longer-dated put option at the same strike.

1.10 MEASURES OF OPTION RISK MANAGEMENT

1.10.1 Delta

The delta of an option looks at how a change in the underlying price will affect the option's premium. In one sense, delta can be thought of as the trader's directional exposure. There are a variety of different definitions of delta:

- The rate of change of the premium with respect to the underlying
- The slope of the price line
- The probability of exercise
- The hedge ratio.

The most traditional way of defining delta is the rate of change of the premium with respect to the underlying. Using this method, with a known delta value the analyst can see how a small change in the underlying price will cause the premium to change. In formula, delta is expressed as:

$$\text{Delta} = \frac{\text{Change in premium}}{\text{Change in underlying price}}$$

So, if an option has a delta value of 0.58, we can say that the premium should change by 58% of the change in the underlying. Using the option parameters presented in Table 1.4, if the underlying price rises by 50 cents from USD 425.40 to USD 425.90, the premium should rise by 29 cents to USD 21.79.

Delta will be either positive or negative depending on whether the option has been bought or sold or is a call or a put. If I have bought a call or sold a put the associated delta value will be positive because, in both instances, the value of the option will rise if the underlying price rises – a positive relationship. If I have bought a put or sold a call the delta value for these options will be negative because a rise in the underlying price will cause the options to fall in value – a negative relationship. Both of these statements can be verified by looking at the payoff diagrams that were introduced in Figure 1.1.

Table 1.5 Premium against underlying price; illustrating delta

Underlying price (USD)	Premium (USD)	Delta value
325	0.11	0.01
350	0.77	0.05
375	3.27	0.16
400	9.60	0.35
425	21.27	0.57
450	38.23	0.77
475	59.16	0.89
500	82.40	0.96
525	106.74	0.99

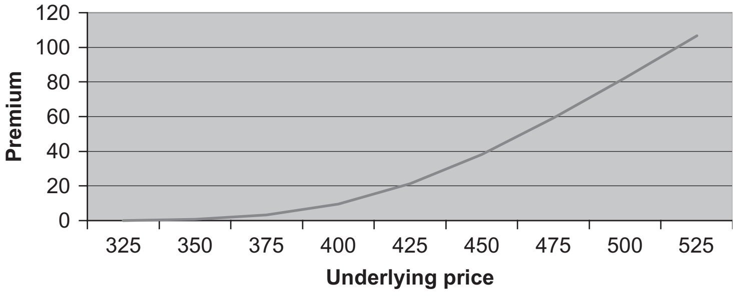

Figure 1.4 Premium against price (prior to maturity)

The second way of illustrating the concept of delta is to see how the value of the premium changes as the underlying spot price changes (all other things being equal). Table 1.5 shows the premium and delta values for an instantaneous change in the spot value.

The result of plotting these values is shown in Figure 1.4.

The second definition of delta is the slope of the price line prior to maturity (e.g. premium against price). Delta is the slope of a tangent drawn at any point on the curve. This also suggests that delta should only be used to measure small changes in the underlying price. As the premium – price relationship is convex in nature, using delta to determine the impact of significant price movements would give an incorrect estimate.

Although it is not strictly true, delta is often interpreted as the probability of exercise. The application of the term in this definition is not uniformly accepted but it certainly provides a useful rule of thumb.

Delta can be measured in a variety of ways:

- As a number whose value will range from 0 to 1.
- As a percentage.
- As a delta equivalent figure.

The most common way of expressing delta is a value ranging between 0 and 1, with an associated positive or negative sign. Out-of-the-money options will have a delta that ends towards 0 while in-the-money options will tend towards 1. It would be incorrect to say that all at-the money options have a delta of 0.50 but the value will certainly be close. An alternative method is to express it as a percentage number, which was also illustrated earlier. The third way is to express it as a delta equivalent value. This technique compares the market risk position of the option with that of an equivalent underlying asset. Thus, if for example we were to buy a call option on a notional amount of 50,000 ounces with a delta value of 0.58, the option has the same degree of market risk (for small changes in the underlying asset) as 29,000 ounces (50,000 × 0.58)

The concept of delta can be extended to option positions. If we were to add another bought call option to the above position with a notional of 60,000 ounces and a delta of 0.40, this individual position would have a delta equivalent value of 24,000. On a net basis, therefore, the entire option position has a delta exposure of 53,000 ounces (24,000 + 29,000) – the option trader's position that would have the same exposure to small movements in the underlying price as an actual physical holding of 53,000 ounces. This technique is referred to as delta hedging and demonstrates the use of delta as hedge ratio.

If the option trader chose to neutralise this exposure, he could take an offsetting position in other options or the underlying instrument. He may wish to do this if he is seeking to benefit from a change in another market factor, such as implied volatility. In the previous example the trader is delta positive, so his option position will rise in value if the underlying market price rises and vice versa. To neutralise the exposure he could therefore sell 53,000 ounces of gold in the underlying market or trade options in different combinations to achieve a net delta exposure with which he is comfortable. If the net delta position is reduced to zero, this is described as delta neutrality.

If the delta neutralising trades are executed in the underlying market, the trader is faced with making a choice between trading in either the spot or the forward market. If the option position is European then the forward market should be chosen, because the European option can only be exercised upon maturity and the appropriate equivalent market is therefore the forward market. However, some traders may choose to hedge in the spot market and then hedge the resultant interest rate risk (recall that the difference between the spot and forward markets is driven by interest rate considerations).

We have seen that delta changes as a function of the underlying price but will also evolve over time. For example, as the option approaches maturity the price line shown in Figure 1.4 above will become more linear in nature and at maturity will resemble the hockey stick shape that was introduced Figure 1.1. As a result, the delta of the option will tend towards 0 or 1.

1.10.2 Gamma

Gamma is correctly defined as the rate of change of delta with respect to the price. By looking at the rate of change, this makes it a second-order function.

$$\text{Gamma} = \frac{\text{Change in delta}}{\text{Change in underlying price}}$$

However, this definition is not really helpful in trying to grasp what gamma is actually doing. Gamma will be studied in detail in Chapter 2, so some of the definitions given here may not seem clear at this point. We will define gamma as:

- The speed with which a delta hedged position becomes unhedged
- The exposure to actual volatility in the market
- The rate of change of a trader's profit or loss
- The exposure to significant changes in the underlying price.

Gamma is positive for option buyers (irrespective of whether the option is a call or a put) and negative for sellers. Using the call option in Table 1.3, the option pricing model returned a value of 0.58 for delta and 0.009 for gamma. If we take an exposure of 50,000 ounces the delta exposure is equivalent to owning 29,000 ounces of the underlying of the asset. We will assume that the trader hedges this exposure by selling an equivalent amount in the underlying market. However, although the position is initially delta neutral a movement in the underlying will lead to a change in delta and the position will no longer be 100% hedged. The gamma measure can be used to predict the magnitude of any move in delta for a small move in the underlying price. So, in this instance we can say that for a small increase in the underlying market price, the delta will move to 0.589 and for a small downward move it would move to 0.571.

1.10.3 Theta

Theta reflects the impact of time on the value of the option and is defined as:

$$\text{Theta} = \frac{\text{Change in premium}}{\text{Change in time to expiry}}$$

Theta is positive for sellers and negative for buyers. This reflects the fact that, for buyers, options will experience time decay. In other words, the buyer will pay a premium and will acquire a contract that has time value and possibly intrinsic value. However, over time, the time value element will fall to zero. Using the original call option example, and holding all the other parameters constant, the effect of the passage of time on the option is reflected in Table 1.6.

Table 1.6 The effect of time on an option

Time to expiry (years)	Premium (USD per troy ounce)	Change in premium
1.0	32.74	–
0.9	30.74	2.00
0.8	28.66	2.08
0.7	26.49	2.17
0.6	24.21	2.28
0.5	21.79	2.41
0.4	19.19	2.60
0.3	16.33	2.86
0.2	13.06	3.27
0.1	9.02	4.04
0.01	2.88	6.14

Table 1.6 shows that, for the buyer of an option, its value will decrease as time passes and that the rate of decay will accelerate as the option approaches maturity. The opposite will hold true for sellers of options. Theta can be thought of as the slope of the price line (column 2 in Table 1.6) with respect to the time to expiry. Theta will be small initially and will increase with the passage of time (column 3 in Table 1.6).

1.10.4 Vega

Vega is defined as the change in the option premium for a 1% change in an options implied volatility. It is expressed as:

$$\text{Vega} = \frac{\text{Change in premium}}{\text{Change in implied volatility}}$$

Although most people are happy with the concept of volatility in general terms – change in the market with no direction suggested – the notion of implied volatility still remains difficult to grasp.

A related concept is historical volatility, which, as the term suggests, measures the magnitude of movement of the underlying asset over a historical period, measured in per cent per annum. The statisticians would describe this as a standard deviation. However, as is often cited in financial markets "that past performance is not a guide to future performance", there is a need for a more forward-looking measure of volatility.

One of the most popular ways of trying to explain the concept is to describe it as the volatility implied in an observed option price. The rationale for this is that if one looks at all of the option pricing model inputs, the only real unknown is the implied volatility. Terms such as the spot rate or strike are either easily observable or negotiated as part of the option contract. Since it is the only unknown factor, the traditional definition suggested taking an observed option premium, inserting the value into the option pricing model, running the model in reverse and backing out the volatility implied by this price. However, this argument traditionally ceases at this point and does not address the obvious circular argument: Where did the market maker quoting the observed option price obtain his volatility input? It also gives no real feel for what the measure is trying to achieve. Tompkins (1994, p. 139) describes it as "the risk perceived by the market today for the period up until the expiration of a particular option series". However, since no one can foretell the future with complete certainty, this is nothing more than a "best guess". As an anecdote, the author recalls teaching a class attended by an experienced option trader who described the interbank options as the "market for guesses". In the interbank market option prices are quoted in terms of a bid–offer implied volatility number. Since implied volatility is the only truly unknown variable, this is the factor that is traded on an interbank basis. Trading strategies that aim to exploit movements in implied volatility will be analysed in the chapter on interbank trading strategies.

The main confusion regarding implied volatility often surrounds the relationship with movements in the underlying asset price. The delta value of an option is sensitive to the magnitude of the implied volatility input, and as the size of the volatility number increases the delta on the option will tend towards 0.50. This is because the range of values that the underlying is expected to take at maturity is wider and the likelihood of exercise of

Table 1.7 The impact of implied volatility on a variety of options

Implied volatility	OTM option		ATM option		ITM option	
	Premium	Change in Premium	Premium	Change in premium	Premium	Change in Premium
0%	0.00		0.00		17.60	
5%	1.00	1.00	6.01	6.01	18.02	0.42
10%	5.48	4.48	12.00	5.99	22.25	4.23
15%	10.98	5.50	17.99	5.99	27.49	5.24
20%	16.77	5.79	23.98	5.99	33.03	5.54
25%	22.68	5.91	29.95	5.97	38.69	5.66
30%	28.64	5.96	35.92	5.97	44.41	5.72

an out-of-the-money option will therefore increase and for an in-the-money option it will decrease. However, it does not follow that as the underlying price starts to move, the implied volatility must also move. In an option model the implied volatility and the spot price are two independent variables and, as such, a movement in one factor can be independent of the other. Although it may be reasonable to assert that current observed movements in the spot price may alter a trader's expectations on the distribution of spot values at maturity, there is no inbuilt dependency. There is no "ruling" that says if spot were to change by $X\%$ that implied volatility must change by $Y\%$.

The confusion over the relationship between spot and volatility is similar to that seen between vega and gamma. Gamma was earlier introduced as a trader's exposure to actual volatility. If we saw a large value for gamma then the difference between two delta values will be substantial for a given change in the underlying market. If the difference between two delta values is significant, this means that the movement of the spot prices was quite large, and one could reasonably describe such a market as being currently volatile.

Table 1.7 shows the effect of implied volatility on the premium for three variants of the call option introduced in Table 1.4:

- An out-of-the-money (OTM) option with a strike rate of USD 450.
- An at-the-money (ATM) forward option with the strike rate equal to the forward rate of USD 432.40.
- An in-the-money (ITM) option with a strike rate of USD 415.

As before, all of the other pricing parameters are held constant.

From Table 1.7 we can draw a number of conclusions:

- The relationship between an ATM option premium and implied volatility is proportional. Doubling the implied volatility will double the premium.
- If the option has no implied volatility, an OTM or ATM option will have no premium, although the ITM option will have its intrinsic value. This also suggests that volatility is arguably the defining factor that distinguishes an option from its underlying asset.
- The relationship between an OTM option premium and the implied volatility is also non-linear. Notice how a doubling of the implied volatility at low levels causes the premium to rise by a much greater magnitude.

Option buyers are said to be long vega (i.e. they will benefit if implied volatility rises), while sellers are short vega – they will suffer if implied volatility rises, but will benefit if it falls. The logic behind this lies in the pricing formula for options. Recall that the premium of an option is composed of the intrinsic and time value. The intrinsic value of the option is driven by the difference between the price of the underlying and the strike price. The main drivers of the time value are time and implied volatility. The buyer of an option acquires both elements and, all other things being equal, an increase in implied volatility will increase the value of the option – a positive relationship.