1 An Introduction to the Major Asset Classes

This introductory, and slightly eclectic, chapter focuses on the liquid investment asset classes in which derivatives and structured products are normally developed, priced and traded, namely equities, fixed income, commodities and foreign exchange. The aim is not to describe in detail the everyday products traded in these markets but more to give a sense of the general price characteristics of these markets, e.g., how they move through time. There are many excellent product books on the market and we would recommend readers interested in the range of products available to pick up a copy of one of them; we have listed some in the References.

In this chapter, we aim to make some observations on how to model underlying asset price behaviour for each asset class.[1] In order to have a real understanding of the value that can be extracted from financial derivatives, one must understand how prices behave in their composite parts. This is not a quantitative finance textbook, however, and we do not aim to be completely rigorous. Instead, in this chapter as in the others, we try to get across an intuitive (rather than an academic) understanding.[2]

On the way, we will also highlight some interesting features of the markets under review so that you, the reader, can feel more personal affinity with them; we think that, if your interest is tweaked, you will enjoy the book (and the chapter) much more. Many readers of this book will have experience in one or two of these asset classes but probably not all. If short on time, you could, of course, read the sections on the least familiar asset classes and skip the rest. That said, if you do have some experience but you are a little rusty, the following sections should be a good review.

1.1 EQUITIES

1.1.1 Introduction

Equities are perhaps the most familiar of the main asset classes. The shareholder capital or capital stock (or, merely, stock) of a company represents the capital paid into or invested in the business by its shareholders. Along with other assets in the firm, shareholder capital serves as security for the creditors of that company. The capital stock is divided into *shares* (also called *equity*). The company may have different types (*classes*) of equity, each class having distinctive ownership rules, privileges and/or prices.

Equity typically takes the form of *common* or *preferred*. Common shares typically carry voting rights, which shareholders can exercise at certain times to influence the company and its directors. Preferred shares typically do not carry voting rights but holders are legally entitled to receive their dividends before other shareholders. Since holders of preferred shares get preferential treatment, they are called *preference* shares in the UK.

[1] There will be a little mathematics but not much. Instead, we hope to explain the issues more intuitively.
[2] If you are looking for a quantitative finance textbook then probably anything by Wilmott is worthwhile.

Convertible preferred shares are preferred shares that include an option for the holder to convert them into a fixed number of common shares, usually any time after a predetermined date.

Preferred shares may sometimes have a hybrid structure, having the qualities of bonds (such as a fixed percentage dividend) and common stock voting rights.

1.1.2 Pricing equities

The price of a single stock/share/equity (the terms are interchangeable) is determined as always by the interaction of supply and demand. Factors that impact only on the single stock price are called *stock-specific*. These factors can include a change in CEO, new patents awarded to the company, new product development, etc. The prices of individual stocks are also influenced by factors that affect the entire equity market, however. These are *market-specific*. In fact, statistical analysis suggests that the majority of the variation in the prices of individual stocks usually comes about from events that impact the overall equity market, for example changes in monetary policy or the risk tolerance of investors, rather than stock-specific factors. Stocks that tend to move by more than the overall market moves are called *high-beta* stocks while those that move less than the market are called *low-beta* stocks. This is essentially a carry-over from the Capital Asset Pricing Model (CAPM),[3] which you do not really need to know much about for the purposes of this book. High-beta stocks tend to be in cyclical sectors, such as consumer durables, property and capital equipment. Low-beta stocks tend to be non-cyclical, such as food retailers and public utilities. Another way of describing low-beta stocks is *defensive*.

To be sure, there is a huge industry dedicated to predicting likely future movements in the price of a single stock or equity market. The three main approaches are *fundamental analysis*, *technical analysis* and *quantitative analysis*. We discuss these techniques below. There are some more esoteric theories of equity price determination, such as those predicated on sun spot activity or the Chinese horoscope, for example. What all these theories have in common, however, is that they have their fans as well as their detractors.

Of course, investors could just ignore forecasting entirely and invest passively, judging that the future direction of price movements is unpredictable in all but the long run. We will say more on this below.

1.1.3 Fundamental analysis

Fundamental analysis attempts to use economics and accountancy to value a single stock or the overall equity market (or sector), thereby obtaining the *fair value price*. Investors can then compare the actual price with the fair value price and then decide whether the stock price or index level is over- or undervalued. Put simply, analysts attempt to assess the likely flow of cash flows from the company or the market as a whole and discount these using a relevant discount rate. The net present value of the future cash flows is then the equilibrium (or fair value) price of the equity or the overall market.

[3] The beta is usually estimated by regressing the historic returns to the stock against the returns to the market over a given sample period (for example, or one business cycle). Alternatively, beta may be interpreted as a stock's contribution to the volatility of the total market portfolio. In this case, it is the covariance of returns between the stock and the "market portfolio", divided by the variance of returns to the market portfolio. Variance is a measure of dispersion and is the square of the standard deviation of returns. One of the problems with stock betas, however, is that they are notoriously unstable and significantly different estimates will be obtained depending on the sample period selected for the price return histories.

Although discredited by many, the Efficient Markets Hypothesis (EMH) suggests that equity markets are essentially rational. That is, the price of any stock (and, consequently, market index) at any given moment represents a complete evaluation of the known information that might have a bearing on the future value of that stock. In other words, stock prices are, at any point in time, equal to the sum of the discounted future cash flows. Furthermore, prices can only change if there is some "news" that has an effect on those future cash flows.[4] In a sense, the EMH supports the idea of fundamental analysis although there is a chicken-and-egg problem here. The EMH says that fundamental analysis to value a stock or equity market is unnecessary since current stock prices are already at their fair value. But, of course, how did the market arrive at the prices in the first place? Presumably some entity somewhere must be cranking out the numbers?

In recent years, even the most die-hard fans of the EMH have come to accept that the markets are not perfectly efficient, especially the least liquid ones. To quote Professor Robert Schiller, "the Efficient Markets Hypothesis is a half-truth".

1.1.4 Technical analysis

Technical analysis attempts to predict future price changes in a single stock or the overall equity market (or sector) by analysing past price movements. Technical analysis is extremely popular among certain market participants although there is considerable controversy attached to its use. Certainly, if you subscribe to the school of thought that says that equity markets are efficient, then there is no place for technical analysis. For a more in-depth discussion of this point, see Chapters 10 and 11 of *Financial Theory and Corporate Policy* (2003) by Copeland, Weston and Shastri.

Whichever approach you prefer, it is irrefutable that stock prices are determined by supply and demand. Although new stock can be issued, for example when a company is wishing to raise new capital for fixed investment, stock prices at the overall equity market level are likely to change from day to day because demand conditions are changing rather than because the net supply of equity is increasing. If equity investors are feeling particularly confident, then they may well allocate more capital to stocks and this will push up prices. On the other hand, if equity investors are feeling much less confident, then they may well take capital out of the stock market and this will push down prices. If changes in the confidence levels of investors are unpredictable, as seems the case empirically, then short-term price changes will tend to be random.[5] Of course, fundamental analysts will argue that the market or single stock has temporarily moved away from fair value and recommend buying (if cheap) or selling (if rich). The technical analysts, on the other hand, will try to explain these short-term movements with reference to past price patterns.

1.1.5 Quantitative analysis

But what about the passive investors, those who judge that the future direction of price movements is unpredictable in all but the long run? Empirically, single equity prices and stock market indices tend to follow a pattern that closely resembles what mathematicians call

[4] Although not recent, a 1988 NBER Working Paper, entitled *What Moves Stock Prices?*, found that "large market moves often occur on days without any identifiable major news releases (which) casts doubt on the view that stock prices movements are fully explicable by news about future cash flows and discount rates".

[5] "I can predict the motion of heavenly bodies but not the madness of crowds", Sir Issac Newton.

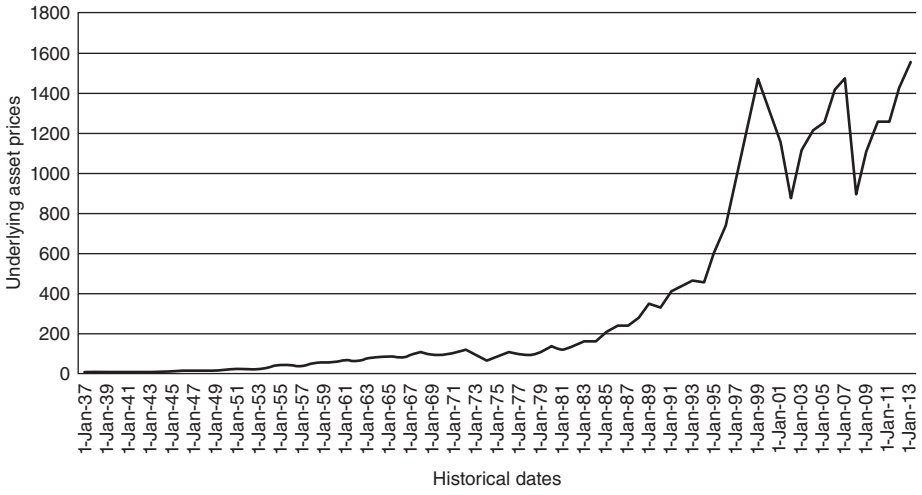

Figure 1.1 S&P 500 historical prices

a random walk (with drift). Predicting future price changes over anything but the medium to long term in such a world is likely to be unfulfilling. Investors can "predict" that they will earn a long-run average return above the risk-free rate (since equities are risky) but the intermediate path will be far from predictable. In this case, a single investor can only systematically "beat the market" by being incredibly lucky. This is the world of the quantitative analyst. Put very simply, quantitative analysts attempt to model the returns from assets, rather than their fair value prices since, empirically, price movements appear random but returns, over time, seem to follow something akin to a normal distribution (Figure 1.1).

More recently, academics have suggested that, in fact, equity prices may more closely resemble a first-order autoregressive process than a random walk. Without wishing to get too mired in mathematical-speak, such a process allows for mean reversion to the trend. If so, then there may be opportunities, after all, for savvy investors to beat the market. They just need to be able to determine the trend!

Neither the random-walk model nor the autoregressive model seems to be an entirely accurate description of the process of equity price determination, however. Stock markets are much more volatile than such models would seem to imply. In particular, they do not seem fully to be able to explain the prevalence of stock market crashes.[6] And, the models do not explain why technical analysis has so many fans. If equity prices change with a random walk, then technical analysis is useless. If equity prices follow a first-order autoregressive process, however, then it may be possible to predict price moves from statistical analysis but the "head-and-shoulder" and "double-bottom" formations popular with technical analysts the world over are likely to be little more than colourful nonsense. As Simon Beninga suggests in his book *Financial Modeling*, "if you are going to be a technician, you have to learn to say these things with a straight face".

[6] Many quantitative analysts argue that, in practice, it is not worth trying to model crashes, for example. While they can create significant damage, there is very little that we can actually do about them. They may appear more regularly than the normal distribution would imply but their prevalence still tends to be relatively limited. So, in practice, one can ignore them for modelling purposes as long as one is also aware of them and tries to account for their effects.

So, why does technical analysis appear to work at some level? Perhaps because it has lots of fans who follow the techniques and take investment decisions based on such. Is this rational, not entirely so but, then, humans are far from rational. How else can one explain the Dot-Com Bubble?

We have described the processes driving actual equity prices in some detail since those same processes are at work in the other asset classes, most especially commodities, foreign exchange and indices. Fixed income products are much less easy to model, however.

1.1.6 The equity risk premium and the pre-FOMC announcement drift

The equity premium is the difference between the "average" return on the stock market and the yield on short-term government bonds. Most academic research into the size of the premium finds that it is too high to compensate for the average riskiness of equities. The *equity premium puzzle* is the name given to this phenomenon. Mehra (2008) provides a review of the literature on this topic.[7]

A recent New York Fed staff report suggested that since 1994 more than 80% of the equity premium on US equities was earned during the 24 hours preceding scheduled Federal Open Market Committee (FOMC) announcements (which generally occur just eight times a year).[8] The researchers called this phenomenon the *pre-FOMC announcement drift*.

The interesting aspect to note is that the pre-FOMC announcement drift is generally earned *ahead* of the Fed's announcement and so it is not directly related to the actual post-meeting monetary policy actions. The historical record suggests that equity prices rise in the afternoon of the day before FOMC announcements and then rise even more sharply the next morning. Following the announcement, equity prices can vary widely but, on average, they tend to finish the day around the same levels prevailing at the time of the announcement. In other words, unchanged. The gain over the 2 days as a whole, however, averages about 0.5%.

On the other hand, since 1994, historic returns from equities are in line with the returns on government bonds if the windows around scheduled FOMC announcement days are excluded.[9] Note that similar patterns are found in other equity markets,[10] although, in these cases, they are reacting to the FOMC announcements rather than announcements from their own central banks.

Finally, while one might expect similar patterns to be evident in other major asset classes, the researchers concluded they were not. In other words, the pre-FOMC drift seems to be restricted to equities.

1.2 COMMODITIES

1.2.1 Introduction

What is a *commodity*? The standard definition of a commodity is "any marketable item produced to satisfy wants or needs" and can include goods and services. For most of us, however, a commodity is essentially a basic resource that we use that comes out of the ground

[7] See http://www.academicwebpages.com/preview/mehra/pdf/FIN%200201.pdf.

[8] See http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1923197.

[9] The return on the 24-hour period ahead of the FOMC announcement was about 3.9% per year, compared with only about 0.9% per year for all other days combined. In other words, more than 80% of the annual equity premium is earned over the 24 hours preceding scheduled FOMC announcements.

[10] The FTSE 100, DAX, CAC 40, SMI, IBEX and the TSE.

(e.g., wheat, corn, sugar, coffee, copper, aluminium, gold, silver, crude oil, natural gas, coal, diamonds, uranium, etc.). Obviously, this description is far from rigorous but it is intuitive. We can go one step further and say that soft commodities are grown, while hard commodities are extracted through mining or drilling.

Commodities generally have actively traded derivative markets while the spot markets tend to be less liquid. The most active derivative markets are the commodity futures exchanges in the USA and the UK, although the first real futures market was probably the rice futures market that began in seventeenth-century Osaka, Japan. Spot markets provide for the immediate delivery of the relevant commodity while the derivative markets provide for or accommodate delivery at some point in the future; in some cases, as much as 10 years out. These markets help to determine "world" prices although most transactions in commodities do not actually take place "on market". Instead, they are bilateral arrangements undertaken directly between a buyer and a seller. The prices at which these trades take place are generally not made public, since they are private transactions, but they will tend to be determined to some extent in reference to the prices traded on the public markets.

Most commodity products will have undergone some form of basic refining before they are traded publically. For example, the copper that is traded on the London Metal Exchange (LME) is 99% pure, is in the form of cathode (a flat sheet of metal) and is quoted in lots of 25 tonnes (if one buys one lot, one is buying 25 tonnes of copper). Other commodities that will have gone through some form of refining include gasoline and fuel oil (both derived from crude oil), orange juice, lean hogs, sugar (refined from raw sugar), etc. Another important commodity is electricity. Electricity is derived from other fuel sources (oil, coal, nuclear, renewables, etc.) and has the particular characteristic that it is uneconomical/difficult to store. Hence, electricity is produced as and when it is demanded/consumed.

The main participants in the commodity markets are producers (mining companies, farmers, refiners, oil companies, etc.), manufacturers (who use the raw materials to make consumer durables and non-durables), households (who use electricity or gasoline), traders (who help the producers and consumers to hedge) and investors. Governments are also involved since they tend to intervene, perhaps to build strategic reserves in the case of the US government or to manipulate the price in the case of OPEC. We discuss the main participants in later sections describing their typical activity.[11]

Another aspect of commodities is that, although they tend to be lumped together under one single heading, there is no such thing as an average commodity. Their prices behave differently and they are influenced by commodity-specific factors. Gold prices will be impacted by the festival season in India, oil prices by the driving season in the USA, Australian natural gas prices by tsunamis in northern Asia, global grain prices by droughts in the USA, Russia and Australia, copper prices by earthquakes in Chile and coffee prices by the frosts in Brazil and Colombia. As the global economy expands, we can assume that commodity prices will rise but the actual path of prices will be heavily dependent on the forces of nature.[12]

1.2.2 Hedging

Producers and consumers of commodities will use the derivative markets to hedge their future production or consumption. Producers will want to lock in the prices of their future output

[11] An excellent book on the real life world of commodities is *The King of Oil, The Secret Lives of Marc Rich* by Daniel Ammann.

[12] In the short run, prices will be volatile and this volatility will tend to outweigh the underlying growth rate. In the medium to long run, however, prices will still be volatile but this volatility will tend to be outweighed by the underlying growth rate.

so that they have certainty of revenue. This certainty is required since their capital spending requirements are enormous and spread out over many years. Likewise, consumers will look to lock in the supply of raw materials over time since they will not want to risk not being able to find the materials they need to manufacture their finished goods. This is hedging. The producers will generally sell their future output using prices determined today. They will deliver their output at predetermined dates in the future and receive the price at the level set today. The consumers will generally buy their future needs at prices determined today. They will receive the raw materials at predetermined dates in the future and pay the price at the level set today.

The hedging actions of the producers and consumers very rarely match. The producers tend to hedge their output 3 to 5 years out while the consumers tend to hedge their demand out to 2 years. Moreover, the producers will tend to want to hedge when average prices are high (locking in high prices), while consumers will want to hedge when average prices are low (locking in low prices). Someone will need to sit in the middle, the *speculators*.

The traders generally get a bad press from governments, regulators, consumer bodies and charities since they are perceived as profiting from some kind of unfair activity in commodity prices. Such complaints are ill-targeted, however. Producers and consumers need to hedge. Without this hedging, producers and consumers would not be able to plan effectively. By stepping in to offset the mismatch between the hedging activities of the producers and the consumers, the traders are providing a publically undervalued service.

1.2.3 Backwardation and contango

Backwardation describes a situation where, for a given commodity, short-dated prices are higher than longer-dated prices. For example, if the price for delivery of copper today is $8,000 per tonne and the price for delivery 5 years into the future is $6,000 per tonne, then the copper curve would be backwardated. In this situation, one can argue that prices are assumed (or discounted) to fall over time (Figure 1.2).

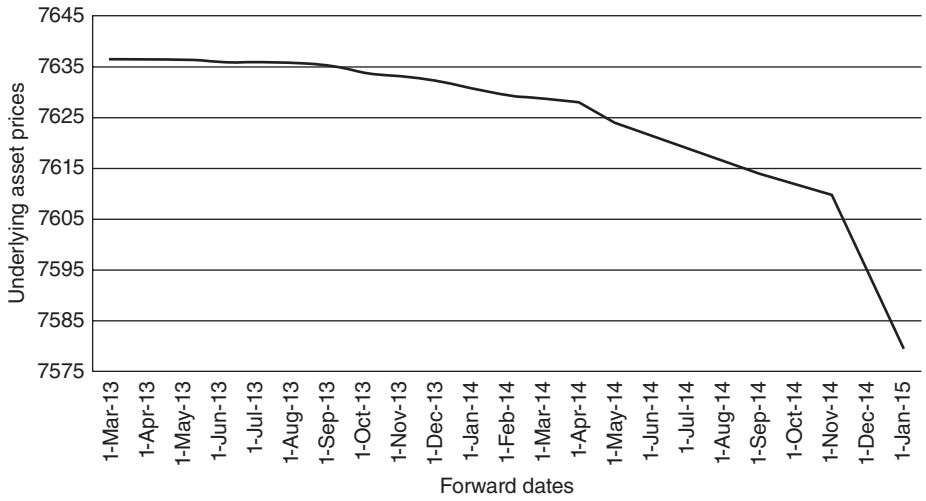

Figure 1.2 Backwardation curve

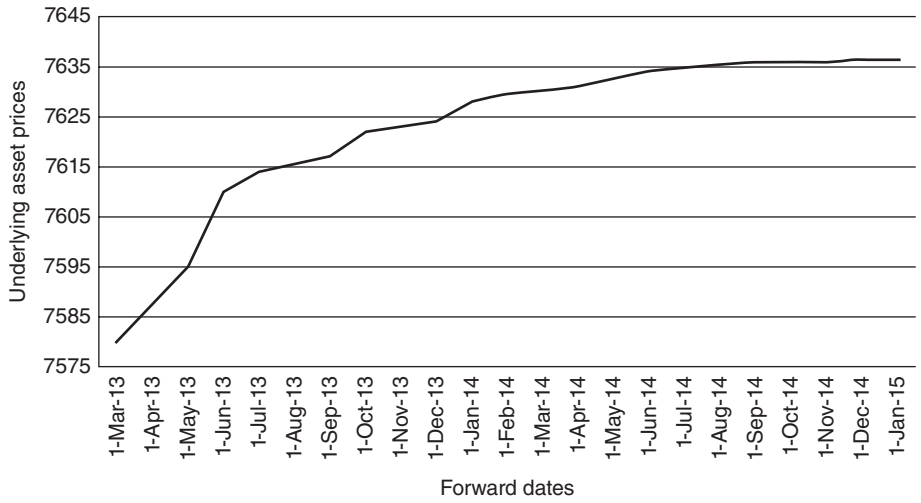

Figure 1.3 Contango curve

In a curve that is in *contango*, short-dated prices are lower than longer-dated prices. For example, if the price for delivery of copper today is $8,000 per tonne but the price for delivery 5 years into the future is $10,000 per tonne, then the copper curve would be in contango (Figure 1.3).

Which curve shape is the natural state of affairs? When there is excess supply of a particular raw material available, perhaps because the world economy is in recession and demand is weak, then the curve for that commodity may be in contango. Prices for immediate or near-immediate delivery will be low since producers will be happy to sell their excess inventories. Over time, however, producers will tend to cut back on their production, especially the high-cost producers. This will reduce the available supply, bringing supply and demand into balance at higher prices. The commodity markets will anticipate this and so longer-dated prices, prices for delivery further into the future, will be higher than short-dated prices.

When there is excess demand for a particular raw material, perhaps because the world economy is booming and demand is strong, then the curve for that commodity may be in backwardation. Prices for immediate or near-immediate delivery will be high since producers will demand high prices for their production. Over time, however, consumers may switch to substitutes or completely alternative products. This will reduce demand over time, bringing supply and demand into balance at lower prices. The commodity markets will anticipate this and so longer-dated prices, prices for delivery further into the future, will be much lower than short-dated prices.[13]

The ability to store or warehouse commodities (with the notable exception of electricity) might suggest to the average market participant that the "normal" state of commodity curves should be gently upward-sloping, i.e., in gentle contango. If a particular commodity curve

[13] Some commodity analysts interpret the long end of the commodity curve as a guide to the long-term marginal cost of production for that commodity. Empirically, when we look at forward curves at different points in time, long-term prices tend to cluster within a smaller range of prices.

were too steeply upward-sloping and prices for the immediate or near-immediate delivery of the commodity were much lower than longer-dated prices, then speculators could buy the commodity, store it in a warehouse and simultaneously sell for future delivery at the higher longer-dated price. Short-dated prices would tend to rise relative to longer-dated prices as the excess supply was taken out of the market and stored. The profit from such a transaction would depend upon the difference between short-dated and long-dated prices, the cost of warehousing the commodity (such as the rental cost of the warehouse and the insurance premium to insure it) and the interest cost of all the capital tied up in the commodity. Theoretically, then, there is a limit to the extent to which a commodity curve can be in contango.

But what if the curve is in backwardation? In this case, traders who are holding inventories may release their inventories in the spot market (i.e., for immediate delivery), selling at the high short-dated prices and buying back the inventory that has sold for future delivery at lower prices. By not having to warehouse the inventories, they would save on warehouse rental, the insurance premium and the interest cost of the capital deployed. Such flows would tend to cause a backwardated curve to move less into backwardation (or more into contango). The problem is that this trade cannot continue indefinitely. Once the available inventories have been sold off, there will be pressure on the curve to return to the original extent of the backwardation. It is easier to amass inventories than it is to keep selling them; once they have gone, they have gone. Hence, there is a limit to the extent of the contango but perhaps not to the extent of the backwardation.

Indeed, it is probably not an unreasonable statement to make to say that near-arbitrage situations are more likely to exist (and persist) in commodity markets than in other asset classes. This can occur for many reasons – for example, the activities of commodity consumers, the actions of index players in the futures markets, the inability to take advantage of the arbitrage opportunity (e.g., is there enough storage or transportation capacity?) and/or the inability to access sufficient finance.

1.2.4 Investment in commodities

The past 5 to 10 years have seen increased interest in commodities as a form of investment. Commodities are now demanded not just for the production of goods, say, but as a potential source of asset returns. Pension funds, insurance companies, sovereign wealth funds and wealthy individuals have been investing increasing amounts of money into commodities partly because they believe that commodity prices are likely to trend higher over time (as the world demands more and more of these scarce resources) and because they are believed to offer portfolio diversification benefits relative to financial assets. Equity-like returns from uncorrelated assets is the lure.

At first, this investment took very simple forms, perhaps investment in a financial product offered by one of the many investment banks involved in commodities. These financial products were, and still are, often based on the average price of a basket of commodities. Later, investors might have moved onto slightly more esoteric products offered by these same banks or they may have invested in one of the many commodity hedge funds that started in the mid to late part of the last decade. Next, they might have invested in the underlying commodities themselves, possibly trading them on a public exchange. Finally, as their experience in commodity products grows, they might have invested in non-listed assets such as commodity-related infrastructure; for example, buying timberland, mines, agricultural land, refiners, etc.

Note that the investment world has long been involved in commodities. Many investors have long held equity in commodity producers and traders, such as the major oil producers (BP, Shell, etc.), the mining companies (Rio Tinto, Xstrata, etc.), commodity traders, electricity producers and the like. For equity investment, however, the main risk tends to be the general level of the world's major stock markets. The prices of commodity-related equities tend to be highly correlated with the prices of equities issued by financials, utilities, producers of consumer goods and services, etc. So, investment in commodity-related equities tends to have returns that are only indirectly related to commodity prices. This may explain why governments often complain about direct investment in commodities while not complaining about investment in commodity companies, although this may be giving governments a little more credit than they perhaps deserve.

The experience of many investors in the commodity sector has not been a rewarding one. The early part of the last decade saw commodity prices rise rapidly, in large part due to the industrial emergence of China and the other BRIC nations. Those investors who were in from the start were able to ride the tremendous growth in prices in the 2000–2007 period. Those investors who became involved from 2007 onwards, encouraged by the gains of the previous years, have lived through a very volatile period for commodity prices. The sought-after diversification benefits of direct commodity investment have not been as evident as predicted since the financial crisis has, for example, hit commodity prices as much as it has hit equity and property prices. A number of high profile commodity hedge funds opened and closed in this latter period as investors first flocked to them and then departed from them, discouraged by the extreme volatility of returns.

Nevertheless, the amounts invested in commodities have continued to rise as more and more asset managers look for ways to increase investment returns in the current low interest-rate environment.

1.2.5 Commodity fundamentals

Producers and consumers of commodities tend to think of commodity prices being determined by demand and supply fundamentals and, on top of this, by investor flows. They tend not to think of investor flows as being "real" supply and demand. One will hear statements such as "copper prices are too high relative to supply/demand fundamentals". Clearly, however, the current copper price is determined solely by supply and demand. That is how prices are determined. Commodity researchers will spend much of their professional lives analysing Chinese commodity imports, Chile's copper output, cheating on production quotas by individual countries within OPEC, maintenance schedules for oil rigs, weather patterns, etc. Then using these supply and demand *fundamentals*, they will forecast prices. In the past, commodity prices were much more stable and researchers tended to have more success in forecasting prices. With the recent increased volatility in prices, however, comes greater propensity for forecasting errors.

The investment community is an easy target to explain those errors. "If it were not for those pesky investors, prices would be lower", might be one refrain from a researcher who has under-predicted prices. Is this fair? Possibly, but perhaps one should ask why these same researchers are not trying to model the impact of investor flows, if investors do indeed have an impact (see below).

One issue with commodity prices is that the demand and supply schedules are inelastic (in an economics sense). Shifts in these schedules will have significant impacts on commodity prices

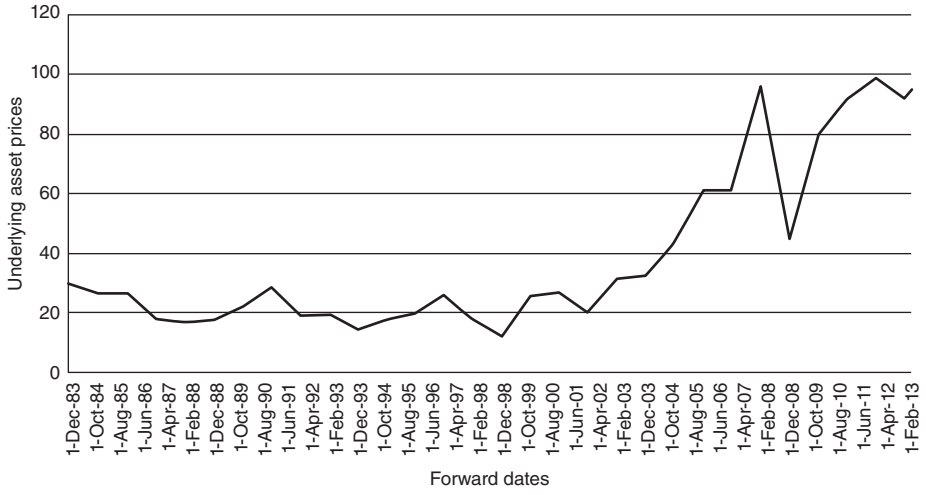

Figure 1.4 Crude Oil WTI

and, unfortunately, these shifts are unpredictable and unobservable except through the prism of price changes. Could it be that unobservable movements in the fundamentals are really to blame for recent volatility of prices? After all, as demand increases with the industrialization of the BRIC economies and the rest of the newly industrialized economies (NIEs), there will be a whole set of demand schedules that researchers know little about. This is especially the case for China where the actual level of inventories is a multiple of the amount recorded in official warehouse statistics.

The truth of the matter is that commodity prices are becoming more and more random in the short term and this is not the fault of the investor. As commodity markets become tighter and tighter, with the growth of the NIEs and the depletion of the more easily accessible commodity supply, one has to accept greater volatility in prices. This does not mean that the fundamentals no longer apply. It just means that forecasting prices will become more and more difficult; as difficult as forecasting FX or equity prices (Figure 1.4).

1.2.6 Super-cycles in commodity prices

A paper published in February 2012 by the *United Nations' Department of Economic and Social Affairs*[14] reviewed the literature around *super-cycles* in commodity prices.

What are super-cycles? They are decades-long periods of significantly above- or below-trend movements in prices. Super-cycles differ from short-term fluctuations that come about from microeconomic factors in that they tend to span "20–70 year complete cycles". These long cycles can see commodity prices varying between 20 and 40% above or below their long-run trend. The authors of the UN report argued that approaching the analysis of commodity prices from a super-cycle framework is "an important innovation over the more traditional analysis

[14] http://www.un.org/esa/desa/papers/2012/wp110_2012.pdf.

of trends and structural breaks (since) it allows us to analyze the gradual change in long-term trends instead of a priori assuming a constant deterministic or stochastic trend".[15]

1.2.7 Future regulation

Although evidence to support the view that direct investment in commodities leads to higher prices is patchy, there is a belief among many that those people who do not directly use commodities in the production process should not be allowed to buy them since these people obviously (sic) push commodity prices higher than they otherwise would be to the detriment of all. As noted earlier, however, many traders, especially those who help producers and consumers to hedge, provide a very valuable service. Hence, there is a begrudging acceptance of such activity although it is still viewed with suspicion by some politicians, regulators, charities and consumer bodies.

Investment in commodities to turn a profit, however, wins few fans outside of the investment business despite there being conflicting evidence as to whether or not investors themselves push up prices. A review by the last Labour government in the UK found that investment flows did not impact commodity prices on average but did increase the overall liquidity of the marketplace while a recent study by the Commodity Futures Trading Commission (CFTC) found that investor flows were the main determinant of the moves in agricultural futures prices in the USA.

Consider the case of commodity futures markets, however. If investors are buying, someone must be selling since futures contracts are a zero-sum game. The buyers always equal the sellers. Are the investors merely hedging producer supply? To be fair, much of the work done in this area tends to be politically motivated and, as such, the results will have a political bias. Interestingly, when prices are rising, the investor tends to get the blame, when prices are falling, however, it is the state of the world economy. Nevertheless, the upshot is that investors in commodity futures will find themselves more and more limited in what they can and cannot do. Some regulators would clearly like to be in a position where only "real" commodity people trade commodities even if liquidity is harmed in the process.

1.3 FIXED INCOME

1.3.1 Introduction

Fixed income products can be split into three broad categories, *rates*, *credit* and *inflation*. For all these, a significant driver of the equilibrium price of the product will be the general level of interest rates, as determined in a large way by the relevant central bank. For rates products (government bonds and interest rate swaps), changes in the general level of interest rates are essentially the only significant driver of equilibrium prices. Credit products, however, are those where there is a second, significant driver of prices, i.e., credit risk – for example, in the case of a bond issued by a highly-geared corporate or by a lowly-rated emerging market sovereign (in the last few years, credit risk has become a major factor also in the sovereign bond market of some Euro Zone countries). By extension, inflation products are those where

[15] As an aside, the authors also note that there is a short-run relationship running from crude oil prices to changes in world output. This finding supports the widely held view that oil price hikes constrain economic growth in the short run. This is in sharp contrast with non-oil prices, which tend to follow world GDP and are thus more demand determined.

the second significant driver of prices is the future development of some inflation index, such as Consumer Price Index (CPI) or Retail Price Index (RPI) – for example, in the case of an inflation-linked bond issued by a sovereign or the case of an inflation swap offered by an investment bank.

Fixed income products may also be thought of as *flow* or *structured*. Flow products are essentially those products that are traded by many market participants and they have simpler cash flow structures. A flow product, for example, would be a government bond, a bond issued by a corporate or an interest rate swap (we discuss swaps in some detail later in the book). Structured products, on the other hand, are essentially those with relatively complex, non-linear cash flows, although, with the passage of time, the distinction between what is flow and what is structured has become a little blurred. Products that were once considered to be structured are now thought of as flow, such as Credit Default Swaps (CDS), Mortgage Backed Securities (MBS) or basic options.

1.3.2 Credit risk

Credit risk is the risk that a borrower will not be able *or willing* to meet its interest or principal repayment obligations. Historically, the risk for many borrowers has been assessed by ratings agencies, such as Standard and Poor's or Moody's. Each of the ratings agencies has a slightly different approach to valuing credit risk, although their qualitative assessments tend to be very similar.

How does one measure credit risk? There are many ways to measure credit risk depending, to some extent, on how the enquiry is framed. In simple terms, however, credit risk can be thought of as the probability of a *credit event* occurring (e.g., a default) multiplied by the loss incurred as a result of that credit event.[16] Credit events include bankruptcy, failure to pay, default or repudiation (a polite term for sovereign default).

As noted above, for all fixed income products, one would expect their prices to be influenced *directly* by the general level of interest rates. However, credit risk (and inflation) will also be impacted *indirectly* by the level of interest rates. As interest rates are reduced, for example, one would expect that the ability of a borrower to meet his or her obligations will improve, i.e., credit risk would decline. On the other hand, if interest rates are increased, one would expect the ability of the same borrower to meet his or her obligations to deteriorate, i.e., credit risk would increase. In a sense, the prices of credit products are influenced both directly and indirectly by the overall level of rates and by factors related to the specific borrower.

1.3.3 The empirical pattern of yield curve moves

Until 2010, the historical record showed that yield curves bull steepened or bear flattened.[17] When yields fell, they tended to fall across the entire yield curve and to be led by moves in short-tenor instruments. Not only that but the absolute moves in short rates were greater than the absolute moves in long rates, hence, the yield curve steepened. Conversely, when

[16] Credit risk (or spread) = probability of default × (1 – recovery rate). What you expect to lose multiplied by the probability that you will lose it.

[17] Since the start of 2010, central banks around the world have held short-term interest rates at very low levels as they have tried to combat the *balance sheet recession*. With short rates held in place, interest-rate volatility that would usually manifest itself along an entire yield curve has been transmitted into the medium to long end of the curve. Consequently, higher yields have tended to be associated with steeper (rather than flatter) curves and vice versa for lower yields. This pattern is likely to remain the case until the central banks refrain from holding short rates at abnormally low levels.

yields rose, they tended to rise across the entire yield curve and to be led, again, by moves in short-tenor instruments. Again, the absolute moves in short rates were greater than the absolute moves in long rates and, hence, the yield curve flattened. At the same time, however, yield curves would see significant changes in their shape. As rates fell, they tended to become more convex (imagine an upside-down bowl). As rates rose, yield curves tended to become less convex and more concave (imagine a normal bowl).

The key driver in most cases was central bank activity. If central banks are in control of anything, it is short-term interest rates. Short rates are raised and lowered in order to control macroeconomic variables. As short rates are changed, the entire yield curve follows but long rates typically move by less; they tend to be more stable.

1.3.4 Modelling interest rate movements

Modelling stock or commodity price movements is relatively easy; assume a random walk, with an element of mean reversion and some capacity for occasional (random) jumps to "explain" market crashes. Modelling interest rates is a lot more complicated, however. In fact, we include this section (and the next) for completeness but most readers will probably prefer to skip it (them) and move straight to the section on foreign exchange.

In *Quantitative Finance*, Paul Wilmott notes, "... there's no reason why interest rates should behave like stock prices, there's no reason why we should use the same model for interest rates as for equities. In fact, such a model would be a very poor one; interest rates do not exhibit the long-term exponential growth seen in the equity markets." A thorough description of the leading interest rate models is beyond the scope of this book but interested readers may want to take a look at Rebonato's classic *Interest-rate Options Models*.

In the jargon, interest rate models either attempt to model the future movements in a liquid, "short rate", such as the Ho & Lee approach, or they attempt to model movements in the entire forward curve,[18] such as the Heath, Jarrow & Morton (HJM) approach. When modelling the path of the short rate, this will be something like a 1-month or 3-month rate rather than the overnight rate since changes in the latter tend not to be a very good guide for other short-term rates. Just as with modelling equity or commodities, the interest rate modeller will need a "drift" parameter, i.e., by how much will interest rates tend to change from period to period, and a "volatility" parameter that measures the random changes in interest rates around the drift. Unlike models for equities and commodities, however, there will also need to be mean reversion to some kind of constant value (rather than to a trend). Finally, if one can model the future path of the spot rate, then one can, of course, create an entire yield curve.

1.3.5 Modelling the risks of default

When modelling default, it is generally assumed that the probability of default is exogenous to the asset under consideration. Such an approach is popular since it is simple to use even if it does not sound very clever. Think about simultaneously tossing three coins every month. If you get three heads, then assume default.

Now think of a corporate bond and a risk-free bond of equal maturity. In the world just described, where the risk of default is exogenous, the incremental yield that the corporate bond

[18] Forward rates are those future short-term interest rates implied by the current yield curve.

needs to earn to compensate for the default risk can be thought of as a fixed spread added to the yield on the risk-free bond.[19]

For example, a 5-year risk-free bond has a yield to maturity of 6% while the yield to maturity of a 5-year risky bond is 7.5%. The extra 1.5% is the compensation for the risk of default. In this example, if we have 100% loss at default, then the implied probability of the company default is approximately 1.5% per year.

A slight refinement to this model is to assume that the probability of default changes with time. For example, one could assume a Poisson distribution so that there is a very small chance of a default initially, rising over time to some maximum value before tailing off indefinitely. This seems to mirror the empirical record. Over the long term, companies either do well or have already gone bust; it is unusual for well-established companies to go into default, they are generally taken over beforehand.

In practice, loss on default is not 100% and there is usually some recovery value and one can assume a recovery value given the historical record. Such data is easily obtained from the ratings agencies. Of course, alternatively, the recovery rate could be modelled but this is probably a digression too far.

Finally, as an aside, we note that there is the possibility of a re-rating, either positive or negative. The ratings agencies publish data on the likelihood of transition from one rating to another and this data could also be incorporated into models to price credit risk.

1.4 FOREIGN EXCHANGE

1.4.1 Introduction

An exchange rate is the ratio of the price of one currency versus another. In most cases, exchange rates between the currencies of the developed economies are freely floating although there are occasional bouts of central bank intervention.

Empirically, the best forecast of tomorrow's exchange rate is today's spot rate. In other words, short-term changes in exchange rates tend to be random. This does not cause too much worry for the FX forecasting profession, however, which makes a great living from making short- and long-term forecasts. As with equities, there is fundamental analysis, technical analysis and quantitative analysis.

There are many fundamental models of exchange rates. Unfortunately, despite the efforts of many academic researchers in particular, such models fail to explain the following real world experiences:

1. Exchange rate movements very often tend to be unrelated to changes in macroeconomic variables over periods exceeding a year.
2. Exchange rates are excessively volatile relative to the underlying economic fundamentals (assuming that we can agree on what these are).[20]

[19] The spread to the risk-free rate is not all credit risk. There are other issues to take into account such as the relative liquidity of the credit bond versus the risk-free bond but, as a means of understanding, it is probably a fairly reasonable approximation.

[20] A 2011 IMF Working Paper argued that inflation and growth rates are the two principal fundamental factors used by foreign exchange forecasters in the formation of exchange rate expectations. The prominence of inflation points to broad acceptance of purchasing power parity. Other oft-mentioned factors, including the current account balance, do not appear to play a common role in the formation of exchange rate expectations. See http://www.imf.org/external/pubs/ft/wp/2011/wp11116.pdf.

Note also that excess returns (i.e., returns over and above the risk-free interest rates) from speculating in currencies tend to be predictable but inexplicable; what academics call the *forward bias puzzle*.

Given the weaknesses inherent in fundamental analysis of exchange rates, it is surprising that it is still popular, although, of course, we all like to have some kind of prop to support our decision-making. When the forecasts support the investment view, they are popular. When they do not support the investment view, they can be ignored.

Technical analysis is also very popular for foreign exchange markets, perhaps because the fundamental models have such a poor forecasting record. Nature abhors a vacuum and it is human nature to try to explain the world through repeated patterns. There are many adherents and some techniques within technical analysis do seem to work for foreign exchange forecasts because of this. If enough market participants expect a certain exchange rate level to hold, then it will very likely hold. The problem is, we tend to remember the successes more than the failures and most testing of the value of technical analysis tends to be undertaken in hindsight.

Surveys of foreign exchange market participants suggest that they generally base their short-term forecasts on recent trends, i.e., they extrapolate recent behaviour. This is why moving averages are so popular as technical analysis tools. Moving averages are essentially extrapolative, i.e., they extrapolate the recent trend. Longer term, however, i.e., 1 to 2 years, market participants tend to focus more on fundamental models in their assessment of value. Nevertheless, foreign exchange trading tends to be short-term trading, given the sensitivity of most players to large mark-to-market swings. As a consequence, while fundamental models may matter, they have little relevance for near-term price developments.

1.4.2 How foreign exchange rates are quoted

Foreign exchange rates are always defined in currency pairs, e.g., EUR/USD. The first currency in the pair is called the *base currency* and the second is called the *term currency*. The foreign exchange rate represents the number of units of the term currency that must be surrendered in order to acquire one unit of the base currency. So, a foreign exchange rate defined as CCYA/CCYB and equal to X means that X units of the CCYB (term currency) can be exchanged for 1 unit of CCYA (base currency). In other words, if the EUR/USD rate is 1.27, it means that 1.27 units of USD can be exchanged for 1 unit of EUR.

When expressing foreign exchange rates against USD, the standard market conventions are as follows:

- When quoting EUR, GBP (Great Britain pound), AUD (Australian dollar) or NZD (New Zealand dollar) against USD, the USD is the term currency. Consequently, the standard market convention is to quote EUR/USD, GBP/USD, AUD/USD or NZD/USD respectively, i.e., number of US dollars per unit of the other currency.
- When quoting the remaining currencies against USD, the USD is the base currency. Consequently, the standard market convention is to quote the number of units of the foreign currency per 1 US dollar, e.g., 1 USD = 0.89 CHF in the case of Swiss franc.
- When quoting EUR crosses, EUR conventionally tends to be the base currency, i.e., one will generally see market prices such as EUR/AUD, EUR/CHF, EUR/JPY (JPY is the Japanese yen).

SUMMARY

In this chapter, we focused on the liquid investment asset classes in which financial derivatives are normally traded, namely equities, fixed income, commodities and foreign exchange. The aim was to give a sense of the general price characteristics of these markets, e.g., how they are quoted and how they move through time. We also aimed to make some observations on how to model underlying asset price behaviour for each asset class since investors must understand how prices behave in their composite parts in order to have a real understanding of the value that can be extracted from multi-asset derivative products.