1

Risk

1.1 What is risk?

It is arguable that risk is the key concept in modern finance. Every transaction can be viewed as the buying or selling of risk. The success of an organization is determined by how much return it can achieve for a given level of risk. Before we can justify these statements, we need to achieve some understanding of what risk is.

In a typical pure mathematical ploy, let us start by trying to understand the absence of risk. A riskless asset is an asset which has a precisely determined future value. Do such assets exist? The fundamental example is that of a government bond. We can buy a government bond for say a £100 today and know that we will receive say £5 a year, (called a coupon payment), until a pre-determined date, when we receive our £100 back. Is this asset truly riskless? There is of course a possibility that the government will renege on its promise to pay. (This is known as defaulting.) But if we pick the right government this possibility is sufficiently remote that we can for practical purposes neglect it. If this seems unreasonable, consider that if the British, American or German government reached such straits, the world's financial system would be in such a mess that there would be precious few banks left to employ financial mathematicians. In fact, the existence of such riskless assets is so fundamental both to financial mathematics and to the modern finance industry, that the fiscal policy of the American and British governments of running budget surpluses, and therefore reducing the number of bonds they have issued, caused great consternation. The reader who is tempted to chuckle at the predicament of the finance industry should consider that financial institutions fund pensions by buying long-maturity government bonds and using the interest coupons to pay the pension. The shortage of long-maturity bonds therefore makes pensions harder to fund and ultimately results in smaller pensions.

We can now define a risky asset to be an asset which is not riskless. That is it is an asset of uncertain future value; risk can be regarded as a synonym for uncertainty. The most basic example of such an asset is a share of a public limited company and

we shall return to this example again and again. However, it is important to realize that almost anything except a riskless government bond is such an asset. For example, we could hold foreign currency and be exposed to the risk that the exchange rate will change against us, or we could buy a flat in London and be exposed to the possibility that there is a property crash, as occurred in the early 1990s.

The sharp reader will have noted that the definition in the paragraph is not quite right, in that an investor would not actually care about the riskiness of an asset if the worst possible future value of the asset was greater than today's value.

However, we have to be slightly careful about what we mean by value here. Unless there is no inflation, £1 a year from now will buy less than £1 today. This means that £1 a year from now is effectively worth less than £1 today. In addition, even in a non-inflationary world, most people prefer jam today to jam tomorrow and so would not be happy to receive the same amount of money back in a year with no compensation. A better view of riskiness is that the asset can return less than the same amount invested in a riskless government bond for the same period. A good example of such an asset is the premium bond. In the United Kingdom, one can buy a government bond, called a premium bond, redeemable at any time which pays no coupon but instead the holder gets a free entry in a prize draw paying up to a million pounds a month. This seems too good to be true at first, but the issue is, of course, that the bond is not very different from investing some money and using the interest to buy lottery tickets. That said, the expected winnings for the amount of interest foregone is much better for premium bonds than for lottery tickets. The investor is effectively buying risk.

1.2 Market efficiency

Before we can understand why risk is so important we first have to understand the concept of market efficiency, which underlies most of financial mathematics and modern economics. This concept roughly states that in a free market, all available information about an asset is already included in its price. Therefore there is no such thing as a good buy – the only value an asset has is its market value and it is meaningless to attempt to think otherwise.

Is this hypothesis correct? To see that it cannot be wholly so, consider the apocryphal story of the two economists who see a ten dollar bill lying in the gutter. The first one goes to pick it up but the second one tells him not to be silly as if it were real, someone else would have already picked it up. However, the true moral of this story is that market efficiency only works if someone does not believe in it – the first person who does not believe in market efficiency picks up the ten dollar bill and thereafter it is no longer lying in the gutter. Warren Buffett is the most famous example of a non-believer who has very effectively made a lot of money through

his disbelief. He has largely done so by buying shares in companies he believes are undervalued by the market. Indeed, until Bill Gates overtook him he was the richest man in the world, and he made his money by beating the market.

So although market efficiency is not wholly correct, there are enough people attempting to be the next Buffett, for it to be sufficiently correct that we can work under the assumption that it is true. What does this mean for us? Well, the first thing it means is that it is pointless to for us to try to predict the future price of a share by looking at graphs of its past prices. All this information is already encoded in the share price. This is sometimes called the Markov property, and is also called the *weak efficiency of markets* as it's a consequence of the strong form mentioned above.

It is interesting to note that the modern white-collar crime of insider trading is really based on the principle of market efficiency. Insider trading is trading based on knowledge which is not publicly available and therefore not included in the share price. For example an employee of a company might know that the company was about to announce unexpectedly large losses or profits which would move the share price in an obvious direction, and take advantage in advance. The perception of this as a crime rather than a natural action is fairly recent, and is based on the ubiquity of the concept of market efficiency.

Given that all assets are correctly priced by the market, how can we distinguish one from another? Part of the information the market has about an asset is its riskiness. Thus the riskiness is already included in the price, and since it will reduce the price, the value of the asset without taking into account its riskiness must be higher than that of a less risky asset. This means that in a year from now we can expect the risky asset to be worth more than the less risky one. So increased riskiness means greater returns, but only on average – it also means a greater chance of losing money. From this point of view, an asset's price reflects the value it is likely to have in the future reduced by a factor depending upon its riskiness.

To illustrate these ideas, let us consider a simple game. Suppose we toss a coin, if it comes up heads I give you £3, if it comes up tails you give me £1. Unless beset by moral qualms, you would consider this game a very good deal and play it – your expected winnings would be

$$\frac{1}{2}3 - \frac{1}{2}1 = 1,$$

and your maximum losses would only be 1. Suppose we play a slightly different game, I pay you £13 on heads, you pay me £11 on tails. Your expected winnings are still £1 but are you still so keen to play? If not why not? If you are still keen, let's take the payment on heads to be £103, and on tails to be £101. At some point, when the stakes become high enough you will stop regarding the game as a good deal. The point where you stop depends upon personal risk preferences; the stopping

point is where the expected gains stop outweighing your aversion to the possibility of losing money.

Now suppose the game is changed a little again. The sum you lose is paid to me today and we toss a coin a year from now. If the coin comes up heads, I return your money to you and pay you my losses, otherwise I keep your money. What has changed? During the year in between, I have put the money on deposit with a bank and earned some interest. If you were not playing the game you could have done so also. The amount of return you would want from the risky game would increase to express the interest foregone. And since you could have made money from the interest payment without taking any risks, you will demand that the expected winnings be greater than the amount of interest you could have earned.

The moral is that there is no such thing as a guaranteed high return. The reader would be well-advised to remember this the next time he sees a guaranteed high return in a newspaper or Internet advertisement.

Let's return to the concept of weak market efficiency. This says that all the past movements of an asset's market price is already expressed by today's price. At this point, the prospects of a financial mathematician could be regarded as being pretty bleak. Why? This tells us that trying to predict the future price from past data is a waste of time – there is no periodicity nor trends to be read. The only mathematical information is today's price which tells us very little. In that case, why is financial mathematics a burgeoning field? The job of a financial mathematician is not to predict prices but instead to relate the movements of price in one asset to that of another. These price movements are viewed as being driven by information arriving in the market and since that information is by definition unknown until it arrives, we can view it as being random.

The key point in mathematical finance is to use market instruments which are affected by the same information in such a way as to cancel out randomness. This process is called hedging. The objective of mathematical finance is to understand how to do this and to understand the consequences.

1.3 The most important assets

We have been discussing an asset rather vaguely so let's look at the basic assets in finance from the point of view of risk.

1.3.1 Bonds

The simplest asset, already mentioned, is a government bond issued by a reliable government. Typically, the government issues a bond of, say, thirty years in length, which pays every year a sum called the *coupon* and gives the investor his original

investment back at the end of the thirty years. The original investment is called the *principal*. The day the investor gets the principal back is called the *date of maturity*, when the bond is said to *mature*. In the meantime, each year the investor receives interest payments to compensate him for the fact the government has his money. From a mathematical point of view however the coupons just confuse things, so mathematicians typically study zero-coupon bonds although they are in fact rather rare. A *zero-coupon bond* is a bond which pays no interest but instead just returns the investor his investment on the date of maturity. Why would anyone buy such a bond? Suppose the principal is one dollar. The point is that the investor does not pay a dollar for the bond, but instead pays a smaller amount which if it was invested in interest-paying bonds would give a dollar in total (including the compound interest) at maturity.

The interesting thing about the riskless bond is that there is some risk in it. Not in the possibility of default, but instead in the possibility that interest rates might change in the meantime. At the time of writing, interest rates are quite low by historical standards so an investor might be wary of buying a bond with a long time to maturity: he has locked in today's interest rate and if rates go up he loses out. Of course, if rates go down he gains. There are well-established markets in the major government bonds so our investor need not hold his bond until maturity instead he can just sell it in the market. But what price will he get? If interest rates have gone up, the market price will be less than he paid for it as other investors will want the fixed sum on maturity to reflect how much money they could have got it by investing the money in a newly issued riskless bond. Thus as well as the interest rate reflected by the coupon payment, the price of a bond reflects today's interest rates and indeed reflects today's expectations of future interest rates. The effective interest rate implied by the market price is called the *yield* of the bond and this can be very different from the coupon. Since the yield of a bond reflects expectations of future rates, bonds of different maturities can have different yields implied by the market.

As the date of maturity of a bond approaches, there is less and less uncertainty left. The principal of the bond is known and will be paid on the date of maturity, and the interest rate is unlikely to move much in a short period of time so there is less and less uncertainty as the maturity date gets closer and closer. A bond of longer maturity will be exposed to more uncertainty, that is risk, than one of short maturity. We can therefore expect long-dated bonds to have higher yields to compensate investors for this additional risk. This is generally, but not always, true. Indeed at the time of writing, gilts (gilts is the financial jargon term for UK government bonds) display a hump: the yields first rise and then decrease. See Figure 1.1. Recall that the yield also reflects expectations of future rates so one explanation is that the market expects UK interest rates to rise in the short term, but to decrease in the long term. An alternative and probably more correct explanation is that there is

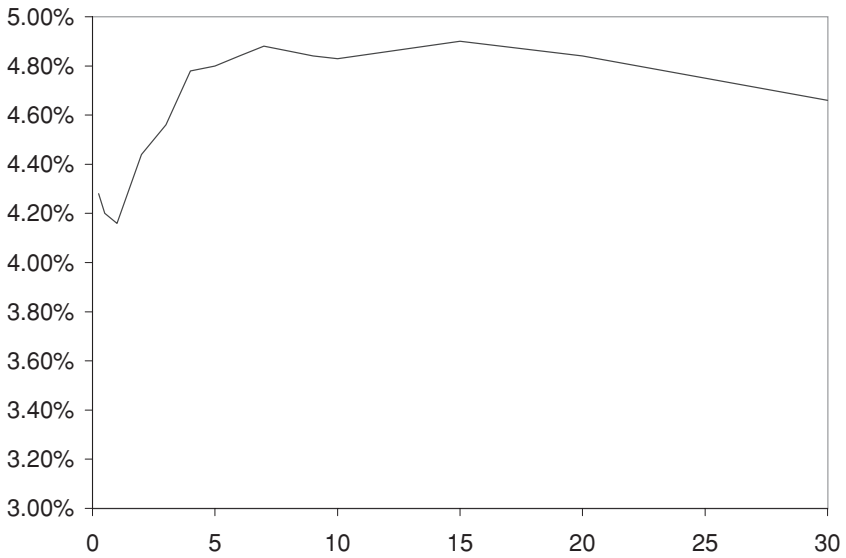

Fig. 1.1. The yields on UK government bonds (gilts) in October 2001 as a function of years.

currently a shortage of long-dated bonds which drives their prices up and hence their yields down.

1.3.2 Stocks and shares

Probably the most ubiquitous sort of traded asset is the share or stock in a company, (share and stock are equivalent terms.) While the reader is almost certainly familiar to some extent with shares, it is worth examining precisely what that term means. The holder of a share of a company owns a fraction of that company. Companies traded on the stock exchange typically have plc after their name reflecting the fact that they are public limited companies. Public just means that anyone can buy shares in them. Limited means that they are of limited liability; the owners of such a company have no liability for its debts if it goes bankrupt. The importance of this fact should not be underestimated. In the author's opinion the existence of limited liability companies is the foundation of modern capitalism. Why? Because it reduces the riskiness of investing in a company by capping the total losses to be the amount invested. If the company is sued for a billion dollars and goes bankrupt, you the shareholder are not liable. If you could be liable would you still buy shares?

To emphasize this point, consider the 'Names' at Lloyds. The Lloyds insurance market worked in the following way. A person would agree to set aside at least

£100, 000 against claims by people insured in the Lloyds' market, in return for which they would receive a lucrative stream of insurance premiums. The £100, 000 could also be happily invested in other assets allowing the Name to do particularly well. The Name got a very high rate of return but high returns equal high risk. Indeed, the particular feature of the Lloyds market which made investing very different from buying shares in an insurance company was that a Name had unlimited liability. Not only could he lose the £100, 000, he could lose everything including his house and his shirt. And in a particularly bad year that is precisely what happened. Perhaps the most interesting aspect of this story is how surprised many Names were; somehow they had failed to appreciate the connection between risk and return.

We have seen that the holder of a share owns part of a company and the sole risk he bears is that the value may drop to zero. The share may bring the investor money in two ways. The first is simply that the value of the share may go up, the other is that the company will generally pay dividends that is payments, often annual, to shareholders dispensing the profits of the company. As one might expect, the total return to shareholders on average is much greater than the rate received on depositing the money in a riskless bond or a bank account to compensate for the danger that the company will go under or just not do very well.

1.3.3 The corporate bond

Another asset commonly traded in the markets is the corporate bond. This lies somewhere between a share and a government bond. A company wishing a loan issues bonds in the market paying some coupon in interest. The coupon is generally higher than that of a riskless bond. The investor's risk is that the company may default on its payments as it has gone bankrupt in the meantime. However, bondholders have more claim on the company's assets than shareholders so the riskiness is reduced. The main disadvantage of a bond is that if the company share price soars the bondholder does not gain at all. Thus both the returns and the riskiness of bonds are lower than those of shares. In order to entice investors to buy bonds, companies sometimes issue *convertible bonds*, that is, bonds that can be converted into shares if the investor so chooses. This allows the investor the upsides of both bonds and shares – of course, typically the coupon or yield on such a bond would generally be less than that of an ordinary bond.

1.3.4 Positivity

All the assets discussed so far have one similarity, they all carry rights without liabilities. This has an important consequence for the mathematician: their values

are always positive. In an extreme case, the value could become zero on bankruptcy but the important point is that they will never be negative. There are plenty of market instruments that do not share this property, and we shall encounter many of them in this book.

1.3.5 The risk paradigm outside the markets

It is important to realize that the concept of risk is inherent in all investment decisions not just those of what to purchase in the financial markets. For example, suppose a company wishes to invest in a new plant to produce a new product. It will estimate the amount of return it will receive on the invested capital. What level of return should it demand? Well to answer that question it needs to assess the riskiness of the project and it should demand the same return as a market instrument of comparable riskiness. Otherwise, it would do better just to buy the relevant market instrument and forget about the plant.

1.4 Risk diversification and hedging

We have treated all risks as being equal but some are better than others. In particular, some risks can be effectively eliminated by judicious trading. There are two main ways to proceed: hedging and diversification.

Consider a contract that pays £100 if a coin flip is heads and zero otherwise. From what we have said so far, we would expect this to trade for less than £50 depending on the risk aversion of investors. We can also consider the complementary contract that pays £100 if the *same* coin flip is tails. We also expect this to be worth less than £50. However, if both these contracts are trading in the market, we can buy both and be guaranteed £100 whatever happens. The risk has disappeared and since we do not expect to be able to make riskless profits, we conclude that the original contracts were worth £50 after all. We have a paradox, after arguing at great length that such a contract would have to be worth less than £50, we conclude that we were wrong and that it is worth £50. What has changed? The addition of a second contract has removed the risk. This process is called *hedging*. As long as only one of the contracts is tradable, its value is less than £50 but as soon as both are, the risk stops being unhedgeable and the risk premium disappears.

A related concept is that of *diversification*. Suppose we can bet a very small amount on a very large number of independent coin tosses. We therefore divide our portfolio into N bets each paying $100/N$ on heads. Our average pay-out will still be £50 but as N gets larger the variance gets smaller and smaller. The risk has therefore been effectively eliminated and we cannot expect the individual contracts to trade for less than 50 pounds.

The lesson of this example is that the market will only compensate investors for taking risks that are not diversifiable. Undiversifiable risk is known as *systemic* risk. The job of an investor is therefore to achieve the maximum amount of systemic risk for a given level of riskiness by diversifying his portfolio.

One paradoxical side-effect of investors' need to diversify is the increasing demand for products that express *purity of risk*. An investor who feels he is overexposed to one sort of risk will want to buy other sorts to offset it. Derivatives are one way to manipulate an investor's risk profile.

1.5 The use of options

Given that risk is inherent to all financial decision making, a bank or company will want to manage its risk carefully. In particular, it may want to buy certain sorts of financial instruments to increase or decrease a certain type of risk. This is where options and related products enter the picture. These products can be used to reduce risk or to increase it. Whether they are a good or a bad thing depends purely on the way they are used.

To give an example before we start making definitions, consider an American company which exports to Japan. The Japanese importer pays the company in yen but the company prefers dollars. The company estimates that it will receive between one and two billion yen next year, which it will need to exchange for dollars. The variability of the exchange rate between dollar and yen means that the company is exposed to some extra risk which it would prefer to avoid. One solution is to enter a forward contract. This is a contract to exchange a fixed amount of yen for dollars at a fixed future date at today's exchange rate (modified slightly to take account of interest rates). The company's problem is that it does not know precisely how many yen it will wish to sell so it cannot remove (or hedge) all the risk. One solution would be to enter a forward contract for a billion yen, since it is sure it will need to sell at least that much and treat the rest separately. To deal with the rest of the yen which is a variable amount between zero and one billion, the company could buy an option.

What is an option? An option is typically an instrument which gives the holder the right to buy or sell a quantity of some fixed asset during a specified period of time at a price fixed today. The important point is that unlike a forward contract there is no obligation to buy or sell. The option carries rights but not obligations, and therefore will have always have a positive value before the time of expiry.

In this case, the company could therefore buy an option to sell a billion yen at today's price a year from now. Buying the option would, of course, cost the company a fee but it would cap the amount of losses it might make if the exchange rate moved in the wrong direction, thus reducing the company's risk. The important

thing to realize is that the company will not decide whether to use the option on the basis of how many yen it needs to convert but instead on the basis of whether the market price is higher or lower than the price which the option guarantees. Any excess yen can always be sold in the market.

The financial derivatives market is full of jargon. An option to buy is called a *call* option. An option to sell is called a *put* option. (The easiest way to remember which is which is that C is close to B for buy in the alphabet.) Using an option is called *exercising* it. The price which is guaranteed by the option is called the *strike* price and the option is said to have been *struck* at that price. There are many different sorts of rules for when the option can be exercised. The simplest sort is the *European* option which can be exercised on one specified date in the future. An *American* option can be exercised on any day before a specified date in the future. Note that since an American option carries all the same rights as a European option and more on top, it will always be worth at least at much as a European option and generally more. One thing we will demonstrate later in the book is the surprising result that under certain quite natural circumstances American and European call options have the same value. The options we have mentioned so far are very much the beginning of the list, and the list of possible options goes on and on, growing every day. The options we have mentioned above are generally called *vanilla* options to express the fact that they are standardized and less interesting than *exotic* options. Banks sometimes have different teams of mathematicians for the pricing of vanilla options and of exotics. Many exotic options are not really options in the sense that the holder does not get a choice but instead receives a payoff, which is possibly negative, dependent on the behaviour of some asset. This asset is generally called the *underlying*. The generic term for all instruments whose value is defined in terms of the behaviour of some other asset is *derivatives*. This name expresses the idea that their value is derivative of the behaviour of the price of the underlying.

Ultimately, an option is a powerful instrument to change one's exposure to risk. The big difference between buying an option on a stock and buying the stock is that if the stock price moves the wrong way the option will be valueless whereas the stock will not, but on the other hand if the stock moves the right way, the option holder will have made much more money for the amount of money spent than the stock holder. One attraction of speculating using options is that the maximum downside is the loss of the initial premium, whereas the up-side is unlimited. Of course, it is important to appreciate that for the option seller the position is reversed.

Another way an investor could use an option to reduce risk is as follows. Suppose he holds a large number of stocks which he knows he might need to sell a year from now, but he is worried about a crash in the meantime. Whilst he knows that he can