Uses And Abuses Of Discount Rates A Primer For The Wary

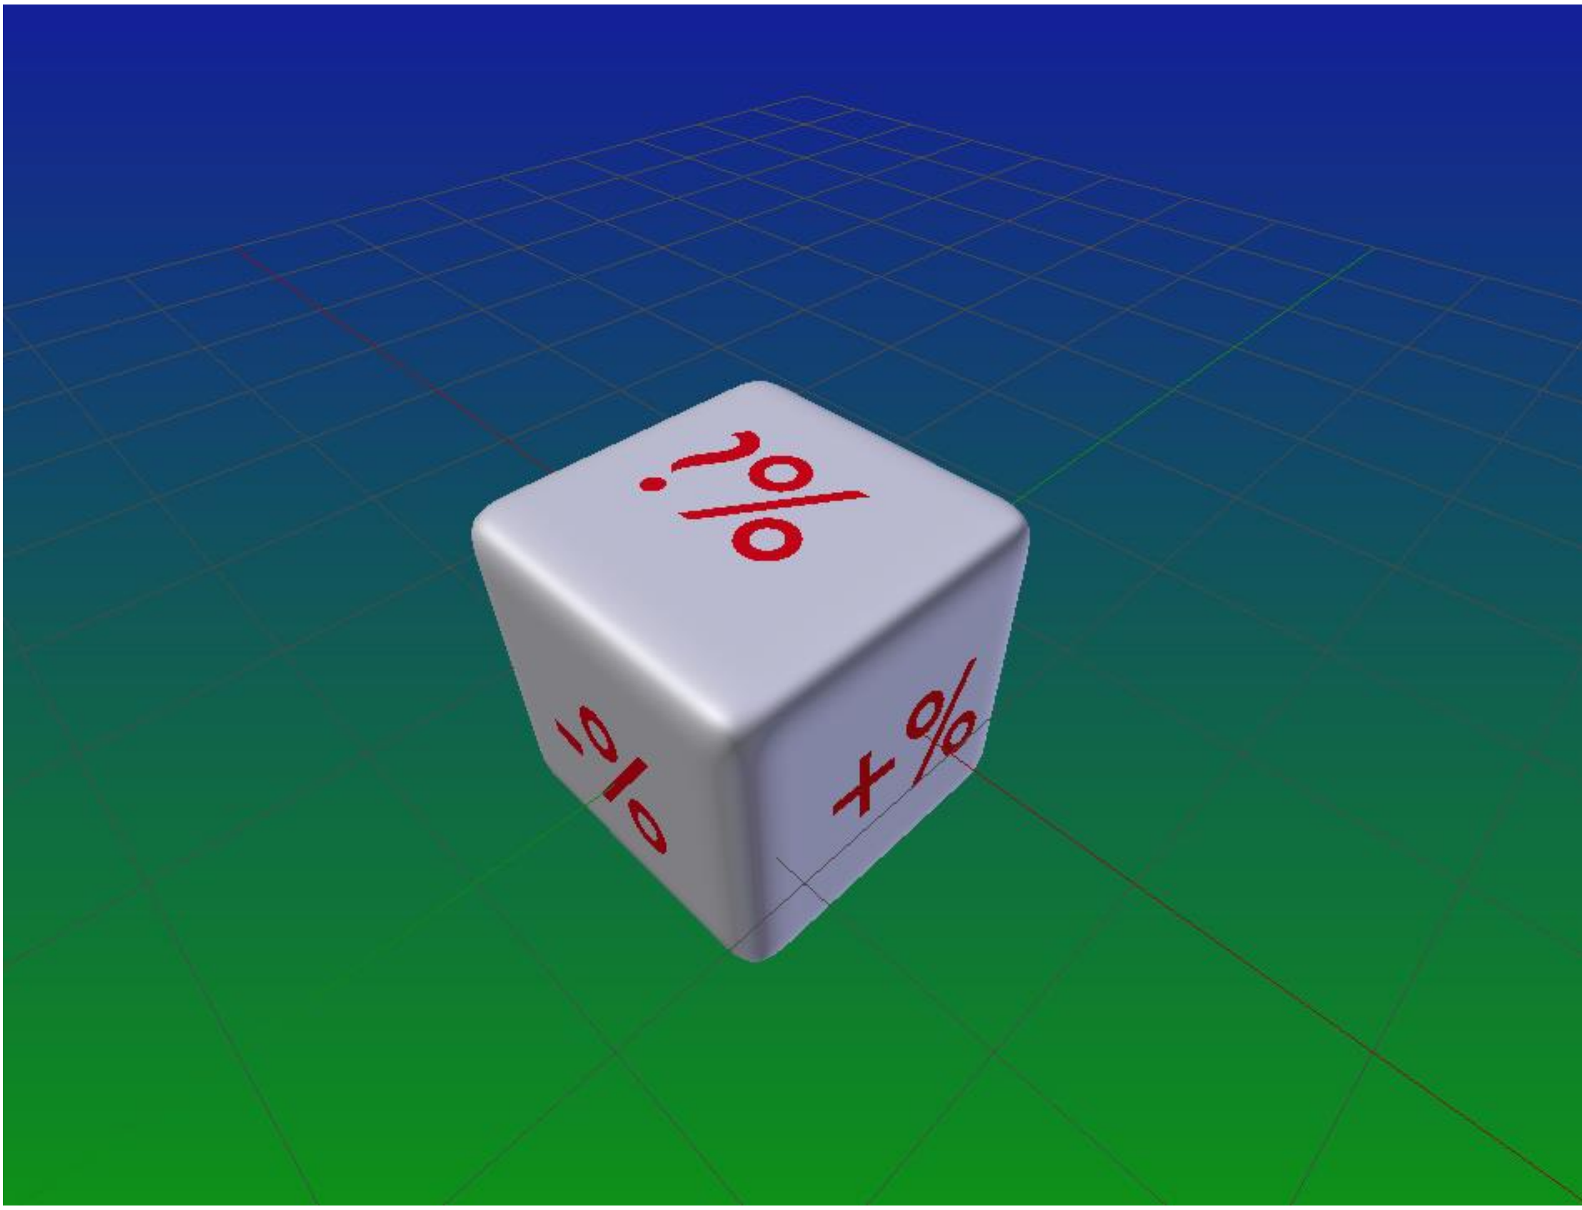

September 2015

A report prepared by Dr Nick Goddard for Long Finance

Uses and Abuses of Discount Rates: A Primer for the Wary

Dr Nick Goddard

2015

Edited by Michael Mainelli and Chiara von Gunten

Contents

Preface

Long Finance aims to *“improve society’s understanding and use of finance over the long-term”*. We do not consider ‘long’ as a synonym for ‘good’. However, we do believe society should improve its ability to finance solutions to long-term social problems.

Discounted cash flow (DCF) and net present value (NPV) analyses have long been part of the financial analyst’s toolbox. In order to use both tools we need to decide on a discount rate and use that discount rate in some exponential equations. Exponential equations lead in turn to infinities and are thus inherently problematic in a constrained world. The use of discount rate tools leads to conflict in values over time.

In 2006 Charles Young posited on the Environmental Economics blog a *reductio ad absurdum* argument about discount rates in general; “an extra glass of wine for Alexander the Great matters more than all today’s capital stock”. In The 2006 Stern Review, *The Economics of Climate Change*, Sir Nicholas Stern used a discount rate significantly below those found in typical financial analyses. The importance of discount rates warranted a technical annex – “Ethical Frameworks and Intertemporal Equity”. When a discount rate is altered for one risk, then discount rate analysis can be altered for any risk. Decision-making moves from the analytical to the political.

The same discount rate setting arguments arise in other inter-generational transfers such as pensions, health, public infrastructure, taxation, biodiversity, or cultural heritage. Under discount rate assumptions, it is easy to argue that future generations will be richer than us and can pay for more in their future, so we can spend wantonly now. Or that something horribly expensive is really quite cheap if you make a very small change in the discount rate. Nick Goddard does us all a favour with his short, easy-to-read primer. By making discount rates easier to understand he shows the utility, and potential for abuse, of these ancient tools over the long term.

Professor Michael Mainelli
Executive Chairman – Z/Yen Group Limited

Preamble

"*Talk to him of Jacob's ladder, and he would ask the number of the steps.*"

Douglas Jerrold, ***A Matter-of-Fact Man*** (1859)

Back in 1996, a financially naïve scientist was parachuted into a group of investment bankers to provide some technology insights. In return, he was inducted into the mysteries of discounted cash flow (DCF) and net present value (NPV) analyses. Was this, he wondered, an unavoidably complex piece of financial wizardry which provided invaluable insights into financial decision-making? Or was it unnecessarily complex obfuscation, useful mainly for conferring an aura of technical rigour to whatever the banker's gut-feel was telling them? As a physicist he at least had the advantage of not being bamboozled by the maths. So he watched and listened, and when he left banking nearly a decade later, he had concluded that it was, in fact, a bit of both. This essay expands on that conclusion by trying to separate the uses from the abuses.

1. Discount rates in layman's terms

1.1. The basics – saving for a car or selling some saplings

It's worrying when any profession claims that its analysis is too esoteric to be explained in layman's terms. Medicine is highly complex, but that does not prevent a good doctor from explaining to a patient the key features of an illness and laying out clearly the various treatment options with their attendant risks. The first potential abuse of discount rates is therefore a generic one: the temptation to say "*don't worry your little head about why this is right – just trust me because I understand things about financial engineering that you don't*". The power to make pronouncements without having to justify them is as corrupting as any other kind of power.

So how do we explain discounted cash flow (DCF) and net present value (NPV) using discount rates in layman's terms? At its simplest level, DCF and NPV seek to answer the following questions:

- If I have a bill to pay at some point in the future, how much money should I have in the bank right now to cover that liability?
- If there is an entitlement to receive a sum of money at some point in the future, what should I be prepared to buy or sell that entitlement for today? What is it worth right now?

Taking a very simple example; Duncan knows that five years from now he will need to buy a new car which will cost him £10,000. Let us assume that he can obtain 3% interest on his savings if he ties them up for this period. On this basis, the NPV of his liability is £8,626, because if he puts this sum into a bank account with 3% compound interest it will grow to £10,000 over five years, and will be enough to pay for the car when the time comes.

The above example considers a future liability. What about a future asset? Suppose that Rupert owns some young trees which will be cut down in 20 years' time and sold for timber, and that a timber merchant has entered into a contract to buy that timber from him for £10,000 when the time comes (£10,000 being the price that would be paid in today's market for a similar quantity of cut timber). If he believes that he could get 4% compound interest by tying up his money for 20 years, then the NPV of his asset is £4,564 because if he put this sum into a bank account with 4% compound interest it would grow to £10,000 over 20 years.

It is probably fairly obvious that the mathematics used to calculate the NPVs in the above examples is based on expressing the interest rate as a multiplying factor for the annual increase in the savings (i.e. expressing 4% as 1.04) and then raising it to the power (or exponent) of the savings period. After twenty years at 4% interest, a sum of money will increase by a factor of 1.04 to the power of twenty (1.04^{20}), which is equal to a factor of 2.19, or an increase of 119%. It should be noted that for this to be true, the interest has to be compounded (i.e. added to the account) rather than paid out each year. There is nothing particularly scary about exponential mathematics other than the fact that these curves tilt upwards at a steeper rate every year. As can be seen from Graph 1, the higher the interest rate, the more immediate the effect of the steepening.

Graph 1: Comparing 4%, 8%, and 12% compound growth rates

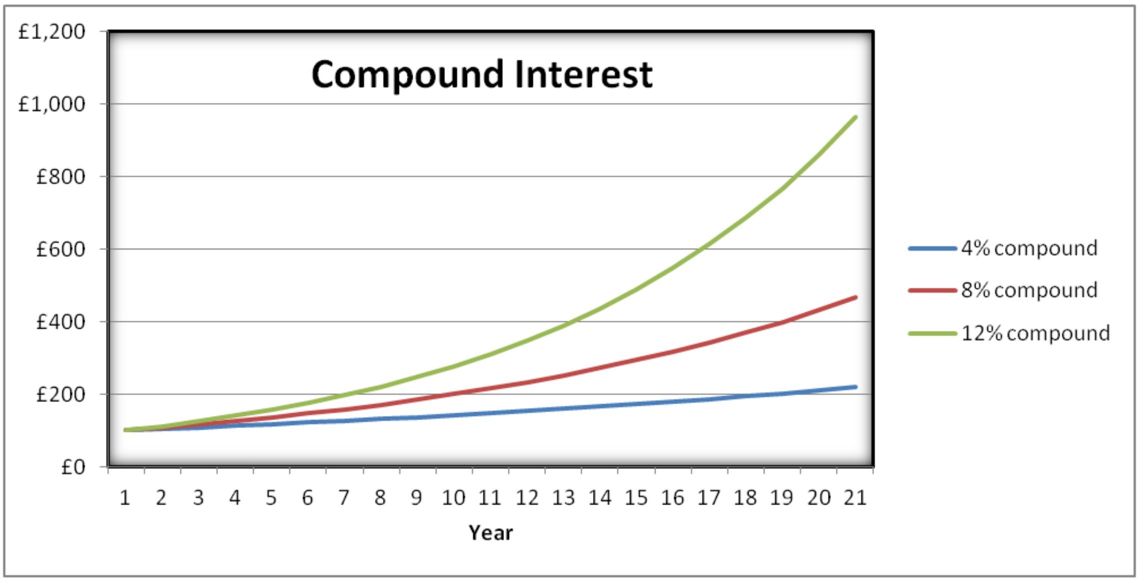

This steepening of the curves is good if you are receiving interest, but not so good if you are paying it (for example, to a payday lender). Indeed, some may argue that payday lenders have systematically abused people's failure to understand the impact of compound interest – effectively an abuse of discount rates. Payday lenders have very few supporters, but it is only fair to point out that high street banks can, on occasion, be just as greedy. The fees they charge on small, non-authorised overdrafts can often amount to several tens of pounds when the overdraft is just single digit pounds for a few days. Expressed as an interest rate on an Annual Percentage Rate (APR) basis, this would be every bit as usurious as payday lending, but the charges are treated as administrative fees rather than interest. While such fees are far from popular, they seem to attract much less odium than the payday lenders with their sky-high headline interest rates. This serves to illustrate that some abuses of discount rates may be semantic rather than mathematical.

In the case of the car, Duncan probably does not have much option other than to accept the liability of £8,626 onto his personal balance sheet, as he cannot function without one. In the case of the trees, the figure of £4,564 represents an asset which is already on Rupert's balance sheet – he is just another an aristocratic landowner counting his wealth. However, it could also represent the fair price which I would be prepared to offer to buy those trees off him, because I have £4,564 of cash in my bank account and my reasonable expectation for a return on my savings is 4% per year over 20 years. These days, I may even feel that my money is safer tied up in trees than being looked after by a bank. Finally, £4,564 would be a reasonable price for Rupert to ask if he sold the trees to someone else with the same expectation as me for the rate of return on their savings over 20 years.

Someone greedier than me, who wanted a 5% return on their savings, would not be prepared to pay Rupert £4,564 for his trees. Instead, they would want to pay him no more than £3,769 because £3,769 would grow into £10,000 over 20 years at a 5% compound interest rate. In fact, they may not necessarily be greedier than me: unlike me, they may still believe that money is safer kept in a bank because trees can fall over or suffer from disease (clearly things which could never happen to the global financial system). So they want a premium rate of interest to compensate them for the higher risk of withdrawing their savings to buy Rupert's trees.

Given that these two examples – saving for a car and selling some saplings – incorporate all the key assumptions of DCF analysis it may seem amazing that such a straightforward tool could be abused. The devil is, of course, in the Byzantine detail, of which the remainder of this essay merely scratches the surface.

1.2. Discount rates reflecting target rates of future return

At the end of the previous section we introduced the idea of using discount rates to evaluate rates of return which someone wants to achieve rather than those which may typically be available from savings accounts or bonds. Such discount rates are subjective and cannot simply be looked up in the Financial Times. Suppose that I set a 'hurdle rate' for the 'Internal Rate of Return' (IRR) on my capital of 10%, simply because that would be nice and I think that I ought to be clever or aggressive enough to achieve it. I would now only be prepared to pay Rupert £1,486 for the trees, having calculated that this will grow into £10,000 over 20 years at 10% compound interest.

The business executive concerned with IRR may not currently own any trees, and right now may not even know anyone who has trees for sale. He or she is probably trying to work out whether forestry is, in principle, a good business for someone who wants to make a 10% IRR on their capital. IRR is a much more rigorous investment comparison approach than Rupert's. Rupert has trees growing on his family estates and simply wants to estimate his wealth. The executive may write down a whole list of costs which will be incurred at different dates in the future – starting with the immediate cost of buying and planting the saplings, including an annual rental payment for the land on which they are planted, and finally putting in a figure for the cost of cutting and transporting the timber. The objective is to see if the NPV of selling the timber is sufficient to offset the NPV of the associated costs.

Most practical applications of DCF analysis involve large numbers of individual credit and debit items put together to determine if a project will be able to deliver a desired IRR. This is discussed in more detail in Chapter 2, but for a simple illustration let's revisit the future liability for buying a new car which was calculated above as £8,626. Suppose Duncan is very lucky and has an aunt who gives him £1,000 every birthday, and that he intends to chip this in to his car purchase fund. As today happens to be his birthday, he will immediately put £1,000 into his bank account with its 3% interest rate. In a year's time, he will put in another £1,000, but this is only the same as putting in £971 today (because by next year £971 would have grown to £1,000). In two years' time he will put in another £1,000, but this is only equivalent to putting in £943 today. The NPV of his 'buying a car' liability can therefore be calculated as follows:

NPV of purchase cost of car:	£8,626
Less NPV of £1000 put into the bank today	-£1,000
Less NPV of £1000 put in the bank one year from now	-£971
Less NPV of £1000 put in the bank two years from now	-£943
Less NPV of £1000 put in the bank three years from now	-£915
Less NPV of £1000 put in the bank four years from now	-£888
Less NPV of £1000 put in the bank five year from now	-£863
Total NPV of Duncan's liability	**£3,046**

So it seems that if Duncan were to put £3,046 in the bank today, along with his aunt's birthday money every year, he would indeed be able to buy a new car for £10,000 five years from now. To double-check this, let's look at how the cash would accumulate in Duncan's bank account:

Table 1: NPV of contributions made towards saving for a car

Item	Today	End Yr 1	End Yr 2	End Yr 3	End Yr 4	End Yr 5
Duncan's contribution	£3,046	£3,137	£3,232	£3,328	£3,428	£3,531
Today's birthday	£1,000	£1,030	£1,061	£1,093	£1,126	£1,159
Future birthday 1		£1,000	£1,030	£1,061	£1,093	£1,126
Future birthday 2			£1,000	£1,030	£1,061	£1,093
Future birthday 3				£1,000	£1,030	£1,061
Future birthday 4					£1,000	£1,030
Future birthday 5						£1,000
Total savings	**£4,046**	**£5,167**	**£6,322**	**£7,512**	**£8,737**	**£10,000**

Typically, NPV calculations are used to model situations where the income and the expenditure for an enterprise are quite widely separated in time. In the case of the car, the income appears first and the expenditure in a lump sum right at the end. For the forestry, the expenditure is incurred up front and then the income is a lump sum at the end. Both situations occur in business, but it is more common to make the investment first and reap the rewards later. The converse situation is, however, also encountered. For example, someone who owns an oil well reaching the end its life will receive income for the first few years and then have decommissioning costs at the end. NPV is therefore essentially a tool for time-shifting unbalanced income and expenditure items.

Interlude – "The Rule of 72"

This essay, and many others, seem to want to make readers reach for their calculators. One of the more fun mathematical tricks is the "look Ma, no calculator" Rule of 72. This mathematical device can be used for quickly estimating a discount rate's 'doubling period'. To estimate the number of periods required to double an original investment, divide the number 72 by the expected growth rate, expressed as a percentage.

A quick example - taking a £100 investment, compounded at a rate of 9% per annum, the rule of 72 gives **72/9 = 8 years** required for the investment to grow to £200 [a more exact answer using logarithms is ln(2)/ln(1+.09) = 8.043 years].

Another example – taken from Luca Pacioli's ***Summa de Arithmetica*** (Venice, 1494. Fol. 181, n. 44) – "When the interest is 6 percent per year, I say that one divides 72 by 6; 12 results, and in 12 years the capital will be doubled." [a more exact answer is ln(2)/ln(1+.06) = 11.896 years]

For those who do reach for their calculator or spreadsheet, Excel has a number of DCF and NPV functions such as 'future value' (FV), 'repayment' (PMT), 'present value' (PV), 'net present value' (NPV) or 'discount rate' (DISC), all worth exploring.

2. Setting the right discount rate

2.1. Taking tax into account

In the simplified examples presented in Chapter 1, the discount rate is deliberately set at an interest rate which might be received (or a targeted rate of return) for money on deposit. It may sound trivial to decide on what this rate should be, but it is not. One thing which has to be considered is tax. For a private individual there may well be tax levied on any interest received and the tax rate is likely to depend on several factors. Taking the example of Rupert's trees, we originally arrived at a NPV of £4,564 because if this sum were put into a bank account with 4% compound interest it would grow to £10,000 over 20 years. But suppose there is 40% tax levied on the interest received in that bank account. This reduces the net interest rate to 2.4% and at this rate the £4,564 will only grow to £7,334 after 20 years. If someone wanted to have £10,000 in such a bank account 20 years from now, then they would need to have £6,223 in the account today. Incorporating taxation arguably changes the valuation of Rupert's trees from £4,564 to £6,223.

However, the correction in value calculated above might not be the full story. It may be that when Rupert harvests the trees, the proceeds from selling them is treated as income and taxed at his marginal rate of 40%, so he only receives a net £6,000 from the timber merchant. He would need to put only £3,734 in his '4% gross & 2.4% net' bank account to arrive at £6,000 in 20 years' time. But supposing all Rupert's income is from various sorts of farming and he also raises calves born to his herd each year and sells them for £10,000 in mid-April. Does he assume that this £10,000 is the first £10,000 he earns that tax year and so can be banked tax free (because he has a £10,000 nil-rate tax band), while the £10,000 from selling the trees comes much later in the year and by then he has crossed the 40% tax threshold. On an NPV basis, this might make poor returns from cows look more profitable than good returns on trees. It is not logical that the discount rate should depend on the sequence in which revenues are earned. However, it may not be practicable to try and work out the blended tax rate for all your activities before you can assign a discount rate to any one of them. By definition, the discount rate is assigned years in advance and the blended tax rate cannot be known until the relevant tax year has ended.

Rather than pay income tax at whatever rate is relevant, Rupert may be able to argue that the growth in value of the trees over the intervening years is a capital gain and should be taxed as such. He now has a nil rate capital gains allowance to use up and for the remainder may be taxed at 10%, 18% or 28% depending on the specific circumstances of his gain. Let's assume that he pays 28% on everything and so his net proceeds from sale of the timber are £7,200. He would need to put £4,481 in his '4% gross & 2.4% net' bank account to arrive at £7,200 in 20 years' time. But supposing his 4% bank account is a UK government-approved Individual Savings Account (ISA) and he receives the interest gross. He would need to put £3,286 in such an ISA to reach £7,200 in 20 years (if the timber proceeds are treated as a capital gain), or £2,788 to reach £6,000 (if they are treated as income at 40% marginal tax).

We have taken perhaps the simplest possible example of a future asset – remember that Rupert did not bother to include the sophisticated schedule of all future costs taken into account by the business executive. And yet simple variations in tax treatment have given estimates for the NPV of those trees ranging from £2,788 up to £6,223. The actual tax implications may be much more complex to model. Suppose Rupert is in debt via a 'current account mortgage' on which he pays 5% interest out of his post-tax income. Such mortgages effectively act as an asset-backed overdraft facility (asset-backed borrowing is discussed in more detail in Section 3.3). If Rupert pays 40% tax, then for every £1,000 which he earns gross each year he will receive a net £600 which will cover the interest on £12,000 of the debt. So if he had a windfall of £12,000 and paid down his mortgage, it would be the equivalent of adding £1,000 a year to his gross salary. This is a gross rate of return of 8.3%, so perhaps this is the gross return he would seek if he were selling his trees to me.

If the tax implications for an individual are complex to model, the tax implications for a multi-national company take complexity to a whole new level. However, many of the issues will mirror those just discussed for Rupert. Most corporations have debt, which is why the rather contrived example of the 'current account mortgage' was given above. They may pay corporation tax on both their trading profits ('income') and capital gains, although they may also find clever ways of avoiding this tax such as using timing differences or inter-company charges. Unlike Rupert, they will get tax relief on interest payments made on money owed but, like him, they will have to pay tax on interest earned. Further, many government tax rules are designed to favour certain industries or certain types of company, for example growing companies being given depreciation advantages or capital allowances to encourage further investment.

For large companies the situation will be complicated further by a treasury function which may borrow in one currency for use in another and engage in aggressive tax minimisation, so the appropriate interest rate to use may be some blended figure based on foreign exchange risk, multiple tax regimes and various local borrowing costs. In our simple example, tax considerations altered the NPV of Rupert's trees by a factor of more than two; for corporations the margin of error is likely to be much greater.

2.2. Taking inflation into account

Another factor which needs to be considered is inflation. In some ways this is similar to a tax on the real value of the returns obtained. In the simplified examples set up in Chapter 1 inflation was ignored but in the real world it can be a significant factor, particularly if a volatile commodity price or currency exchange rate is involved in the DCF calculation.

Taking the example of saving for a car, if car prices are experiencing inflation of 5% per year, then a car which costs £10,000 today will cost £12,763 in five years' time. Based on the original assumption that Duncan can get 3% compound interest on his bank account (and adding the corollary that this is the figure net of tax), then the NPV of his liability is now £11,009 rather than the £8,626 calculated previously. Interestingly, the NPV of the liability is now higher than the cost of buying the car today. In theory, this means that it could be better for Duncan to draw £10,000 out of savings, buy the car today and to keep it pristine in his garage for five years than to buy it in five years' time. Obviously in the case of a car this would not be ideal, but it serves to illustrate the point that in some project finance situations the effect of inflation is not merely to adjust NPV calculations, but also to alter the timing of purchases.

The 28% increase in the NPV of the car liability just calculated is not huge, but five years is not a long period for the effects of inflation to be felt. In the example of Rupert's trees we assumed a 4% compound interest rate (again, let's now treat that as net of tax) and a 20-

year period before receiving proceeds of £10,000. Now let's assume that the price of timber is increasing at 6% per year and that instead of entering a contract now to sell the timber for today's market price of £10,000, Rupert waits and charges the prevailing market price 20 years hence – which should be £32,071. The NPV of his proceeds now increases from the previous figure of £4,564 to a new figure of £14,637. The promise of timber in 20 years' time is apparently worth considerably (46%) more than an actual pile of timber today. And, unlike a car, timber bought today may even turn out to be of superior quality if it is left to season for 20 years.

Clearly, it now makes sense to buy timber speculatively and to sit on it for future use. This drives up demand, which is likely to drive up the price. But since the future estimate of the price of timber is based on today's price lifted by inflation, this will not close the gap. In fact, if the price of timber spikes up today, analysts might be tempted to lift their long-term inflationary assumptions above 6%, making the gap even wider. This feedback loop arises whenever an asset is projected to increase in value over extended periods faster than the return on capital from bank interest. It will often act to inflate asset price bubbles.

Now it is unlikely that someone would buy a car five years before it was needed and put it in a garage for the intervening period. Despite my speculation on the benefits of seasoning timber, it is unlikely that it would be practicable to store wood for 20 years. But it would be practical to buy a much bigger house than was needed and to live in it for 20 years (or, more likely, to continue living in a big house for 20 years after the children have left home). So the British mania for speculative investing in housing ("it's my best pension"), with all its attendant market distortions, could in fact be described as an abuse (or, at least, a misuse) of discount rates.

2.3. Taking risk appetite into account

Section 1.2 introduced the notion that discount rates could diverge from the prevailing rate of credit interest and instead become an arbitrary hurdle rate for the IRR someone wants to achieve. It may be hard to prove than an arbitrary choice of IRR is objectively 'wrong' but it may well be inappropriate. For example, it may be appropriate to use a different rate when putting money aside to cover a liability (as with the car in Chapter 1) than when deciding whether to liquidate an asset (as with the trees in Chapter 1). For the first time, we are thinking about the influence of risk on the discount rate, and this will be considered in more detail in Chapter 4.

In the case of Duncan saving up to buy his new car in five years' time, there is considerable downside risk from failing to put aside sufficient funds because he really cannot function without one. In this case, he might use a discount rate which is a very conservative estimate of the minimum interest which he will actually be paid in a savings account, he will remember to take into account tax deductions from that interest, and he may assume a relatively high rate of car price inflation. If he manages to save tax free in an ISA or inflation is lower than expected, he will have made inefficient use of his capital but he is happy to pay this 'opportunity cost' as an insurance premium against the risk of not being able to afford a car when the time comes.

In the case of selling some trees, Rupert may not need the money right now and so he might prefer to accept the downside risk of receiving a disappointing price for the harvested timber in 20 years' time rather than the opportunity cost of selling the young trees too cheaply today. Alternatively, he might be more inclined to sell the boring old trees at a knock-down price and put the money to work in a riskier, more exciting investment. There have been various versions of the tree saga, so let's stick with the one where we ignore inflation and assume that 20 years from now Rupert would probably be able to sell the wood at today's market price of £10,000. If he can get 4% interest in a savings account, we concluded that he should be prepared to sell the trees for £4,564 today and put the proceeds on deposit for 20 years – this would be another way to generate a £10,000 asset 20 years from now.

If Rupert thinks that in 20 years' time the trees are more likely to be worth £15,000 than £10,000, then he should hold out for £6,846 today. This may make the trees hard to sell if he is more optimistic (or greedy) than other participants in the timber market. There is a risk that he will turn down £4,564 today, hang around for 20 years and then sell the trees for £9,000. He'll wish that he had taken the original offer of £4,564 because effectively the value of his trees has increased at a compound rate of only 3.5% (from £4,564 to £9,000). There was an opportunity cost of not selling for £4,564 and putting the money in the bank at 4% interest. In this case, he has not paid the opportunity cost as an insurance premium – he has used it as the stake in a 20-year bet against the market which he subsequently lost.

However, a higher risk appetite may have precisely the opposite effect on Rupert's behaviour. Since he does not need to sell his trees to fund the necessities of life, he may not be very interested in seeing their value grow by a measly 4% per year, regardless of whether this is the value of the trees in a field or the value of the proceeds in a savings account. He may decide that he wants to put this money into opportunities which will grow at 10% per year. He may have been convinced (perhaps by his banker friend Dominique) that such opportunities are abundant. He would now be happy to sell the trees for as little as £1,486 so that he can free up some money to put into one of Dominique's 10% investment schemes – because £1,486 will grow into £10,000 over 20 years at an interest rate of 10%. But Rupert is not stupid – he knows that the current market price of saplings is £4,564, so he will initially put them up for sale at that price. Unfortunately, if lots of other tree-owners know bankers who can offer them 10% returns, they may all be trying to sell their trees at once. If someone makes Rupert a 'take it or leave it' offer of £3,000, he may be inclined to take it just to get the deal done.

So a heightened appetite for risk may make Rupert demand £6,846 for his trees, or it may cause him to accept £3,000 for them (once again, the variation is by a factor of more than two). This illustrates an important point: discount rates can model the consequences of a person's decision on risk – they cannot tell that person what their decision should be. There is a tendency to use an industry benchmark discount rate (say 10%) because it is the 'right rate' for a particular type of project. In effect, you are assuming that the risks in your project, and your appetite for risk, are the same as everyone else in a similar line of business. As with the tree owner and the car buyer, the correct approach is to consider your personal circumstances and risk appetite. Factors you should consider include: how much you might lose (in business called the 'downside value at risk') relative to what you can afford to lose; what else you might do with the finite amount of money you have ('opportunity cost'), and; the existence of insurance against such risks and your willingness to pay the relevant premium ('underwriting'). Taken in the round, these factors should inform your decision on discount rate.

The end of this section is a good place to reiterate the important general point that discount rates are inferred from human decisions, not fundamental constants of the economic universe derived by clever economists in a manner analogous to physicists working out a value for acceleration due to gravity. In fact, to be rigorously empirical, they should ideally be derived from human actions (the actual investment of real money) rather than just a popular consensus of disinterested parties. This point could have been made when discussing the fundamental principles of NPV analysis in Chapter 1. However, it seemed better to present a few behavioural and attitudinal examples first which might make clearer why it must be so. The dangers of relying on an *a priori* 'right discount rate' (or, worse still, a 'one size fits all' rate) are discussed further in Chapter 5.

2.4. Taking past experience into account

Whenever risk has to be taken into consideration the problem immediately arises of how it should be measured. In some very simple situations (e.g. the 'risk' of not winning the lottery) it can be modelled using probability theory. But most real world situations (e.g. the risk of a plane crash) are far too complex for theoretical modelling and so estimates have to be based on historical or empirical data. The problem in finance is that "past performance is not a reliable indicator of future results" – a mantra so important that it has become a mandatory health warning. Past performance may not be a *reliable* indicator, but what else do we have? When Rupert sells his trees for £3,000 he probably has faith in his banker friend Dominique because in each of the preceding three years her investments have indeed returned a 10% profit – there are plenty of hedge funds which can point to such performance. This is not a reliable indicator that such funds can deliver 10% every year for 20 years, but that is what gets fed into the DCF analysis, the results from which confer an aura of mathematical precision on what is basically just a hopeful hunch.

Historical data can be equally harmful when it is too pessimistic – in finance it is not always without risk to "err on the side of caution". After the 2007 credit crunch, interest rates fell to negligible levels, equity values dived and government bond values rose. Government bond yields (the interest paid on the now much higher bond values) consequently dived to levels which in some cases were lower than the rate of inflation over extended periods. Duncan may feel that while saving up for a car he cannot take risks with his money, but his predicament pales into insignificance when compared to someone saving up for their pension. A cautious person using historical performance as a guide to their pension investment allocation in 2008 would have been sorely tempted to put all their assets into government bonds. This is especially true for someone within ten years of retirement age – often described as "too old to risk waiting for the equity market to recover". And as their existing pension pot was unlikely to have been held in cash, "putting all their assets into government bonds" probably means selling equities which have just dropped in value and using the proceeds to buy government bonds.

DCF models of the 'prudent' behaviour just described may well have suggested that locking in the low returns from government bonds for a decade would indeed remove uncertainty, only to replace it with the near certainty that on retirement the existing pension pot would be smaller, in real terms, than its current value. For many people, this in turn would imply that there would be absolutely no 'risk' of it being large enough to generate a pension any higher than the level of means-tested benefits paid to those who had saved nothing (a form of 'moral hazard'). In such a situation, the logical conclusion may be either to stop making any pension contributions, or else to make what would effectively be a one way bet that over the next ten years interest rates would rise, bond prices would fall and equity markets would recover. It's easy to state this dilemma with the benefit of hindsight, but some pundits did put forward similar analyses at the time. Setting the wrong discount rate is not always an abuse – it can sometimes be an accident arising from the best of intentions.

2.5. Applying a range of different discount rates

So far we have assumed that the discount rate for credit and debit items is the same. There is some logic in this because if I leave cash on deposit I receive a certain amount of interest, but if I take it out and put it to work in another investment (i.e. 'lend' it to someone) my 'opportunity cost' of the capital is the same amount of credit interest which I now won't receive. It is even more logical when we are not trying to reference any actually paid real world interest rate but instead considering an aspirational IRR hurdle rate. It surely makes sense for a company with diverse operations to set a 'level playing field' figure for its target Return on Capital Employed (ROCE) in order to ensure that it places its bets on the best operational and geographic opportunities. However, we have also noted briefly (in the reference to 'current account mortgages' in Section 2.1) that this assumption may be called into question by tax considerations. We will now consider in more detail whether in some situations, setting the right discount rate should involve different rates for assets and liabilities.

The first thing to note is that real world credit and debit interest rates are not generally the same. In fact the whole banking edifice depends on them being different. Bankers need to make money and they achieve this by lending at rates higher than the rates they pay depositors. Long before the development of sophisticated DCF analysis there was an adage that commercial bankers operated on the 'three threes' principle – arrive at work in the morning, take a deposit at 3% interest, lend it to someone at 3% margin, then go and play golf at 3 o'clock. The natural order of things is that individual savers and corporate treasurers expect to receive a lower interest rate on their deposits than they pay on their borrowings. And yet companies investing in business opportunities and individuals buying homes both hope to borrow at a lower rate than their long-term return. If I can borrow at 4% and invest at 7% then I do not need sophisticated DCF analysis to guide my behaviour – I simply need to know how much I can borrow (or 'leverage up'). If DCF analysis plays any role, it is likely to be the post-rationalisation of a decision to borrow too much. The abuse of discount rates has therefore played a role in the over-leveraging of western economies, the other great economic threat (alongside asset price bubbles) to financial system stability.

The difference between real world credit and debit interest rates can have important practical consequences when actually financing a project rather than simply using DCF as a ranking tool to determine which of various candidate projects might theoretically be the more profitable. Let's consider the case of an ambitious young man called Tim, who has decided that he would like to build his own private nuclear power station. Before we do so, I should make it clear that this essay is not a treatise on the economic viability of nuclear power, so please treat the following cost estimates as entirely arbitrary. They are used simply to illustrate a possible flaw in DCF analysis. No doubt those actually advising Tim on the planning for his nuclear power station project would take care to avoid such errors.

For simplicity, we will assume the station costs £25bn to build, is paid for up-front and is completed at the end of Year 1. Next, we will assume that it will generate electricity which

Tim can sell initially for £2.5bn per year and that this will increase at 3% per year to account for energy inflation (he was canny enough to negotiate that before a recent collapse in the price of oil). Then we will assume annual running costs of £0.6bn in the first year, also rising at 3% inflation. After completion the plant will run for 49 years (giving us a 50-year time horizon in total) and will then be decommissioned at a cost of £10bn paid entirely up-front (i.e. at the end of Year 50). This last figure may be particularly controversial but remember that we have 50 years of inflation to take into account before that cost is incurred. Finally, let's assume a 10% discount rate because that is a nice round number and may be an attractive rate of return for Tim's commercial partners. The table below sets out the total NPV of the annual cost and income items and, for illustration, shows the contribution to that total from Years 1, 2, 10, 20, 30, 40 and 50 (there is no room for a table with 50 columns).

Table 2: Project finance DCF for a nuclear power station

Item	**Total NPV (£bn)**	**Contribution to total NPV from particular years (£bn)**						
		Year 1	**Year 2**	**Year 10**	**Year 20**	**Year 30**	**Year 40**	**Year 50**
Construction	25.00	25.00	0	0	0	0	0	0
Operations	8.23	0	0.55	0.32	0.17	0.09	0.04	0.02
Decommission	0.09	0	0	0	0	0	0	0.09
Total cost	**33.32**							
Electricity sales	34.29	0	2.27	1.34	0.70	0.36	0.19	0.10
Total income	**34.29**							

The first thing to do is to remind ourselves what the total NPV values mean. The 'total cost NPV' of £33.32bn means that if Tim could put that sum in a bank account yielding 10% net interest, then he could draw out of that bank account all the money needed to cover all future costs of the power station, and it would be empty just after he had paid the decommissioning bill (he is a young chap, so there is every chance he will be around in 50 years' time to sign the final cheque). Of course, he would draw out most of it (£25bn) to pay for building the power station during Year 1. The 'total income NPV' of £34.29bn means that if Tim were to put all his future revenues into a bank account yielding 10% net interest, then at the end of 50 years the total would be the same as if he had put £34.29bn into that account today and left it to grow at 10% over 50 years. He probably doesn't have £33.32bn hanging around in his current account, but if he were to borrow £33.32bn it would be as valuable to him as putting £34.29bn into a bank account delivering 10% per year. The conclusion is that this project actually returns a little over 10% on the capital employed (i.e. it passes the 10% hurdle rate) and if Tim can borrow £33.32bn in the market at an actual interest rate of <10%, then he certainly should do so.

One thing which may be surprising is how little is contributed to the NPV totals from the period beyond Year 20. The NPV of the £600m of operational costs in Year 2 is £545m while the NPV of the operational costs in Year 20, by which time they will have risen to £1,021m due to inflation, is only £167m. This is reassuring really, as it is very difficult to predict what the price of electricity and the costs of running a nuclear power station will be 20 years from now. It tells us that quite big errors in those predictions will make only minor differences to the viability of the project. Perhaps the biggest surprise is that the £10bn of decommissioning costs in Year 50 contribute only £94m to the NPV of the project – a negligible 0.3% of the total project costs. This is because 10% compound interest left to run over 50 years will increase a sum of money by a factor of 117 (11,700%). In view of this, why do people fuss so much about the cost of decommissioning?

We know that these days people do not receive or pay 10% interest in practice. Tim will probably borrow the first £25bn at about 5% and be delighted that he can 'put it to work' at 10% return. In each future year he will have revenues, out of which his costs will be paid, leaving him a tidy profit. He discounts that profit at a rate of 10% when he considers its notional worth today. He won't actually have to put that figure in the bank and see it grow at 10% because he is counting future profits rather than providing for future liabilities. Except in the case of the decommissioning costs. Here, the DCF analysis imagines that Tim really does put £85m into a bank account yielding a net 10% compound return. He can then forget about it until he draws out the full cost of decommissioning in 50 years' time.

But we know that such bank accounts do not exist. Suppose Tim can only get a net 2.5% compound return. The amount he would have to put in the bank account today rises from £94m to a whopping £2,909m, meaning that the project no longer meets its 10% IRR hurdle rate. And what if he has underestimated the cost of dealing with nuclear waste over thousands of years? The figure of £10bn was arrived at by assuming that today the cost would be about £2.25bn and that inflation in decommissioning costs would run at 3% per year. But what if actual inflation in these costs was 5%? The estimated cost in 50 years' time would rise from £10bn to £26bn. At a 2.5% savings interest rate it would be necessary for Tim to put aside £7.5bn today to cover that liability.

Hopefully what is clear from the above is that while DCF may be a good tool to test out a 10% IRR hurdle rate for different power generation options (nuclear, fossil, wind, solar) and thereby produce a 'Levelised Cost of Energy' for each, it may not offer a very good picture of the actual project finance economics. To be fair, Tim could say that, instead of a single up-front deposit, he will put aside most of his profits from Year 30 onwards to cover the cost of nuclear decommissioning. This will remove very little from the DCF total income value but would be plenty to cover the decommissioning cost. However, that is not what is being modelled by a simple DCF analysis and so it may not appear in the contractual obligations of the operator. It may be that Tim is tempted to maximise profits for 30 years and then conveniently go bankrupt. So, perhaps, after all, people are right to make a fuss about the cost of decommissioning.

A key lesson from all this is that DCF shows we do not have to worry much about things which will happen 20 to 50 years from now. At a human level this is true – roughly half the population will not even be alive in 50 years' time (even young Tim will be collecting his pension by then). For economic evaluations across society, discount rates are typically lower. When people attempt to estimate intergenerational transfers, they attempt to estimate the 'pure rate of time preference' - the rate people would ethically use to evaluate transfers to future generations. To get there, analysts sample the population using ethical questions about saving lives versus costs to try and find these utility functions, often arriving through a thicket of contradictions, as did HM Treasury, at an estimate around 1.5%. Because people die, and the average annual death rate for adults is about 1.5%, this is not a surprising number. Individuals, at this point in time, quite rightly want to see payback in their lifetimes. No wonder old people are crotchety about long-term investments.

However, what makes selfish economic sense may not make social sense. Is it right for us as a society to commit irrevocably to a project which makes economic sense during our life-times but becomes an inevitable and massive liability for our grand-children? Suppose humankind now had to decommission the Great Pyramid of Cheops and that this turned out to cost hundreds of trillions of dollars – leaving us nothing for the basics of life. Would it be any comfort that Cheops had provided for this by investing a tiny sum of money in 3,000 BC? At this point in the essay it is worth reminding readers of Michael Mainelli's quote from Charles Young in the Preface: "an extra glass of wine for Alexander the Great matters more than all today's capital stock".

2.6. Taking account of externalities

In economics, an 'externality' may be defined as a cost or benefit which affects a party who did not choose to incur that cost or benefit. In the case of a cost, this is what a military person might refer to as 'collateral damage'. For DCF analysis, externalities may be thought of as the indirect costs and benefits which get left out of the NPV calculation. In many situations of practical importance what gets left out is collateral damage to the environment. If such damage is ignored, then the costs of preventing it happening cannot be properly compared to the costs of remediating it at a later date, even though this would be an obviously appropriate use of NPV analysis. My grandmother was fond of saying that "a stitch in time saves nine" – but if you have to pay for the first stitch yourself while others would have to pay for the nine, you may not be inclined to make that investment. And, playing devil's advocate, with a suitably high and accurate discount rate, the cost of one stitch now might actually be higher than nine stitches a decade from now.

To illustrate this point, let's consider the case of Sophie who owns some juvenile fish which by coincidence are all the same size and currently weigh 200 grams each. Sophie is the sister of Rupert (the owner of the trees) and the fish live in a pond on the family estate. They are of a particularly tasty variety and can be sold to a fishmonger for £5/kg, making each fish worth £1.00 right now. If left in the pond, they will put on weight at a rate of 5% per year. Sophie has just had dinner with Dominique, the banker friend of her brother who in Section 2.3 was flogging "sure fire" investment opportunities yielding 10% per year. She wonders whether to drain the pond, kill all the fish and sell them to provide funds for this investment. Clearly she should do so – a year from now each fish will weigh 210g and be worth £1.05 whereas the financial investment will be worth £1.10. Of course, none of the fish will now grow up and produce offspring, which would have provided additional returns and a sustainable source of income. But the disparity in discount rates means that, on an NPV basis, it simply is not worth waiting for that to happen.

What 'did for' for the poor fish was that the financial returns promised by industrialisation (if Dominique invested in the equities of manufacturing industry) or short-term financial wizardry (if she chose hedge funds) significantly exceeded those which could be returned from the sustainable harvesting of nature. This is often the case in real life – nature seems to be prepared to offer single digit returns while humans are eager for double digits. The tragedy is not that one fictitious lady owned some fish in an ornamental pond, but that the human race lays claim to a large number of real fish in the world's oceans and is treating them in a similar manner. What makes economic sense for Sophie will probably also make economic sense for everyone else (we look at the perils of everyone acting on identical models further in Chapter 5) and so entire species of fish end up almost extinct. This will have a cost for humankind which is not captured in the NPV model and is a classic ‘externality’.

The economist Jeffrey Sachs has addressed this issue in his 2008 book ***Common Wealth***. He points out that “the market price of a species will generally not reflect its societal value as part of the earth’s biodiversity’ and that ‘if the value of the resource is likely to grow more slowly than the market rate of interest, the blaring market signal is to deplete the resource now and pocket the money”. Michael Mainelli and Ian Harris, the founders of the Long Finance initiative, pick up on this discussion in their appropriately named 2011 book, ***The Price of Fish***. As well as furnishing the above quote from Jeffrey Sachs, they point out that pricing biodiversity (i.e. putting a value on the collateral damage caused by extinction) is a similar problem to arriving at a carbon price which is ‘fair’ to the environment and properly reflects the costs of global warming. There is no market price for species extinction but there is a market price for carbon and many commentators believe that it is too low. If it was set at a more realistic level, it might help to make the case for new nuclear power stations on an NPV basis without needing to duck the issue of waste.

Another group which has grappled with this issue is the International Institute for Sustainable Development (IISD). Their January 2014 report on ‘Value for money infrastructure procurement: the costs and benefits of environmental and social safeguards in India’ includes a section on ‘How environmental and social risks affect discount rates and risk sharing’. This suggests that, rather than try to put a market price on the future ‘good’ of carbon not emitted or a species not driven to extinction, such externalities should be captured by increasing the discount rate used when modelling returns generated from the option which does not address such aspects. The reasoning is that the increased discount rate captures a risk of future environmental legislation eating into those returns.

Oscar Wilde famously described a cynic as “a man who knows the price of everything and the value of nothing”. A similar charge might well be levied at DCF modellers who leave out externalities. However, the temptation to ignore or understate them is not always driven purely by the desire to make a project look more profitable. Furthermore, as we saw with nuclear waste, any adverse impacts of shoddy analysis tend to be felt in the distant future – how much do I care about costs incurred long after I have died? It is actually very hard to quantify the ‘economic cost’ of environmental damage. Like much of economics, DCF can tell you what is expedient – it can never tell you what is morally or socially acceptable. 2014 marked the 100th anniversary of the extinction of the Passenger Pigeon – exactly what price do each of us pay today for its loss? We will now consider in more detail the effect of applying DCF analysis over trans-generational periods of time.

3. Understanding the effects of timescale

3.1. Long timescales discount future income and expenditure to neglectable values

The example of the decommissioning costs for Tim's nuclear power station illustrates a general aspect of exponential mathematics – that the numbers start to run away with themselves when extrapolated over extended periods. The fact that the NPV calculated for this key aspect of the project may range from £94m to £7.5bn does not provide confidence limits for the required budget – it indicates that we are flying entirely blind. As a schoolboy I was intrigued to be told that if, when Jesus told the Pharisees that the silver penny should be "*rendered unto Caesar*", he had added that this was because Caesar was paying 3% compound interest, then it would have grown to be worth £472 trillion today. Two thousand years is a long time for a curve to become ever steeper.

It does not, in fact, take thousands of years for exponentials to run away with themselves, even at relatively low discount rates. In the case of nuclear decommissioning, we saw that a 10% discount rate reduces the NPV of a liability to less than 1% of its nominal value in 50 years. The table below should help provide a feel for when liabilities will be reduced to negligible (in its true sense of 'neglectable') proportions.

Table 3: Diminution in nominal value at various discount rates

Discount Rate	Period to reduce to 10% of nominal value (years)	Period to reduce to 1% of nominal value (years)
1%	231	463
3%	78	156
5%	47	94
7%	34	68
10%	24	48
15%	16	33
20%	13	25

3.2. Interest rates cannot actually be fixed over long timescales

When choosing discount rates, it is important to consider whether they apply over the entire timescale of the project. Graphs 2, 3, 4, & 5 take a Long Finance view of the situation. You can see that UK and US real interest rates, i.e. Real Interest Rate = Base Rate – Inflation, have exhibited two characteristics over the centuries (a) they have grown more volatile, and (b) they have been decreasing, as indicated by the moving averages dropping. One can see by inspection that 'real is actually well below 2% for the past century on a 50 year moving average (green line on Graphs 3 and 5), not the 4% or 5% frequently talked about as 'normal' by many people and economic commentators. Volatility has been so significant in the past century it's odd that people even think normal is 4% or 5%. In addition, you can see that quantitative easing since 2008 has already had a significant downward effect on the 50 year moving average.

Graph 2: UK Real Interest Rates

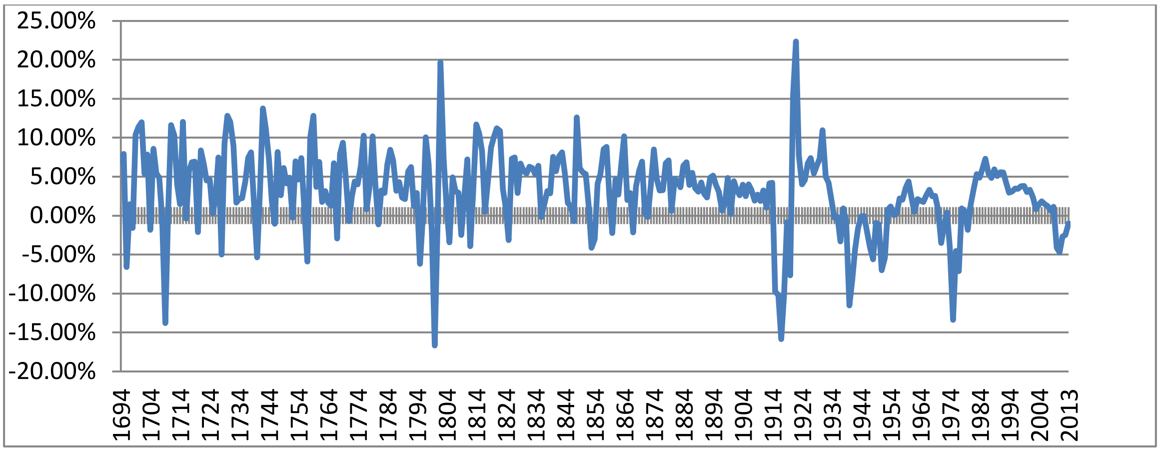

Graph 3: UK Real Interest Rate Moving Averages

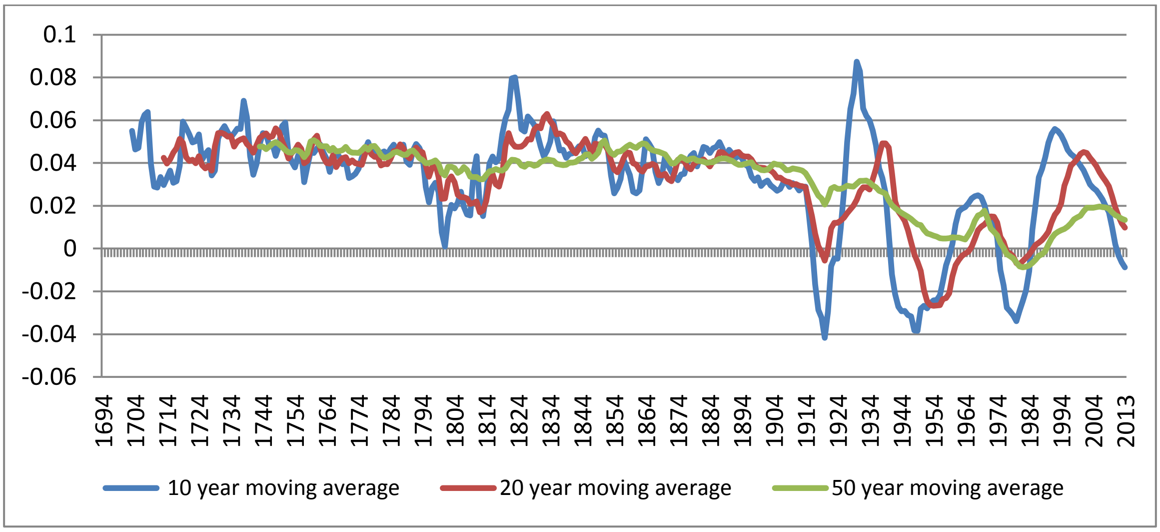

Graph 4: US Real Interest Rates

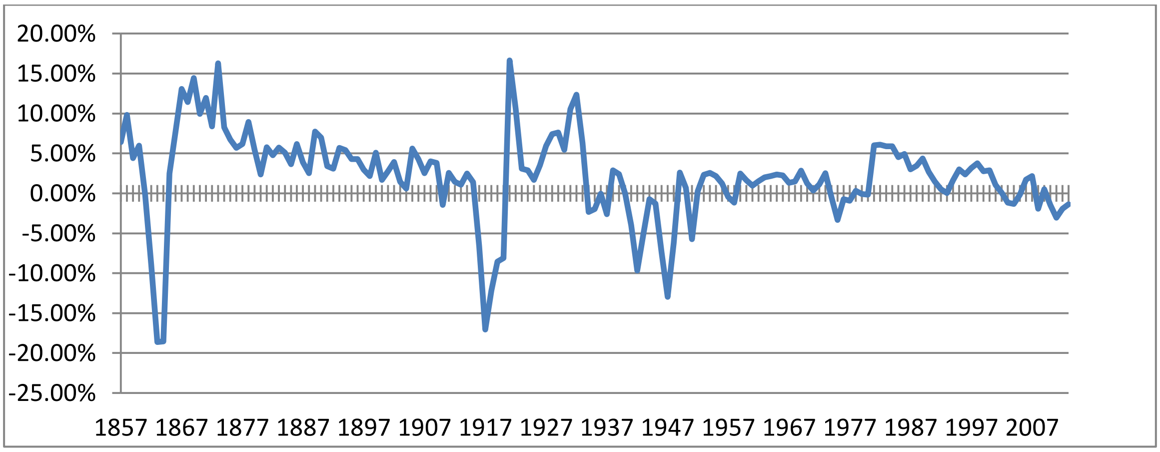

Graph 5: US Real Interest Rate Moving Averages

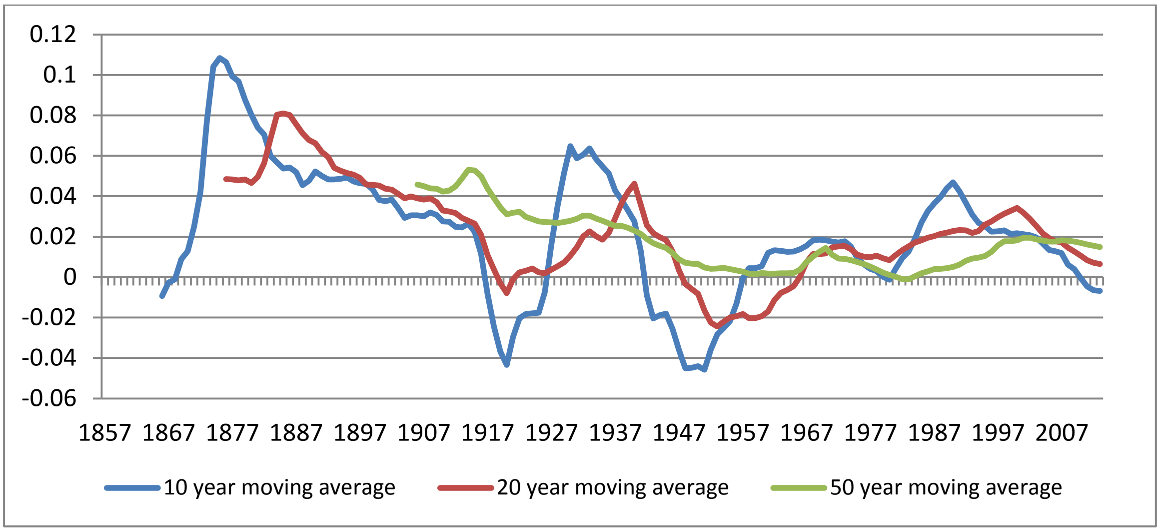

In Section 2.4, we touched on the issue of risk and the potential pitfalls from using recent past experience as a predictor for future performance. We noted that this could potentially have profound consequences for anyone with a defined contribution pension scheme approaching retirement. In some ways, the key economic indicators (and for DCF these will be interest rates, inflation rates, and, to some extent, tax rates) are rather like weather indicators. In the case of weather patterns, the last few days will be a reasonable (but not entirely reliable) indicator of what will happen for the next few days. However, they will tell us next to nothing about the weather on a particular date ten years from now. Economic indicators are similar except that the coherence timeframe is a few years rather than a few days. Interest rates for the past three years may be a useful guide to those for the next three years, but tell us very little about what they will be in 2050. This is illustrated in Graph 4 on the previous page which shows the historical real interest rates in the US since 1857. If this was held to be similar to weather pattern data, then it would be for a region with periodic and very dramatic storms.

Over recent years, interest rates have, in fact, been very low – with a 0.5% base rate in the UK since March 2009. However, putting 0.5% into a 30-year DCF would not be the best estimate for future base rates – it would probably be one of the worst possible estimates. This is because when base rates are this low they will almost certainly have to increase (unless, perhaps, you live in Japan). If the base rate is an important parameter and the timeframe is 30 years, then the best approach is probably to carry out multiple simulations based on the typical patterns of base rate movement over similarly extended periods. You might weight the outcome from these different simulations according to the likelihood of them happening and use this to arrive at a best estimate for a 'central case' outcome. At the same time, it would be possible to estimate the probability of different deviations from this central case, which could be used to make contingency provisions. This approach is sometimes called 'Monte Carlo' simulation. There are similarities to the Long Finance initiative's 'Confidence Accounting' proposal, a plea to use ranges in balance sheet valuations. Monte Carlo analysis and Confidence Accounting are outside the scope of this essay, but the point is that DCF models may often be best used to provide inputs to further 'probabilistic' models rather than to provide 'oven ready' conclusions. We pick up the treatment of risk again in Chapter 4.

As well as touching on risk, Chapter 2 made the point that DCF analysis may not offer a very good picture of actual project finance economics when contrasting an IRR hurdle rate with actual credit or debit interest rates. We noted that in many cases the modeller has no intention of using real world interest rates – citing the example of using a 'Levelised Cost of Energy' tool with a uniform 10% IRR to compare the merits of different power generation technologies. When timeframes extend beyond about five years, the modeller in fact has very little ability to use real world interest rates even if she wants to - if by 'real world interest rates' we mean the rates that would actually be paid throughout the project. This is because it is difficult to fix these rates for very long periods in advance.

Duncan, the car buyer we first met in Section 1.1, may just possibly be able to find a five-year fixed rate savings account offering compound interest. But Rupert, the tree seller who lives in the big house down the road from him, would certainly not be able to find a similar account with a rate fixed for 20 years. There are debt instruments with 20-year tenors, government bonds being a good example, but these pay simple rather than compound interest. Nor, indeed, is it really likely that Rupert would be able to find a counterparty prepared to enter into a fixed price contract to buy his timber 20 years hence (and by Section 2.3 he was anyway having doubts about the wisdom of doing this).

Given historical levels of volatility, it is not even very easy to estimate what the level of future interest rates will be several years ahead, let alone secure them contractually in advance. Over long timescales, the choice for a DCF modeller is therefore not between 'real interest rates' and 'arbitrary IRRs' but between 'guessed interest rates' and 'hoped for IRRs' – there may be equal amounts of aspiration in each. So DCF is good for illustrating that a certain financing plan *could* be attractive – while not guaranteeing that it *will* be attractive as events unfold in the real world.

3.3. Asset-backed borrowing may provide some certainty on debt interest rates

Although long periods of fixed rate compound credit interest and pre-agreed commodity prices ('forward prices') for transactions in the distant future are very difficult to secure, the situation is a bit easier for simple annual interest paid on debt. Pension funds are looking for stable long-term streams of regular interest payments on low risk deposits (which is why pension funds are so fond of government bonds) and building societies will look at relatively long-term fixed rate mortgages on property. Mortgaging a property is a form of 'asset-backed borrowing' and is popular among private individuals. In the case of companies, a similar strategy to mortgaging physical assets is the 'securitisation' of future revenue streams – for example borrowing to build a student accommodation block against the security of the future rental income. Sometimes the distinction between individuals and corporates becomes blurred – David Bowie famously issued a $55m bond securitised by the anticipated stream of future royalties on his back catalogue of songs.

Where a project has up-front investment in a tangible asset which is paid back directly from the returns made operating that asset, it may be best to model it based on the assumption that the investment is paid for under a fixed rate repayment mortgage or some other form of asset-backed borrowing. For example, a 'repayment mortgage' type of financing for an asset other than a house can often be arranged as a hire purchase (HP) agreement, which in the case of a corporate is referred to as a 'financing lease'. This is particularly appropriate if that asset is (unlike most houses) destined to depreciate steadily in value. The period of the loan should reflect the useful life expectancy of the asset and the repayments will be covered by the income it generates before reaching the end of its useful life.

Let's imagine that Margaret develops a business plan to buy a caravan for £10,000 and then rents it to a select group of friends for their holidays. This is not a bad idea, as many caravans are only used for a couple of weeks a year and take up space for the remainder, quite apart from tying up capital. She has five friends who will pretty certainly hire it every summer for two weeks each at £200 per week, bringing in £2,000 per year of revenues. This income will increase at 3% per year in line with inflation. Because it will be used for large numbers of trips, the shared caravan will be worn out after ten years, at which point it will be worthless and simply towed away for scrap. We shall assume that it is bought on a ten-year hire purchase agreement with an interest rate on outstanding balances of 7%. The economics of the business will then be as shown below.

Table 4: Economic model for funding a 'buy-to-let' caravan using hire purchase

Item	Year 1	Year 2	Year 3	Year 4	Year 5	Year 6	Year 7	Year 8	Year 9	Year 10
Rental income	£2,000	£2,060	£2,122	£2,185	£2,251	£2,319	£2,388	£2,460	£2,534	£2,610
Payment to HP company	£1,424	£1,424	£1,424	£1,424	£1,424	£1,424	£1,424	£1,424	£1,424	£1,424
Interest component	£700	£649	£595	£537	£475	£409	£338	£262	£180	£93
Capital repaid that year	£724	£774	£829	£887	£949	£1,015	£1,086	£1,162	£1,244	£1,331
Debt remaining at end of that year	£9,276	£8,502	£7,673	£6,786	£5,838	£4,823	£3,763	£2,574	£1,331	£nil
Retained profit	£576	£636	£698	£762	£827	£895	£964	£1,036	£1,110	£1,186

It turns out that if Margaret pays exactly £1,424 to the HP company at the end of each year, by the end of Year 10 the entire principal has been repaid. A proportion of the £1,424 covers interest on the balance outstanding at the start of the year and the remainder pays down the capital. As the outstanding balance drops, the interest proportion decreases and the amount repaid increases. This will be familiar to anyone who has taken out a repayment mortgage. Those who have taken out endowment-linked mortgages will, conversely be familiar with the notion of leverage introduced in Section 2.5. The idea, justified by a different NPV calculation, is that if the return on an endowment fund will be consistently higher that the cost of borrowing against my home, then I should remain in debt for as long as possible and put the money to work in the stock market. At the end I supposedly could pay off my mortgage and have cash left over. It might be argued that selling endowment mortgages was an abuse of discount rates while working out the cost of a repayment mortgage is a good use of them.

Returning to Margaret and her caravan, we can see that the HP agreement has allowed her to lock in a fixed interest rate on her debt over a ten-year period, thereby removing some of the uncertainty in her business model. She does not have to base her planning on a DCF analysis using a theoretical interest rate, which may change over the course of ten years, or an arbitrary IRR. There is a risk on the revenue side that she won't be able to find enough people to hire the caravan, but let's assume she knows her five friends well enough to be pretty confident in that respect. This leaves her with a stream of near-certain and steadily increasing profits. In the ten years before the caravan is scrapped, these will add up to £8,690. However, this money will not be available up-front and the use of an arbitrary IRR would in this case be entirely appropriate for working out its NPV. Margaret may, for example, want to compare the relative merits of a buy-to-let caravan with that of a buy-to-let sailing boat in order to decide which business offers the most attractive opportunity. At a 10% discount rate the NPV of the profit stream is £5,521. At a 3% discount rate, which is more representative of the returns made in actual savings accounts, the NPV becomes £7,491.

Hopefully, the above example illustrates the advantages of using the actual project finance cash flows to model some aspects of a business and a simple discount factor to address other aspects. It also illustrates the way that asset-backed borrowing can help to mitigate the risk from fluctuating interest rates when the project runs over extended periods. In Section 2.2, Duncan was worried that the combined effects of car price inflation and falling interest rates may mean that putting £8,626 in the bank today would not be sufficient to cover the cost of his vehicle in five years' time. He could potentially have avoided this risk if he bought the car today for £10,000 on a HP scheme. DCF would have been a good way to find out if this was actually a sensible idea based on the interest rate offered by the finance company.

All the examples used in this essay are rather contrived (no, really!). They are only useful if they provide simplified explanations of more complicated situations which might reasonably occur in the real world – albeit more often for businesses than for private individuals. We have considered the merits of asset-backed lending with fixed interest rates to reduce the uncertainty in DCF project finance projections stretching over extended periods. It is reasonable to ask whether such borrowing facilities actually exist in the real world. For individuals, fixed rate mortgages tend to run for three to five years and HP agreements for two to four years. Corporates make use of financing leases to benefit from what is effectively HP over slightly longer periods (perhaps up to seven years). For the UK government, building schools and hospitals under the Public Finance Initiative (PFI) resembles a property HP arrangement and similarly makes sense only if the discount rates are correct (especially the discount rate of government debt).

As an alternative to mortgaging their properties, companies seem to prefer sale-and-lease-back arrangements. This involves selling the premises which the company currently owns to a financial partner and then leasing it back from them under a long-term arrangement. To the uninitiated, this might seem to be a somewhat desperate measure akin to a private individual using a pawn shop. But, it can be justified on the basis of DCF analysis. If the company can put the cash raised to work at a higher IRR than the annual cost of not owning its building (which comprises the rent it now pays plus the capital appreciation from which it no longer benefits) then it makes sense to do so. However, the cost to the company is equal to the return (from rent received and capital appreciation) to the finance company. So sale and lease back is really based on the assumption that the risk-adjusted IRR achievable by the business selling the property is higher than that of the finance company buying it.

It might be argued that such out-performance is a problematic risk to take, especially since property has been one of the best performing asset classes in the UK over long periods. There are some similarities with endowment mortgages, where private individuals delayed paying off the debt on their homes in order to invest the cash in assets offering a superior return (this is discussed in more detail in the next section). The cynical might point out that many companies have properties on their books at much less than their current market value, and that sale and lease back allows a profit to be recognised, perhaps masking unexpected losses in the operational side of the business. If this is true, then the use of NPV modelling to expose the faulty rationale would be a good use of discount rates.

With the partial exception of the property financing schemes just discussed, there is not, in fact, a highly developed corporate market for asset backed repayment mortgage type borrowing over periods of ten years or more. However, it is possible to imagine creating such a facility synthetically by issuing a series of long-dated fixed interest bonds, for which there is a very clear demand. Thinking back to Margaret, our caravan entrepreneur, it is clear that she could achieve the same project finance arrangement by issuing a 1-year bond for £724, a 2-year bond for £774, a 3-year bond for £829 and so on up to a 10-year bond for £1,331. When each bond matures, the amount that needs to be repaid mirrors the 'capital repaid that year' figure in Table 4.

There will be complexities to take into account with this approach. The so-called 'yield curve' for lending shows that shorter dated bonds tend to have lower interest rates than long-dated ones, so it may not be appropriate to issue all the bonds at the same interest rate, but the interest rate on all the bonds will at least be fixed up-front. Also, the largest bond (for £1,331) is collateralised during its final year by a virtually worn-out caravan. Notwithstanding these issues, the idea of synthetic repayment mortgages for financing assets such as wind and solar farms is worth investigating. Furthermore, use of the PMT function in Excel to model such financing structures may be a better approach to calculating the levelised cost of the energy produced than an arbitrary 10% discount rate.

3.4. The infinite timescales implied in equity valuation

Sometimes the implied timeframe for a DCF is infinite (or, to be pedantic, 'unbounded'). The market value of a company is often described as the "net present value of all its future dividends". But we don't know how long the company might exist. Mathematically, this long stream of future dividends is not as frightening as it sounds because a geometric series can be summed to infinity, and if the discount rate and inflation in the nominal value of the dividends both stay constant, the stream of all future earnings does indeed form a geometric series when reduced to their NPV. If the maths doesn't appeal, then look back to Table 3 and note that after about 100 years the NPV of any sum is reduced to a neglectable amount, so a good estimate for the NPV of our stream of future dividends can be obtained from the value of those delivered over the next 100 years, with those delivered from Year 101 through to the end of time safely being ignored. The problems are not mathematical – they arise from scepticism that any company will be around for the next 100 years and the fact that, if it is, I won't be alive to collect my dividends.

Of course, the near-term dividends which will dominate the eventual value of the DCF should themselves form a 'diminuendo' of decreasing individual NPVs. The sum of 100 dividends might, for example, prove to be about 15 times the value of the first number in that series (i.e. the dividend being paid this year). This explains why price/earnings (P/E) ratios are useful as a 'quick and dirty' method of valuing shares. (I know that earnings and dividends are not the same thing, but the principle holds good.) Buried inside a P/E ratio is therefore the notion of a DCF calculation. And that's why stocks which are growing their earnings (and presumably therefore also their dividends) are deemed to merit a higher P/E ratio than those which have static earnings. The growth in the nominal value of the dividend each year offsets some of the decrease in its NPV caused by an additional year of discounting. If the P/E ratio is dependent on the growth rate of the earnings, then perhaps the ratio between them should provide a useful guide to whether a share is over- or under- valued. This leads to interest among some investors in the P/E ratio divided by the growth – the 'PEG' ratio.

However, if the anticipated growth rate becomes higher than the discount rate for an extended period, we start to encounter a mathematical instability. This is illustrated in Table 5 for some ten year streams of dividends which are growing at different rates but are all discounted backwards at 5%.

Table 5: NPV of stock dividends growing at various rates

Item	Year 1	Year 2	Year 3	Year 4	Year 5	Year 6	Year 7	Year 8	Year 9	Year 10
Nominal value of dividend, 0% growth	£1,000	£1,000	£1,000	£1,000	£1,000	£1,000	£1,000	£1,000	£1,000	£1,000
NPV of dividend at 5% discount rate	£1,000	£952	£907	£864	£823	£784	£746	£711	£677	£645
Nominal value of dividend, 3% growth	£1,000	£1,030	£1,061	£1,093	£1,126	£1,159	£1,194	£1,230	£1,267	£1,305
NPV of dividend at 5% discount rate	£1,000	£981	£962	£944	£926	£908	£891	£874	£857	£841
Nominal value of dividend, 5% growth	£1,000	£1,050	£1,103	£1,158	£1,216	£1,276	£1,340	£1,407	£1,477	£1,551
NPV of dividend at 5% discount rate	£1,000	£1,000	£1,000	£1,000	£1,000	£1,000	£1,000	£1,000	£1,000	£1,000
Nominal value of dividend, 7% growth	£1,000	£1,070	£1,145	£1,225	£1,311	£1,403	£1,501	£1,606	£1,718	£1,838
NPV of dividend at 5% discount rate	£1,000	£1,019	£1,038	£1,058	£1,078	£1,099	£1,120	£1,141	£1,163	£1,185

The NPV of the dividend streams growing at 0% and 3% do indeed form a diminuendo and so their sum over an infinite period will be a finite number. For reasonable combinations of growth rate and discount rate this sum is typically 15 to 25 times the value of the first dividend. If a company pays out half its earnings in dividends, then the NPV of all future dividends is 8 to 12 times the current year earnings. In other words the appropriate P/E ratio for such stocks lies in the range 8 to 12. However, when the dividend growth rate and the discount rate are both 5%, the NPV of the dividend is the same each year. This is logical because every time you step forward a year the discount rate knocks an additional 5% off the NPV but the growth in dividends adds that 5% back on. The sum of the NPV of these dividends over ten years would be £10k, over 20 years would be £20k, and over an infinite period would be an infinite sum. A similar argument applies if the dividend growth rate is greater than the discount rate, except that the sum now reaches infinity somewhat faster (a mathematically dodgy concept, but you know what I mean).

The idea of a 'mathematical instability' makes it sound like the problem lies in the maths but in fact the model is entirely accurate – it is the input assumptions that are dodgy. Jesus' silver penny really would have grown to be worth £472 trillion today if the Pharisees could have found a bank which paid 3% compound interest every year for 2,000 years. The idiocy is to assume that such a bank could stay in business over such a long period. Endowment mortgages really would have made people much richer if they could borrow against their homes at an average of 5% interest, paid annually for 25 years, and then invest the money for up to 25 years at a compound net rate of return of 7%. That 7% would, of course, have been paid only after the deduction of some cumulatively eye-watering management fees, so let's assume 9% gross returns would be necessary to make this work.

The problem with endowment mortgages was not in the maths, but in believing that such high returns could be sustained over such a long period. Most of the investment was in equity – owning the means of wealth creation. If the growth in an economy averages 3% per year, then the value of owning the means of wealth creation presumably increases at more or less the same rate. Over 25 years, a 3% compound growth would double the value of the economy whereas 9% growth would require the value of owning the means of wealth creation to increase just over eight-fold. OK, endowment mortgages did not actually involve borrowing money against the house and putting it straight away into the stockmarket – the money was used to buy the house and then additional sums over and above the interest payments, which would traditionally have been used to pay down the principal, were instead invested in the market, but the general fallacy remains the same.

3.5. Valuing the equity of high growth companies

We have just highlighted the idiocy of assuming that all shares can continue to grow their earnings by 'super normal' rates for indefinite periods. However, some companies will clearly grow their earnings very rapidly for fairly sustained periods. This is particularly true of technology companies riding a wave of changing behaviour – one only has to think of Microsoft, Google, Apple, or Facebook. These companies typically start life with no earnings – they make a loss (which they refer to as 'burning cash') while they invest their shareholders' money, preparing their product offering so it is ready for the wave to rise (or, perhaps, so that it can create the wave on which it will then ride).

Price-earnings (P/E) ratios are clearly of no use in valuing these high growth companies. If the earnings are taken to be zero then assigning any price whatsoever to the shares implies an infinite P/E ratio (another mathematical instability), whereas plugging in negative earnings would create a negative share value – presumably implying that the company has to pay people to take ownership of its shares. In this situation, the 'quick and dirty' shorthand DCF calculation implied in a P/E ratio has to be replaced by working out the NPV manually. During the 'dot com' bubble, carrying out such calculations turned out to be very lucrative for a number of star equity analysts. It was not always so lucrative for investors who bought the shares they recommended.

Let us imagine a typical technology stock that currently has enough money in the bank to complete its journey to profitability. To use the parlance, it can "cover its cash burn up to the point where it becomes cash flow positive" (this is somehow more reassuring than the thought that you are paying good money for a loss-making company). Three years from now it makes a profit for the first time, and pays a modest dividend of £1 per share. As it rides a wave of technological innovation it is able to double this dividend every year for the next five years. It then experiences three years of transitional growth rates as its market matures – lifting the dividend by 70%, 40% and 10% in successive years. Thereafter, it is fully mature and only able to grow its dividend in line with wider economic growth – say at 3% per year (the 'terminal growth rate'). The NPV of these dividends at a 5% discount rate is shown in the table below (note that the first column of the table is Year 3 and not Year 1).

Table 6: NPV of dividends from a hypothetical technology start-up company

Item	Year 3	Year 4	Year 5	Year 6	Year 7	Year 8	Year 9	Year 10	Year 11	Year 12
Dividend payment	£1.00	£2.00	£4.00	£8.00	£16	£32	£54	£76	£84	£86
NPV of dividend at 5% discount rate	£0.86	£1.65	£3.13	£5.97	£11	£22	£35	£47	£49	£48

The total value of the first ten dividends (paid in Years 3-12) is £223. For a company which was delivering no growth in its dividends, the payments received over the first 12 years would account for around half of the NPV of all payments likely to be received through to infinity. However, for this technology company, the sustained period during which dividends grew much faster than the discount rate means that the most valuable dividend in the NPV series is that paid in Year 11. Year 12 is the first year of 3% growth and as this is lower than the 5% discount rate, it is the first year that the NPV of the dividends starts to fall. Because we have only continued our summation for one year beyond the peak of the NPV'ed dividend stream, we have ignored a considerable amount of value.

We can attempt to capture this additional value in two ways. The first is to extend the NPV model for a further period of about 40 years (referring back to Table 3 tells us it takes 47 years for a 5% discount rate to reduce the NPV to 10% of its nominal value). Alternatively, we can say that in Year 12 our technology company is just like any other mature company and can be valued on a P/E basis. If we assume that a P/E of 15 is appropriate, then its value in Year 12 will be equal to 15 x £86 = £1290. This has an NPV of £721. So a better estimate of the current value of shares in our currently loss-making but soon to be profitable technology company would be £944, being the NPV of the dividends paid in Years 3-12 (£223) plus the NPV of the likely value of the mature company share in Year 12 (£721).

An explicit NPV model has ridden to our rescue when seeking to value a company that is not amenable to the implicit NPV treatment underlying a P/E valuation. While the historical earnings of a company can be known fairly accurately (always assuming there are no skeletons in the accounting closet), and these are often a reasonable guide to future earnings over the medium term, our NPV is now based on some pretty major assumptions, which may be little better than guesses. For example, if the company can lift it earnings by only 50% during its five-year growth spurt, then the NPV of the shares becomes £225. But if the company achieves a 100% growth spurt and sustains it for six years instead of five, the NPV of the shares would be £1,803. If it has a six-year growth spurt and then subsides to a 10% terminal growth rate rather than 3%, its value is theoretically infinite (because 10% is greater than the discount rate). However, investors will reckon on this falling back after a few more years and so rather than write a cheque for infinity, their response may be to put the shares on a P/E of 30 instead of 15 in Year 12 – this would give them an NPV of £3,360.

Our estimate of £944 for the value of our technology company shares sounded reassuringly accurate with its three significant figures, but in fact some very minor flexing of our assumptions shows the value might lie in the range from <£250 to >£2,500. Furthermore, technology start-ups tend to be either spectacularly successful or very unsuccessful – generally they do not deliver slightly different versions of average. So if we really stress-tested our model to include a more realistic range of outcomes we could probably widen our valuation range to be anything between £100 and £10,000 – two orders of magnitude rather than three significant figures. Has the NPV model really ridden to our rescue, or has it burnished what is effectively a complete guess with a gloss of apparent precision? Even if I stick with £944 as my central case scenario (sounds better than ‘best guess’), at the very least I will want to take into account the risk that this is wrong. In the next chapter we consider how DCF modelling offers to do this, and whether it would be wise to accept that offer.

4. Including high levels of risk in the discount rate

4.1. The exponentials used in DCF analysis amplify errors in the input assumptions

The preceding chapters have shown that complexities are introduced into DCF modelling not by mathematically complex relationships between well-known inputs, but by mathematically simple relationships between uncertain inputs. There is a well-known principle in modelling of ‘GIGO’ (Garbage In, Garbage Out). The nature of exponents means that for DCF it should perhaps be amended to ‘GIMAGO’ (Garbage In, Massively Amplified Garbage Out).

This amplification effect means that if inappropriate input assumptions are used in a DCF model it can produce catastrophically wrong conclusions. This could lead us to build a nuclear power station, take out an endowment mortgage, put our entire pension into long dated gilts during a period of quantitative easing or buy shares in dot-com companies - when really we should not (or indeed, stop us from doing so when really we should). Other UK-specific examples which we do not have the time to discuss here could include funding hospitals through the PFI, education through the Student Loans Company, or energy-saving home improvements through the Green Deal. Indeed the ‘deal’ implicit in the Green Deal could effectively be expressed as “we can assure you that the NPV of the work you are having done to your house under this government scheme will always be positive”.

The use of inappropriate input assumptions can arise in two ways – deliberate or accidental. The wide range in the possible outputs from a DCF model offers scope for deliberate abuse – the modeller can adjust the inputs to obtain an output that is most expedient to whatever proposition they are trying to sell. The proposition may be commercial, political or an unholy nexus of the two. The expediency is often to leave you unaware of the risks you are taking, or to misprice that risk. The only way to guard against being bamboozled by deliberately wrong input assumptions is to understand DCF at least as well as your counterparty and force them to explain their proposition in detail. This essay may or may not put you in a better position to do so. Where the risk is that the wrong input assumptions have been used accidentally (‘inadvertently’ might be a better word), then DCF modelling itself, repeated more rigorously, can be used to mitigate that risk.

4.2. Relationship between risk and rate of return

It was suggested in Section 1.1 that the price offered to Rupert for his trees might depend on the risk appetite of the buyer. If the buyer thought the risk was higher than that estimated by Rupert then he would use a higher discount rate, driving down the NPV of the trees. In Section 2.3, we revisited Rupert’s situation and speculated that he might have a different risk appetite when investing the proceeds from selling his trees than that of Duncan who was saving up for a car. Being quite wealthy, Rupert would be happy to play with his spare cash (or ‘hobby money’) by putting it into a higher risk and higher reward investment. So we have already been introduced to the well-known principle that higher risk lending or investing needs to be compensated by higher rates of return. Before going any further, we should remind ourselves of why this is and quantify the relationship between interest rate and risk.

Interest rates vary over time, but at any given moment for a particular currency and economy, there should be a ‘risk free rate’. This is the return on an investment which is 100% secure (or 0% risk). The benchmark for such investments is often taken to be a government bond, if only because governments can always repay the principal by printing more money. Let's assume the risk-free rate is 3%. If I invest £1,000 for a year, I can expect (with a very high degree of certainty) to have £1,030 at the end of that period.

Now suppose that any given company has a 10% chance of going bankrupt in any given year. If I were to invest £100 in the shares of each of ten companies, then at the end of a year I can expect nine of those companies still to be in business and one to have gone bankrupt, wiping out my investment. I would be crazy to invest in equities if I *expected* to be left with less money than investing in government bonds. Therefore, my *expectation* (albeit without complete certainty) should be that the shares in my nine remaining companies are worth £1,030 – or £114.44 per company. I had started by investing £100 in each, so the appropriate rate of return on their shares to compensate for the added risk is 14.44%. The 'equity risk premium' over the 3% risk free rate is thus 11.44%.

Calculated below are the theoretical interest rates required to lend at different levels of default risk when the risk-free rate is 3%. Of course, in reality it may not be a binary choice between receiving your entire investment back with interest or else losing the lot.

Table 7: Theoretical interest rates at different levels of default risk

Annual risk of default	Implied Interest Rate
Zero	3.00%
1%	4.04%
3%	6.19%
5%	8.42%
7%	10.75%
10%	14.44%
15%	21.18%
20%	28.75%
30%	47.14%
50%	106.00%

There are two important points to note from the above. The first is that relatively modest levels of risk require levels of interest which would, in practice, 'squeeze the borrower until the pips squeak'. The second is that the mitigation of risk by seeking higher levels of interest only protects the lender if they take a portfolio approach. By investing £100 in the shares of ten companies each having a 10% annual risk of default and all paying me a 14.44% yield, I can reasonably expect to be left with £1,030 at the end of the year. This is by no means the only possible outcome, but it is more probable than any other. But if I invest £1,000 in one company paying £14.44% interest, then it is actually impossible for me to be left with £1,030: I will end up either with nothing or with £1,144.40, with the latter outcome being by far the more likely (i.e. having a 90% probability of happening).

So anyone seeking to mitigate risk needs to take a portfolio approach – the more investments in the portfolio, the more certainty of the outcome - unless the fate of those investments is somehow correlated. 'Portfolio theory' suggests that investors in equities should hold 20-50 stocks in order to smooth out their 'unsystematic' (i.e. company specific) risk. Fund managers offer to choose these stocks for you, and in return cream off a slice of your profits. If you decide to forego the costs of portfolio investment and switch from investing £1,000 in a single risk-free investment to investing £1,000 in a single risky investment, you should be much more concerned about choosing a lucky company than in squeezing out a couple of extra percentage points of risk premium. Sophisticated risk analysis, such as correlations within the portfolio or value-at-risk, is beyond the scope of this essay but adjusting the discount rate to incorporate risk is a fundamental start.

4.3. Reflecting high levels of risk in discount rates

Having quantified the relationship between interest rate and risk, how should this read-across to the relationship between discount rate and risk? In the case of our technology company, assuming we were looking for 5% growth in our money, and that the company delivered on its promise for the next 12 years, we calculated that we should value the shares today at £944. But technology companies have a habit of not delivering on their promises. Let's model that by saying that there is a 50% chance that the company will go bankrupt in the first year after I have invested, but that if it survives this period, it will go on to deliver its plan. In other words, there is a 50% chance of default in the first year.

I could mitigate this by buying a portfolio of ten such companies and paying £472 for each – costing me the same as buying five companies at £944. I can expect five of the companies to fail, but the remaining five will deliver their business plans and return to me a flow of dividends with a NPV at 5% discount rate (the discount rate I apply to money I fully expect to receive) of 5 x £944. In order for the flow of dividends set out in Table 6 to have a NPV of £472 (the price I have just agreed that I am prepared to pay), the discount rate needs to be increased from 5% to 11.6% - which is therefore the discount rate I should apply to money I have doubts about receiving. So when people bring me their technology companies and ask me to invest, I will apply an 11.6% discount rate to their projected cash flows.

Like all the examples used in this essay, the above is a gross over-simplification. For a start, it does not model two equivalent situations. In the case of lending to a portfolio of borrowers with a 50% risk of default we charge 106% interest and expect (probably) to turn £100 into £103 after one year (as per Table 7). At the end of this period we are free to invest the money in whatever else we like. In the case of investing in a portfolio of technology companies with a 50% risk of default in the first year and zero risk for the 11 years thereafter we expect (again, probably) to turn £100 into a stream of dividends which would have an equivalent value to putting the £100 into a bank account paying 5% interest and drawing it out progressively over a 12-year period. In this case we have to consider the 'illiquidity risk' of tying up our money in an investment which we cannot easily 'realise' or 'liquidate'. It is also obvious that the risk profile of real technology companies is not concentrated into the first year of their existence, after which they become as safe an investment as a government bond. This residual default risk adds a further layer of anxiety to the illiquidity risk.

Despite being over-simplified, the example given does serve to illustrate an important point. Because the returns from exponential growth accelerate with time, a relatively modest discount rate can, if left long enough to do its stuff, compensate for relatively high levels of risk. We saw for the technology company that at a 50% risk of default the interest rate would be 106% and the discount rate only 11.6%. In fact, 50% risk of default is actually modest for a start-up company. If it was 90%, then I would need to apply an interest rate of 930% but a discount rate of 28.8%. Clearly the interest rate is absurd – who would borrow £100 knowing that the interest would reach £100 after just 39 days? However, issuing shares which might one day provide an investor with a 28.8% compound rate of return does not seem to be a particularly 'big ask' for a high growth company. As risks mount, so interest rates become an ever less appropriate way to deal with them and debt becomes ever less appropriate than equity as a means of financing the enterprise.

4.4. High risk projects should be funded by equity rather than debt

Not only are high interest rates an inappropriate way to deal with significant levels of risk, they actually have a tendency to amplify the risks they are trying to guard against. Let's consider a major tunnel built by a private company at a cost of £10bn. It is debt financed at 5% interest which requires servicing by annual payments of £800m if calculated on a 20-year repayment mortgage basis. The company's modellers assumed that the gross profit from tolls would amount to £1bn and this would be ample to service the debt. However, over the first couple of years of operation the income settles at level of ca. £800m per year. The company is perilously close to becoming incapable of servicing its debt – the risk of default is clearly rising. So the creditors decide they must raise the interest rate to 6% to compensate for this. At 6% interest the cost of servicing the debt increases to £870m per year. The company now has real problems servicing its debt and rolls up some of the interest in the hope that it can grow its revenues. The risk of default rises further and so therefore must the interest rate.

There is clearly a feedback loop operating which means that the tunnel operator which might or might not have had problems servicing its debt had the interest rate remained at 5% will soon be driven into bankruptcy. The risk anxieties of the lenders have turned into a self-fulfilling prophecy. If the tunnel constitutes a critical part of the national infrastructure, it may have to be rescued by the government. Could this have been avoided? Well, if the government had guaranteed the debt of the tunnel operator as soon as it became clear that the initial revenue projections were over-optimistic, but still viable, then arguably the interest rate should have dropped to the risk free rate of, say, 3%. This would reduce annual interest payments to £670m which could readily be serviced from £800m of revenues. The risk of default has almost disappeared and so the government guarantee is never likely to be called. The company makes a profit and pays tax to the government rather than additional interest to a bank. Another self-fulfilling prophecy and certainly a less disruptive one. Lenders ramping up interest rates can be like a driver breaking the speed limit to mitigate the risk of arriving late and thereby creating the risk of a crash which will mean not arriving at all.

So it appears that debt funding might be appropriate if interest rate risk premiums are limited to about 5% which, for a risk free rate of 3%, corresponds to an annual risk of default of about 4.5%. When the risk levels are greater than this, equity should be used to finance some or all of the enterprise. Evaluating the impact of different interest rates and debt to equity ratios is a good use of DCF modelling. For high risks, the enterprise will need to be funded entirely by equity and it will be essential (remembering the second key point made in Section 4.2) to take a portfolio approach. This is precisely the model of a venture capital (VC) fund, which invests in a portfolio of (usually) 10 to 30 very risky companies. Most of them will, like tadpoles, fail to grow up but hopefully those which do will more than repay the sum invested in the entire portfolio. In other words, in love you have to kiss a lot of frogs to get a prince but in finance you need to fund a lot of tadpoles to get a frog.

By investing in a portfolio of companies, the VC fund acts as a risk smoother for investors who do not have the opportunity or the inclination to evaluate hundreds of different opportunities and select the best 20 or so in which to make an investment. If it can be shown by a combination of modelling and track record that the portfolio is likely to deliver, say, a 10% IRR even though most of its constituent companies will fail, then the VC fund ought to be able to borrow at, say, 5% interest rate, thereby leveraging the returns to its equity investors. The fund is therefore not only smoothing the returns, but decoupling the high risk of failure in individual companies from the exorbitant interest rates which would only serve to promote that failure (VC funds do not tend to be leveraged only because most of them do not have a track record of delivering stable returns). The money raised by the VC fund is invested in the portfolio companies entirely in the form of equity, consistent with the level of risk.

4.5. The pitfalls of using high discount rates to model high risks

Since the VC fund manager is addressing risk via discount rates on equity returns rather than interest rates on debt, it seems logical that they should use DCF modelling when selecting the most promising investment opportunities for inclusion in their portfolio. To see why this might not be such a good idea, let's revisit for the final time our technology start-up company. The table below repeats the DCF modelling presented in Section 3.5 using a 5% discount rate and then shows what the NPV figures would have been if a discount rate of 30% had been used (calculated above as the sort of figure that would be appropriate for a failure probability of around 90%).

Table 8: NPV of dividends from a hypothetical technology company at different discount rates

Item	Year 3	Year 4	Year 5	Year 6	Year 7	Year 8	Year 9	Year 10	Year 11	Year 12
Dividend payment	£1.00	£2.00	£4.00	£8.00	£16.00	£32.00	£54.00	£76.00	£84.00	£86.00
NPV of dividend at 5% discount rate	£0.86	£1.65	£3.13	£5.97	£11.37	£21.66	£35.07	£46.76	£48.98	£48.05
NPV of dividend at 30% discount rate	£0.46	£0.70	£1.08	£1.66	£2.55	£3.92	£5.13	£5.52	£4.67	£3.70

If you recall, the fair value of the company was treated (at least in one version of the analysis) as the NPV of the dividends from Years 3-12 plus the NPV of the terminal value calculated at 15 times the Year 12 earnings. The two different discount rates lead to the following valuations:

- 5% rate: £223 for stream of dividends plus £721 for terminal value = £944
- 30% rate: £29 for stream of dividends plus £56 for terminal value = £85

Now it is entirely reasonable that the higher discount rate should have such a dramatic effect on the value of the shares. Most of the tadpoles will die and those that grow up have to yield returns that more than compensate for the money invested in their dead brothers and sisters. Most technology entrepreneurs accept that they will have to sell their birth right at a huge discount. They are often pathetically grateful to find anyone interested in investing in early stage technology rather than inclined to argue over the price.

The key issue is rather that the higher discount rate changes the relative importance of the different contributions to the total NPV. At a 5% discount rate the dividend stream accounts for 24% of the NPV while at a 30% discount rate it accounts for 34%. The Year 9 dividend by itself accounts for 6% of the 30%-NPV but only 3.7% of the 5%-NPV. If the fund manager is looking to invest in companies with a clearly positive NPV (which is presumably the only reason to carry out NPV analysis in the first place), then they might be tempted to pay more attention to the Year 9 dividend payment than to the Year 12 growth rate. However, in those all-important companies which survive through to maturity, the assignment of a P/E of 30 rather than 15 when calculating the terminal value (on the basis that they are still growing quite rapidly in Year 12) will in reality be much more important than the value of any interim dividends. This is, in effect, an example of Goodhart's law that 'when a measure becomes a target, it ceases to be a good measure'. If VC fund managers spend too much time calculating individual NPVs, the risk is that they will start to target them rather than to focus on overall portfolio returns over the medium term.

There is another reason why VC fund managers should not become too fixated on NPVs: they do not model the actual return on the equity investment made. The 'fair value' of a share is, as we noted in Section 3.4, the NPV of all its future dividends stretching out to eternity. The 'market value' is, of course, whatever you can sell it for on the day. VC funds do not hold shares for the rest of eternity: they buy them at a cheap price while there is still substantial technology and execution risk, and hope to sell them a few years later at an expensive price when that risk has dropped away (or at least when the public markets can be persuaded to believe that it has).

Notwithstanding the synopsis put forward above when modelling our pet technology company, in almost all cases a VC investor will sell its stake in a company ('exit the investment') before it has received a single dividend. Microsoft, founded in 1975 and floated on NASDAQ in 1986, did not start to pay dividends on its shares until 2003. The key driver of returns for a VC fund manager is the percentage of companies in the portfolio which can be sold successfully – either to a corporate buyer or to public investors via an Initial Public Offering (IPO) – and the size of the windfall achieved at the point of exit.

4.6. Alternative valuation models – lottery tickets and intuition

In view of the potential pitfalls described above, a more realistic approach to valuation may be to view early stage investments as lottery tickets. To describe something as a 'lottery' may sound pejorative, implying that any profits made are simply a matter of luck. However, lotteries are particularly amenable to mathematical analysis. If I pay £1 for a lottery ticket and have a 1% chance of winning £10 then I am a fool. If it gives me a 1% chance of winning £10,000 then I would be a fool not to buy it, and as many other tickets for the same lottery as I can get my hands on. The fair value of a lottery ticket is the size of the prize times the probability of winning it. A good VC fund manager who has spent years assessing technology investments should develop an intuitive sense for the value a company might have on exit and the chances of achieving a successful exit. In other words, each investment can be valued as a lottery ticket.

Different companies will require different amounts of investment, will expose that money to different levels of risk, will pay out different sizes of prize, and will require the investment to be tied up illiquidly for different periods of time. And the best estimates for these parameters made in advance will vary by different amounts from what actually happens in practice. This is ideal territory for the Monte Carlo type of probabilistic modelling mentioned in Section 3.2. DCF analysis is not particularly useful for deciding the size and payout date of the various prizes, but balancing the exposure of the fund to the different lotteries on offer may well be a good use of it. It will also be essential when choosing between a smaller prize won sooner and a larger prize won later.

There is a saying that “if the only tool you have is a hammer, then every problem will look like a nail”. Notwithstanding the potential benefits of applying probabilistic modelling to the analysis of early stage investments, the ‘old familiar’ DCF is still much more popular. An extreme example of the contortions undertaken to use DCF for the purposes of start-up company valuation came with the various internet incubators that sprung up in the last days of the dot-com bubble. These were pooled investment vehicles designed to offer exposure to the growth of dot-com hot-stocks during the period prior to IPO. They were often run by young, untried fund managers – perhaps on the basis that only such people could truly understand young, untried companies. Their unique selling point was that they would provide not just capital, but also company nurturing (‘incubation’) services.

Incubator stocks with trendy names proved very popular with investors and their shares were often valued at multiples of the cash raised before they had made a single investment. This was hard to justify on a DCF basis, as not only were there no obvious cash flows to discount, there were not yet any companies to generate such cash flows. It was therefore argued by some that the genius of the investment manager would very likely secure an attractive return on the capital, say 40%, and that the risk of this not happening could be captured by a modest discount rate, say 10%. Now a pound which grows for seven years at 40% becomes worth £10.54, and if this sum is discounted back at 10% it has a NPV of £5.41. So obviously (not) it is reasonable for the yet-to-be-invested pound in the incubator fund to be worth £5.41 today. While it may or may not be true that “a pound in the hands of an internet wunderkind is worth more than a pound in my pocket”, or that (more plausibly) “a pound in the hands of Warren Buffett is worth more than a pound in my inexpert hands”, these are instinctive judgements and DCF brings little to their quantification, though it may highlight some apparent contradictions.

The first conclusion from this chapter is that small adjustments to risk can be captured appropriately by small adjustments to the discount rate in a DCF model. This had anyway already been covered to some extent when discussing ‘choosing the right rate’ in Chapter 2. The sort of business situations modelled in this way will be predominantly debt-financed. The second conclusion is that large risks must be predominantly equity-financed and that these are best modelled using a probabilistic approach, particularly if the timescales are fairly long and the investments are illiquid. DCF calculations may well be useful to provide inputs to these probabilistic models, but they are not suitable for the whole job.

5. One size fits all discount rate

5.1. A universal discount rate does not necessarily ensure a level playing field

Hopefully we have seen that while DCF modelling is good for some things, it is not good for everything. And in those situations where it *is* the most appropriate tool, it can never offer particularly high precision. Even when approaching a DCF analysis with the best of intentions (and many practitioners don't), there will always remain a risk of making input errors which become grotesquely amplified by the power of exponential growth (remember GIMAGO!). Analysts who take a 'single club' approach, reaching for DCF as the answer to all problems, are dangerous. Even more dangerous are those who always use DCF and, when doing so, always apply the same discount rate. In this final chapter we will look at why this might be so.

Ironically, those with good intentions may be more prone than rogues to this error of applying a monotonous discount rate. The latter relish the ability to adjust discount rates to suit their nefarious purposes. The former may see the use of a single universal discount rate as the fairest thing to do – thereby ensuring a level playing field. It was noted in Section 2.5 that companies may feel that this approach is the only way to identify the best investment opportunities across their competing geographical and sectorial operations. Governments obviously feel obliged to treat all proposals fairly so as not to be accused of massaging the figures. And yet Section 2.5 suggested that a single discount rate, divorced from actually achievable rates of return, may be an unfair method of comparing a project with significant end-of-life costs (the nuclear power station) and a project with most of its costs up-front (e.g. a wind farm).

A single headline discount rate may, in fact, provide a false sense of security with regard to fairness and objectivity. Factors such as inflation may influence the outcome more significantly than the use of different discount rates and if these are not treated with equal objectivity, the end result may nonetheless be distorted either deliberately or accidentally. In the case of power generation, one of the most important factors in any comparison between nuclear power and renewables is the assumed deflation in the cost of PV solar panels and energy storage systems during the multiple decades over which the nuclear power station will be operating. If solar energy can be generated for 5p/kWh and stored for a further 5p/kWh anytime soon, then the index-linked electricity strike prices guaranteed for some nuclear power plants may look wildly over-generous on a DCF basis.

5.2. Management by discount rate is lazy management

In the corporate sector there is a 'managerialist' school of thought that you don't need any detailed domain knowledge to run a business – it is possible for a senior executive to move from a food company to a car company and do an equally good job. This is often coupled to the belief that the key to success in any business boils down to minimising the Weighted Average Cost of Capital ('WACC') while maximising the Return on Capital Employed ('ROCE'). Perhaps this is why 'MBA' is sometimes taken to mean 'Management by Acronym'. The cost of capital has to be 'weighted' because it is assumed that there will be different costs for the debt and equity components. The problems of assigning a future-looking cost of equity based on historical trends in the stockmarket could be the subject of an entire new essay. The role of a CEO who subscribes to this school of thought sounds somewhat similar to that of the old school bankers we met in Section 2.5. They call in the finance team and tell them to drive down the WACC by 3%, then call in the operations team and tell them to increase the ROCE by 3%, and then presumably head off to the golf course (or fly to Davos) at 3pm.

In Section 4.5 we came across Goodhart's law that 'when a measure becomes a target, it ceases to be a good measure'. ROCE is a useful measure, but when it becomes a target problems start. The temptation is to drive up ROCE not by doing everything a little bit better, but by weeding out operations with a low perceived return on capital. A similar situation arose in UK schools where the focus switched to the percent of pupils passing their exams: this can be increased by coaching marginal candidates, but also by not allowing them to sit the exam. If senior management are gunning for operations with low ROCE, then they will also be keen to demonstrate that they are operating a level playing field. This means that they are likely to set a 'one size fits all' hurdle rate for the ROCE on existing activities and a similar universal hurdle rate for the IRR on proposed new activities. This will favour mature sectors using fully-depreciated assets: it will lead to the manager keeping his television tube plant open but closing his flat screen R&D project.

Superficially, a single discount rate for all and sundry sounds fair and reasonable. It can even work well for a company with a single business activity – and one which is amenable to modelling. For example, a property company looking at the rental yields from commercial property, student halls of residence and residential letting, and comparing new building with the acquisition of existing properties, may well use a single hurdle rate to identify the best opportunities. But the devil is usually found in the detail. How does a supermarket compare the ROCE from selling clothing and selling food in the same store? How is the tied up capital (e.g. in the brand) allocated between the two functions and how are operational costs shared out when calculating the return on that capital? Any attempt to divide up costs will be fraught because the people best able to make the calculation will all have a very significant vested interest in its outcome. Often this means bringing in management consultants who have no vested interest in the outcome (but a strong vested interest in the process being as complex and protracted as possible).

Things can get worse when the time comes to use the ROCE data to make operational decisions. To take an absurdist example, imagine that a separate ROCE was calculated for every step in a car production line and it was found that the highest return (or Economic Value Added, EVA) came from putting the wheels on right at the end. The pneumatic nut-tightener turns out to be much less expensive than a welding robot (less capital employed) and over the period the analysis was carried out by the consultants, the guy who operated it was hugely experienced and so worked very fast (he has since retired and been replaced by a youngster who is both lazy and incompetent, but data will always lag reality).

In view of this, the CEO may decide that it would make sense to restructure the business. All upstream portions of the production line can be sold to other car makers and the money invested in a huge wheel-attaching factory. The promising young guy who currently puts on wheels and has replaced the more expensive older chap with his defined benefit pension, can be promoted to run the entire complex. All other car makers will cease to put on their own wheels because they cannot match the economies of scale of our huge dedicated factory. They will clearly play their allotted parts in this plan (i.e. will acquire from us the facilities we wish to divest and then procure from us the services we have decided it suits us to sell). The acronyms are all aligned for us: what can possibly go wrong?

At least the car company is only in one business and the entire production line is in the same location (so the same taxes, regulatory environment, wages, utility costs, and currency zone). Many companies will be trying to assign ROCE figures to operations in different business sectors and different geographies. Most businesses experience some sort of cycle and the attractiveness of countries can wax and wane as currencies fluctuate and politics evolve. A CEO who closes any operation which has had a disappointing ROCE over the previous five years and keeps open any which has had an impressive ROCE is no different from an equity fund manager who buys any stock on the market which has gone up by 50% in value and sells any stock in the portfolio if it goes down by 50%.

It is obvious that sometimes it will be right to reinforce success and sometimes it will be right to invest ahead of recovery. Sometimes it will be right to sell at the top of the market; sometimes it will be right to exit from a perennial under-performer. Management is about judgement and an acronym-assisted look-up table cannot safely replace that judgement. If all companies end up being run by people using the same analytical framework, then they will all move in lockstep. They will all add capacity at the same time in sectors which have historical supply shortages and have therefore recently been very profitable, for example fibre optic networks in the late 1990s, and they will all close capacity at the same time in sectors which have recently suffered from over-supply and low margins, for example photovoltaic solar panel manufacturing a few years ago.

5.3. A universal discount rate ignores risk

Perhaps the most complex chapter of this essay was Chapter 4 on risk, and risk clearly affects the expected return on capital. Setting a single hurdle rate for ROCE will make middle managers favour more risky projects, particularly if the specified rate is a ‘stretch goal’ (i.e. it is much higher than the prevailing risk-free rate). This is because they have a similar incentive environment to the one which was so corrosive in the financial services sector – if they succeed they get a big bonus but when they fail they lose someone else’s capital. The astute manager will set a target in terms of Risk Adjusted Return on Capital (RAROC). However, this is another idea which sounds good in principle but is difficult to implement in practice. Once again, the people best able to carry out the risk quantification needed in order to perform the calculation will all have a very significant vested interest in its outcome. All of the above arguments about retrospective ROCE or RAROC will apply equally well to the prospective IRR on planned future projects. Except that uncertainty about the future is even more pronounced than uncertainty about the past. Those who have dreamed up a project, or will be promoted to run it, will be both best placed to calculate its IRR and most interested in massaging this to meet the hurdle rate.

Worse than those who always use DCF, and than those who use a single discount rate, are CEOs who set ‘chest beating’ ROCE or IRR targets to impress financial analysts - what might be called a ‘Big Hairy Audacious Goal’ of bean counting. Apart from being a hostage to fortune, this simply tells your competitors how much they have to cut prices over the short term in order to precipitate your departure from the market. At the very least, it can lead to special dividend payments on the grounds that “there are no opportunities for us to put this capital to work” (i.e. there are currently no 10%+ IRR opportunities). So the money is given back to investors who have to pay some of it in tax and put the rest in a savings account at 1% interest, when they would have been quite happy to leave it in equity at 5% return. In some cases the tax position can be ameliorated by a share buy-back programme rather than a special dividend. This should have the effect of driving up the share price which may help to trigger the CEO's bonus, but that would, of course, be entirely serendipitous.

Management by acronym is poor management – it is an abuse of the trust placed in appointing the manager and the huge salary paid to them for exercising their superlative judgement. Many of the favoured acronyms have buried within them a DCF analysis. Management by acronym therefore often represents another abuse of discount rates, and is the final one we will discuss in this essay.

5.4. Conclusions

DCF modelling cannot make the answer anything you want it to be, but can all too often make it anything you want it to be within one or two orders of magnitude. This provides plenty of scope for abuse – be it intentional or accidental. Getting the wrong rate accidentally can have a huge impact on the outcome. Remember the one contribution we may have made to your acronym vocabulary – GIMAGO! This in turn means that it is often possible to get a desired outcome by deliberately choosing a discount rate that supports that outcome.

Accidental errors are more accurately a misuse of discount rates rather than an abuse. Sometimes they give an answer which is both wrong and, *ex post facto*, turns out to have been convenient for one party in a transaction. This can lead to accusations of foul play, though incompetence is much more common than conspiracy. However, whenever someone starts with the answer and works backwards to find the inputs needed to support that answer, then you are always justified in suspecting abuse. But we must not forget that DCF analysis can also be a blessing when used appropriately. DCF is a simple but elegant concept which can, for example, be used to compare two different finance packages via annual percentage rates (APRs) or to illustrate the effect of management fees on the eventual size of a pension pot.

Like priests, guardians of the mystical truths of discounting can use the special insights with which they have been blessed either to serve their fellow human beings, or to bamboozle the rest of humanity into serving them. Like priests, there is a mix of good and bad practitioners. It is, perhaps, interesting that the use of APRs on credit agreements and fee deduction modelling on endowment policies were in both cases forced on the financial services sector by the regulators, when they might more honourably have been offered proactively as obvious best practice. Bankers and their DCF acolytes were trusted to practise a certain amount of self-regulation, and this may have been a mistake. Thus, like every tool ever invented by humans, from the flint knife onwards, discount rates can be used or abused.

Acknowledgements

Both the genesis of this essay and its realisation owe a huge debt of gratitude (not to be discounted!) to Professor Michael Mainelli of Long Finance and Z/Yen. Several of the ideas presented were suggested by Michael over one particularly nice lunch at the Guildhall – others emerged from his meticulous and insightful editing of the draft. I am also grateful to Jan-Peter Onstwedder, Chiara von Gunten, and Arjuna Sittampalam, of Long Finance, and Tim Mercer, Head of Economics at the Ashcombe School Dorking, for reading the manuscript. Finally, I would like to thank my son, Alex Goddard, for checking the calculations.

About the author

Dr Nick Goddard qualified as both a scientist and an engineer followed by ten years of experience in the oil & gas, power generation, aerospace and defence industries. During this period Nick migrated from his original role as a technologist at BP into strategic planning, marketing and business development, his final post being within the Ministry of Defence. Leaving government, he worked for eight years in the City as an investment banker, exiting with inadvertently good timing in 2005. This was followed by a further eight years of independent consultancy advising corporates and government agencies on technology commercialisation. During this period he occasionally got his hands dirty by joining the technology companies he was advising as an interim manager. In 2013 he found an engineering start-up company in the energy sector which he liked so much that he joined it full time, his career thereby completing a full circle. Rightly or wrongly he believes that his share options have an attractive NPV and hopes to live long enough to benefit from this.

Nick received a BA in Natural Sciences (Physics) from Jesus College, Cambridge, and a PhD in Materials Engineering from Imperial College, London. He is a Chartered Physicist and a Chartered Engineer.

About Long Finance

Long Finance aims to improve society's understanding and use of finance over the long term through four programmes, research, publications and events. The initiative began with a question – “when would we know our financial system is working?” – and seeks to challenge a financial system that revolves around short-term thinking and practices. Long Finance is a community which can be explored and joined at http://www.longfinance.net

LONG
FINANCE

www.longfinance.net