IN-DEPTH

Investment Treaty Arbitration

THE DISCOUNTED CASH FLOW METHOD OF VALUING DAMAGES IN ARBITRATION

LEXOLOGY

Investment Treaty Arbitration

EDITION 10

Contributing Editor

Barton Legum

Honlet Legum Arbitration

In-Depth: Investment Treaty Arbitration (formerly The Investment Treaty Arbitration Review) fulfils an essential function. Updated every year, it provides a current perspective on a quickly evolving topic. Organised by topic rather than by jurisdiction, it allows readers to access rapidly not only the most recent developments on a given subject – from jurisdictional and procedural issues to damages and much more – but also the debate that led to and the context behind those developments.

Generated: June 30, 2025

The Discounted Cash Flow Method of Valuing Damages in Arbitration

Ronnie Barnes
Cornerstone Research

Summary

Introduction

In the chapter 'Choosing the Appropriate Valuation Approach for Damages Assessment' in last year's edition of this title,[2] the authors note that while '[t]here are four general approaches to valuing an asset. . ., [t]he [discounted cash flow (DCF)] approach is the fundamental basis of valuation. . .' and go on to explain: 'There are two major components of a DCF analysis: (1) projected free cash flow (estimates of the cash available to debt and equity holders after all cash expenses of the business have been met); and (2) the discount rate (a measure of the risk of the projected cash flows, with higher rates for riskier assets and lower rates for less risky assets).' The objective of the current chapter is to expand on this high-level description of the DCF valuation methodology and, hopefully, equip the reader with a sufficient grasp of certain key issues and concepts relating to the methodology so that – for example – when faced with an expert's valuation report, they are able to identify the questions they should be asking.

The basic premises that underpin the DCF methodology are straightforward. First, the value of any investment – be it a company, a project or an individual asset – at a given point in time is determined solely by the future cash flows that the investment is expected to generate from that point on. Second, the sooner a cash flow occurs, the more valuable it is. For example, $100 today is worth more than $100 a year from now, which in turn is worth more than $100 two years from now. Third, the less risky a cash flow is, the more valuable it is. For example, $100 a year from now is worth more if it is 'relatively safe' than if it is 'relatively risky'. However, operationalising these basic premises in such a way that a DCF valuation leads to a robust estimate of the value of the investment in question may be far from straightforward – for example, what exactly do we mean when we refer to 'the future cash flows that the investment is expected to generate', how do we measure the 'risk' of a future cash flow, and how do we translate this measurement of risk into a measurement of value? The remainder of the chapter is aimed at providing some answers to these questions.

Discounting risk-free cash flows

We begin our analysis with the simplest possible example, that of the valuation of a single 'risk-free' cash flow. While in common English usage, the word 'risk' typically has negative connotations, in the valuation context, the meaning is somewhat different. Here, risk is essentially synonymous with uncertainty – for example, a future cash flow that could be either $100 or $50 with equal likelihood is risky (since there is uncertainty as to what the actual cash flow will be), while a future cash flow that will with certainty be equal to $100 is risk-free (since there is no such uncertainty).

By means of a concrete example, suppose that today is 1 July 2025, and a bank offers you a 4 per cent interest rate on a one-year deposit subject to a government guarantee. If you deposit $100 today, on 30 June 2026, the amount in the account will have increased, with no risk or uncertainty, to $100 × 1.04 = $104. Similarly, if you could convince the bank of your creditworthiness, you would be able to borrow $100 today and repay $104 a year later. As a matter of terminology, $104 is said to be the *future value* of $100 one year from today, while 4 per cent is the *one-year risk-free (interest) rate*.[3] Another piece of terminology that is important for the following discussion is *return* – for an investment that does not generate any intermediate cash flows, the return over a given period is defined as

(end of period value - start of period value) / (start of period value) [1]

so that, in our example, the return from the bank deposit is (unsurprisingly)

($104 - $100) / ($100) = 4%

Now suppose that you know that you will receive, again with no risk, $800 on 30 June 2026. However, you have a pressing need for cash today and decide to borrow from the bank as much as you possibly can, subject to the limit that however much you borrow, together with interest on the borrowed amount, must be able to be repaid out of the $800. Given that borrowing $P from the bank will lead to you repaying $P × 1.04 one year from now, the amount you will be able to borrow is the solution to the equation $P × 1.04 = $800, that is, $P = $800/1.04 = $769.23.[4] This amount – $769.23 – is said to be the *present value* or *discounted value* of $800 one year from today and may be thought of as the amount of cash today that is 'equivalent' to $800 in a year's time. Symbolically, we may write this as

PV[X] = X/(1 + r1)

where X is the future cash flows, PV[X] is its present value, and r*1* is the interest rate available on a one-year bank deposit.

While this is a simple numerical example, it does provide a road map for valuing future cash flows that will be useful as we tackle more complex examples:

1. determine the risk of the cash flow you are seeking to value (in this case, risk-free);
2. determine the return that is available to you from an investment with the same level of risk (in this case, the interest rate on a risk-free bank deposit); and
3. use this return to calculate the present value of the cash flow.

We now tackle the question of how to calculate the present value of a future cash flow that occurs other than one year from today. Suppose again that today is 1 July 2025, and the same bank offers a 6 per cent interest rate on a two-year deposit, that is, the two-year risk-free rate is 6 per cent. If you deposit $100 today, on 30 June 2026, the amount in the account will have increased, with no risk or uncertainty, to $100 × 1.06 = $106, while by 30 June 2027, this amount will have increased further to $106 × 1.06 = $112.40.[5] Note that this is equal to $100 × 1.06*2*, which enables us to generalise as follows: if the t-year risk-free rate is r*t*, then the present value of a risk-free cash flow of X to be received or paid t years from today is

PV[X] = X/(1+ rt)t.[6]

To calculate the present value of a series of cash flows that will occur at different future points in time, simply calculate the present value of each individual cash flow and add these up. So, for example, if the one-year and two-year risk-free rates are 10 and 12 per cent respectively, the present value of $100 to be received one year from now and $500 to be received two years from now is equal to

$100/(1.10) + $500/(1.12)*2* = $489.51.

While the bank deposit example is a convenient means of introducing the concepts in this section, in practice, the source of the information needed to determine risk-free rates is the government bond market. The basic idea is that information regarding the risk-free rates that are appropriate for discounting future cash flows that are denominated in a specific currency can be extracted from the prices at which investors are currently buying and selling bonds issues by a country that has control over the supply of that currency – the assumption being that a government will never default on bonds that it issues in its own currency, that is, the bonds are genuinely risk-free.

Introducing risk

So we now have a framework for determining the present value of a series of future risk-free cash flows denominated in a particular currency – simply calculate the present value of each individual cash flow and add these up. The same approach works when the future cash flows are no longer risk-free, although the way in which the individual present values are calculated does change. To motivate the discussion, consider the following extremely stylised example. Suppose that an investor is offered an investment that, in one year's time, will pay either (1) $111 with 88 per cent probability or (2) $11 with 12 per cent probability. Assuming a one-year risk-free rate of 10 per cent, how much would the investor be willing to pay today in order to acquire the investment? Equivalently, what is the present value of the future (risky or uncertain) cash flow that the investment will generate?

One possible solution might be the following. Although we do not know exactly what the cash flow in a year's time will be, we can calculate what is referred to as its expected value – this is nothing more than a probability-weighted average of the possible values of the cash flow, that is

0.88 × $111 + 0.12 × $11 = $99

Having calculated this expected value, we might then be tempted to discount this at the one-year risk-free rate of 10 per cent to arrive at a present value of $90 (since $99/1.10 = $90).

However, this fails to take account of the fact that individuals are *risk-averse* – that is, they would prefer to receive a risk-free cash flow with a guaranteed value of $99 rather than a risky cash flow with an expected value of $99. The implication of this is that the investment will be valued at something lower than $90. To determine this lower value, we continue to discount the expected cash flow of $99, but use a discount rate that is *higher* than the risk-free rate:

present value = (expected future cash flow) / (1 + risk-free rate + risk premium) [2]

Before considering how we might assess the risk premium that an investor will require from this investment, it is useful to consider the question of the return that the investment offers. Because we do not know what the actual cash flow in a year's time will be, we similarly do not know what the return over the year will be. But we do know what the *expected* return is – this is (analogous to [1]) defined as

(expected future cash flow – initial price paid) / (initial price paid) [3]

Clearly, the lower the initial price paid, the higher the expected return will be. For example, if the investor pays $75, the expected return is ($99 - $75)/$75 = 32 per cent, whereas with an initial price paid of $70, the expected return increases to ($99 - $70)/$70 = 41.4 per cent. Critical for the discussion that follows is the fact that if the investor pays an initial price that is equal to the investment's present value, then combining [2] and [3] yields:

expected return = risk-free rate + risk premium

In other words, if the amount invested is equal to the investment's present value, the investor's expected return is equal to the risk-free rate (compensation for the time value of money) plus a risk premium that compensates the investor – no more and no less – for the risk that they are assuming in making the investment. Put simply, the higher the risk that an investment carries, the higher the expected return and risk premium that an investor will need to be offered in order to make the investment an attractive one and the lower the present value of the investment will be. An obvious, but important, corollary is that given two investments with the same risk, an investor will require the same risk premium, that is, will use the same discount rate when determining the present values of the investments.

The obvious question is how the appropriate risk premium should be quantified. If we are willing to take on faith the proposition that the riskier a particular investment is (holding fixed the expected cash flow), the less attractive and valuable it will be to investors – that is, the higher the risk premium they will use in their analysis – then it becomes clear that what we need is (1) a way of measuring risk and (2) a means of converting this measurement of risk into an appropriate risk premium. Consequently, we now address the question of how risk should be measured for the purpose of assessing the risk premium to be incorporated into a discount rate.

While a detailed discussion of this question is beyond the scope of this chapter, the basic idea can be summarised as follows. The uncertainty in the future cash flows from a risky investment can essentially be decomposed into two elements. The first – referred to as *systematic risk* – is that which is related to ('correlated with') the overall state of the economy, while the second – referred to as *unsystematic risk* – is that which is unrelated to the overall economy. Investors are compensated (in the form of a risk premium) for bearing the former, but not for bearing the latter – that is, risk premia, expected returns and discount rates are a function of systematic, but not unsystematic, risk.

To see the relevance of this point, consider the following not atypical argument that an expert may advance:

> *I am trying to value a particular project. I am unable to measure the risk of the project directly, but I have identified a number of companies that are in the same industry as the project and that I believe have broadly the same operating characteristics – the average level of risk across these companies is 'X', which translates into a discount rate of (say) 12 per cent. However, I believe that there are certain risks that the project is exposed to that these 'comparable' companies are not, and so, rather than valuing the project using a discount rate of 12 per cent, I am going to use – based on my experience and subjective judgment – 15 per cent to reflect these additional risks.*

When faced with an argument of this nature, it is critical that to ask the question 'Are these additional risks systematic in nature?' In other words, do they lead to the future cash flows of the project that is being valued being more highly correlated with the overall economy than the future cash flows of the comparable companies are? If the answer is yes – these additional risks are systematic in nature – then using a higher discount rate may be justified (although this would at least potentially suggest that the set of comparable companies might have been chosen better). It is much more likely, however, that the answer is no and that the additional risks are specific to the project. In this case, it is contrary to the basic principles of financial economics that underpin the DCF methodology to increase the discount rate to reflect these risks – the additional 3 per cent in the example above is nothing more than a 'fudge factor', and – as explained in one of the pre-eminent corporate finance textbook – this approach is fundamentally misguided:

> *Avoid Fudge Factors in Discount Rates . . . Fudge factors in discount rates are dangerous because they displace clear thinking about future cash flows.*[7]

This remains the case even if the proponent of increasing the discount rate labels the increase not as a 'fudge factor' but as a 'project-specific risk premium'. This is a contradiction in terms – as discussed above, project-specific risks do not command a risk premium.[8]

The capital asset pricing model (CAPM)

We now have a framework for valuing a project from which the future cash flows are risky or uncertain: (1) for each year in the future, estimate the cash flow that the project is expected-[9] to generate in that year; (2) calculate the present value of this expected future cash flow as of today (the 'valuation date') by discounting it a rate that is equal to the risk-free rate plus a risk premium that serves as compensation for systematic risk; and (3) sum the resulting present values. But how exactly do we determine what that risk premium (equivalently, the discount rate) should be? In this section, we introduce what might be referred to as the 'workhorse' model for determining discount rates, namely the capital asset pricing model (CAPM). We start by presenting the equation that defines the model, explaining exactly what each term means and then go on to consider how it is typically implemented.[10]

The equation defining the model is as follows:

E[r] = rf + ß × EMRP

where *E[r]* is the expected return on a given stock over the next year, *rf* is the one-year risk-free rate, *ß* is the 'beta' of the stock, and *EMRP* is the 'expected market risk premium'

As above (see equation [3]), the expected return on a stock over the next year is defined (assuming that the stock is not expected to pay any dividend) as

(expected end of year value – current value) / (current value)[11]

and the CAPM posits that this is equal to the risk-free rate (extracted from the prices of default-free government bonds) plus a risk premium which is equal to the product of the stock's beta and the expected market risk premium. In slightly loose terms, the beta for a given stock is a quantitative measure of the extent to which returns on the stock are (1) correlated with returns on the aggregate stock market (as proxied, for example, by the S&P 500 index) and (2) more or less volatile than returns on the aggregate stock market. Previously, we introduced the general concept of systematic risk (correlation with the overall economy) and explained why the risk premium component of an expected return should be a function of the systematic risk, rather than the total risk, of an investment. The CAPM simply makes this more concrete by telling us that beta is the appropriate measure of systematic risk for a stock. But how do we actually calculate the beta for a particular stock? Skipping over the technical details, we simply take the actual returns on the stock over some time period and actual returns on the aggregate stock market over the same time period and apply to these a statistical technique known as *regression analysis*. Experts may – and will often – disagree regarding questions such as the appropriate time period, the appropriate data frequency (do we use returns measured over daily, weekly or monthly intervals?), and even the appropriate proxy for the aggregate stock market, but there is likely to be relatively little disagreement over the approach of using regression analysis in this way.

That leaves the question of the expected market risk premium, and how to quantify this is one of the most controversial questions among finance academics and practitioners. It is easy to define from a conceptual point of view – it is simply the expected return (over the next year) on the aggregate stock market in excess of the current one-year risk-free rate. But how should this be estimated? It is impossible to provide anything approaching a meaningful discussion of this question in the confines of this chapter, but suffice it to say that there are three broad categories of approach[12] and that results can differ dramatically depending on the approach chosen and, even within a given approach, the details of how it is implemented – for example, an analysis of historical excess returns on the aggregate US stock market is extremely sensitive to the time period chosen. Given that, by definition, there is no right answer (at least *ex ante*) to this question, we would suggest that an expert looking to determine the expected market risk premium as an input to the CAPM should look to support their choice by reference to a range of contemporaneous evidence. Statements such as 'the expected market risk premium, based on single source X, is Y per cent' should be viewed with a certain amount of caution.

Let us now return to the question of project valuation. Armed with our projections of the future cash flows that the project is expected to generate,[13] we next identify a set of companies whose stock is publicly traded and that are as closely comparable to the project being valued as possible.[14] Using regression analysis, we then estimate the beta for each of these comparable companies and calculate the average beta. Finally, we input this average estimated beta, along with an estimate of the expected market risk premium and the current risk-free rate,[15] into the CAPM, the result of which is the discount rate that should be used to discount the project's expected future cash flows.[16]

Bringing debt into the equation

In the previous section, we assumed that all projects and companies are entirely financed by equity. This conveniently simplified the discussion. For example, there has been no need to distinguish between the cash flows generated by a project and the cash flows that accrue to the holders of the equity in the project – these are one and the same thing. Similarly, the value of the project and the value of the equity in the project are entirely synonymous. However, once we bring debt into the equation – as indeed we must if we are to value real-world projects and companies that are typically financed by a mix of debt and equity – things become somewhat more complicated. At this point in standard corporate finance textbooks, a plethora of formulae would start to abound, dealing with the quantitative relationships between, for example, 'levered' and 'unlevered' betas, and the 'unlevered cost of capital' and the 'weighted average cost of capital'. Given the objectives of this chapter, we restrict ourselves to discussing – at a high level – various key issues that are essential for a conceptual understanding of how to value projects and companies that are partially financed by debt.

The first such issue is that the value of a project that is partially financed by debt is higher than the value of the same project if it is entirely financed by equity. The reason for this is straightforward, namely the tax deductibility of debt interest. To illustrate this point, consider a project that is expected next year to generate operating profits of $500 and that is subject to a corporate tax rate of 20 per cent, so that the expected tax payment is $100. If we assume for simplicity that profits and cash flow in this example are the same, then if the project is entirely financed by equity, the cash flow that is expected to accrue to the providers of the project's financial capital[17] (in this case, the holders of its equity) is simply $500 - $100 = $400. Now suppose that the project is partially financed by debt – specifically, next year $200 of interest is expected to be paid to the holders of this debt. In this case, the project's expected taxable profit is $500 - $200 = $300 (because interest on debt is tax-deductible), and the expected tax payment is 20 per cent × $300 = $60. Consequently, the cash flow that is expected to accrue to the holders of the project's equity is $300 - $60 = $240, which together with the expected debt interest of $200 yields a total cash flow that is expected to accrue to investors of $240 + $200 = $440, that is, $40 higher than when the project is all equity financed. The additional $40 is referred to as the 'debt tax shield' and is simply equal to 20 per cent (the corporate tax rate) × $200 (the expected debt interest). To understand the implications for the valuation of the project, note that we can summarise the previous calculation as follows:

expected cash flow to investors = expected cash flow to investors if project is all-equity financed + debt tax shield

In value terms, this becomes:

value of project = value of project if all-equity financed + value of debt tax shield

This is the 'adjusted present value' (APV) approach to valuation. In this approach, we first calculate the value of the project assuming that it is financed entirely by equity, then adjust this value to reflect the additional value that is created through the debt tax shield.

To tackle the first piece of this – the value of the all-equity project – we essentially use the framework set out above, but with one important modification. This arises from the fact that because debt is senior to equity – that is, the debt holders have the first claim on the cash flows generated by a firm or a project – the equity in the firm or project is riskier than it would be in the absence of debt. Consequently, if we use regression analysis to estimate the beta for a given company, and that company has debt in its capital structure, what we are actually estimating is what is referred to as the 'levered' equity beta. This may be thought of as the beta that we would have estimated had the firm been all-equity financed (its 'unlevered' equity beta), plus 'something' to reflect the additional risk to equity created by the debt financing. Thus, if we want to get at the unlevered equity beta, we need to take this 'something' out of the beta that we have estimated, a process that is referred to as 'unlevering'.

The implication for the 'playbook' is as follows. We are attempting to estimate the project's unlevered equity beta, so that we can input this, inter alia, into the CAPM in order to derive the project's unlevered cost of capital, which we can then use to calculate the unlevered value of the project. Economic reasoning tells us that 'similar' firms and projects should have 'similar' unlevered equity betas – however, two identical firms will have different levered equity betas if they have different financing mixes, and it is levered equity betas that we are able to estimate using regression analysis. Consequently, the discount rate components of the playbook now read as follows:

1. identify a set of companies whose stock is publicly traded and that are as closely comparable to the project being valued as possible;
2. using regression analysis, estimate the levered equity beta for each of these comparable companies;
3. 'unlever' each of these estimated equity betas (using information on the relevant comparable firm's financing mix), then calculate the average unlevered equity beta; and
4. plug this average estimated unlevered equity beta, along with an estimate of the expected market risk premium and the current risk-free rate, into the CAPM, in order to derive the project's unlevered cost of capital.

Having calculated the project's all-equity or unlevered value, we now need to calculate the value of the debt tax shield by discounting each year's expected tax saving at a discount rate that reflects the risk in those tax savings. Typically, this discount rate will be either (1) the project's unlevered cost of capital (as calculated above) or (2) the project's cost of debt (as discussed below), although this choice typically has relatively little impact on the overall valuation.

Note that we are valuing the project as a whole, which is equal (tautologically) to the aggregate value of the securities (debt and equity) that are used to finance the project. Often we are interested in the value of the equity in isolation. As we discuss below, valuing the debt is typically relatively straightforward, so we can determine the value of the equity by using the APV approach to value the debt and equity together and then subtracting the value of the debt.

As the above discussion hopefully makes clear, the APV approach to valuation has the advantage of clearly and separately identifying the additional value that is created by having the project partially financed by debt. An alternative approach – which for reasons that will shortly become apparent is referred to as the weighted average cost of capital (WACC) approach – addresses this additional value through an adjustment to the discount rate. Specifically, in the example above, using the WACC approach, we would continue to discount the $400 of cash flow that is expected to accrue to investors in the all-equity scenario but, rather than discounting at the project's unlevered cost of capital, we would discount at a lower, adjusted rate.

At the outset, it is important to note – and this is something that is often overlooked – that this alternative approach may only validly be used under certain restrictive conditions. Broadly speaking, these conditions may be characterised as relating to the stability of the project's financing structure. Specifically, if the project is to be financed with a constant debt (or 'leverage') ratio,[18] then using a lower, adjusted discount rate – which can be obtained formulaically from the unlevered cost of capital – is possible. Otherwise, the APV approach has to be used. For example, if a 10-year project has $100 million of debt outstanding as of the valuation date and this is expected to be paid off in equal instalments of $10 million, adjusting the discount rate will simply not work.[19]

However, if these conditions are met, then it can be shown that the adjusted discount rate is equal to the project's WACC, defined as

$WACC = (1 - L)rE + LrD\ (1 - T)$

where L is the project's leverage ratio (the percentage of debt in its financing mix), rE is the project's cost of equity, rD is the project's cost of debt in turn, and T is the project's marginal tax rate.[20]

Let us now discuss the cost of equity and the cost of debt in turn. Starting with the cost of equity, this is the answer to the question 'what is the appropriate rate at which to discount the expected cash flows to equity holders, properly taking account of the risk of these cash flows?' Importantly, the answer to this question is not the unlevered cost of capital. As we discussed earlier, the unlevered cost of capital is equal to the cost of equity only when the project is all-equity-financed. Once we introduce debt into the financing mix, the equity becomes riskier and the cost of equity therefore increases, since the equity holders demand a higher expected return for bearing this additional risk – the more debt there is, the riskier the equity is, and the higher the cost of equity will be. Consequently, the playbook needs further modification. Specifically, having calculated an average unlevered equity beta across the comparable firms as an estimate of the project's unlevered equity beta, this needs 'relevering' in order to arrive at an estimate of the project's levered equity beta, with the details of the relevering depending on how much debt is in the project's financing mix – all other things equal, the more debt, the higher the levered equity beta. This can then be input, inter alia, into the CAPM to generate an estimate of the project's cost of equity.[21]

Insofar as the cost of debt is concerned, this is often measured by reference to the coupon rate on the debt.[22] However, a word of caution is in order. The coupon rate is set when the debt is issued and therefore reflects economic conditions at that time, including the general level of interest rates in the economy along with then perceived risk of default on this debt. In other words, the coupon rate represents the cost of debt at the time of issue and may be thought of as the expected return that investors in the debt demand in return for the risk they are assuming. As economic conditions change, the cost of debt will change – for example, if the perceived risk of default has increased, so will the cost of debt. Again, a detailed explanation of how to take these factors into account is beyond the scope of this chapter, but it is important to be aware that they do need to be taken account of.

Endnotes

1 Ronnie Barnes is a vice president in Cornerstone Research's London office, head of the firm's international arbitration and litigation practice, and co-head of its valuation, M&A and bankruptcy practice. The views expressed in this chapter are solely those of the author, who is responsible for the content, and do not necessarily represent the views of Cornerstone Research.

2 Jessica Resch, Maja Glowka and Tim Giles, 'Choosing the Appropriate Valuation Approach for Damages Assessment' in Barton Legum (ed.), *In-Depth: Investment Treaty Arbitration*, 9th edn., Law Business Research Ltd, 2024.

3 The government guarantee ensures that there is no possibility that the bank will default on its obligations, thereby justifying the assumption that the deposit really is a risk-free

investment. To avoid complicating matters unnecessarily, we skip over any discussion of exactly how you would convince the bank of your creditworthiness.

4 For simplicity, we ignore the fact that in practice, the bank is likely to charge you a spread over the risk-free interest rate to cover your perceived (lack of) creditworthiness.

5 This assumes that interest is compounded once a year.

6 Note that the *subscript* t in these equations denotes that r_t is the t year risk-free rate, while the *superscript* t denotes that $(1 + r_t)$ should be raised to the power t.

7 Richard A Brealey et al., *Principles of Corporate Finance* 241 (13th edn., 2019).

8 As a concrete example, suppose that the project being valued is heavily dependent upon a single customer for its revenues – the loss of this customer would have a catastrophic impact on the project. It seems obvious that the project is less valuable than it would be if it had a diversified customer base (as do the comparable companies that are being used for benchmarking the risk of the project), so surely it is appropriate to increase the discount rate? The short answer is no – unless the risk of losing the customer is systematic in nature (which seems unlikely), there is no justification for such an increase. That is not to say that the possibility of losing the customer should be ignored – clearly it should not. Rather, it should be dealt with (as the textbook quote above suggests) when estimating the expected future cash flows from the project. We do not dispute that determining exactly how to adjust these expected future cash flows to reflect the possibility of losing the customer is difficult. Our point is that increasing the discount rate by an arbitrary amount does not obviate this difficulty, it simply sweeps it under the carpet.

9 Here, as throughout this chapter, the term 'expected' should be interpreted to mean a probability-weighted average across all possible scenarios (so in the example in footnote 7, this would include a scenario in which the customer is lost). In practice, this averaging will be done implicitly, i.e. the expected value will be forecast directly – this is fine, providing that what is forecast is in fact an average, rather than a 'best-case' or 'aspirational' scenario.

10 To simplify matters, for now we assume that all projects and companies are financed entirely by equity, that is, there is no debt financing.

11 For a stock that is expected to pay a dividend, the definition of expected return is *(expected end of year value + expected dividend – current value) / (current value)*.

12 These are (1) historical analysis, which looks at the average actual returns (in excess of the then-prevailing risk-free rate) on the aggregate stock market over some past period; (2) survey analyses; and (3) the 'implied' approach, whereby the current level of the aggregate stock market, together with assumptions regarding expected future dividends, are used to infer the discount rate that market participants are implicitly using.

13 This is not to downplay the fact that these projections will often be the subject of fierce debate – issues that often arise include questions as to the extent to which is it valid to use contemporaneous management projections, in particular whether such projections genuinely represent 'expected values' or rather are 'upside' projections that are 'aspirational' in nature.

14 While this is somewhat subjective, a typical starting point will be to identify companies within the 'same industry'.

15 In principle, given that risk-free interest rates differ by maturity (meaning that the one-year risk-free rate is different from the two-year risk-free rate, and so on), we could look to apply a different discount rate to the expected cash flows at different future points in time. In practice, this is rarely, if ever, done. Typically, an expert will use a risk-free rate with a maturity that corresponds to the 'average' maturity of the cash flows being valued as an input to the CAPM. Exactly what is meant by 'average' may require some subjectivity, but in today's economic environment where – at least for developed market currencies – risk-free rates show relatively little variation by maturity, whatever choice is made will likely not have much in the way of an impact on the overall valuation. In other economic environments, where interest rates differ significantly by maturity, there is much more of an argument for using different discount rates (at least insofar as the risk-free component is concerned) for different future cash flows. However, even if the decision is made to use different risk-free rates, the beta for the project will still be estimated in exactly the same way. Finally, it would be rare indeed to use anything other than a single estimate of the expected market risk premium.

16 As a matter of terminology, this discount rate is referred to as the project's cost of capital or, because all of the project's financial capital is in the form of equity, its cost of equity (although strictly speaking, as we explain in the next section, this discount rate is actually the project's *unlevered* cost of capital or *unlevered* cost of equity – a project or company that is all-equity financed is said to be 'unlevered', whereas if it is financed by a mix of debt and equity, it is said to be 'levered').

17 For ease of exposition, we shall henceforth refer to the providers of the project's financial capital simply as 'investors'.

18 This means that the amount of debt issued against the project is adjusted over the life of the project so that it is always a fixed percentage of the overall value of the project.

19 Ignoring the details, this is because the formula that relates the adjusted discount rate to the unlevered cost of capital depends on how much debt is in the project's financing mix. Essentially, if the project's financing mix is expected to change over the life of the project, the formula essentially breaks down.

20 The word 'marginal' means that this is the rate at which each additional dollar of taxable income is taxed – typically, this will be taken to be equal to statutory corporate rate.

21 Note that applying this cost of equity to the cash flows that are expected to accrue to the project's equity holders allows us to calculate the value of the equity directly, rather than as the difference between the project as a whole less the value of the debt.

22 The coupon rate is what is applied to the face, or notional, value of the debt in order to calculate the periodic interest payments to the debt holders. For example, a coupon rate of 8 per cent means that the holder of $100 face value of debt will receive an interest payment of $8.

Ronnie Barnes rbarnes@cornerstone.com

Cornerstone Research