4 Setting the social discount rate

What does the foregoing theory, together with the available empirical evidence, say about the appropriate discount rate to use? What is the appropriate way to adjust for taxes, foreign capital flows and risk?

4.1 Adjusting for taxation and foreign capital flows

Accounting for the effects of personal taxes and foreign investment flows reduces the social discount rate by (at most) about 1 percentage point. For example, it is plausible that the before-tax risk-free real rate of return (on government bonds) is around 4 per cent and the after-tax risk free rate of return around 1 per cent.[1] Even if 30 per cent of a project's capital came from sources that earned the after-tax return, the weighted-average riskless rate would be above 3 per cent.[2] Yet the before-tax rate is the appropriate forgone return for not only for investment, but also much consumption and foreign borrowing. The before-tax risk-free return in the private sector may be higher than the bond rate because of property and other taxes. Further, the marginal cost of foreign investment lies above the average cost, and so the weighted average rate is likely to lie closer to the before-tax return.[3]

4.2 Adjusting for risk

It follows from the foregoing section on risk that governments should only discount with the risk free return if either:

- the project is risk free, or
- the market is able to spread all the risk associated with the project, or
- the government spreads all risk so that the project does not impose risk on beneficiaries and taxpayers, or
- the expected values of cost and benefit flows have been converted into 'certainty equivalents' — see box 4.1.

[1] See appendix G.

[2] 0.7*4 + 0.3*1 = 3.1.

[3] See appendix G.

Box 4.1 Using certainty equivalents instead of a risk premium in the discount rate

The certainty equivalent of a risky payment is the smallest certain payment which the household would accept in exchange for the risky one. When a risky payment increases overall consumption risk for a risk averse household, its certainty equivalent is less than the expected value of the payment. Some analysts reject adding a risk premium to the discount rate and endorse converting expected values of benefits and costs into certainty equivalents to account for risk and then discounting with the risk-free rate.[4]

For example, Abelson describes the certainty equivalent approach as 'theoretically attractive', but impractical, and adding a risk premium as practical, but 'crude'.[5] Yet, as shown in appendix H 'Asset pricing', in general the two approaches give exactly the same answer, expressed in different ways. The discount rate approach is more transparent, easier to compare with market data and more practical. When they differ, the certainty equivalent approach is the theoretically correct approach; it is more flexible and makes it easier to make adjustments needed to reflect departures from the standard model. But it is complex to apply, and lack of relevant information often makes it totally impractical. Using a risk-adjusted discount rate is practical, and gives close to correct answers.

The more important distinction is not whether a certainty equivalent or risk premium approach is undertaken, but whether the calculation is based on market data or directly calculated from assumptions about preferences and perfect markets.[6] As pointed out earlier, using the direct calculation approach gives a negligible risk premium, far below the market price of risk. Using the same assumptions in a certainty equivalent approach would also give a much smaller risk correction than using market parameters.

Source: See appendix H.

Using the riskless rate when these conditions do not hold ignores the cost of the risk the project imposes on the private sector and makes it likely the taxpayer will be saddled with undesirable risks.

What adjustments need to be made to account for risk?

---

[4] For example, Boardman et al. (2006, p. 238; footnote 3, p. 271) and Abelson (2000, pp. 130-31). Bureau of Transport and Regional Economics (2005, p. 30) recommend it when 'the welfare of a small number of individuals affected by the project varies greatly across the states of nature'.

[5] Abelson (2000, pp. 130-31). Bureau of Transport Economics (1999, p. 77) also describes adding a risk premium as 'a crude allowance for risk' and 'second-best' and the certainty equivalent approach as 'rigorous' but 'problematic' (pp. 74-76).

[6] As done in Bureau of Transport and Regional Economics (2005, p. 55).

Adding a risk premium to the risk-free weighted average rate

One way to adjust for risk is to calculate the risk-free weighted average discount rate and then add an appropriate risk premium which accounts for the cost of risk imposed by the project. The before-tax market price of risk is the appropriate social price of risk.[7]

But it is problematic to calculate the parameters in that equation. The theory has run far ahead of our empirical knowledge.

The information to calculate the quantity of risk for government projects is seldom available. For example, government projects do not have the historical market data used to calculate market beta risks of private projects, much less that needed to calculate the influence of other risk factors. Further, the liquidity risk of government projects is difficult to judge.

Moreover, it is difficult to calculate an appropriate price of risk. The risk premium is the extra return needed to compensate investors for expected risks and cannot be observed, as expected returns cannot be measured. The usual approach is to use the past realised price of risk, along with the assumption that expectations are correct on average to calculate the price of risk in the past.

But such assumptions are problematic. The enormous variability in realised risk premiums makes it difficult to pin down a precise estimate of the average (their standard deviation is around 20 percentage points). More years of data helps (some studies go back over 100 years), but the average risk premium may not have been constant over these long time periods. For example, the risk premium in the 1890s may not tell us much about today's market.

The high variability in returns means the average return is very sensitive to the starting and finishing dates of the measurement period. For example, whether estimates include or exclude the 1970s, when realised returns were negative, makes a significant difference to the average. The 1970s experience also demonstrates that expectations can be wrong for long periods of time.

The risk premium is not constant over the business cycle, and so

> standard cost-of-capital calculations featuring the CAPM and a steady 6% market premium need to be rewritten, at least recognizing the dramatic variation of the initial

---

[7] See appendix G. For example, if the CAPM held the appropriate social discount rate for asset k is $w_k = w_f + \gamma\beta_k$ where γ is the social price of risk, β_k is the quantity of market risk in asset k, $w_f = ai_f + (1-a)r_f$, i_f is the before-tax risk free rate of return, r_f is the after-tax risk free rate of return, a is the proportion of government investment sourced from investment.

premium, and more deeply recognizing likely changes in that premium over the lifespan of a project and the multiple pricing factors that predictability implies.[8]

Another issue is the so-called 'peso problem': where investors take into account some probability of a rare disaster when forming expectations. If the disaster is not observed in the sample, realised returns seem inexplicably high, but they are compensating investors for the risk of rare disasters that happen not to have materialized.[9] To complicate matters, the expected probability of disasters may change over time.

Some authors conclude that a substantial part of the long term 6 to 8 per cent equity premium was luck, such as avoiding a world war or depression for the past 60 years and that the expected probability of these events has fallen. If the true risk premium falls, that would boost share prices and increase the realised share returns — meaning the measured equity premium overstates the expected future premium. They believe the true average expected equity premium is closer to 4 per cent, a figure the earnings and dividend growth based estimates support.[10]

The risk premium is not constant over the business cycle, but even the average premium over the cycle may have changed systematically over time, because of (for example) declining transactions costs of trading in shares, changes in the proportion of the population owning shares, changes in market risk, changes in companies' average gearing ratios and greater international capital flows. Also the relationship between bond and share market returns could change over time with changes in tax, such as the introduction of imputation and capital gains taxes, changes in tax rates, and with changes in expected inflation.[11] Further, the effects of these changes are controversial.[12]

Another complication often ignored in the risk free rate plus risk premium approach is the effects of taxation. Estimates of the risk premium invariably use the after-company tax return to equity and the before-personal tax bond return. Yet the company and property taxes that firms pay affect the opportunity cost of forgone investment. These taxes could be responsible for at least part of the gap between the before-all taxes return on capital and the bond return. The before-tax risk premium

8 Cochrane (undated).

9 Weitzman (2007, p. 1104). See appendix H.

10 See appendix G.

11 Gordon and Gould (1984) find this with Canadian data. Merton (1980) finds that risk has changed substantially over time.

12 For example, see Lally (2002) for a review of different approaches to the effect of imputation on estimates of the risk premium. Jones (2008a) disputes the common practice of approximating the economic effects of taxes on realised capital gains with accrual based taxes set at lower rates.

companies need to earn on their investment will tend to be the after-tax premium grossed up by the company tax rate. When nominal returns are taxed and losses are not fully deductible, the gap between the before- and after-tax risk premium would be even greater. To the extent that companies pay property and other taxes on their risk-free returns, they would have to earn more than the government bond rate on risk-free investments. The effects of taxes on the risk premium are both complex and controversial.[13]

Calculating the expected real risk free rate of return from nominal bond rates depends on expected future inflation. Again, expectations cannot be directly observed and the usual approach is to infer the expected real return from realised real returns, assuming that expectations are unbiased so that average realised inflation equals expected inflation.

The taxation of nominal receipts means that even if the underlying after-tax real rate of return were constant, the expected before-tax real return would not be stable, and would increase with expected inflation and the tax rate. Average realised real returns would depend on the historical pattern of inflationary expectations and the structure of taxes.[14] As there is no guarantee that future patterns will be the same, that makes it difficult to predict future rates, especially for the long term. On the other hand, the variability in realised risk free real interest rates is low (its standard deviation is around 3.5 per cent), reducing the impact on the social discount rate.

Using the marginal return on capital

Given all these problems, the better approach is to start with the marginal rate of return on capital, such as the national accounts measures of the before-all tax real rate of return on private capital. It averages 8.9 per cent over long time periods and is more stable than share market returns (see section 3.4 'The marginal return to capital').[15] The measured real marginal return to capital is understated because the CPI used to deflate the nominal figures overstates inflation.

The marginal rate of return to capital should be adjusted to reflect the impact of tax distortions and foreign borrowing. That reduces the rate of return by around 1 percentage point to around 8 per cent.

[13] See appendix G for details.

[14] See appendix A, section A.3.

[15] The discount rate derived from the marginal return to capital is consistent with that derived from the risk free rate plus market risk premium approach. See appendix G.

The resulting weighted average market rate includes the market risk premium. It is appropriate for a government project which has the same risk as the average private sector investment. The rate reflects the opportunity cost – the funds used in the government project would have produced this return if left in the private sector, with no greater risk cost.

In the absence of information on the quantity of risk in a government project, it is reasonable to assume that the average government project is no less risky than the average private investment (in terms of covariance with aggregate consumption). The consensus view is that most government projects are highly correlated with returns to the economy as a whole. If society receives an 8 per cent rate of return on the average private sector project, they would also require at least that rate of return on the average government project.

The presumption in favour of the market risk premium can be varied if there is a clear argument to the contrary, such as evidence that the amount of market risk in project flows is likely to be low. For example, some projects may offer insurance benefits.

If the risk the government project imposes differs from the average private sector project, then the risk premium should be adjusted. If, for example, the government can successfully spread some diversifiable risks that the market cannot, then the discount rate should be lowered to reflect the lower risks imposed.

Conversely, there may be reasons why the risk premium should be higher for government projects. For example, government projects impose high liquidity costs and governments may do a worse job of spreading risks than the private sector. If so, using the market weighted average discount rate is conservative.

As pointed out above, choosing how much to adjust the discount rate to reflect lower or higher risk is likely to be difficult. The best approach is to do the adjustment together with sensitivity testing (see below).

Likewise, the adjustment for taxes could be varied if there is strong evidence that project funds come from sources with low rates of return. If the costs of a government project mainly came at the expense of consumption, then the discount rate would be closer to the consumption rate. But as many consumers have a high consumption rate, that would not necessarily lower the discount rate much.

In the absence of further information, 8 per cent is a reasonable default discount rate, but there is real uncertainty about the appropriate discount rate to use. The weights and returns to use in the weighted average discount rate, or the appropriate risk premium are not clear. The use of sensitivity analysis is recommended as an important tool to identify the relevance of uncertainty over the value of the discount rate.

Sensitivity testing

Accounting for risk means that discount rates should be project specific. The uncertainty, or variability, of flows, and how the amount of market risk differs across projects affects how the recipients value them. A single discount rate for all projects would not reflect their varying degrees of risk. Likewise, accounting for taxes would make rates differ across projects if capital sourcing differed. Because the appropriate discount rate varies across projects, depends on factors analysts often know little about, and for any particular project is uncertain, analysts should conduct sensitivity testing with the discount rate.

For the reasons summarised in table 4.1, the net present values should be calculated with the discount rates 3, 8 and 10 per cent real — representing the weighted average riskless rate of return, the weighted average rate of return and a rate of return for a riskier asset or that reflects the marginal productivity of capital during the 2000s. The effect on present value of reducing the discount rate by 1 percentage point is larger than increasing it by 1 percentage point.[16]

If the sensitivity analysis reveals that the choice of discount rate is important (changes the sign of the project's net present value), then more consideration should be given to the choice of an appropriate rate — such as the risk characteristics of the proposal (for example, the extent of fixed costs and how costs and benefits vary with the state of the economy). Project flows that are more sensitive to market returns and other factors should have a higher discount rate, while projects that are less sensitive should have a lower one.

Further sensitivity testing can be used to help determine the appropriate rate, such as calculating the project rate of return (the rate at which the net present value is zero). If the plausible level of discount rate is below the rate of return, then the project improves efficiency.

[16] For a constant perpetual flow of net benefits, an x per cent increase in the discount rate reduces the present value by x/(1+x) per cent. A given percentage point increase in the discount rate represents a smaller percentage increase the higher the discount rate, and so the smaller the percentage change in net present value. For example, increase the discount rate from 3 to 8 per cent is a 167 per cent increase, which reduces present value by 62.5 per cent. Increasing the discount rate from 8 to 13 per cent is a 62.5 per cent increase, which reduces present value by 38.5 per cent.

Table 4.1 **Deriving a range of discount rates for sensitivity testing**

Conceptual guidance		*Quantifying the concept*
(i)	Setting a default rate:	
	Start from the marginal return to capital (which contains a risk premium reflecting average market risk)	9 per cent — see p. 37, details at appendix G
	Adjust for tax and foreign capital flows (the weighted average of project financing costs)	say 8 per cent — see p. 55, details at appendices E and G
(ii)	Sensitivity testing for risk	
	A higher rate suitable for high risk projects	10 per cent — see pp. 45, 57, details at appendix H
	A lower rate representing the weighted average risk free rate	3 per cent
(iii)	If sensitivity testing reveals that the choice of discount rate is important	Changes the sign of the project's net present value or its ranking against alternative projects
	Consider the project's exposure to undiversifiable risks.	Correlation with market risk and other risk factors. See p.42-45 and appendix H
	Use a risk premium as large as applies to equally risky private projects	