3 The discount rate and market benchmarks

The capital market is where inter-temporal trades take place — such as investment and saving, borrowing and lending. Market interest rates indicate the opportunities forgone when funds are withdrawn from the private sector and invested in government projects. People's inter-temporal preferences, how they value costs and benefits spread out over time, are revealed through the trades they make in the capital market. The interest rate is the price people pay to have resources now rather than later.

3.1 The discount rate in an undistorted capital market

Efficiency requires that the social discount rate be set at the social opportunity cost of funds allocated to public investment. If there were a perfectly competitive capital market, with no taxes, transaction costs, or risk, determining the social discount rate would be relatively straightforward. The capital market would determine a single interest rate. Dollar flows of costs and benefits should be discounted using this market interest rate, because it would be the rate at which dollars today can be turned into dollars tomorrow.

There are many reasons why future dollars are valued less than current dollars, including: impatience, the expectation that wealth will grow over time, opportunities for productive investment, and uncertainty (one reason people value future consumption less than current is that they may die before receiving the future consumption). All these are reasons for the market interest rate to be positive.

The marginal rate of time preference is the rate that people are willing to trade present consumption for future consumption. If someone is willing to give up \$1 today in return for \$1.04 a year from now, then the marginal rate of time preference is 4 per cent.

Consumers will tend to adjust their savings (or borrowing) until the rate at which they are willing to give up current dollars for future dollars just equals the market price of future dollars. That is, their marginal rate of time preference equals the interest they receive on their savings (the consumer rate of return).

Likewise, capital investment will adjust until the interest paid by investors (the investment rate) equals the marginal return to investment in capital.

In the absence of taxes, the consumption rate of interest equals the investment rate equals the market rate. Individuals trade consumption through time at the market interest rate, which is equivalent to saying that they discount future consumption at the market rate of interest. In addition, the market rate of interest is also the rate of return on capital investments, so this is the rate at which consumption can be translated through time via private sector investments.[1]

If a government project has a positive net present value when discounted with the market rate, then its benefits can compensate private investors and consumers for the consumption they forgo because of the project. The market rate shows the opportunity cost of the capital used by the project.

In reality the capital market is distorted, for example by taxes. There are different returns on different assets, depending on their risk characteristics and how they are taxed. Even for a given before-tax return, after-tax returns differ between consumers (with different marginal tax rates). Which of these numerous market interest rates should be used as the social discount rate? What is the appropriate measure of opportunity cost in the presence of capital taxes and risk?

3.2 The effects of capital taxes

Capital taxes drive a wedge between the investment and consumption rate of interest

The efficiency effects of a project depend on its interactions with existing taxes and other distortions, such as monopoly, externalities and capital market imperfections. Tax is one-third of GDP, and seems likely to impose a much larger distortion than monopoly profits or externalities.

Capital taxes distort inter-temporal choices through taxing the rate of return to savings. Capital taxes drive a wedge between the before–tax rate of return that

[1] Lind (1982a, pp. 25-28).

investments earn, or the investment rate of interest, and the after–tax rate of return that lenders receive, or the consumption rate of interest.[2]

The wedge is substantial. Income tax rates in Australia range up to 46.5 per cent. Abatement of welfare and family tax payments can increase the effective tax on income, adding 30 percentage points or more to statutory tax rates. Moreover, Australia's income tax system taxes the nominal rate of return, part of which only compensates the asset holder for inflation. The effective tax rate on an asset's real rate of return is much higher than the nominal income tax rate. In particular, the nominal payments on government bonds (usually considered the risk-free return) are taxed.

For example, if the annual nominal interest rate were 6 per cent and the expected inflation rate 3 per cent, the real before-tax interest rate of return would be 2.9 per cent (=1.06/1.03 – 1). If the tax rate were 50 per cent, the nominal after–tax rate of return is 3 per cent and the real after–tax rate of return 0 per cent. The effective tax rate on the asset's real rate of return is 100 per cent, double the tax rate on nominal income.[3] Further, the effective tax rate would change over time with expected inflation.

Is the investment or consumption rate the social discount rate?

There are two intuitions when it comes to determining the social discount rate, focusing on investment and consumption. Those who argue for the investment rate emphasise opportunity cost, indeed they often call it the opportunity cost of capital. If government investment comes at the expense of private investment, the opportunity cost to the economy is measured by the social returns that would have been generated by that investment. Generally, the before-tax rate of rate of return, i, measures the value of output that the funds would have generated for society (this ignores any positive externalities associated with private investment). It is also called the investment rate of interest, producer rate of interest, marginal rate of return to investment or capital, the marginal efficiency or product of capital, or the

---

[2] If a borrower can deduct interest payments, then capital taxes reduce the after-tax investment interest rate the borrower pays.

[3] More generally, when an asset's nominal rate of return, ρ, is taxed at τ per cent (or nominal interest payments deducted) and the expected inflation rate is π per cent, the after-tax real rate of return is: $r = [\rho(1-\tau) - \pi]/(1+\pi)$, as by definition $1 + \rho(1-\tau) = (1+r)(1+\pi)$

The before-tax real rate of return is $i = (\rho - \pi)/(1+\pi)$, as $1+\rho = (1+i)(1+\pi)$

The effective tax rate on the real before-tax rate of return is given by $(1 - t)i = r$ where:

$t = (i - r)/i = \tau\rho/(\rho - \pi)$ which is higher than the tax rate, τ, on the nominal rate of return.

social opportunity cost of capital. If a government project does not have a positive net present value using the investment rate, then society could get a higher return by investing in the private sector.

On the other hand, a cost-benefit analysis of a project values the stream of costs and benefits that accrue to consumers. The consumption rate of interest determines the consumer's valuation of current relative to future consumption (the consumer's marginal rate of time preference). The consumption rate of interest, r, is usually measured by the after–tax real rate of return on savings — the supply price of savings. It is equal to the consumer's rate of time preference.

In the absence of distortions, the consumption rate equals the investment rate, and the consumer's rate of time preference equals the marginal rate of return to capital — which is the appropriate discount rate. But when capital taxes drive a wedge between the two, should the social discount rate be the before-tax investment rate or the after-tax consumption rate?

Both intuitions are correct. A cost-benefit analysis must value future consumption flows using the consumer interest rate. It must also take account of the social opportunity cost of the capital it uses — the amount of consumption forgone when the capital is withdrawn from the private sector.

The opportunity cost of investment should account for the future consumption that displaced private investment would have produced. Because taxation means the investment rate exceeds the consumption rate, every dollar of private investment creates consumption with a present value greater than \$1. A unit of private investment produces a stream of social returns at a rate greater than that at which they are discounted by individuals. In other words, the social value of \$1 of private investment is greater than \$1. The social opportunity cost per dollar of private investment forgone is greater than a dollar.[4]

One way to proceed is to use shadow prices to convert project cost estimates into consumption equivalents, which account for the full opportunity cost of capital used, and then discount with the (after-tax) consumption rate.

An alternative approach is to discount the project flows with a discount rate that is a weighted average of the consumption and investment rates, with the weights being the proportion of project costs sourced from consumption and from investment.

4 See appendix D.

Weighted average or shadow price approach?

Appendix D, 'The shadow price of capital and weighted average discount rate', sets out both approaches and establishes they are equivalent under reasonable assumptions. The most practical approach is to assume the conditions for the weighted average discount rate to be valid hold, and to use it. If they do not hold, then the shadow pricing approach is impractical because the information it requires is seldom available (which explains why it is rarely done).

For example, different consumers face different marginal tax rates and so have different consumption rates of interest. Moreover, even working out the relevant consumption rate for any particular individual is difficult. A consumer may save into superannuation, pay off a mortgage, run up credit card debt and make personal investments, all at the same time. What is the appropriate consumption rate?

When different groups have different consumption rates, the shadow pricing approach becomes totally impractical. Each group's benefits should be discounted at its own rate, which requires the analyst to know the project benefits that accrue to each group, knowledge that is seldom available.[5] Working out the social value of private investment would be even more difficult, requiring knowledge of each group's share in the benefits from private investment. Applying the shadow pricing approach in even an approximate way is extremely difficult, and seldom done.

In contrast, the return on investment is a familiar and conventional way to think about government projects and the weighted average approach has an intuitive meaning — the project must earn enough to compensate all the losers from the project's effect on the capital market. That is, the project must cover its opportunity costs. When private investment is displaced, the cost to the economy is the social yield those investments would have had (measured by their before-tax return). When consumption is forgone, the cost to the economy is measured by the consumption rate.

A further advantage of the weighted average approach is that it can be readily adapted to incorporate foreign capital. Capital is internationally mobile so that increased demand for capital can also be met from increased capital funds from abroad.

[5] See Jones (2005, pp. 195-97) for details on how to evaluate projects when consumers have different discount rates because marginal tax rates differ.

The weighted average discount rate would then be the weighted average of the investment rate, the consumption rate and the marginal cost of foreign funds, with the weights equal to the amount of capital sourced from each.[6]

The sourcing of capital used in government projects[7]

The weights used to calculate the weighted average discount rate come from the amount of consumption and investment displaced and extra borrowing from foreigners when the government project increases the demand for capital in the capital market (or equivalently, reduces the supply of capital to the private sector). For example, when the government borrows (sells bonds) $100 million to finance an infrastructure project, it directly draws on the capital market, driving up interest rates which displaces private investment and increases private savings (reduces consumption).

But other government projects that impose costs and benefits in different time periods will also increase the demand for capital, even if the government does not draw directly on the capital market. When a regulatory proposal imposes present financial costs on firms and individuals, it increases the demand for capital. For example, if a construction firms is required to purchase scaffolding to improve safety, it may need to borrow, or forgo other investment. Any capital required by a regulatory proposal must be sourced from displaced investment or newly stimulated savings (that is, decreased consumption). For example, if regulation imposes large set-up costs on producers, such as requiring electricity producers to provide smart meters, then producers would have to finance their higher costs through borrowing in the capital market. For the regulatory proposal to increase efficiency, its benefits must exceed those to be had from these alternative uses of the capital.

Even policies that directly impose costs and benefits on the consumer are likely to increase the demand for capital and draw on the private capital market. For example, if the regulation instead forced the consumer to pay for a smart meter, would the capital market effect be any different?

Consider a policy that imposes a cost on a consumer in order to reap a future benefit, say regulation only allows energy efficient air-conditioners to be sold. Air-conditioning units now cost more, but consumers can expect lower future electricity bills.

---

[6] That is, $w = ai + (1–a–f)r + fMC_f$, where MC_f is the marginal cost of foreign funds and f is the proportion of investment sourced from foreign funds. Edwards (1985).

[7] For more details, see appendix B, which considers the effect of project financing on savings and investment, and appendix E, which deals with the implications of foreign capital flows.

Consumers may well pay the higher air-conditioning price out of savings, rather than reducing other consumption, to restore their desired consumption path. This reduces the supply of savings and crowds out investment in exactly the same way as government borrowing would.

Policies that directly impose costs on consumers (which include tax financed expenditure) can lead to the same type of capital market effects as in the standard project analysis, depending on how people perceive the value of future project output and future tax liabilities.[8]

There may be cases where the costs of a government policy mainly act to directly reduce consumption, and they would be discounted at a lower rate (as the relevant weighted average discount rate is close to the consumption rate). But these exceptions are likely to be where the discount rate makes little difference, for example, where a regulation imposes net costs in every period and so would be undesirable at any discount rate.

Determining the weighted average discount rate

Different investments receive different tax treatment and have widely varying effective tax rates. For example, income retained within companies is subject to company tax. The returns to investment in owner occupied housing are not taxed at all. Different forms of saving receive different tax treatment and marginal tax rates differ across individuals. As a result, the marginal productivity of investment differs across the various sectors of the private economy. Moreover, the marginal rate of time preference differs across various groups of savers. The cost of foreign borrowing may also differ across lenders and between different types of finance because of, for example, different tax treatment.

Ideally an estimate of the social discount rate would break down each source to allow for differences in the social yields of different classes of alternative investments, differences in the marginal rate of time preference for the different sources of forgone consumption, and differences in the marginal costs of borrowing from the various foreign sources, attaching an appropriate weight to each. Just to state these requirements is to realise the magnitude and complexity of pursuing an ideal measure of the social discount rate.[9]

In practice, analysts have little information about the weighting of each source, or even its cost. Nevertheless, to use any particular social discount rate is to implicitly

---

8 See appendix B for details.

9 Harberger (1992, pp. 28-9).

assume something about these variables, and it is better to make the assumptions explicit. What is needed is a reasonable rule of thumb.

There is a reasonable presumption that the weighted average rate will lie close to the before-tax investment rate because capital used in government projects comes mainly from activities than earn the before-tax investment rate of return or more.

The academic consensus is that increased demand for domestic capital comes mainly at the expense of investment rather than consumption.[10] There is evidence that investment is sensitive to interest rates, but little evidence that consumption responds to changes in interest rates.

Further, many consumers borrow and borrowing to finance consumption or owner–occupied housing is not tax deductible in Australia. Their marginal interest rate is the before-tax rate. Consumers that face borrowing constraints (credit–rationing) would have an even higher rate of time preference (value an extra dollar in the future even less). Stiglitz (1982) points out that plausible capital market imperfections (such as imperfect annuity markets) raise the social discount rate. Rather than society saving too little for the future, it may be saving too much.[11] Direct estimates of consumers' rate of time preference usually find it to be extremely high.[12]

Foreign capital is an important source of investment funds, and fairly responsive to interest rate changes. Its cost to Australia depends on knotty details of international tax law. The cost of debt is likely to be the before-tax return. The cost of foreign equity investment is likely to be the after-tax return. Further, if extra borrowing by Australians bids up the price that foreign lenders charge (that is, the supply of

10 See for example, Harberger (1969, p. 108); Boardman and Greenberg (1998, p. 305); Abelson (2000, p. 129); Department of Finance and Administration (2006, p. 66) and Zhuang et al. (1997, p. 10). Mulligan (2002) disputes this, arguing that the relevant consumption interest rate is not an after-tax bond yield (as used in many studies) but the after-tax rate of return on a representative piece of capital. He finds that consumption is interest elastic and forecastable when the interest rate is measured by the after-tax rate of return on capital from the national accounts. On the other hand, this rate is higher than the after-tax bond yield, offsetting the effect on the weighted average discount rate.

11 See Stiglitz (1982, p. 134).

12 For example, Warner and Pleeter (2001) find the vast majority in their sample had discount rates of at least 18 percent. Frederick et al. (2002, see table 1, p. 379) survey estimated rates and find they vary from negative to several thousand percent per year. Viscusi (2007, p. 228) finds people discount their own lives with real interest rates in the range 11 to 17 per cent.

foreign funds is upwards sloping) then the marginal cost of funds is greater than the average cost.[13]

Where the opportunity cost of the major sources of project capital is the before-tax return or more (i.e when project capital mainly comes from forgone private investment, forgone household borrowing and extra foreign debt), the weighted average discount rate would be close to the before-tax investment rate. In this (plausible) case, the consumption rate is of little relevance, and it is reasonable to ignore all the complications from different consumption rates. That would be impossible to with the shadow pricing approach, because its whole foundation is discounting consumption flows with a consumption rate.[14]

3.3 The Ramsey formula

One influential approach to determining the discount rate is to use a model of consumer behaviour to build the discount rate up from estimates of underlying parameters. Under common assumptions about consumer behaviour,[15] consumers will adjust their holdings of the risk free asset and their consumption pattern over time until the following equation for the risk-free rate holds: $r^f = \theta + \eta g$ where $\theta >$ 0 is the 'pure rate of time preference', which is used to discount future utility.[16] It arises from impatience and the chance of death. g is the rate at which consumption is expected to grow and η is the (absolute value of) the elasticity of the marginal utility of consumption (the percentage fall in the marginal utility when consumption increases by one per cent). The equation must hold for each consumer with access to the risk free asset.[17]

The Ramsey equation is sometimes used to estimate a social discount rate, and is often used to justify a relatively low rate.[18] Its use is one reason for the trend of declining recommended discount rates. It is also used in the context of inter-

13 See appendix E and also Stiglitz (1982, p. 128). Makin (2006) has estimated the real realised cost of foreign capital from 1995-96 to 2004-05 varied from -0.4 to 5.4 per cent, averaging 3.6 per cent (with a standard deviation of 1.6 per cent). The risk characteristics of foreign investment are unclear, such as how much goes into risk free assets like government bonds.

14 See Harberger (1987, pp. 174-75).

15 Technically, the equation holds for a consumer with a power utility function — which is a first approximation to any standard utility function

16 Some authors call this the rate of time preference, but that is also used to refer to the social discount rate (which is used to discount consumption, not utility).

17 See appendix F.

18 For example, by Stern (2007); Quiggin (2006); Garnaut (2008) and HM Treasury (2003).

generational choice to estimate long-run discount rates for the economy as a whole. It can also be derived from the optimal long-run growth model with a representative consumer, which can take account of population growth, often used to examine these issues. In this application, θ becomes the pure social rate of time preference, which is the discount rate on the utility of future generations.[19]

Sometimes a descriptive approach to the use of the Ramsey formula is taken, with empirical evidence on each of the parameters used to estimate the population's rate of time preference, and so the discount rate.[20] As marginal tax rates differ across individuals, they have different after-tax discount rates. Presumably the single Ramsey rate is meant to be some sort of average of the different consumption rates, although such uses of the Ramsey approach do not explicitly consider the relevant weights. When people face different marginal tax rates, taxation of nominal rates of return means that even if interest rates in the absence of inflation were stable, expected inflation changes the after-tax return, with the effect depending on the person's marginal tax rate. The effect of expected inflation on interest rates would depend on the proportions of savers and borrowers in each tax bracket and the responsiveness of their saving and borrowing to changes in the interest rate.[21]

In the descriptive approach, the Ramsey formula approach can be interpreted as a positive model of how the economy works, which estimates the equilibrium market interest rate — the marginal rate of return to capital and the rate at which consumers trade consumption over time.

Others adopt a prescriptive approach to use of the Ramsey formula: what the population's preferences should be when decision making is ethical. The Ramsey equation is used as a framework to guide the ethical choice of a social discount rate. For example, such approaches often specify a pure social rate of time preference, θ, equal to zero (or very close to zero).[22] The other variables may also be specified from the analyst's value judgements.

If the Ramsey rate specified by a prescriptive approach is below observed market rates, as argued earlier (in section 2.6), that is an argument for increased savings and investment, not for lowering the discount rate used to evaluate any particular project. At best it only addresses the issue of the appropriate transfer to future generations, rather than the policies and projects to achieve that transfer. As argued

19 See appendix C.

20 For example, see HM Treasury (2003, p. 97-98).

21 See appendix A, section A.3.

22 For example, Garnaut (2008); Quiggin (2006) and Stern (2007).

earlier, project choice should be based on discounting with an efficiency-based social discount rate.

Zhuang et al (2007) survey empirical estimates of θ (the pure rate of time preference) and η (the elasticity of the marginal utility of consumption or the co-efficient of relative risk aversion). They find that θ ranges from 0 to 3 per cent and η from 0.2 to 4 (with most between 1 and 2).[23] Average annual per person consumption growth is usually in the range 1 to 2 per cent. For example in recent applications, the Treasury has assumed 1.2 per cent, whereas the Productivity Commission has used 1.75 percent (further examples are given below). Annual GDP growth per capita in Australia averaged 2.1 per cent from 1960 to 2007. Peak to peak labour productivity 1969-70 to 2003-04 averaged 1.71 per cent. Using these parameter ranges in the Ramsey formula give estimates of a risk free discount rate ranging from 0.24 to 11 per cent.

The wide range in empirical estimates of the relevant parameters means that the descriptive Ramsey approach does not resolve disagreement about the appropriate discount rate. It is consistent with a wide range of risk free market rates.

Nor does the prescriptive approach resolve disagreement. Table 1.3 summarises different parameters that have been suggested. The parameters differ greatly, resulting in recommended discount rates that range from 1.3 to 8 per cent. As the differences reflect different judgements about values as well as about the empirical literature, they cannot be resolved objectively. Further, that the Ramsey formula gives this wide range even with an assumed zero pure rate of time preference (θ) confirms that the debate over the size of that variable is not decisive to discount rate choice.[24]

[23] See Zhuang et al. (2007, p. 6-7).

[24] See appendix C for more details.

Table 3.1 **Different discount rates derived from the Ramsey formula**

Source	*Pure rate of social time preference, θ*	*Elasticity of marginal utility of consumption, η*	*Growth rate in consumption, g)*	*Discount rate = θ + ηg*
	per cent		per cent	per cent
Stern (2007)	0.1	1	1.3	1.4
Quiggin (2006)	0	1	1.5	1.5
Cline (1993)	0	1.5	1	1.5
Garnaut (2008)	0	1-2	1.3	1.3-2.6
HM Treasury (2003)	1.5	1	2	3.5
Nordhaus (2007)	1.5	2	2	5.5
Weitzman (2007)	2	2	2	6
Arrow (2007)	0	2-3	If 1-2	2-6
Dasgupta (2006)	0	2-4	If 1-2	2-8
Gollier (2006)	0	2-4	If 1.3	2.6-5.2
Empirical evidence	**0–3**	**0.2–4**	**1.2–2.1 (for Australia)**	**0.24–11 (given range)**

Source: The empirical evidence comes from the summary in Zhuang et al. (2007).

Although the proponents of applying the Ramsey formula ignore tax, recognising that there are large tax wedges in capital markets has important implications for the use of the Ramsey formula discount rates.

The Ramsey formula for the risk free rate is based on the conditions for a consumer to be happy with his asset holdings and consumption patterns. It determines a consumption rate and is appropriate for discounting future consumption. For a consumer subject to capital taxes (or who can deduct interest payments), this would be an after-tax rate. But it is incorrect to discount with the consumption rate unless the project flows have been shadow priced to convert them into consumption equivalents. Ordinary project flows should be discounted with a weighted average of the consumption and investment rates, which accounts for the consequences of capital taxation and the full change in consumption from displaced private investment.

Even if low parameter estimates are used to produce a low Ramsey rate, inflation and high tax rates on nominal rates of return mean that the before-tax investment rate is likely to lie substantially above the Ramsey consumption rate. For example, if the tax rate was 50 per cent and expected inflation 3 per cent, it would take a before-tax rate of 9 per cent nominal, or 6 per cent real, to give a Ramsey rate of 1.5 per cent real.[25]

As most of the weight in the weighted average discount rate lies on the investment rate, accounting for capital taxation means the social discount rate is far above rates produced by the Ramsey formula. What is the evidence on the investment rate in Australia?

3.4 The marginal return to capital

One way to estimate the overall social yield of capital is to derive a direct measure of the real return to capital from the national accounts by dividing the total income from capital generated in the private sector by an estimate of the private sector capital stock.[26] The average return to investors over a long time period provides a reasonable estimate of the return to capital.

Dolman (2007) estimates the nominal rate of return to capital in the market sector. He calculates it by expressing capital earnings before interest and direct tax in each year as a percentage of the net capital stock at the beginning of the year. Earnings are calculated based on the gross operating surplus of corporations (and a portion of the gross mixed income of unincorporated businesses), after subtracting the cost of physical depreciation of the capital stock during the year and adding back the carrying gain that firms make because the price of capital increases over time. Indirect taxes and subsidies on production are also removed.

The rates of return are set out in table 3.2, which expresses them in real terms, deflated by the consumer price index. The returns are high – with an (arithmetic) average annual real rate of return of 8.9 per cent since 1965 and 11.7 percent since 2000. The standard deviation is 3 per cent giving a 95 per cent confidence interval of plus or minus just under one percentage point. The low volatility means the geometric mean is only slightly less than the arithmetic (it is 8.6 per cent).[27] As the CPI is thought to overstate the inflation rate, true real returns would be higher.[28]

25 If the tax rate is τ, the expected rate of inflation rate π and the before-tax nominal rate ρ, then the real after-tax rate of return $r = \rho(1-t)-\pi$ and so the real before-tax return is $i = \rho - \pi = (r+\pi)/(1-t) - \pi$. It increases with expected inflation and the tax rate. See appendix A, section A.3.

26 As suggested by Harberger (1987, pp. 177-78; 1992, p. 29), and by Poterba (1997, p. 10–11).

27 See appendix G for a discussion of arithmetic and geometric rates of return.

28 See appendix A, section A.2.

Table 3.2 **Real Rate of Return to Capital, Market sector Australia**

Per cent. Year to 30 June

1960s		*1970s*		*1980s*		*1990s*		*2000s*	
		1970	11.7	1980	8.2	1990	4.6	2000	9.0
		1971	8.8	1981	9.2	1991	5.2	2001	5.2
		1972	7.7	1982	6.9	1992	6.4	2002	10.0
		1973	10.9	1983	2.6	1993	7.6	2003	10.6
		1974	11.1	1984	8.4	1994	8.6	2004	13.5
1965	11.6	1975	7.9	1985	8.2	1995	5.8	2005	15.2
1966	9.9	1976	8.2	1986	8.2	1996	6.9	2006	14.7
1967	12.2	1977	4.1	1987	5.0	1997	9.9	2007	15.4
1968	9.5	1978	6.4	1988	6.4	1998	11.1		
1969	14.0	1979	8.7	1989	6.1	1999	10.5		
Average	**11.4**		**8.5**		**6.9**		**7.7**		**11.7**

Source: Dolman (2007).

The Dolman (2007) estimates are consistent with other national accounts based estimates of the before-tax rate of return to investment in Australia and the United States and with estimates of the cost of capital in Australia.[29]

One limitation of these estimated rates is that they are the average rates of return to private capital investment, whereas the cost of displaced private sector investment is the marginal rate of return. If the capital market were competitive and production constant returns to scale, then factors are paid their marginal product and factor payments would exhaust output. Capital would be paid its marginal product. Capital income would be the marginal product of capital times the capital stock and so the measured average return would equal the marginal return. Further, the same would be true if capital was heterogeneous, so long as marginal rates of return were equalized across different capital investments.

If firms have some monopoly power (that is, can set prices above marginal cost), then the measured average return to capital could exceed the true marginal return because some measured profits are not a return to capital. On the other hand, Summers argues one interpretation of the empirical evidence is that monopoly power is combined with increasing returns and the threat of entry and so the private return to increased capital investment is below the social return and the measured return to capital understates the true productivity of new investment.[30] Further, monopoly power in other sectors may mean measured profits understate the social yield from capital — for example, if unions extract some of the return from capital

29 See appendix G, section G.3, which also discusses possible biases in the measures.

30 See Summers (1990, p. 114, 127-30).

investment (through increasing wages). Some evidence suggests labour captures between 20 and 40 per cent of the return on incremental capital investments, with effects similar to a tax on the rate of return to capital.[31]

Any positive externalities from capital investment would raise the social return further. Summers suggests the evidence on the relationship between national investment rates and rates of economic growth means the level of capital accumulation is linked to technical change and has substantial external benefits.[32]

A reasonable estimate of the marginal rate of return to capital (or the opportunity cost of foregone private investment) is 9 per cent real. But this market return includes a risk premium which compensates investors for the risk they bear. What adjustments should be made for risk to determine the opportunity cost of capital?

3.5 Risk

'Nothing is more certain than risk', Thomas Sowell, *Applied Economics* (p. 129)

Investment decisions generate real risks. No one knows with certainty the future costs and benefits from any multi-period investment project: benefits may not be realised or costs may blow out. One financial market expert writes 'risk is at the centre of all investment decisions … investing is a bet on an unknown future'.[33]

Most people are risk averse (at the margin), and need to be compensated for bearing risk. In the private sector, the quantity and price of an asset's risk are key determinants of an asset's value. The risk premium is a major part of the market rate of return. The historical equity market risk premium in Australia is usually estimated to average around 3 to 8 per cent, but its high volatility (a standard deviation of around 20 percentage points) makes it difficult to estimate with confidence.[34]

Much financial market activity is about evaluating, pricing and managing risk. Risk is pooled and transferred, usually from those who are prepared to accept a lower return to get rid of it, to those who are less averse to risk (but as illustrated in the

31 See Summers (1990, pp. 130-32).

32 See Summers (1990, pp. 133-36).

33 Bernstein (2007, pp. ix, xii).

34 See appendix G. The standard deviation is 18.3 per cent. If normally distributed about one-third of observations should be one standard deviation away from the mean. In the last 39 years in Australia, 10 annual returns lay more than 18.3 percentage points away from the mean.

recent financial crisis, sometimes to those who imperfectly understand what the risk is).

Governments must manage the risks from public projects. Here too, failures of risk management are plentiful in history, with common examples in many countries arising in defence procurement, infrastructure investment (such as building stadiums and dams), industry policy and social policy.[35]A key analytical issue is how to value the risk from a proposed government project. The cost of risk a project imposes depends on the quantity of risk and its price.

Assessing risks

Risk assessment should be a standard component of the evaluation of any major proposal. The estimates of future costs and benefits should be expected values. 'Expected value' is not the anticipated or most likely outcome. It is not 'plausible "comfortable scenarios" about how the project will evolve'.[36] It is the probability-weighted average of all possible outcomes — including possible disasters and windfalls. It should consider what could happen as well as what should happen.

The expected value estimates should be based on a realistic risk analysis, using a reasonable range of possible outcomes. Ideally the analyst carefully assesses potential risks, determines the size of the cost or benefit in each possible outcome and estimates the probability that outcome will occur. The expected value is the probability-weighted sum of the values in all possible outcomes. For example, if there is a 90 per cent chance of a payment of $100 next year, but a 10 per cent chance that a recession will occur and the payment would only be $40, the expected value of the payment is $94 (that is, 0.9*$100 + 0.1*$40).

Harberger recommends calculating expected values by developing four or five scenarios keyed to the variables that are most important to a project's success or failure. For example, scenarios that are quite unfavourable, unfavourable, somewhat unfavourable, normal, somewhat favourable, quite favourable, etc. For each scenario the analyst works out the profile of costs and benefit, attaches a probability and generates expected values.[37]

35 For example, government risk management of even the largest projects leaves a lot to be desired. See Flyvbjerg, Bruzelius, and Rothengattern (2003) who identify the main cause of the 'strikingly poor performance records' (p. 3) of megaprojects as 'inadequate deliberation about risk and lack of accountability in the decision making process' (p. 6)

36 Harberger (1997, p. 21).

37 Harberger (1997, p. 23).

One common problem with cost-benefit studies is optimism bias, a pervasive tendency to underestimate costs and overestimate benefits. Downside risks should not be ignored, but should be accounted for in the expected value estimates.

Expected values are the risk-weighted average of all possibilities. They do not account for the cost of risk a project imposes. Risk arises because project flows depart from their expected values. In general, a risky flow is valued differently from its expected value. For example, a 90 per cent chance of a payment of $100 and a 10 per cent chance of a payment of $40 has the same expected value as a certain payment of $94, but the risk imposed, and the payment's value, may be very different.

Modern financial theory suggests what matters to individuals is the risk to their consumption, so the cost of project risk to an individual depends on how project costs and benefits contributes to his overall consumption risk. That is, it is how a payoff co-varies with consumption that matters, not its variability. Assets whose returns co-vary positively with consumption make consumption more volatile and investors need a higher expected return to be induced to hold them. That is, their rate of return includes a risk premium to compensate for the cost of the risk they impose on investors. Their net benefit flows contain non-diversifiable or aggregate risk. For example, the popular capital asset pricing model (CAPM) assumes future consumption is funded solely from returns to portfolios of securities and it prices assets based on the covariance of their returns with the market portfolio (the market risk they contain, which so called 'beta' measures). In practice, other non-diversifiable risk factors, such as recessions, also appear to be important.[38]

Households receive the benefits and costs of public projects through expected increases in real consumption and through expected tax changes. The benefits and costs of government projects accrue to particular private households and firms whom the project affects and it is how the project affects their consumption risk that matters.[39]

The adjustment made for risk depends on the variability of the project net benefit flows and how they are correlated with individuals' consumption. In practical terms, it depends how the project flows are correlated with factors that determine the marginal utility of consumption and give rise to costly risk. Projects with market risk need to earn a higher return than those without, and so the future expected value of their costs and benefits should be discounted at a higher rate (one that includes a risk premium). Alternatively, but to equivalent effect, the expected

38 See appendix H.

39 Jensen and Bailey (1972, p. 16-17).

values could instead be converted into certainty equivalents, which should then be evaluated using the risk free discount rate. If the project contains market risk, the certainty equivalents in the second approach are less than the expected values. If the risks were negatively correlated with the market, the risk premium would be negative and the certainty equivalent would be greater than the expected value [40]

Projects contain more market risk (higher beta), the more net benefits co-vary with the state of the economy, and should be discounted with a higher discount rate. Projects with large, up front fixed costs (often called high operating leverage) tend to be riskier than projects where costs are variable year-by-year. (For example, some government projects such as infrastructure investments and regulations have high set-up costs, whereas others such as welfare programs can be adjusted year by year if circumstances change.) The high level of fixed cost means a small percentage change in benefits will produce a much larger percentage change in net benefits and greater co-variance with the factors that affect benefits.[41]

Valuing risky flows: the traditional approach

How should we determine the present value of uncertain income flows? The traditional approach is to use observations of market behaviour – how a private sector asset of equivalent risk is priced in financial and capital markets. An asset gives its owner future net income flows and the price of an asset is the present value of those flows, discounted at an appropriate rate. For example, the price of a share in an efficient market tends towards the present value of its expected dividend flow.

Asset pricing theory tries to understand the prices or values of claims to uncertain payments.[42] Theories of asset pricing can be used to value a stream of risky payments from a government project and to account for the cost of risk imposed on people. If a government project gives a particular flow of benefits, the recipients would value the benefits in the same way as a private project that produced the same benefit flow.

The traditional approach to dealing with risk is to calculate the net present value of expected values with a discount rate that includes a risk premium based on the market price of risk.

In a complete market, all diversifiable risk is spread, and non-diversifiable or ‘market risk’ is the only risk worth paying to avoid. For example, in the CAPM, the

[40] See chapter 4, box 4.1.

[41] Brearley and Myers (1991, p. 199-200).

[42] See Cochrane (2005) for a comprehensive treatment.

return on an asset does not depend on the variability of its returns, but on the contribution it makes to overall portfolio risk (and ultimately, to consumption risk). The return on each asset is the risk free rate of return plus a risk premium which compensates risk averse consumers for bearing any market risk in the asset's net cash flows. Assets that only have diversifiable risk earn the riskless rate of return (a zero risk premium). Assets with returns that co-vary positively with the market return must pay a higher rate of return. Assets that co-vary negatively with the market reduce portfolio risk (they provide insurance) and earn less than the riskless return.[43]

The Arrow-Lind theorem

The issue is then to determine the quantity of risk the project imposes. For example, in a widely cited, classic contribution to public sector discounting theory, Arrow and Lind (1970) showed that if a government project was 'small' (in relation to the total wealth of taxpayers) and 'the returns from a given public investment are independent of other components of national income', then the social cost of the risk for project flows that accrue to taxpayers tends to zero as the number of taxpayers tends to infinity.[44] That is, government investments with diversifiable risks spread over many households should be evaluated using the riskless rate to discount expected benefits (that is, with no adjustment for risk).

This result, known as the Arrow-Lind theorem, is consistent with the CAPM approach. If a project contains only diversifiable risk and no market risk, an efficient private sector would spread the risk and would also use the riskless rate to discount expected project returns.

If a project contains aggregate risk, then a risk premium should be used

Aggregate risk is an irreducible social risk that cannot be diversified, even by government. It is caused by shocks such as recessions and variations in the market return. As Bailey and Jensen point out:

> the 'private' (and 'social') risk of even a small project which is perfectly correlated with the average returns on all other assets cannot be reduced one iota by transferring it from the private to the public sector.

43 See appendix H.

44 Arrow and Lind (1970, p. 171).

> The question regarding the size and sign of the covariances of returns on prospective projects is an empirical issue. However, some brief consideration of the problem seems to indicate (contrary to Samuelson et. al.) that the vast majority of government projects will have outcomes correlated with national income. For instance, any government investment that facilitates ordinary commerce will produce more benefits when national income is high than when it is low. Electric power, highways, waterways, airports, and postal service, for example, all have this character.[45]

When a project imposes risk on individuals, it should be evaluated with a discount rate that reflects the risk premium they demand for bearing risk.

Most government projects involve aggregate risk

Some have applied the Arrow-Lind theorem to argue that government projects are riskless and should be discounted using a risk-free rate. But most government projects involve market risk. The demand for the services of infrastructure projects is linked to the state of the economy. Less obviously, so are the benefits of regulation. For example, the valuation of 'statistical lives' saved by safety regulations depends on the level of wages.

Further, Foldes and Rees (1977) show the Arrow-Lind assumption that project net returns are statistically independent of each person's disposable income in the absence of the project is very stringent. It is true when people are identical, but when people pay different amounts of tax or if they receive different amounts of taxable benefits from the project, the assumption fails, even when the project's social returns and gross income are uncorrelated.

They also set out why the assumption that the share of net benefits accruing to any one person becomes negligible as population tends to infinity is unacceptable for three cases: public goods (non–rivalry); where project scale is adjusted in proportion to the size of the population; and projects whose benefits in part accrue to a small section of the population.

Lind himself argued that a risk adjusted rate, which reflects the covariance between the investment returns and market returns, should be used to evaluate public investments.[46] He agrees that the Arrow–Lind theorem does not apply to projects that contain market risk or where the risks are not spread, and he concludes that

> in most cases in would be extremely difficult to estimate the probability of benefits and costs over time and measure the covariance of the net benefits of a public policy or project with the market portfolio. … As a practical alternative, I suggest that the best

45 Jensen and Bailey (1972, p. 7).

46 Lind (1982, pp. 69-71).

assumption is that returns to most government policies or projects are highly correlated with returns to the economy as a whole, unless there is a clear argument to the contrary.[47]

Some government projects may warrant a higher risk premium than the private sector average

Experience suggests governments can run some projects badly and increase risk. For example, governments may find it more difficult than the private sector to wind up or restructure poorly performing projects, because of political pressures and lack of incentives. As these hard decisions are more likely to be necessary in a recession, this difficulty adds to market (non-diversifiable) risk.

It may be difficult to identify, ex ante, which government projects are particularly vulnerable to such risk management problems. This is a reminder of the value, in areas where significant risk is likely and risk management is suspect, of sensitivity testing with a range of discount rates that includes a ‘top end’ with a risk premia above the market average (see section 4.1 below.)

Arrow and Lind also point out that the government may be worse at allocating aggregate risk than the private sector because ‘the government does not discriminate among taxpayers according to their risk preferences’.[48] In contrast, risk is voluntarily assumed in private markets, ensuring market risk is generally traded to those most willing to bear it.

For those persuaded that government projects often contain non-diversifiable risk, use of a discount rate that has an appropriate market risk premium in it follows, with the practical difficulty then being selection of that premium. But analysts approaching the issue from other starting points have raised various arguments why a risk premium should not be used. These arguments are briefly reviewed in the remainder of this section.

Discounting losses

Some analysts have argued that a risk premium in the discount rate is unsatisfactory because it implies ‘an increase in uncertainty about a future cost makes that cost

---

47 Lind (1982b, p. 447).

48 Arrow and Lind (1970, p. 172). Stapelton and Subrahmanyam (1978) also argue the government finds it difficult to achieve an efficient allocation of risk bearing.

less important as viewed today'[49] and would 'make the higher risk project appear preferable to a lower risk project'.[50]

But as noted earlier in this section, what matters for the discount rate is not the variability in cost and benefits, but the amount of aggregate risk the project contains — the covariance with the market return (beta) and other risk factors. The same principles apply to negative cash flows as any other cash flows. A negative cash flow has the same beta as the equal (and opposite) positive cash flow, and should be discounted with the same discount rate (see box 3.1).

Box 3.1 Risky cost flows should be discounted in the same way as risky benefit flows

Define a risky cash inflow CF_{in}. The negative of this cash flow is $CF_{out} = -CF_{in}$. Let r_{in} and r_{out} be the correct risk adjusted discount rate for CF_{in} and CF_{out}. Therefore the cash flows have present values defined by:

$V_{in} = E(CF_{in})/(1+r_{in})$ and $V_{out} = E(CF_{out})/(1+r_{out})$.

But by definition, $CF_{in} + CF_{out} = 0$. Each period the flows exactly cancel each other, so a person who receives both flows would receive 0 each period. Therefore,

$V_{in} + V_{out} = 0$ or $V_{in} = -V_{out}$. But $E(CF_{in}) = -E(CF_{out})$. Therefore $r_{in} = r_{out}$.

Source: Ariel (1998).

Using the same discount rate to discount costs as would be used for a positive cash flow of the same magnitude makes intuitive sense. A project with a positive net benefit flow that is high in good times and low in bad times increases portfolio risk and should be discounted with a rate above the riskless rate.

The negative of that flow is high (more negative) in good times and low in bad times, reduces portfolio risk and makes the project more valuable. It should be discounted with a discount rate above the riskless rate, which makes the present value of the cost flow less negative and makes the project more attractive.

49 Stiglitz (1982, p. 150).

50 Partnerships Victoria (2003, pp. 23, similar statements are also at pp.29-30). The same claim is made, in Bureau of Transport and Regional Economics (2005, p. 7, 104; 1999 p.74-75); Victorian Competition and Efficiency Commission (2007, p. 2);, Bazelon and. Smetters (1999, p. 217); Arrow and Lind (1970, p. 177); Victorian Department of Treasury and Finance (2007, p. C-9) and Boardman et al. (2006, p. 271).

Conversely, a project with a positive net benefit flow that is low in good times and high in bad times, reduces portfolio risk and makes the project more valuable. It should be discounted with a discount rate below the riskless return, which increases the projects net present value and makes the project more attractive.

The negative of this flow, which is high (more negative) in bad times — just when you least want high costs — increases portfolio risk and so lowers the project's net present value. It should be discounted with a discount rate below the riskless rate, making its present value more negative.

Adding an appropriate risk premium to the discount rate to compensate for risk gives the correct net present value even when there are risky negative net benefit flows.

Market imperfections

Some analysts have argued that since capital markets are imperfect and various asset pricing puzzles (such as the equity premium puzzle) persist, observed risk premiums can be ignored in discounting government projects. They argue that analysts should use the risk free rate (in practice, the government bond rate) because in a perfect market, the market risk premium would be negligible.[51]

Arrow and Lind reject this argument that when markets are imperfect, the government should nevertheless act as if markets were perfect:

> if we are to measure benefits and costs in terms of individuals' willingness to pay, then we must treat risks in accordance with these individual valuations. Since individuals do not have the opportunities for insuring assumed in the state–preference model, they will not value uncertainty as they would if these markets did exist. … The critical question is: What is the cost of uncertainty in terms of costs to individuals?[52]

Even imperfect markets convey information about the cost of risk. The market price of risk shows the cost of extra non-diversifiable risk to the private sector, whether the market is efficient in handling risk or not.[53] For example, people require a risk premium for holding equities that are likely to fall in value during recessions, at the time they are most at risk of lower returns from other sources of income. If a

51 See for example, Grant and Quiggin (1998, p. 4; 2003). Quiggin (1996; 1997); Spackman (2001; 2004) and Bureau of Transport and Regional Economics (2005, p. i, ix, vii). See appendix H for a discussion of the pricing puzzles and the cost of risk.

52 Arrow and Lind (1970, p. 167).

53 Sandmo (1972) and Stiglitz (1982, p. 152).

government project imposes aggregate risk (which cannot be diversified, so is costly no matter how it is spread), it should not be ignored or assumed away.

Further, it is not necessarily true that the risk premium would be negligible in a perfect market. To the extent various explanations of the equity premium puzzle are true, they give a significant risk premium even in a perfect market.[54]

A practical approach is to use the risk margin observed on private sector investments in the same risk class as the government project. A project that imposes similar risks on consumers should have those risks priced in the same way.

The government should only price risk differently if it has some advantage that allows it to improve on an imperfect market. The advantage should be specified, because the government cannot correct some market imperfections.[55] Any comparison of how the government and private sector manage risk should take account of how the government in fact operates under the incentives of the political process.

A liquidity advantage?

Another argument for using a low discount rate without significant risk premium is that 'by virtue of its superior ability to issue a liquid security the government enjoys a cost advantage relative to issuers of private equity. Hence, the appropriate rate of discount for public projects is the bond rate'.[56]

The liquidity of government bonds ultimately comes from the government's sovereign power to tax.[57] That is, it can draw on the resources of taxpayers to ensure the debt is repaid. That does not reduce the risk and illiquidity costs of committing funds to any particular government project. Those risks arise because funds are put into that project, rather than alternative uses. How the funds are raised is immaterial. The cost of those risks are borne by taxpayers and programme beneficiaries, who are essentially equity holders in the government project: 'the taxpayers will have assumed a contingent liability for which they are not remunerated, that is, there is no risk premium in the sovereign borrowing rate. In fact, the taxpayers play the role of investors bearing the risk of failure'.[58]

---

54 See appendix H.

55 Such as moral hazard. See for example, Grant and Quiggin (2003, p. 264).

56 Bureau of Transport and Regional Economics (2005, p. 110).

57 See appendix G for a discussion of liquidity risk.

58 Klein (1996, p. 1; 1997, p. 30). The same point is made in Congressional Budget Office (2004, p. 4-5); Currie (2000) and Economic Planning Advisory Commission (1995, p. 37).

If anything the liquidity premium for government assets should be higher than for equity because from the point of view of the taxpayer, a government investment is less liquid. Unlike equity it cannot be traded or disposed of by an individual taxpayer.

An advantage in spreading risk?

A further argument that has been advanced for government use of a lower discount rate (reflecting a lower price of risk) is that the private market might be imperfect in trading diversifiable risk and the government has some advantage over the private sector in dealing with the imperfection and spreading diversifiable risk.[59]

If governments do not spread the risk, then it is borne by individuals and discounting should be done at a rate that reflects the risk premiums they demand for bearing risk. For example, when governments introduce regulation, usually the resulting costs and benefits are imposed on the producers subject to that regulation and their customers and not spread more broadly.

A government might spread all project risks if it appropriated all benefits and paid all costs so that all risk is borne collectively and is then diversified. In fact, governments seldom appropriate project benefits. As Bailey and Jensen point out:

> In most cases, the government gives away the services of the project without charge, so that its contribution to national income flows into the hands of the persons and firms whom it particularly benefits. … The project risk is borne by the specific households that receive the (uncertain) benefits.[60]

Arrow and Lind agree, arguing that in the 'typical case where costs are borne publicly and benefits accrue privately'[61] analysts should discount the expected value of costs at the riskless rate (as their costs are spread) and the expected value of benefits at a higher rate (reflecting the preferences of the people who receive the benefits), a procedure that would qualify fewer government projects. They conclude that if:

> some benefits and costs of sizeable magnitudes accrued directly to individuals so that these individuals incurred the attendant costs of risk bearing … it is appropriate to discount for risk … As a practical matter, Hirshliefer's suggestion of finding the marginal rate of return on assets with similar payoffs in the private sector, and using

59 Grant and Quiggin (2006, p. 264).

60 Bailey and Jensen (1972, p. 4, 17).

61 Arrow and Lind (1970, p. 177).

this as the rate of discount, appears reasonable for discounting those costs and benefits which accrue privately.[62]

The cost of using the stock market

In a variant of the foregoing argument, Spackman (2004) argues that the high equity premium is specific to the share market and is not relevant for government investments. On this view, irrational changes in market sentiment cause share prices to be excessively volatile and creates risk unique to the share market.[63]

However the broad national account measures of the rate of return to capital indicate it is high (around 9 per cent real) and stable — and so high premiums relative to the risk free rate are not limited to the share market, but seem to reflect broader attitudes to the cost of risk.[64] Government projects that come at the expense of investment have a high opportunity cost. The government should only invest in a project if can reap a better return from it than from investing in private sector project of equivalent risk. That requires using the private rate of return to capital as the discount rate (or, more accurately, using the weighted average discount rate, which adjusts for taxes).

Further, if the government can directly invest in the private sector, the return on private capital is the opportunity cost of investing in government projects, even if it reflects an irrationally high price of risk that the government has an advantage in dealing with (see box 3.2).

3.6 Discounting the distant future

It is difficult to observe market interest rates for the far future. For example, the maximum term for bonds is around 30 years. Even in an efficiency analysis, there may be reasons for adjusting the social discount rate used to evaluate costs and benefits received in the far future away from current market rates. There are valid arguments to lower the discount rate for evaluating costs and benefits received in the far future, but there are also valid arguments to raise it.[65]

---

62 Arrow and Lind (1970, p. 175-76).

63 See Spackman (2004, pp. 480-81, 503).

64 See section 3.4.

65 See appendix I for more detail.

Box 3.2 Opportunity cost when direct investment in the private sector is possible

Over recent decades, governments around the world have created sovereign wealth funds which include in their investment portfolios significant investments in the private sector (for example, through company bonds and national and foreign equities).[66] Australia's Future Fund is an example of a sovereign wealth fund.

Even if the government can avoid the risk costs from private sector investment, the market risk premium may still be part of the appropriate discount rate. The government could invest in the private sector, use its advantage to reduce risk costs (for example. by spreading risk the market cannot) and reap the market return.[67] For example, say the expected market return was 9 per cent, of which 5 per cent was a risk premium to compensate individuals for the cost of risk contained in the market portfolio. If the government could diversify this risk, it does not need the risk premium. But it could still reap the 9 per cent return from investing in the market portfolio. The market return would be the appropriate starting point discount rate for evaluating alternative projects.

Mishan concluded in 1972, well before the upsurge in sovereign wealth funds, that if:

> the public agency is permitted to use funds appropriated for investment purpose, either to invest directly in the private sector, so availing itself fully of an actuarial rate of return prevailing in the private sector, or else to buy out private investment projects already undertaken– buying them out at (not less than) the capitalized value of the certainty–equivalent of their expected stream of benefits — the Arrow–Lind amendment does not apply. For under either of the latter options the opportunity rate of return open to the public funds will again be the full actuarial rate of return on private investment, which private rate of return can then be adopted as the appropriate rate of discount for public investment projects. Moreover, inasmuch as permitting public agencies these additional options of investing in, or buying from, the private sector promotes optimal use of investible funds, the economist should recommend them.[68]

For example, one popular reason for lowering the rate (associated with Weitzman) is that when there is uncertainty about what discount rate to use, the appropriate discount rate to calculate present value is lower the further in the future the payments are received, so called hyperbolic discounting.[69]

Another reason provided for using a lower real interest rate to discount costs and benefits received further in the future is that uncertainty about the path of interest

66 See Devlin and Brummitt (2007).

67 See Mishan (1972).

68 Mishan (1972, p. 163).

69 Weitzman (2001). See appendix I for details..

rates means the yield curve for real interest rates tends to slope down.[70] Further, the appropriate risk premium may also fall with the length of project.[71]

Even if the net present value of the project is positive, delaying the project may increase the net present value. Although hyperbolic discounting increases the net present value of long term projects, it may also make delaying them attractive.[72]

Further, starting a project gives up the possibility of waiting for new information to arrive that may change the desirability or timing of the expenditure. This lost option value is an opportunity cost that must be included as part of the cost of investment. For a project to go ahead, the net present value must exceed the value of keeping the option alive. As long as time is likely to reveal relevant information, an option is more valuable the greater the uncertainty over net benefits and the higher are sunk costs. The option value of delay increases the appropriate discount rate to apply to government projects.[73] But other kinds of real option, such as expansion and abandonment options, may lower the appropriate discount rate.

The effect of uncertainty over multiple periods is complex. It depends on the characteristics of the uncertainty, such as exactly what is uncertain, how it is correlated with aggregate risks and how it interacts with investment decisions. The appropriate discount rate is often difficult to determine, yet crucial for long-lived projects. For example, sources of risk include the risk associated with next period's actual cash flow, the risk associated with revision of expectations, a changing riskless rate of return and changes in investment opportunities.

Consider global warming policy.[74] The effect of greenhouse gas emissions on the welfare of the future population is highly uncertain, in several dimensions: uncertainty over timing, uncertainty over magnitudes and uncertainty over its effect on the returns on investment. For example, the current willingness to pay to avoid a permanent loss of 1 per cent of GDP in 100 years time is much less than willingness to pay to avoid a 1 per cent annual probability of the same loss (which has the same expected timing). This, in turn is much less than willingness to pay to avoid a small probability of a large catastrophic loss (with the same expected loss). A small chance of a large loss with very uncertain timing (even if expected waiting time is large) can generate substantial willingness to pay.

70 See appendix I.3 for details.

71 See appendix I.5 for details.

72 See appendix I.2 for details.

73 See appendix I.4 for details.

74 The following summarizes Murphy (2008).

Murphy concludes that the gains from mitigation against modest climate change are greatest in the highest GDP states and so have a positive beta. For example, where losses are proportional to GDP (as the in the first two examples above), the appropriate beta is 1 and the market risk premium should be used in the discount rate. Mitigation against very bad climate outcomes pays off in the worst states and is likely to have a negative beta, and a discount rate below the riskless return — favouring policies that focus on avoiding extreme outcomes.

Determining the appropriate discount rate for the very long term is difficult, yet makes a huge difference to a project's value. It is worth carefully considering the relevant discount rate on a case by case basis, accounting for the risk and other characteristics of the particular project. If the discount rate in year i is r_i then a payment received in n years should be discounted by: $(1+r_1)(1+r_2)(1+r_3)\ldots(1+r_n)$. The starting point is always the current discount rate r_1. If the current discount rate is expected to persist for a few years, it will have the highest weight in the discount rate used for future payments.