Australian Government
Productivity Commission

Valuing the Future: the social discount rate in cost-benefit analysis

Visiting Researcher Paper

April 2010

Mark Harrison

The views expressed in this paper are those of the individual involved and do not necessarily reflect the views of the Productivity Commission.

ISBN 978-1-74037-313-5

This publication is available in hard copy or PDF format from the Productivity Commission website at www.pc.gov.au. If you require part or all of this publication in a different format, please contact Media and Publications (see below).

Publications Inquiries:
Media and Publications
Productivity Commission
Locked Bag 2 Collins Street East
Melbourne VIC 8003

Tel: (03) 9653 2244
Fax: (03) 9653 2303
Email: maps@pc.gov.au

General Inquiries:
Tel: (03) 9653 2100 or (02) 6240 3200

An appropriate citation for this paper is:

Harrison, M. 2010, *Valuing the Future: the social discount rate in cost-benefit analysis*, Visiting Researcher Paper, Productivity Commission, Canberra.

JEL code: D9, H4

The Productivity Commission

The Productivity Commission is the Australian Government's independent research and advisory body on a range of economic, social and environmental issues affecting the welfare of Australians. Its role, expressed most simply, is to help governments make better policies, in the long term interest of the Australian community.

The Commission's independence is underpinned by an Act of Parliament. Its processes and outputs are open to public scrutiny and are driven by concern for the wellbeing of the community as a whole.

Further information on the Productivity Commission can be obtained from the Commission's website (www.pc.gov.au) or by contacting Media and Publications on (03) 9653 2244 or email: maps@pc.gov.au

Contents

Preface

Governments face a bewildering array of policy choices, ranging from regulatory interventions with little direct cost to government and potentially quick payoffs (such as regulating product safety), through to complex investments in social and economic infrastructure costing billions and taking years to implement and decades to deliver their full benefits. Some policies have extremely long-term effects: for example, the costs and benefits of policies to address the uncertainties of climate change span centuries.

Cost-benefit analysis is a powerful tool to inform this range of decisions. People generally expect more than a dollar in the future to compensate them for forgoing a dollar today. The selection of discount rates is important to bring the estimated dollar impacts over time to a common point so as to establish whether a particular project has a present value of benefits greater than its costs, and to rank viable alternatives.

In August 2007, the Office of Best Practice Regulation (then part of the Productivity Commission) published its *Best Practice Regulation Handbook,* containing guidelines including use of a social discount rate of 7 percent real, with sensitivity testing over the range of 3 to 11 per cent (OBPR 2007, pp 129-132).

This Visiting Researcher Paper was initiated, as foreshadowed in the *Handbook*, to examine further the evidence on the parameters influencing the choice of discount rate. It was prepared by Dr Mark Harrison when a visiting researcher at the Commission. The paper reviews conflicting views on the issues influencing discount rate selection, and examines recent evidence of possible market benchmarks for discount rate derivation.

The Visiting Researcher papers

This publication is the second in an occasional series presenting the work of the Commission's Visiting Researchers.

In support of the Commission's core function of conducting public inquiries and studies commissioned by the Government on key policy and regulatory issues, the Commission conducts supporting research into diverse issues concerning productivity and its determinants, environmental and resource management, labour markets, and economic models and frameworks to aid policy analysis.

To assist the research effort, the Visiting Researcher Program seeks to attract established researchers in areas closely related to the Commission's primary research themes. Visiting Researchers share their expertise and knowledge with Commission staff and contribute to the work of the Commission during appointments which are usually for terms up to one year.

During their stay with the Commission, Visiting Researchers take an active interest in the work of the Commission and its staff. It is common for Visiting Researchers not only to conduct their own research in conjunction with Commission staff but also to contribute to the research work of others in their areas of expertise. The views expressed in this series are those of the Visiting Researchers themselves, and do not necessarily represent those of the Commission.

More details on the Visiting Researcher Program are available on the Commission's website at www.pc.gov.au/employment/researcher

OVERVIEW

Key points

- The discount rate is a critical parameter in cost-benefit analysis whenever costs and benefits differ in their distribution over time, especially when they occur over a long time period.
- Approaches to selecting real discount rates fall into two broad groups, both of which have given rise to a wide range of recommended rates:
  - a ‘descriptive’ approach based on the opportunity cost of drawing funds from the private sector; and
  - a ‘prescriptive’ approach that derives from ethical views about intergenerational equity.
- The arguments here support the descriptive approach and provide a starting point for the discount rate — the marginal return to private capital in Australia over the past four decades, which has averaged almost 9 per cent real.
- Market rates reflect the opportunity cost of investing in public projects, and there is no case for allocating resources to low return investments when higher returns are available. Using an artificially low discount rate for project evaluation can make future generations worse off. Ethical arguments for a low discount rate are more a reason to increase savings and investment.
- Government projects for which cost-benefit analysis can assist decision making cover a huge range: regulatory changes; infrastructure investments with significant gestation periods and long benefit streams, whose magnitudes are positively related to general economic conditions; and climate change policies with cost and benefit streams extending over centuries, but with high uncertainty.
- No single discount rate could meet the precise financing and risk characteristics of each project in this wide range of applications.
  - Taxes make a big difference between the before-tax ‘investment rate’ that investments earn and the after-tax ‘consumption rate’ that lenders receive, and a project’s discount rate choice should ideally reflect the extent to which its financing reduces investment and consumption.
  - Government projects are not in general free of risk; some have expected net benefits inversely related to aggregate consumption, but many have expected net benefits positively correlated with aggregate consumption. Discount rates should embody an appropriate compensation for risk. The rate should be equal to the rate of return on private projects with similar levels of risk. The market price of risk is what people have to be paid to bear risk and reveals attitudes to risk even where markets are imperfect.
- The appropriate adjustments for taxes and risk cannot be precisely estimated — one reason why sensitivity testing is important. A base rate of 8 per cent, and testing over a range of 3 to 10 per cent is proposed.

Overview

Cost-benefit analysis is used to improve decision making by systematically comparing the social costs and benefits of government policies, with the emphasis on valuing them (to the extent possible) in monetary terms. It provides decision-makers with quantitative information about the policy's likely effects and encourages them to take account of all the positive and negative effects and the linkages between them. Quantifying the impact of government policies in a standard manner promotes comparability, the assessment of relative priorities, and consistent decision making.

Moreover, the process of trying to describe and measure costs and benefits is valuable in itself. By examining what determines the costs and benefits and how they are likely to vary, policy makers are encouraged to consider different approaches and determine the best way to achieve objectives. Identifying and measuring costs and benefits encourages close examination of the factors that influence them and assists in minimising costs and maximizing benefit, helping decision makers increase net benefits to society.[1]

Cost-benefit analysis can be used to analyse and strengthen a wide range of government choices, including whether to undertake an infrastructure project, provide a service, pass a regulation, produce a public good, change a social welfare programme or adjust a tax.

Most government policies or projects give rise to a stream of costs and benefits over time. A key element of the cost-benefit analysis framework is the use of a discount rate to compare costs and benefits received at different points in time. Yet there is little agreement about the appropriate discount rate, with cost-benefit guides, academics and textbooks giving conflicting advice. A wide range of discount rates has been recommended, with the average and the bottom of that range falling over recent years.

The choice of discount rate can make a significant difference to whether the present value of a project is positive, and to the relative desirability of alternative projects, especially when costs and benefits accrue at different times and over long periods.

[1] See Australian Government (2007, p. 115) for the benefits of using cost-benefit analysis.

As table 1.1 illustrates, the higher the discount rate, the smaller the present value of future costs and benefits. The further in the future the payments are received, the greater the effect of the discount rate. A high discount rate favours projects with benefits that accrue early.

Table 1 **How present value of $1000 varies with when it is received and the discount rate**

Discount rate	*Years in the future*		
Per cent	10	50	100
1	$910	$608	$370
3	$744	$228	$52
8	$463	$21	$0.45
10	$386	$9	$0.07

Source: Author's calculations. The present value of $1000 received n years in the future with a discount rate *r* is $\$1000/(1+r)^n$.

Aims

This paper is about discount rate choice: how to discount estimated cost and benefit flows. It is not about whether cost-benefit analysis should be used, or how well costs and benefits can be quantified. These are separate debates.

Putting a dollar value on a project's costs and benefits may be a difficult task — it is not easy to estimate the benefits flowing from a public good, environmental improvements or safety measures that that save lives. There may be considerable uncertainty about predicted impacts and their appropriate monetary valuation. But that is a reason for conducting sensitivity analysis with the disputed variables and improving the estimates, and does not affect the case for discounting.[2] Defects in the measurement of costs and benefits should be directly addressed and do not justify adjusting the discount rate.

2 Sensitivity analysis provides information about how changes in values assigned to different variables will affect the overall costs and benefits of the policy proposal. It shows how sensitive predicted net benefits are to different values of uncertain variables and to changes in assumptions. It tests whether the uncertainty over the value of certain variables matters, and identifies critical assumptions. See Australian Government (2007, p. 122-23) for more on sensitivity analysis.

Why discount?

The case for discounting future costs and benefits arises from opportunity costs. The cost of investing a dollar in a project is what it would have produced in its alternative use. Because invested capital is productive, an extra dollar invested today in the private sector will (on average) grow to more than a dollar tomorrow, a fact reflected in positive market interest rates. Tying up a dollar today in any project requires a return of more than a dollar tomorrow to ensure the project covers its costs and produces benefits greater than from leaving the money in the private sector.

Dollars are valuable for what they can buy. Inflation is one reason a dollar in the future is worth less than a dollar now. The usual approach in cost benefit analysis is simply to express all costs and benefits in real or constant-price dollars, which avoids having to estimate the future course of inflation. But that requires the analyst to convert past nominal flows into real dollars and to specify a real discount rate.[3]

Expressing cost and benefit estimates in terms of constant purchasing power and using a real discount rate does not mean the prices of individual products and factors are fixed at current prices. Anticipated changes in the relative price of important outputs and inputs into the project (such as possibly rising relative values of environmental benefits), as distinct from general price level changes, should be reflected in the estimates of future costs and benefits.[4] The discount rate should be consistent with the dollar flows being discounted. If costs and benefits are measured in nominal (or current) dollars, they should be discounted with a nominal discount rate; costs and benefits measured in real terms (that is, adjusted for inflation), should be discounted with a real discount rate.

Different approaches to the discount rate

Academics, cost benefit guides and textbooks give widely conflicting advice on discount rate selection, with recommended rates varying from 1 to 15 per cent, with the rates recommended in most developed economies trending down over recent decades. The two major schools of thought are the prescriptive and descriptive approaches to discount rate selection (chapter 2).

The prescriptive or normative approach directly specifies a discount rate influenced by ethical principles (sometimes literally deriving from the assessments of

[3] See appendix A, section A.2.

[4] Harberger (1969, footnote 16, p. 122).

philosophers). It mixes efficiency and equity considerations, and is frequently advocated when projects affect future generations. The prescriptive approach gives a wide range of suggested discount rates, reflecting different value judgements that cannot be resolved objectively.

The descriptive approach to the social discount rate is based on the opportunity cost of capital used in the project: what benefits to society the funds would return if left in the private sector. It is based on the efficiency criterion.

Chapter 2 makes the case that project choice should be based on discounting with an efficiency-based social discount rate, even for projects implying benefits or costs for future generations. But focusing on an efficiency-based social discount rate still leaves plenty of room for disagreement. Complications include the effects of capital taxes, capital market imperfections and uncertainty. Theory about ideal adjustments for these factors is far ahead of our empirical knowledge. Selection of a social discount rate depends on parameters that can only be imperfectly estimated, and reasonable people can make different judgments about them. Recommendations about the social discount rate should be practical and account for this imperfect empirical knowledge.

The discount rate and market benchmarks

Chapter 3 explores the utility of market benchmarks for deriving practical estimates of the discount rate applicable to various projects.

Capital taxes drive a substantial wedge between the before-tax investment return and the after-tax consumption rate of interest, making it important to distinguish between the project's financing impacts on investment and consumption.

If government investment comes at the expense of private investment, the cost to the economy is measured by the social returns that would have been generated by that investment. This has been variously labelled the investment rate of interest, the producer rate of interest, the marginal rate of return to investment or capital, the marginal efficiency or product of capital, or the social opportunity cost of capital.

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 is usually measured by the after–tax real rate of return on savings — the supply price of savings. The discount rate that reflects the opportunity cost of capital is a weighted average of the consumption rate, the investment rate, and the marginal cost of foreign funds with the weights being the proportion of project costs sourced from consumption, from investment and from foreign funds. In practice, there is little information about the weighting of each source, or even the precise cost of each source. Nevertheless, to use any particular social discount rate is to implicitly assume something about these variables, and it is better to make the assumptions explicit. What is needed is a reasonable rule of thumb. For reasons elaborated in Chapter 3, there is a reasonable presumption that the weighted average rate will lie close to the before-tax investment rate.

The risk free before tax rate of return in Australia is around 4 per cent, and so the risk free social discount rate is significantly greater than rates usually derived from the Ramsey equation approach (chapter 3) – which produces a consumption rate (often around 1 to 2 per cent).

Estimates of the rates of return on capital in the market sector in Australia are also presented in Chapter 3. A reasonable estimate of the marginal rate of return to capital (or the opportunity cost of forgone private investment) is 9 per cent real. This market return includes a risk premium that compensates investors for the risk they bear.

Estimates of uncertain future costs and benefits should be a risk-weighted average (or 'expected value)' of all possible outcomes — including possible disasters and windfalls – accounting for what could happen, not just what should happen. The most practical way to account for the costs of project risk is to discount these expected values with a discount rate that includes a risk premium, whether discounting costs or benefits. Some have argued that government projects should be discounted using a risk-free rate. However there is an element of risk — aggregate risk or irreducible social risk — that cannot be diversified, even by the government. It is caused by shocks such as recessions and variations in the market return, which affect consumption. This undiversifiable aggregate risk should be reflected in a risk premium in the discount rate for policies or projects whose net benefits are likely to be positively correlated with the performance of the overall economy. 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'.

Taxpayers and program beneficiaries are essentially equity holders in the government project. When a risky stream of payments from a government project accrues to individuals, they incur the costs of risk. It is their valuation of the stream that is relevant. The market price of risk is what people have to be paid to bear risk and reveals attitudes to risk even where markets are imperfect.

Even if the government can spread risks the market cannot, if the government has the option of investing in the private sector, say through a sovereign wealth fund, the private market return is the opportunity cost of investing public funds and is the appropriate discount rate for public investment projects.

A few policy decisions have potentially significant impacts many generations hence. Climate change policies are a topical example. Some arguments point to lowering the rate used to discount the distant future, while others suggest raising the rate. For example, uncertainty about the long term path of interest rates suggests that, on the one hand, we should use a lower real interest rate to discount costs and benefits received further in the future (the yield curve for real interest rates slopes down). Moreover, the appropriate risk premium may fall with the length of project. On the other hand, the option value of delay may in some cases favour a higher rate.[5] Uncertainty about how to adjust the rate makes it particularly important for sensitivity testing a range of discount rates for analysis of very long-lived projects or policy choices.[6]

Practical implications

The arguments developed in this paper suggest using a discount rate based on the marginal rate of return on capital, such as the national accounts measure of the before-all-tax real rate of return on private capital. This has averaged 8.9 per cent over long time periods, and is more stable than share market returns.

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.

The resulting weighted average market rate includes the market risk premium. It is the appropriate discount rate for a government project which has the same risk as the average private sector investment. The rate reflects the opportunity cost — the

[5] Chapter 3, section 3.6 and Appendix I, discusses further how the value of delaying a project may alter the appropriate discount rate to use in evaluating it.

[6] PC (2008) develops this example in the context of the Stern Review of the economics of climate change.

funds used in the government project would have produced this return if left in the private sector, with no greater risk. 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 whose value is inversely correlated with economic conditions.

In the absence of further information, 8 per cent is a reasonable default discount rate, but there is still considerable imprecision in empirical estimates. The weights and returns to use in the weighted average discount rate, and the appropriate risk premium, are not clear. As with any uncertain variable, sensitivity analysis provides valuable insights.

Sensitivity testing using real rates of 3, 8 and 10 per cent is proposed in this paper — 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 suggested dispersion around 8 per cent is not symmetrical, because there are fewer arguments for much higher rates than for much lower rates.)

If the sensitivity analysis reveals that the choice of discount rate changes the sign of the project's net present value or changes the ranking of alternatives, 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 rate of return is above plausible discount rates, then the project is likely to improve economic efficiency, an important consideration that should inform the decision about whether to adopt the project.

1 The concept of discounting in cost-benefit analysis

Cost-benefit analysis is a tool to improve decision-making. Identifying and measuring costs and benefits encourages close examination of the factors that influence them and assists in minimising costs and maximizing benefit, helping decision makers increase net benefits to society.[1]

1.1 Introduction

Cost-benefit analysis uses willingness to pay to measure benefits and opportunity cost to measure costs.[2] The opportunity cost of resources is their value in the alternative use to which they would have been put.[3] Identifying the costs and benefits of a policy change involves comparing outcomes with the proposed change to outcomes without the change. Analysts can measure the value people place on something by observing how much they are willing to pay. Market behaviour often reveals people's valuations, or is at least a guide to them. Wherever possible, money values of benefits and costs should be based on tradeoffs that individuals would make in markets.[4]

Most government policies give rise to a stream of costs and benefits over time. To evaluate them requires us to compare costs and benefits received in different time periods. That requires choosing a discount rate, which determines the value of future costs and benefits relative to current ones. The choice of discount rate can make a huge difference to the desirability of government projects, especially when their costs and benefits occur over long periods.

---

[1] See Australian Government (2007, p. 115) for the benefits of using cost-benefit analysis.

[2] See Department of Finance and Administration (2006); Australian Government (2007, appendix B, pp. 115–34); and Weimer (2008).

[3] Not the best alternative use to which they *could* have been put, which is how opportunity cost is sometimes defined in textbooks. The two definitions only coincide in a perfectly functioning economy, where if resources are not used in one activity they *would* be used in the most valuable alternative they *could* be used, but that is not generally true in a distorted economy. See Feldstein (1972, pp. 319–320).

[4] Arrow et al. (1996, principle 15, p. 11).

Yet there is little agreement about the appropriate discount rate, with cost-benefit guides, academics and textbooks giving conflicting advice. What influences discount rate choice? Should we use the same discount rate in all applications? If not, what factors make a higher or lower rate appropriate for particular uses? But first, what is a discount rate? Why is it used? Why is it important?

1.2 The discount rate in cost-benefit analysis

How do we weigh future costs and benefits against current costs and benefits? A convenient way to compare, and add up, costs and benefits that accrue at different times is to calculate their present value, which expresses them as an equivalent amount of today's dollars.

Discounting converts the dollar value of costs and benefits received in different time periods to present value. The term discounting refers to the fact that a dollar received in the future is worth less than a dollar now. The value of future dollars relative to current dollars is expressed in terms of a time preference rate or a discount rate and is usually expressed as a percentage rate for a period of one year. If the discount rate were constant at ρ per cent per year, a benefit of B_t dollars received in t years has a present value of $B_t/(1+\rho)^t$ dollars.[5] To calculate a project's net present value, we add up all the present values of all its costs and benefits over time, which measures how much the project increases wealth — how much extra consumption it generates for society. The higher the net present value, the more valuable the project.

Any evaluation of policies with future costs and benefits must specify a discount rate. Choosing a discount rate is the same as choosing the value of future dollars. Failure to discount implies a discount rate of zero, which means a future dollar, however distant, counts as much as one received immediately.

Cost-benefit analysis presents estimates of the costs and benefit of a policy in dollar terms. Choosing a discount rate is about putting relative values on estimates of costs and benefits received in different time periods. This paper is not about whether cost-benefit analysis is reasonable or how well costs and benefits can be quantified. It is

[5] More generally, if the discount rate in period j is r_j then a benefit of B_j dollars received at the end of period j is worth $B_j/[(1+r_1)\ (1+r_2)\ (1+r_3)\ \ldots .\ (1+rj)]$ at the beginning of period 1. There is no problem with having discount rates varying over time: we can have a term structure of interest rates (see, for example, the appendix to Harberger 1969). In fact, if credit market conditions are temporarily unusual (or likely to change), it is desirable to have the discount rate changing towards its normal (or predicted) level over the next few years (Harberger 1969, p. 117).

about discount rate choice: how to discount the estimates that are made. Whether cost-benefit analysis should be used is a separate debate.

Putting a dollar value on a project's costs and benefits may be a difficult task — it is not easy to estimate the benefits flowing from a public good, environmental improvements or safety measures that that save lives. There may be considerable uncertainty about predicted impacts and their appropriate monetary valuation. But that is a reason for conducting sensitivity analysis with the disputed variables and improving the estimates, and does not affect the case for discounting.

The discount rate should not be adjusted because of defects in the measurement of costs and benefits. A better approach is to directly adjust the cost and benefit estimates. A change in the discount rate will have a larger effect the further into the future the cost and benefits are received, which generally would not be the appropriate adjustment for measurement problems.

1.3 Why discount?

An important distinction in economics is between investment and consumption. Investment activities primarily affect future well-being; the main impact of consumption activities is in the present. An investment project involves bearing a current cost that is expected to give future benefits. That is, current output is directed to producing future benefits rather than current benefits. Investment is any expenditure on increasing, maintaining and improving capital. Capital is an asset that produces future benefits.

Although the common view of capital is confined to physical capital, such as buildings and machines, this accounts for a small part of the total stock of capital in countries that have achieved a high level of average income. Other forms of capital include the stock of skills and productive knowledge embodied in people (known as human capital) and the framework of laws and regulations (sometimes referred to as institutional capital).

The term government project, therefore, is a broad one that applies to any government decision or policy that affects future consumption. It may be providing a public good, building infrastructure or the introduction of a new regulation. It may involve public investment or it may be a regulation that requires the private sector to invest, with government costs being limited to monitoring and enforcement costs. For example, hiring labour to build a stadium is an investment in a capital good. Requiring more scaffolding to reduce future deaths or injuries or equipment to reduce future pollution emissions is also an investment. So is spending time training to learn new skills that increase future productivity or resources spent implementing or complying with a regulation that increases future safety or prevents an endangered species being wiped out. All involve current costs, borne for expected future benefits.

The case for discounting depends on opportunity costs. The cost of investing funds in a project is what it would have produced in its alternative use. Because capital is productive, an extra dollar invested in the private sector will (on average) grow to more than a dollar. A future dollar, therefore, is worth less than a current dollar, a fact reflected in positive market interest rates. For example, if $1 invested in the private sector would grow to $1.07 next year, the annual rate of return to investment (the marginal rate of return to capital) is 7 per cent.[6] One dollar received in one year's time is worth $1/1.07 = $0.93 now. A public project would only cover its opportunity cost if it earned an annual rate of return of at least 7 per cent.

Equivalently, if a project's net benefits had a positive net present value using a discount rate of 7 per cent, then its benefits are greater than its costs and the project would increase wealth.[7] For example, if a project cost $50 and produced benefits of $60 in one period (a rate of return of 20 per cent), its net present value is $60/1.07 – $50 = $6.07. If the $50 was left in the private market it would produce $50*1.07 = $53.50 of benefits in period 2 dollars. The project produces enough benefits to fully compensate people for the cost of the resources it uses and have $6.50 (= $60 – $53.50) left over (which has a present value of $6.50/1.07 = $6.07).

The analysis generalises to any stream of costs and benefits from a project. If the net present value is positive using a discount rate that reflects the return capital would have earned in the private sector, the project covers its opportunity costs: the benefits from using capital in the project are greater than from leaving it in the private sector. When a government project draws on savings to make an investment, it is sometimes loosely said that it uses capital, because capital can refer to the asset created by investment or the savings that finance it.

Cost-benefit analysis converts all costs and benefits to money equivalents based on willingness to pay. The project is exactly equivalent to receiving that flow of dollars over time. The benefit estimates represent the sum of money those who would benefit from the policy are willing to pay to receive the benefits. The cost estimates are the amount of money that needs to be paid to compensate for the costs the

[6] If the rate of return on an asset is i, then the return on the asset is $1+i$. If the rate of return is taxed at rate τ then the after-tax rate of return on the asset is $i(1-\tau)$ and the after-tax return is $1+i(1-\tau)$.

[7] A project produces a stream of costs and benefits. The project's benefit less cost in period t is the net benefit in period t. The net benefit produced in any period may be positive or negative.

project imposes. It is argued here that these streams of money payments should be discounted using a rate derived from the return in the private sector. That compares project benefits with the benefits the capital invested in the project would have produced if left in the private sector.

1.4 Real and nominal discount rates

The discount rate should be consistent with the dollar flows that are measured. If costs and benefits are measured in nominal (or current) dollars, they should be discounted with a nominal discount rate. Costs and benefits measured in real terms (that is, adjusted for inflation), should be discounted with a real discount rate. Both methods should result in the same net present value.[8]

Market interest rates are usually nominal rates — showing the rate at which dollars today are traded for dollars in the future, revealing how savers and investors value future dollars. If savers are willing to lend *$1* in return for $\$1+\rho$ next year, then the annual nominal interest rate is ρ per cent.

Market interest rates are usually positive, indicating people generally value a dollar in the future less than a dollar now. When money is given up now for money one period from now (the act of saving), you give up the right to use or consume the money during the period, which could include investing it. That right is valuable, so dollars today are worth more than dollars tomorrow.

Dollars are valuable for what they can buy. Inflation is one reason a dollar in the future is worth less than a dollar now. A general rise in the price level means a dollar buys fewer goods. Future costs and benefits can be valued in nominal or real dollars. In the nominal (or current) dollars approach, the impact of expected inflation is explicitly reflected in the projections (the cost and benefit streams grow faster when expected inflation increases). The real (or constant price) approach expresses all variables in terms of the price level of a given year, usually the present year. A real dollar has the same purchasing power at any time (can buy the same bundle of goods).

The usual approach in cost-benefit analysis is simply to express all costs and benefits in real dollars, which avoids having to estimate the future course of inflation. But that requires the analyst to convert past nominal flows into real dollars and to specify a real discount rate.[9]

---

[8] See appendix A, section A.1.

[9] See appendix A, section A.2.

The real interest rate is the rate at which real dollars (or goods) today are traded for real dollars (or goods) tomorrow — either directly or through converting goods into money, trading money this year for money next year (through lending it at the nominal interest rate) and then buying goods next year.[10]

As people consume — and get utility from — goods, not dollars, it is the real interest rate that is relevant when people decide how much to consume now and in the future.

Again, because present command of resources expands opportunities, people are willing to pay for that command, and people who give it up need to be compensated for doing so.[11] The real interest rate, therefore, is usually positive — goods received now are worth more than goods to be delivered in the future.

Expressing cost and benefit estimates in terms of constant purchasing power does not mean pricing individual products and factors at current prices. Anticipated changes in the relative price of important outputs and inputs into the project, as distinct from general price level changes, should be reflected in the estimates of future costs and benefits.[12] For example, cost-benefit studies that deal with environmental and safety issues need to account for the changing value of non-marketed goods and amenities and of the value of statistical life. Both are likely to grow at least as fast as average incomes. For example, if climate change damages the environment, the relative price of environmental amenities will increase, raising the value placed on marginal units. When the relative price of an important output or input (including the discount rate) is currently out of line with its long-term norm, this should be identified, and the path by which the price is expected to move back to its norm should be specified.[13]

1.5 The discount rate's impact on project viability

The net present value of any project with future costs and benefits crucially depends on the discount rate chosen, especially when the costs are borne in a different time frame than the benefits are received. As illustrated in table 1.1, when the discount rate is higher, future costs and benefits count for less. A high discount rate favours projects with benefits that accrue early.

[10] See Friedman (1990, pp. 333–37) for a discussion of real and nominal interest rates.

[11] See Heyne (2002)

[12] Harberger (1969, footnote 16, p. 122).

[13] Harberger (1992, p. 14).

The discount rate has an important effect on investment decisions. A typical project involves upfront costs, with the benefits coming later. If so, the lower the discount rate, the more attractive is the project (the higher its net present value). If the discount rate is set too high, desirable projects may be rejected. If it is set too low, undesirable projects may be approved. The size of the discount rate makes a huge difference to policies where benefits occur in the distant future, such as many environmental policies. Lower discount rates encourage investors to adopt projects that offer returns at distant dates. It was a lower discount rate that drove the differences between the policy conclusions of the Stern report and the consensus view of previous cost-benefit analyses of global warming. Stern's cost-benefit analysis of global warming assumed a real discount rate of 1.4 per cent and concluded there was a case for an immediate imposition of a high, and increasing, carbon price. Nordhaus assumed a 5.5 per cent discount rate, and favoured a modest carbon price, increasing over time. The recommended polices differed because of the different discount rate assumptions. When Nordhaus ran his computer model using Stern's discount rates, he got similar results to Stern.[14]

The difference is not surprising because most of the effects of global warming take place decades in the future. At Stern's discount rate of 1.4 per cent per year, $1 grows into $4 in 100 years and $16 in 200 years. At Nordhaus's 5.5 per cent, it would grow to $211 in 100 years and $44 719 in 200 years. Put differently, the present value of $1 of damages in 100 years with Stern's discount rate is 25 cents, more than 50 times greater than with Nordhaus's. A dollar in 200 years is valued at 6 cents now, almost 2800 times more than Nordhaus's value.

[14] See Nordhaus (2007, pp. 694, 698-700); Baker et al. (2008, pp. 63-64) and Weisbach and Sunstein (2008, p. 11).

2 Approaches to discounting

While everyone agrees that the choice of discount rate is a crucial determinant of the value of public projects, there is less agreement on the appropriate discount rate to use to calculate present value. Academics, cost-benefit guides and textbooks give widely conflicting advice.

2.1 Different views on the social discount rate in practice

International practice is summarised in table 2.1 (countries are ranked by the discount rate), with recommended rates varying from 1 to 15 per cent.

The table shows that the highest rates are used in developing countries.[1]

Low discount rates are often used in environmental applications, especially when benefits accrue in the distant future. The United States Environmental Protection Agency recommends a discount rate of 2-3 per cent, 'the consumption rate of interest', and also using 7 per cent.[2] It recommends no discounting for inter-generational projects, and sensitivity testing with 2-3 per cent and 7 per cent.[3]

[1] Views that there is a higher opportunity cost in developing countries are not reflected in the realised returns to capital. Recent evidence suggests that the marginal rate of return to capital in developing countries is no higher than in developed countries, Caselli and Feyrer (2007, p. 555). The preferred estimates average 8.4 per cent real in rich countries and 6.9 per cent in poor. These are gross rates of return for 1996. They do not include any capital gains from holding capital.

[2] United States Environmental Protection Agency (2000, p. 48).

[3] United States Environmental Protection Agency (2000, p. 52).

Table 2.1 **Current real discount rates in practice**
International

Country	*Agency*	*Discount rate (per cent)*
Philippines		15[a]
India		12[a]
Pakistan		12[a]
International Multi-lateral Development Banks	World Bank	10–12[a]
	Asia Development Bank	10–12[a]
	Inter-American Development Bank	12[a]
	European Bank for Reconstruction and Development	10[a]
	African Development Bank	10–12[a]
New Zealand	Treasury and Finance Ministry	8[g]. From 1982 to 2008 it was 10[abf]
Canada	Treasury Board	8[c]. From 1976-2007 was 10 (and test 8–12 per cent)[ab]
China (People's Republic)		8[a]
South Africa		8 (and test 3 and 12 per cent)[d]
United States	Office of Management and Budget	7 (and test 3 per cent). Used 10 per cent until 1992.[a]
European Union	European Commission	5 From 2001–2006 was 6 per cent[a]
Italy	Central Guidance to Regional Authorities	5[a]
The Netherlands	Ministry of Finance	4 (risk free rate).e
France	Commissariat General du Plan	4. From 1985-2005 used 8 per cent[ab]
United Kingdom	HM Treasury	3.5 (declining to 1 per cent for costs and benefits received more than 300 years in the future) from 2003.[a] From 1969–78 used 10 per cent[a]
Norway		3.5. From 1978–98 used 7 per cent[ab]
Germany	Federal Finance Ministry	3. From 1999–2004 used 4 per cent[ab]
United States	Environmental Protection Agency	2–3 (and test 7 per cent)[a]

[a] Zhuang et al. (2007, table 4, pp. 17-18, 20). [b] Spackman (2006, table A.1, p. 31). [c] Treasury Board of Canada (2007, p. 37, 1998, p. 45). [d] South African Department of Environmental Affairs and Tourism (2004, p. 8). [e] van Ewijk and Tang (2003, p. 1). [f] Use of the 10 per cent rate by New Zealand government departments is confirmed by Young (2002, p. 12); Abusah and de Bruyn (2007, p. 4). [g] New Zealand Treasury (2008) recommends a default rate of 8 per cent (after adjusting the market risk premium of 7 per cent for gearing).

Some, following suggestions in the academic literature, use a lower rate to discount costs and benefits that accrue decades in the future. For example, the UK Treasury uses 3.5 per cent real, which it identifies as the ‘social time preference rate’. The rate slowly declines to 1 per cent for costs and benefits accruing more than 300 years in the future.[4] The Stern Report used 1.4 per cent real to discount the benefits from greenhouse gas emission abatement policies.[5] The Garnaut Report used 1.35 per cent and 2.65 per cent.[6] Both would be at the bottom of the range in table 2.1.

Not that academics agree on the appropriate discount rate, with numerous symposiums of papers from leading economists demonstrating their conflicting views.[7] There is academic support for the investment (or producer) rate i (the before-tax rate of return), the consumption (or consumer) rate r (the after-tax rate of return) and for a weighted average of the two, $w = ai + (1 - a)r$.[8] Others oppose the use of the weighted average.[9] Some conclude that the appropriate discount rate need not even lie in the range between r and i.[10]

Academics also disagree on whether to include a risk premium.[11] Some reject adding a risk premium, but recommend converting expected values into certainty

---

[4] HM Treasury (2003, pp. 98-100).

[5] Cole (2007, p. 10) and Weitzman (2007, p. 708). It is difficult to find an explicit statement of the discount rate used in the Stern report itself or in the Technical Annex which sensitivity tests the discount rate.

[6] Garnaut (2008, p. 19).

[7] For example, the Intergenerational Equity and Discounting Symposium, *University of Chicago Law Review* (2007, vol. 74, Winter); Lind (1982) and Portney and Weyant (1999a).

[8] Those recommending the investment rate include Kaplow (2006b) and Stiglitz (1982). The consumption rate (usually the after-tax return to savers) is recommended by Arrow et al. (1996) and Feldstein (1972). Harberger (1969); Sandmo and Dreze (1971) and Sjaastad and Wisecarver (1977) recommend a weighted average of the consumption and investment rates.

[9] Feldstein (1972); Bureau of Transport and Regional Economics (1999); Abelson (2000, pp. 129) and United States Environmental Protection Agency (2000, pp. 41–2).

[10] See Stiglitz (1982) and Bradford (1975).

[11] For example, Arrow (1966); Spackman (2004); Viscusi (2007) and Grant and Quiggin (2003) reject using a market risk premium.. Those who argue for a market risk premium to be part of the discount rate include Sandmo (1972); Jensen and Bailey (1972); Lind (1982a); Hirshliefer (1964); Klein (1997); Hathaway (1997); Bazelon and Smetters (1999); Currie (2000); van Ewijk and Tang (2003); Grout (2003); Kaplow (2006b); Brealey et al. (1997); Nordhaus (2007); Dixit and Williamson (1989) and Wright et al. (2003).

equivalents.[12] Some guides explicitly reject adding a risk premium,[13] some implicitly reject it by recommending the government bond rate, usually considered to be the risk free rate. Others explicitly recommend a risk premium,[14] or implicitly by recommending a private sector rate that includes a risk premium over the risk free rate.

Weitzman (2001) surveyed professional Ph.D. level economists about their 'professionally considered gut feeling' for the discount rate that should be used to evaluate the costs and benefits (measured in real dollars) of mitigating climate change. He received 2 160 responses, with a sample mean at around 4 per cent per year, a standard deviation of around 3 per cent, a median of 3 per cent and a mode of 2 per cent. The suggested rates varied from -3 per cent to 27 per cent.

Textbooks also vary in their recommended rates. For example, Perkins (1994) suggests 7 to 13 per cent real. Boardman et al. (2006) recommend 3.5 per cent, with a lower rate for benefits more than 50 years in the future.[15]

The difference reflects a trend towards using lower discount rates. Many agencies have reduced the discount rate they use (as set out in table 1.2). The UK Treasury recommended 10 per cent real from 1969-78. The German Federal Finance Ministry uses 3 per cent real, down from 4 per cent in 1999. Since 1998 Norwegian authorities use 3.5 per cent, down from 7 per cent real (used from 1978). The French Commissariat General du Plan uses 4 per cent real, changed in 2005 from 8 per cent. The US Office of Budget Management specified 10 per cent real until 1992, but now uses 7 per cent real (with sensitivity testing for 3 per cent).[16] From 1974 the Canadian Treasury recommended 10 per cent, but reduced it to 8 per cent in 2007.[17] From 1982 the New Zealand Treasury recommended 10 per cent, reduced to 8 per cent in 2008. And these are only the agencies known to have reduced their rates.

---

[12] Boardman et al. (2006, p. 238; footnote 3, p. 271); Bureau of Transport and Regional Economics (2005) and Abelson (2000, pp. 130–31). Certainty equivalents are defined in chapter 4, box 4.1.

[13] Bureau of Transport and Regional Economics (2005, 1999, p. 74, 78) and the Victorian Competition and Efficiency Commission (2007).

[14] Partnerships Victoria (2003); Department of Finance (1991) and Infrastructure Australia (2008a).

[15] Perkins (1994, p. 310) and Boardman et al. (2006, p. 270). Throughout the chapter on the social discount rate, Boardman et al 'assume benefits and costs are measured in terms of certainty equivalents' (p. 238).

[16] Spackman (2000, table A.1).

[17] Treasury Board of Canada (2007, p. 37).

In Australia, some agencies (including the Commonwealth's Office of Best Practice Regulation) recommend rates around 7 per cent real, usually justified as being approximately the before-tax rate of return on private investment (the investment or producer rate). For example, the NSW Treasury recommends using a real rate of 7 per cent (with sensitivity tests using 4 and 10 per cent).[18] The Australian Department of Finance used to recommend 8 per cent real, and still supports using the before-tax rate of return on investment, but does not prescribe a specific rate because it may vary over time and between proposals.[19] Likewise, the Queensland Treasury used to recommend 6 per cent,[20] but now requests that it be consulted over the appropriate rate (mainly to determine the appropriate risk premium).[21] The Office of Best Practice Regulation uses 7 per cent real (with sensitivity testing at 3 and 11 per cent).[22] Infrastructure Australia recommends cost-benefit studies submitted to it should use 'real risk free' discount rates of 4, 7 and 10 per cent.[23] The Commonwealth Department of Health and Aging and enHealth Council recommends evaluating environmental health policies with a discount rate of 5 per cent, with sensitivity tests ranging from 3 to 7 per cent.[24]

Another group tends to choose around 3 to 3.5 per cent, usually based on the before-tax rate of return on government bonds. The justifications given vary; the rate is said to represent the social rate of time preference, the consumers' rate of time preference (the consumption rate of interest), the risk free rate, or the government's cost of funds. For example, the Victorian Competition and Efficiency Commission recommends 3.5 per cent, 'a recent average of the ten year Commonwealth bond rate to determine the risk free opportunity cost of capital'.[25] The Victorian Department of Treasury and Finance endorses 3.5 per cent (although it adds a risk premium based on a market risk premium of 6 per cent when assessing private sector bids for public-private partnerships).[26] The South Australian Treasury also uses the long-term government bond rate as a risk free rate, which it estimates to be

[18] New South Wales Treasury (1997, p. 52).

[19] Department of Finance and Administration (2006, p. 66-68).

[20] Department of Health and Ageing and enHealth Council (2003, p. 21).

[21] Queensland Treasury (2006, p. 27).

[22] Australian Government (2007, p.120).

[23] Infrastructure Australia (2008b) p.14. Infrastructure Australia (2008a, p. 5, 36) recommends the risk free rate should be based upon a long-term government debt instrument (issued by the relevant government) and then the CAPM model be used to price risk, using a market risk premium of 6 per cent to value the market risk transferred to the private sector.

[24] Department of Health and Ageing and enHealth Council (2003, p. 2).

[25] Victorian Competition and Efficiency Commission (2007, p. 2).

[26] See Victorian Department of Treasury and Finance (2007, p. 5-16) and Partnerships Victoria (2003, p. 8).

5 per cent real.[27] The Tasmanian Treasury recommends the long-term Commonwealth bond rate plus 1 per cent as 'the long term cost of funds to the Government'.[28]

One overview of the literature concludes that 'those looking for guidance on the choice of a discount rate could find justification for a rate at or near zero, as high as 20% and any and all values in between.'[29]

Clearly there is no professional consensus on what discount rate should be used. The appropriate response to the uncertainty about the appropriate discount rate is to conduct sensitivity analysis with it. For example, this paper advocates calculating the net present value of the project with varying discount rates, say 3, 8 and 10 per cent real.

It may be that the project is so bad (or good) that varying the discount rate does not affect evaluation outcomes. For example, if the project results in a stream of net costs each period, it will have a negative net present value at any discount rate.

When a project gives rise to a string of negative and positive net benefits over time, the issue is whether the it is worth bearing the negative flows (costs) in order to reap the positive flows (benefits) at other time periods.

If the sensitivity analysis 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), then more consideration should be given to the choice of an appropriate rate. The purpose of this paper is to provide guidance on how to choose the rate.

Careful consideration of observed data and of the theoretical literature can help set the appropriate discount rate, clarify what judgements or circumstances might warrant selection of higher or lower discount rates and narrow the range of disagreement. The first step is to recognise that cost-benefit analysis uses the yardstick of efficiency.

2.2 Cost-benefit analysis and the efficiency criterion

Cost-benefit analysis focuses on the efficiency effects of a change: the net dollar value of the gains and losses for all people the change affects. It is based on willingness to pay: how many dollars individuals would, if necessary, pay to obtain

27 South Australian Department of Treasury and Finance (2007, p. 57).

28 Department of Treasury and Finance Tasmania (1996, p. 16).

29 Portney and Weyant (1999, p. 4).

(or avoid) a change, measures how much it is worth to them. The amount could be positive or negative, depending on whether the change makes them better or worse off. Summing these amounts across all affected people gives the community's total willingness to pay for the change. If the sum is positive, the benefits exceed the costs and the change would increase efficiency.

In standard efficiency-based cost-benefit analysis, the costs and benefits to all people are added without regard to the individuals to whom they accrue: a one dollar gain to one person cancels a one dollar loss to another person. This 'dollar is a dollar' assumption separates resource allocation from distribution effects — or efficiency from equity effects. That does not mean distributional considerations are unimportant or should be neglected. It means that they should be brought into account as a separate part of the overall analysis of the policy proposal in question — which may be more important than the resource allocation assessment, but should be distinct from it. This separation of efficiency and equity is useful because, in practice, a general consensus about the weight that should be attached to the welfare of different groups is elusive. By contrast, the efficiency or resource allocation effects of a policy has a clear meaning and can, in principle, be objectively measured. Dollar values can be estimated from observed behaviour.

While efficiency is an important goal of public policy and should be given due weight, it is only one consideration. Working out the efficiency effects is not to say that they should determine the whether the project should go ahead, but it is one factor that is worth knowing and should inform the decision. A policy could fail an efficiency test and a government might still judge it worth doing (for example, if it particularly valued the distribution of benefits and costs that it entailed, and thought that impact could be better delivered by the project than by other means such as the tax and transfer systems). Ultimately, it is up to governments, rather than technical analysts, to decide the trade-off between efficiency and other objectives, such as equity. Ideally, these other objectives should be made explicit. A cost-benefit analysis can help to inform the decision and clarify the trade-offs when comparing alternative policy proposals, such as how much income may need to be sacrificed to achieve other objectives.

The cost and benefit estimates should involve a comprehensive and systematic evaluation of all the impacts of a project, accounting for all the effects on all the people in society: not just the immediate or direct effects, not just financial effects, and not just the consequences for one group. All the effects of a policy proposal that are considered desirable by those affected are benefits, all undesirable effects are costs. Cost-benefit analysis requires analysts to identify explicitly the ways in which the proposal makes individuals better or worse off. The relevant costs and benefits are not limited to financial or pecuniary benefits, but are all social benefits including cultural, environmental and other non-market losses and gains. Sometimes non-monetary benefits are difficult to identify and evaluate, but they should not be ignored.

The impacts should be quantified for each time period over the life of the policy proposal. When does a flow of costs and benefits that accrue in different time periods increase efficiency?

2.3 Defining the social discount rate

The social discount rate extends the efficiency criterion to the case where costs and benefits occur over time. If the social discount rate is used to calculate the net present value of a project's social costs and benefits over time, a positive net present value indicates the project increases efficiency or raises wealth: it produces enough benefits to fully compensate individuals for the forgone benefits of the resources it displaces from alternative uses.

The present value of a stream of payments is called its capital value. When capital is traded in capital markets, the price paid for the rights to a stream of payments provides a market valuation of its net present value. Examples of capital markets include: the housing market (the value of the house is the present value of future housing services); the share market (the value of a company's shares is the present value of future dividends); and the bond market (the value of a bond is the present value of its stream of principal and interest payments). Much capital is not traded in any market, including the capital created by many government projects. Cost-benefit analysis tries to estimate the capital value of government projects. If the present value is positive, the project adds to wealth.

The efficiency-based approach to the social discount rate, which dates back at least to Harberger (1969), boils down to determining the opportunity cost of capital used in the project: what benefits to society would the funds have returned if left in the private sector. This 'opportunity cost' is the appropriate discount rate to determine a project's capital value.

But even focusing on an efficiency based social discount rate leaves plenty of room for disagreement. Complications include the effects of capital taxes, capital market imperfections and uncertainty. It turns out that the social discount rate depends on parameters that little is known about or that are difficult to estimate, and reasonable people can make different judgments about them. The theory is far ahead of our empirical knowledge. Recommendations about the social discount rate should be practical and account for our ignorance.

But the first line of disagreement is about what the social discount rate should be measuring. Some label the efficiency-based social discount rate derived from opportunity cost as the descriptive approach, and instead argue for a prescriptive approach,[30] especially when considering policies that affect future generations. One reason for the increasing use of lower discount rates is the explicit or implicit use of the prescriptive approach.[31]

The descriptive versus prescriptive debate has also been described as the positivists versus the ethicists.[32] Users of each approach apply the label 'social discount rate' and often it is not clear which has been adopted. Indeed, the Weitzman survey reported on page 12 did not specify which type of social discount rate those surveyed were to recommend.

2.4 The prescriptive approach

The prescriptive or normative approach directly specifies a discount rate that society 'should' use to discount future consumption flows, based on ethical principles (sometimes literally appealing to the ideas of philosophers). It mixes efficiency and equity considerations, and is often advocated when projects affect future generations.

Specifying the rate at which society is willing to trade present for future consumption is bound to be controversial. 'Society' is not a decision maker (not even the government controls the whole of society). Inevitably, the analyst imposes a specific discount rate (or the parameters that determine it). Yet economists have no particular expertise about how the future should count.

It is not surprising that there is little agreement about the appropriate prescriptive approach. What is ethical depends on value judgements and there is no way to reconcile the different value judgements that people may possess. Equity issues involve trading off the welfare of one group against another's (for example, the present generation's against future generations') and there is no general consensus about the weight that should be attached to the welfare of different groups. The prescriptive approach makes the discount rate a matter of opinion, and provides no basis for determining which opinion is correct.

---

[30] Terminology suggested by Arrow et al. (1995) and see Portney and Weyant (1999, p. 4).

[31] For example, the HM Treasury (2003); Stern (2007) and Garnaut (2008) explicitly adopt a prescriptive approach.

[32] See Weisbach and Sunstein (2008, p. 4).

The implications of the choice of discount rate for equity are not clear. The poor tend to have high discount rates, so using an average discount rate, especially a low one, could make them worse off. For example, if energy efficient appliances pay off at a low discount rate, a policy to force all consumers to use energy efficient appliances would make people with high discount rates worse off.[33] Adjusting the discount rate in project evaluation is not an effective way to achieve equity objectives.

Others criticise high discount rates as being bad for the environment.[34] But a low discount rate could make policies that damage the environment more attractive. Viscusi gives an example of where a higher discount rate made dams less attractive, deterring efforts to build them, which prevented environmental damage.[35]

2.5 Inter-generational issues

Discounting costs and benefits received in the far future involves valuing the effects of policies on future generations, raising ethical issues. For example, the current generation can adopt policies which harm future generations — not only a different group of people, but one that is not around to defend its interests.

Does the opportunity cost logic that underlies the descriptive approach apply to costs and benefits received in the distant future, decades or even centuries away, to be received mainly by people not alive yet? For example, would it be worth spending $100 million now to avert a catastrophe that would cause $50 billion of damage in 100 years? If the discount rate were 8 per cent it would not — the future damage has a present value of $23 million, less than the cost of averting it.[36] The rate of return on the expenditure is less than 8 per cent.

Some people call this the 'tyranny of discounting' and claim discounting is biased against long-term projects.[37] It is, but there is no artificial bias. It merely reflects the power of compounding the benefits of alternative current investments over long periods.

---

33 The Productivity Commission (2005) found this.

34 See for example Ackerman and Heinzerling (2002).

35 See Viscusi (2007, pp. 215-216).

36 To relate to table 1.1, $1000 in 100 years has a present value of $\$0.45 = \$1000/1.08^{100}$ with a discount rate of 8 per cent. $50 billion, therefore, has a present value of $50*0.45 million or $23 million.

37 See, for example, Allsop (1995, p. 375) and Pearce et al. (2003, p. 123).

The logic of rejecting the catastrophe prevention is that if society received a rate of return of 8 per cent on money in the private sector, $100 million dollars would grow to $220 billion in 100 years, more than enough to compensate for the catastrophe.

Alternatively, the present generation could invest $23 million, which gives $50 billion in 100 years, enough to deal with the catastrophe and have $77 million left over for themselves.

> The conclusion, following this logic, is that it is not sensible to invest in any project unless it has a return at least equal to the return available elsewhere. … The long time period does not change the method of analysis; it only makes the issue more important. We should, therefore, discount projects at the otherwise available return — the market rate of return. Only projects that pass discounted cost-benefit analysis should be undertaken. Any other choice throws away resources, which is not good for future generations.[38]

The discount rate should reflect the yields on the alternative use of funds: the return on private investment or the marginal return on capital. A negative net present value with this discount rate means the project is inefficient — the loss to the present generation is greater than the gain to the future generation. It could mean that the future generation is worse off. There is an alternative policy that can make all generations better off.[39]

But for dollars invested today to make it through to 100 years time, the intervening generations must keep investing. The appropriate discount rate does depend on future generations' expected re-investment. The intervening generations could be more, equally or less altruistic to the future than the current generation — but assuming they would continue to care about their children in the same way the current generation does, appears reasonable. The historical evidence is that each generation has continued to invest and improve the next generation's standard of living.

An intervening generation could terminate or offset any investment project, including the catastrophe prevention. Unless there is a reason to believe that transfers to the future from one project are less easily misdirected than transfers from others, the chance that an intervening generation will break the chain of investment does not favour adopting low return projects.

---

[38] Weisbach and Sunstein (2008, p. 14).

[39] The future generation receive the benefits from the government project. They lose the benefits from the higher return private investment that the projects crowds out. It is likely the larger portion of capital used in projects comes at the expense of private investment. Further, future generations may bear the burden of any reduction in current consumption (see appendix B for details).

In fact not discounting, or using a very low discount rate, leads to a ‘tyranny of the future generations’. Not discounting means a benefit to any future generation counts as much as the same benefit now — it is the same as adopting a zero discount rate. A dollar to someone one hundred years from now, or even a thousand years from now, counts as much as a dollar to someone today. That could require the current generation to impoverish itself. For example, a policy that costs $100 million now, but gives $1 million benefit to each future generation — would be undertaken with a zero discount rate because there are a potentially infinite number of future generations. Indeed, a zero discount rate implies that even a policy which costs $100 million now and gives just $1 to each future generation should be adopted. Without discounting, so long as investment has a positive rate of return, the current generation should reduce its consumption to invest for the benefit of future generations.

Further, using a low, or no, discount rate may harm future generations. First, if the current generation adopts low return projects at the expense of investments with higher returns, it can make future generations worse off. Second, if a low discount rate requires the current generation to sacrifice its welfare to make future generations better off, the same argument applies to each generation. Each generation becomes worse off in order to help generations still further in the future.[40]

2.6 Why the efficiency standard is relevant for inter-generational project evaluation

When a project affects the well–being of different generations in the present and future, standard efficiency–based cost-benefit analysis 1) converts future utility into dollars; 2) discounts them at the market return on capital to transform the future dollars into present dollars; and then 3) converts present utility into dollars to compare. Steps 1 and 3 involve standard cost–benefit principles and are within the same time period. Step 2 transforms dollars between time periods. That is, in standard cost-benefit analysis the appropriate discount rate for inter-generational applications is determined in exactly the same way as within a generation: ‘… a dollar invested, in whatever manner, over a period of time does not know or care whether it is an intra-generational dollar (part of someone’s lifecycle savings) or an inter-generational dollar’.[41] The appropriate way to convert future dollars into

[40] Pearce et al. (2003, p. 125).

[41] Kaplow (2006b, p. 3).

current dollars is to discount with the market return on capital because it is the rate at which dollars today can be turned into dollars tomorrow.

In contrast, advocates of the prescriptive approach usually conclude that ethical obligations require discounting costs and benefits that accrue to future generations at a low rate, below the marginal return to capital. That has strong implications. It is implying that the current generation does not save and invest enough for the future and makes a case to reduce private- and public-sector consumption and to increase savings and investment instead. The appropriate response is to increase savings until the rate of return on private capital falls to the advocate's prescribed ethical discount rate. That would make many more potential public and private investments desirable than at a higher discount rate.

The response should not be to accept projects with low rates of return — the high return ones should be adopted first. When evaluating a project, the opportunity cost of invested funds is relevant. Are there better investments that could be made with the capital?

> The market rate of return measures the returns from currently available projects, so as an initial matter, the market rate is a measure of the opportunity costs of this choice. Once again, therefore, we should using discounting at the market rate to choose projects. Project choice and ethical obligations to the future are, to a large extent, separate.
>
> Seen this way, the ethicists' criticism of the positivists' opportunity cost argument is simply irrelevant. It does not matter whether the current market rates of interest are ethically correct because they still represent the opportunity costs of investment.[42]

Put differently, the prescriptive approach implies that the current generation is leaving too little to the future. Even if that is correct,

> it says nothing about the particular choice of projects or policies. If we are going to increase the amount we leave for the future, it is incumbent on us not to do so in a way that wastes resources. Therefore, even if the ethicists' argument is entirely correct, we still must carefully consider the opportunity costs of projects and pick those with the highest return.[43]

If too little is being done for the future, the answer is to increase overall savings and investment rates, not to use below market discount rates and invest in projects with low returns. If an analyst reduces the discount rate when evaluating one particular type of project, it may be adopted over alternative investments with a higher return, which wastes resources.

[42] Weisbach and Sunstein (2008, p. 25).

[43] Weisbach and Sunstein (2008, p. 26).

> Given any level of saving, the interest of future generations is best served by making the highest return on investments in both the public and private sectors. Using a competitive market based rate in evaluating public policy decisions will support that objective.[44]

How much to leave to future generations is an ethical question about how much to transfer. Once some overall allocation between the present generation and a future generation has been selected, then whatever amount that is to be invested on behalf of the future generation should be invested to maximize efficiency —giving the greatest benefit to the future generation from the amount invested. There is no case for wasting resources on low return investments when higher returns are available.

There are many ways to transfer resources from the current to future generations, including by building infrastructure, improving the environment or repaying national debt. The given amount of resources to be transferred to the future should be allocated across the different possible investments so as to equate the marginal return. If one investment had a higher return, funds should be switched to it. As private capital is one possible investment, the marginal return will equal the market return on capital.[45]

At best, the prescriptive approach provides an argument for expanding savings and investment. But whatever the level of savings being allocated, the private return to capital indicates the opportunity cost of government projects.

If the allocation to the future is determined, all intergenerational considerations are about efficiency and projects should be evaluated with the private investment return. Using a below market rate to evaluate public projects could result in worse investments and future generations receiving less.

A policy that discounts benefits to future generations differently from future benefits to the current generation would create time inconsistencies, especially when generations overlap. For example,

> Nuclear waste disposal produces a future benefit of 100 utils, in the year 2020, for a person who will not be born until 2010. Old age insurance produces a future benefit of 110 utils, again in the year 2020, for a person who is alive now and will continue to be alive when the benefits arrive. Which policy is of greater value?
>
> If we treat intragenerational discount rates as exceeding intergenerational discount rates, it becomes possible that nuclear waste disposal has a higher present value, at least if the gap between time preference rates is sufficiently large.[46]

[44] Lind (1995, p. 387).

[45] Kaplow (2006b, pp. 6-7).

[46] Example from Cowen (2001, p. 7).

Yet, when 2010 arrives, old age insurance gives the bigger benefit.

The current generation already transfers resources to future generations in many ways. Examining the efficiency effects of different government projects allows us to maximize the benefit to future generations for a given resource transfer, potentially making all generations better off. Even if it is decided to transfer more resources to future generations, the efficiency of particular methods of doing so is still relevant. If an efficient project was combined with changes in debt levels left to future generations, it could make all generations better off.[47]

These considerations lead to the conclusion that project choice should be based on discounting with an efficiency-based social discount rate. At best the prescriptive approach only addresses the appropriate transfer to future generations, whereas the efficiency approach tells us the best way to make that transfer.[48]

47 Kaplow (2006b, p. 31).

48 Intergenerational issues are considered in more detail in appendix C.

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.

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	

A Accounting for Inflation

A.1 The relationship between real and nominal interest rates[1]

If in period j, \$1 is invested in a bond which pays \$1+ρ_j in period $j+1$, then the nominal interest rate is ρ_j. The expected increase in real purchasing power from buying the bond is: $1+i_j = \frac{E\left(1/P_{j+1}\right)\left(1+\rho_j\right)}{1/P_j}$ where P_j is the current level of the price index, $E(P_{j+1})$ is the expectation of the price level in period $j+1$ formed in period j, and i_j is the bond's expected rate of growth in purchasing power (its expected real return). \$1 has purchasing power $1/P_j$ in period j. Let z_j be the expected rate of growth of purchasing power of \$1: $z_j = \frac{E\left(1/P_{j+1}\right) - 1/P_j}{1/P_j}$ (note that if the price level was expected to increase, z would be negative).

Substituting this into the first equation gives: $(1+i_j) = (1+z_j)(1+\rho_j)$.

Now $1+z_j = \frac{E\left(1/P_{j+1}\right)}{1/P_j} \approx \frac{P_j}{E\left(P_{j+1}\right)} = \frac{1}{(1+\pi_j)}$ where π_j is the expected rate of inflation in period j (the expected percentage increase in the price index over the period). The relationship, therefore, between the real interest rate i and the nominal interest rate is: $(1+i)(1+\pi) = (1+\rho)$. This gives: $i = (\rho - \pi)/(1+\pi)$

Alternatively, taking logs gives $i^c = \rho^c - \pi^c$, where a ^{c} denotes the continuously compounded equivalents. For example, $i^c = ln(1+i)$. If an amount grows at rate i^c, continuously compounded, in one period it grows to $e^{ic} = 1+i$. $i^c < i$.

This formula slightly understates the real interest rate (as in general $E\left(1/P_{j+1}\right) > 1/E\left(P_{j+1}\right)$ so that $(1+z_j)(1+\rho_j) > (1+\rho_j)/(1+\pi_j)$) but the understatement is small compared with the estimation error in inflation forecasts. Further, the formula is exact for realised values of the variables (that is, given actual π and ρ, the formula gives the realised real return, i).

1 Based on Sieper (1981, pp. 49-50).

Public policy analysis usually uses real dollars, with projected costs and benefits measured at today's prices. Real dollars adjust for inflation, they are nominal dollars deflated by a price index to account for changes in the general price level. Real dollars are really bundles of consumption (the basket of goods that goes into the price index). An inflation-adjusted dollar always has the same purchasing power (can buy the same bundle of goods).

The real rate of interest determines the relative value of goods received at different times in the future. The relative price of goods next period to goods now is *(1+i)*. More accurately, *(1+i)* is the relative price of the basket of goods that goes into the price index used to measure inflation. Giving up one basket of goods today gives you *(1+i)* baskets next period. Strictly speaking, when the relative prices of goods change, each good has its own real interest rate.

Using real dollars for future variables avoids having to estimate the future course of inflation. (which would be especially difficult if inflation varies).

Analysts should discount real flows (measured in one period's dollars) using a real rate of interest, and discount nominal flows using the nominal rate of interest. Both methods should result in the same net present value.[2]

[2] For example, the present value of a perpetual stream of nominal payments (received at the beginning of the period) of amount \$I in the first period, which grows at (the inflation rate) π per year is:

$I + \$I(1+\pi)/(1+\rho) + \$I(1+\pi)^2/(1+\rho)^2 + \$I(1+\pi)^3/(1+\rho)^3 + \ldots$

$= \$I[1 + (1+\pi)/(1+\rho) + (1+\pi)^2/(1+\rho)^2 + (1+\pi)^3/(1+\rho)^3 + \ldots]$

$= \$I/[1 - (1+\pi)/(1+\rho)] = \$I(1+\rho)/(\rho - \pi)$

using the nominal interest rate to discount nominal flows. But we can also write

$\$I + \$I(1+\pi)/(1+\rho) + \$I(1+\pi)^2/(1+\rho)^2 + \$I(1+\pi)^3/(1+\rho)^3 + \ldots$

$\$I + \$I/(1+i) + \$I/(1+i)^2 + \$I/(1+i)^3 + \ldots$

But this is the an infinite stream of real payments of amount \$I per period, discounted at the real interest rate which has present value:

$\$I[1 + 1/(1+i) + 1/(1+i)^2 + 1/(1+i)^3 + \ldots] \quad = \$I/[1 - 1/(1+i)] = \$I(1+i)/i$

$= \$I(1+i)(1+\pi)/(\rho - \pi) = \$I(1+\rho)/(\rho - \pi)$

A.2 Converting nominal into real variables

Data on past monetary flows are usually in nominal (or current) dollars. Deflating (or inflating) with a price index (such as the Consumer Price Index or GDP deflator) converts nominal dollars into real, or inflation adjusted, dollars. For example, to convert last year's nominal dollars into this year's real dollars, multiply by the price index for this year and divide by the price index for last year.

A widely used price index is the Consumer Price Index (CPI), which measures the cost of purchasing households' average consumption bundle over time.

More generally, to express an amount $\$A_j$ of period j dollars in terms of period t dollars using the CPI we multiply by the CPI in period t and divide by the CPI in period j. That is, $\$A_j$ in period t dollars is $\$A_j(CPI_t/CPI_j)$.

For example, if period t occurs after period j ($t > j$) and inflation between period j and t (π_{jt}) is positive, then $CPI_t/CPI_j=1+\pi_{jt} > 1$. We inflate the dollar amount because a period j dollar has a greater purchasing power than a period t dollar. It takes $CPI_t/CPI_j = 1 + \pi_{jt}$ period t dollars to buy the same amount as \$1 in period j.

If period j is after period t, then $CPI_t/CPI_j=1/(1+\pi_{tj}) < 1$ if inflation is positive. That is, we deflate period j dollars when we express them in period t dollars, as period j dollars buy less than period t dollars.

If period j is in the future, and period t is now, then we need to multiply period j dollars by $E(CPI_t/CPI_j) = CPI_tE(1/CPI_j) = E(1/(1+\pi_{jt}) \approx 1/[1+E(\pi_{jt})]$ to convert them into period t dollars.

Economists estimate the CPI overstates the actual rate of increase in the cost of living by over 1 percentage point per year.[3] It gives too great a weight to the goods whose prices have risen the most because consumers substitute away from goods whose relative prices have increased towards goods whose relative prices have decreased. New goods are often only included once their price has fallen substantially, and they are widely consumed, so much of their price fall is not reflected in the CPI. Further, the CPI fails to fully allow for improvements in product quality.

The overstatement of the rate of decline in the purchasing power means that estimates of the real interest rate using the CPI to convert nominal into real measures tends to underestimate the real rate of interest.

3 See Boskin et al. (1998) and Gordon (2006)

A.3 The effect of expected inflation on real returns when nominal receipts are taxed[4]

If the variables are expressed as continuously compounded rates then the expected real before-tax interest rate is:

$$i^c = \rho^c - \pi^c$$

To simplify the equations, we will use the continuously compounded rates, but drop the superscripts. The rates are approximately the same as the per period rates.

When nominal returns ρ are taxed at rate τ, the after-tax expected real interest rate is: $r = (1 - \tau)\rho - \pi$

As set out in Sieper (1981), if the underlying after-tax rate of return is given[5], $r = \hat{r}$ and there is a single tax rate τ, and all interest receipts are taxable and all interest payments are tax deductible, then the nominal interest rate is $\rho = \dfrac{\hat{r} + \pi}{1 - \tau}$ which is the tax-adjusted Fisher effect, where nominal interest rates adjust to expected inflation so as to leave the expected after-tax interest rate constant. A 1 percentage point increase in expected inflation increases the nominal interest rate by $1/(1-\tau)$ percentage points. The interest rate rises by more than the increase in expected inflation.

Then the before-tax expected real rate of return is: $i = \rho - \pi = \dfrac{\hat{r} + \pi}{1 - \tau} - \pi = \dfrac{\hat{r} + \tau\pi}{1 - \tau}$

Even if the underlying after-tax real rate of return were constant, the expected before-tax real return is greater, would not be stable. It would increase with expected inflation and the tax rate. Every percentage point increase in expected inflation would increase it by $\tau/(1-\tau)$ of a percentage point. For example, if $\tau = 30$ per cent, $\tau/(1-\tau) = 0.43$. If $\tau = 50$ per cent, $\tau/(1-\tau) = 1$.

As shown in table A.1, the taxation of nominal receipts results in a large gap between the before- and after-tax real interest rate, especially at high tax rates.

For example, some authors suggest applying the empirical evidence to the Ramsey equation gives a real interest rate of 1.5 per cent. But the Ramsey equation produces a consumption rate, an after-tax rate. The before-tax investment rate would be higher.

4 Based on Sieper (1981)

5 For example, it may be tied down by the Ramsey equation, which is derived in appendix F and discussed in section 3.3.

Using the above equation for the before-tax return, for a tax rate of 50 per cent and with expected inflation of 3 per cent, the corresponding before-tax real return would be 6 per cent (see table A.1). The before-tax nominal interest rate would be 9 per cent.

Table A.1 **How the before-tax real return varies with taxes, expected inflation and the after-tax return**

Assuming all nominal returns are taxed at rate τ

	Tax rate					
	T=30 per cent			*T=50 per cent*		
	Expected rate of inflation					
After-tax return	2	3	4	2	3	4
1	2.3	2.7	3.1	4.0	5.0	6.0
1.5	3.0	3.4	3.9	5.0	6.0	7.0
2	3.7	4.1	4.6	6.0	7.0	8.0
4	6.6	7.0	7,4	10.0	11.0	12.0
6	9.4	9.9	10.3	14.0	15.0	16.0

Source: Author's calculations.

If v is the realised rate of inflation, then the realised real interest rate is: $\rho - v = \frac{\hat{r} + \pi}{1-\tau} - v$. If inflationary expectations are unbiased, then the average realised rate of inflation over a large number of observations would equal expected inflation and so the average observed before-tax real interest rate is:

$$\tilde{\rho} - \tilde{v} = \frac{\hat{r} + \tilde{v}}{1-\tau} - \tilde{v} = \frac{\hat{r} + \tau\tilde{v}}{1-\tau}$$ where ˜ represents an average over many observations.

Average realised real returns would depend on the historical pattern of inflationary expectations and the structure of taxes.

We can express the after-tax return as:

$$r = (1-\tau)(i+\pi) - \pi = (1-\tau)i - \tau\pi = (1-\tau)(\rho - \pi) - \tau\pi$$

That is, taxation (or deduction) of nominal returns (payments) is equivalent to a tax on the real return plus an additional term that arises because tax is paid on that part of the nominal return that compensates for expected inflation, equivalent to a capital levy per dollar lent (or subsidy per dollar of deductible borrowing).

In the Australian tax system, people face different marginal tax rates. For example, the marginal income tax rate is increased at higher levels of income, corporate rates may differ from personal rates and effective tax rates differ because people are subject to different benefit withdrawal rates. For a given before-tax interest rate, those with higher marginal tax rates receive a lower after-tax return. Further, some borrowers and savers face a zero marginal rate. Some savers are tax exempt. Although business enterprises and people borrowing for investment purposes can deduct interest payments, borrowing to finance personal consumption, the purchase of consumer durables or owner-occupied housing are not tax deductible and those borrowers pay the before-tax rate.

When tax rates differ across individuals, not only do their after-tax returns differ, but when expected inflation changes there is no adjustment in the before-tax nominal interest rate that would preserve the expected after-tax real interest rate of all lenders and borrowers. The implicit capital levy on savers (subsidy to borrowers) varies with their tax rate.

The effect of expected inflation on the nominal interest rate 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.

Because a large portion of interest receipts are taxed (and some payments deductible), before-tax nominal interest rates would tend to rise by more than the expected inflation rate, but by less than the tax-adjusted Fisher effect for the highest marginal rate. That is, if the before-tax real return in the absence of inflation $\hat{i}$ was fixed, we would expect: $\rho = \hat{i} + \alpha\pi$ where $1 \leq \alpha \leq 1/(1-\tau^M)$ where τ^M is the maximum tax rate. This gives $i = \rho - \pi = \hat{i} + (\alpha - 1)\pi$, so the expected before-tax real return would tend to increase with expected inflation.

The expected after-tax real return for a person with tax rate τ_j is $r_j = \rho(1-\tau_j) - \pi = (1-\tau_j)\hat{i} + \pi[\alpha(1-\tau_j) - 1]$. It would stay the same when expected inflation increased for a person with tax rate $\tau_j = (\alpha - 1)/\alpha$. For people with tax rates smaller than this, the after-tax real return would rise with expected inflation. For people with tax rates higher, it would fall. Even if interest rates in the absence of inflation, $r_j = (1-\tau_j)\hat{i}$, were stable, expected inflation would, in general, change the after-tax return and may increase it for some groups of taxpayers and decrease it for others, depending on their marginal tax rate.

Even if the assumptions that underlie the Ramsey equation hold, the tax system ensures the after-tax real rate of return will vary across individuals and will vary with expected inflation.

It would be good to have an estimate of α, the inflation adjustment factor for nominal interest rates. If we assume that inflationary expectations are unbiased, realised inflation v will be expected inflation plus an unbiased error term: $v = \pi + e$.

If the before-tax real return in the absence of inflation was a constant $\hat{i}$, then from our expression for the nominal interest rate $\pi = \frac{\rho - \hat{i}}{\alpha}$. Substituting this into our expression for realised inflation gives: $v = \frac{\rho}{\alpha} - \frac{\hat{i}}{\alpha} + e$. If we run a regression:

$v = \beta + \gamma\rho + e$ then $\alpha = 1/\gamma$ and $\hat{i} = -\beta/\gamma$

Although some empirical work has been done on the relationship between expected inflation and nominal interest rates in Australia, the researchers usually test whether the classic Fisher effect ($\alpha = 1$) holds.[6] For example, Inder and Silvapulle (1993) test whether realised real interest rate are constant over time, and find they are not. Hawtrey (1997) does the same and finds real interest rates were not constant before deregulation but that since deregulation, after-tax nominal rates (using the company tax rate) rise one for one with expected inflation. Olekahns (1996) has a more sophisticated expected inflation estimate than the one suggested here, and finds that a 1 percentage point increase in expected inflation increases the nominal interest rate by less than 1 per cent (from September 1969 to September 1993) (i.e. $\alpha < 1$). For the period after deregulation, he could not reject that $\alpha = 1$. If Australia were perfectly integrated with the international capital market, the classic Fisher effect would hold for debt instruments (bonds), and perhaps that is reflected in the findings after de-regulation.[7] Nevertheless, even if Australia could borrow as much as it liked at a given world interest rate, changes in world inflation levels could still affect real interest rates.

None of these studies run the regression here. In order to get a rough idea of the size of these parameters we ran this simple regression using quarterly data, from September 1970 to June 2008 and from March 1980 to the June 2008. Realised returns were negative throughout the 1970s, implying that inflation was consistently under-estimated.

Over the full period, we estimated that $\alpha = 1.65$ and $\hat{i} = -0.9$ per cent. For the second period $\alpha = 1.64$ and $\hat{i} = 1.7$ per cent. Both regressions imply the tax rate at which the after-tax return is invariant to changes in the rate of expected inflation is 39 per cent. If $\hat{i} = 1.7$ per cent, the after-tax real return for someone on that tax rate would be 1.0 per cent.

6 See Atkins (1989); Inder and Silvapulle (1993); Olekahns (1996) and Hawtrey (1997). These were all the papers about Australia referred to in the survey Coorey (undated).

7 See, appendix E, box E.1.

For those on lower tax rates the after-tax return would be higher and would increase with expected inflation. Those on higher tax rates receive a lower after-tax return, and it decreases with expected inflation.

All these results apply when inflation-indexed bonds are available, because the inflation adjustment payments are taxed annually as they accrue.[8] The after-tax realised real return on indexed bonds with a return i is $(1-\tau)(i+v)-v=(1-\tau)\hat{i}-\tau v$ which varies with the individual's tax rate and with realised inflation.

8 See Australian Office of Financial Management (2002, p. 6).

B How project finance affects present and future consumption

The effects of a government project on present and future consumption depends on the general public's perceptions of, and responses to, the value of public project output and future tax liabilities. As a result, the effects of the project may depend on how it is financed.

For example, consider the effects of a two period government project. The investment is made in the first period. It is financed either through increasing first period taxes or selling bonds, to which the general public may react. Savings and investment adjust. If savings increase, current consumption has fallen. A fall in investment decreases the private capital stock, which reduces future income.

In the second period, the government project produces output, any bonds sold in the first period are repaid (which may require extra revenue from taxes), the lower private capital stock produces income, and the general public may alter their savings and investment behaviour in reaction (for example, by re-investing the depreciation in the government capital stock — see appendix D).

This two-period model can be used to examine inter-generational equity: in the first period the initial generation makes an investment which provides benefits for the next generation in the second period. Even in this simple framework, the inter-generational equity effects are complex and depend on the form of financing and private sector reactions.

For example, consider who bears the burden of government programmes financed by borrowing from the initial generation – selling them bonds, or issuing debt, that is left for future generations to repay. For the moment, assume that the initial investment comes at the expense of period 1 consumption, just as it would with tax finance. That is, the initial generation reduces its period 1 consumption to purchase the government bonds.

It may seem reasonable to suppose the current generation bears the burden of financing the project. The project uses resources that are drawn from alternative uses in the private sector, reducing period 1 consumption. In the second period, the taxes are levied to pay interest on and repay the debt, but that is a transfer among those living in the second period (from taxpayers to bond-holders). Since the debt servicing does not divert resources to the public sector and does not affect private sector output of goods and services, no net burden is imposed.[1]

The 'burden of the debt' literature, starting with Buchanan argues this is incorrect — the future generation does bear the burden of debt finance.[2] Debt finance differs from tax finance — it is voluntary. In the first period, revenue is raised from those who lend voluntarily (buy bonds). They lend in exchange for promises of future period interest and amortization payments. 'These purchasers of bonds do not, in any sense, "pay for" the benefits of the public spending programme.'[3]

In the second period, taxes are raised to repay the bonds. The repayment just compensates the bond holders for the consumption given up in the first period. The taxpayers whose taxes service the debt bear a burden, their lifetime consumption is reduced. Alternatively, if the debt is rolled over to be repaid by a generation further in the future, that generation bears the burden.

> With government borrowing (debt issue) the ultimate fiscal liability made necessary by the spending programme in the initial period is postponed. This liability is placed, in the aggregate, on taxpayers in periods subsequent to that in which the debt is issued. … Taxpayers in later periods are faced with claims against their incomes that must be met, and which exist only because of the initial-period debt issue.[4]

That is, the appropriate definition of a generation is not all those alive at a point in time but all those born around the same time (say a 20 year period). Then generations overlap and the burden of debt can be transferred across generations.

The rest of the appendix considers the effect of government projects on present and future consumption in more detail in a closed economy – taking account of different forms of financing, of private sector reactions and of capital market effects (such as the crowding out of investment), first in a simple undistorted capital market, and then in a more complicated, but more realistic, distorted capital market.

Appendix E considers the implications of foreign capital flows. If a project is financed through borrowing from foreigners, then the burden is clearly borne by the future generations who repay the foreign debt principal and interest.

---

1 This ignores the efficiency costs of raising tax revenue. Those costs should be accounted for in a costbenefit study. For example, if the project requires extra tax revenue (whether in the first or second period), the cost of raising a dollar in taxes in greater than a dollar. The benefit of a dollar of revenue to the government from its project is greater than a dollar. See Harberger (2008).

2 See Buchanan (2008); Browning and Browning (1994, pp. 438-48) and Rosen (1999, pp. 430-35) on which this section is based.

3 Buchanan (2008).

4 Buchanan (2008).

B.1 In an undistorted market

In a modern economy, the capital market co-ordinates saving and investment decisions. The price that clears the capital market is the interest rate. Figure B.2 depicts a simple, undistorted capital market. The supply curve shows the amount of savings supplied to the private sector at different interest rates. The demand curve shows the private demand for funds to invest at each interest rate.

Figure B.1 **Effect in an undistorted capital market of an increase in bond-financed public investment**

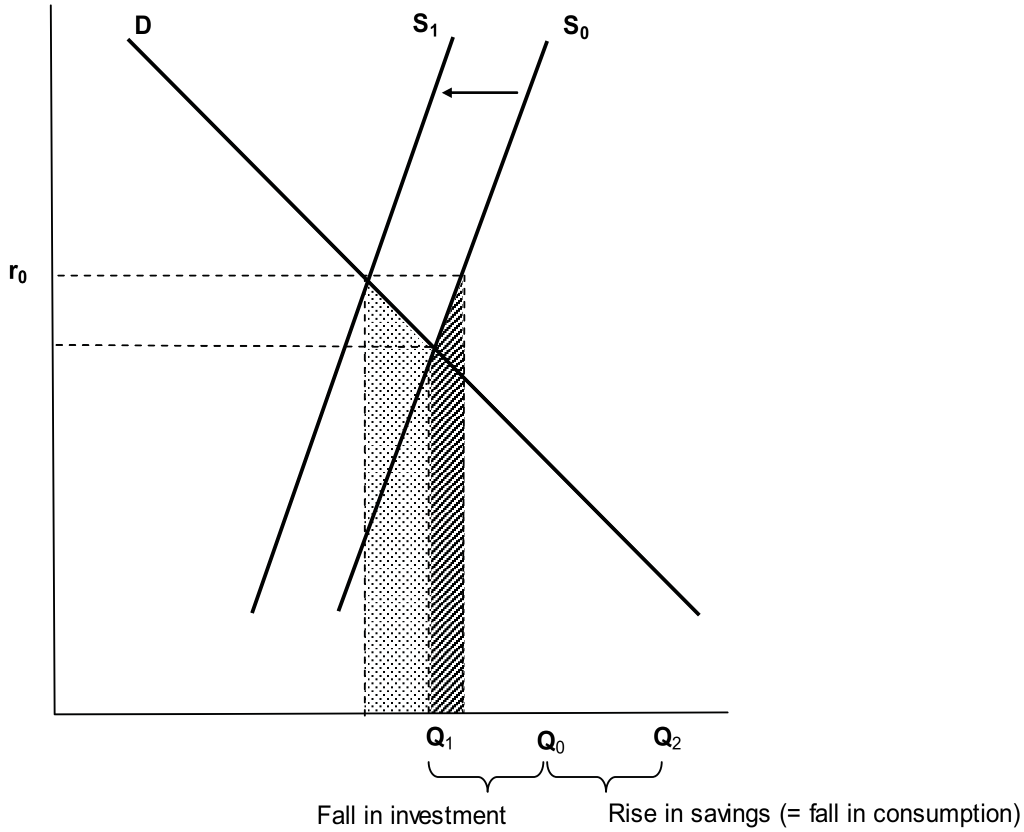

The standard approach to working out the effect of a project is to assume it is financed through borrowing (selling bonds). Bond finance involves selling bonds in period 1 and using the proceeds to increase government investment. In period 2 the project produces output and the bonds are repaid (with higher taxes, project output or both).

The effect in the capital market in period 1 is illustrated in figure B.2. The initial equilibrium is at Q_0, with an interest rate r_0. In the absence of distortions, r_0 is both the consumer rate of time preference and the marginal return to capital.

Extra public investment must come from reduced private investment and consumption. An increase in public investment demand decreases the supply of capital to the private capital market – shifting the supply curve to the left by Q_1Q_2 units from S_0 to S_1. The result is to drive up interest rates enough to increase private savings (reduce private consumption) and reduce private investment by a total of Q_1Q_2 units.

The proportion that comes from each depends on the slopes of the supply and demand curves. In figure B.2, savings rise by Q_0Q_2 units and private investment falls by Q_1Q_0 units. The more inelastic the savings schedule, the greater the reduction in investment. If it were vertical, then every dollar the government borrowed would come at the expense of investment. It is generally agreed that the responsiveness of savings to changes in the interest rate is fairly inelastic, and investment relatively elastic, and so most of the government borrowing would come at the expense of investment.[5]

The decrease in investment reduces the private capital stock, which reduces future consumption. To the extent that the project finance comes from investment, the burden of financing it is automatically shifted to the future.

The right hand area in figure B.2 is the return that must be earned on the extra savings to compensate savers for their forgone consumption. The left hand area in figure B.2 is the return that must be earned on the forgone investment to compensate for the returns it would have earned. For a small marginal change, the project must earn a return of r_0 to cover its costs (the triangles are vanishingly small for marginal changes). Savers are willing to trade \$1 of consumption now for \$$(1+r_0)$ of consumption next period. Each \$1 of forgone private investment would have produced \$$(1+r_0)$ of consumption next period. The opportunity cost of the capital used by the project is $(1+r_0)Q_1Q_2$.

The effect of the government project also depends on how the public:

1. Takes account of future tax liabilities.

2. Perceives the value of future public output.

5 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.

Both will affect the supply of savings to the private sector and directly affect consumption. For example, let b be the cost of the project, which produces $(1+\delta)b$ in period 2. The project is financed through selling bonds, which shifts the supply curve to the left by b units. Let φb be the value placed on future tax liabilities to service the debt. Then the supply of savings will shift to the right by φb units as households increase their savings (reduce their consumption) to pay perceived future tax liabilities.

Let λb be the general public's current valuation of the project's future output. The savings supply curve will shift to the left by this amount as households decrease their savings (increase current consumption) to restore their consumption path over time to the level they desire. The government project reduces current consumption in order to increase future consumption. If people anticipate, and value, the future project output, they will reduce savings. For example, the current generation could cut its bequests by the amount it perceives the project benefits the future generation. The current valuation people place on future public sector output depends on the type of output provided and whether people anticipate receiving it or value benefits received by a future generation.

The net effect is to shift the supply of savings to the left by $b + \lambda b - \varphi b$. That raises interest rates, which reduces private consumption (increases savings) by $(1 - a)(b + \lambda b - \varphi b)$ and investment by $a(b + \lambda b - \varphi b)$ where a is determined by the elasticities of the supply and demand curves.

The total effect on first period consumption is to decrease it by:

$$\varphi b - \lambda b + (1 - a)(b + \lambda b - \varphi b) = b - a(b + \lambda b - \varphi b)$$

In the second period, the project produces output $b(1+\delta)$, the bonds are repaid (with higher taxes, project output or both), private output is $ra(b + \lambda b - \varphi b)$ lower because of the decline in the private sector capital stock and $(b + \lambda b - \varphi b)$ is saved for re-investment which increases consumption by $(1 - a)(b + \lambda b - \varphi b)$ and investment by $a(b + \lambda b - \varphi b)$, which restores the capital stock to its initial position (see appendix D). The total effect on second period consumption is to increase it by:

$$b(1+\delta) - ra(b + \lambda b - \varphi b) - (b + \lambda b - \varphi b) + (1 - a)(b + \lambda b - \varphi b)$$

$$= b(1+\delta) - (1+r)a(b + \lambda b - \varphi b)$$

The project is efficient if the present value of the extra period 2 consumption exceeds the decrease in period 1 consumption:

$[b(1+\delta) - (1+r)a(b + \lambda b - \varphi b)]/(1+r) > b - a(b + \lambda b - \varphi b)$ or $\delta > r$. That is, when the project earns a return sufficient to cover its opportunity costs.

The effect of the project is summarised in table B.1.

Table B.1 **The effect of a government project in a closed economy**

Project costs b, λb is the current private sector valuation of the project's future output, *φb* the value placed on future tax liabilities to service the debt.

Item	*Period 1*	*Period 2*
Consumption	$-[b - a(b + \lambda b - \varphi b)]$	$b(1+\delta) - (1+r)a(b + \lambda b - \varphi b)$
Project output	-b	$b(1+\delta)$
Investment	$-a(b + \lambda b - \varphi b)$	$a(b + \lambda b - \varphi b)$

The inter-generational equity effects of the project are much more difficult to work out — and depend on the particular parameters a, λ and φ and on the extent to which the burden of the debt is passed on and on how project output is distributed between the different generations in the second period. Any statement about the equity effects of a government project must make assumptions about these parameters, yet often we know little about them. Usually the information needed to work out the equity effects is simply not available.

If $\lambda = \varphi$, the public exhibits equal perceptiveness of the value of future public output and the value of future tax liabilities arising from bond finance. The effect on the savings schedule of each offset and the supply of savings falls by the full amount of the bond sale. The savings supply curve shifts to the left by the project cost, which is the traditional assumption.

If the entire burden of financing the project is passed to future generations and if they receive the entire project benefit, then the equity effects of the project depend on its efficiency effects. If the project is efficient, the future generation is better off, if inefficient then it is worse off.

The φ parameter is the extent to which Ricardian equivalence holds. For example, if the project is an inter-generational one, it depends on the extent to which the current generation cares about future generations, is linked to them by bequests and perceives future tax burdens.

If $\varphi = 1$, then people fully account for future tax liabilities, Ricardian equivalence fully holds and the public perceive no difference between current and future taxes (they are equivalent). Although the bond sales reduce the supply of savings to the private sector, they are fully offset by an increase in private savings — as people increase their savings to meet the future tax liability. The net shift in the supply curve (and the amount of investment crowding out) would be the same with bond-financing as with tax financing. Borrowing to finance government expenditures today has the same effect as taxing people today to finance those expenditures, people recognise that their lifetime disposable income has fallen by the same amount in either case.

For example, if a project is bond-financed with $\varphi = 1$ and $\lambda = 0$ then the above analysis shows that the supply of savings does not shift – when the bonds are issued the current generation increases its savings by exactly the same amount, in effect purchasing the bonds to cover the second period taxes needed to repay the bonds. The project is financed entirely from reduced first period consumption: the same outcome as the traditional view of tax finance.

With tax finance, if the public anticipates and values future project output, then $\lambda > 0$, decreasing savings and shifting the supply of savings curve to the left by λb units. If $\lambda = 1$ then the supply curve shifts to the left by b and the situation is the same as the traditional bond finance case. The same outcome would occur under bond financing if $\lambda = 1$ and $\varphi = 1$ since it is an example of equal perceptions of future taxes and future project output, which means that bond-finance gives the traditional bond finance outcome. The supply of savings to the private sector falls by the investment in the project. Where Ricardian equivalence holds fully, bond-financing is the same as tax-financing.

B.2 In a distorted market

Now consider the effects of a government project when capital taxes distort the market. Figure B.3 depicts a simple, distorted capital market. The supply curve shows the amount of savings supplied to the private market at each interest rate lenders receive (that is, at each consumption rate). The demand curve shows the private demand for funds to invest against the interest rate that must be paid (that is, at each investment rate).

The initial equilibrium is at Q_0, with a consumption rate r_0 and an investment rate i_0. An increase in public investment demand decreases the supply of capital to the private capital market – shifting the supply curve to the left by Q_1Q_2 units from S_0 to S_1, which increases interest rates, increasing savings and decreasing investment. Because capital taxes drive a wedge between the consumption and investment rate, it now matters how much the project decreases consumption and how much it crowds out private investment.

In a distorted world, the actual opportunity cost of any resources is their value in the alternative use to which they would have been put. A government investment needs to earn r per cent on funds that come at the expense of consumption and i per cent on funds that come at the expense of investment to cover its opportunity cost (illustrated by the shaded areas in figure B.3).[6]

Figure B.2 **Effect in a distorted private capital market of an increase in public investment**

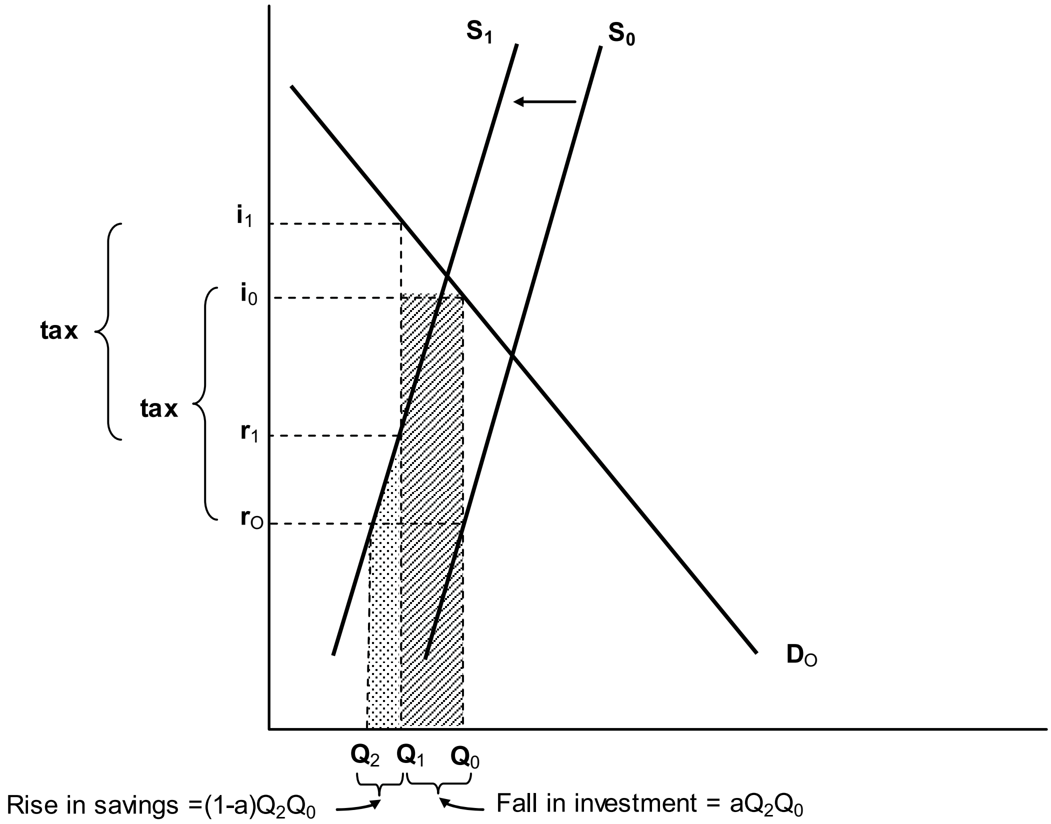

Using the same example as in the previous section, a bond-financed project that cost b in period 1 shifts the supply curve to the left by $b + \lambda b - \varphi b$. Consumption in period 1 falls by $b - a(b + \lambda b - \varphi b)$ and investment falls by $a(b + \lambda b - \varphi b)$ and the capital stock is this much less than otherwise.

In the second period, the project produces output $b(1+\delta)$, taxes are raised to repay the bonds, private output is $ia(b + \lambda b - \varphi b)$ lower because of the decline in the private sector capital stock and $(b + \lambda b - \varphi b)$ is saved for re-investment which increases consumption by $(1 - a)(b + \lambda b - \varphi b)$ and investment by $a(b + \lambda b - \varphi b)$, which restores the capital stock to its initial position (see appendix D). The total effect on second period consumption is to increase it by:

$$b(1+\delta) - ia(b + \lambda b - \varphi b) - (b + \lambda b - \varphi b) + (1 - a)(b + \lambda b - \varphi b)$$

6 For a marginal change, the shaded areas in figure 1 become rectangles and the weighted average discount rate is $w = (1-a)r + ai$.

$= b(1+\delta) -(1+ i)a(b + \lambda b - \varphi b)$

The project is efficient if the present value of the extra period 2 consumption exceeds the decrease in period 1 consumption (discounting with the consumption discount rate because the flows have been converted to consumption equivalents). That is, if:

$b - a(b + \lambda b - \varphi b) < [b(1+\delta) - (1+i)a(b + \lambda b - \varphi b)]/(1+r)$ which reduces to:

$\delta > r + (i - r)a(1 + \lambda - \varphi)$ or

$\delta >(1 - a)r + ai + a(i - r)(\lambda - \varphi)$

If $\lambda = \varphi$, the project is efficient if its rate of return exceeds the weighted average discount rate $w = (1 - a)r + ai$. This is the traditional case where the supply of savings to the private sector decreases (the curve shifts to the left) by b, the cost of the project (and the amount of bond sales).

If $\lambda \neq \varphi$, then the required rate of return is affected. For example, if $\lambda > \varphi$, then the supply of savings decreases more than in the traditional case (the curve shifts by $b + \lambda b - \varphi b > b$). Compared with the traditional case, more investment is crowded out, raising the cost of the capital used by the project (as the tax distortion makes the social value of private investment higher than its private value). The project must earn a higher rate of return to cover its costs – above the weighted average rate. For example, if it is believed that Ricardian equivalence is unimportant (φ is low), that raises the rate of return required for the project to be efficient.

The appropriate discount rate to evaluate whether projects are efficient is not unique, it depends on the difference between λ and φ. We have little information on each, and they may vary by project. Although λ and φ are different concepts, they both depend on similar factors, such as how forward looking the population is and whether they are linked through bequests to future generations. The most reasonable approach is to assume $\lambda = \varphi$ for the purpose of defining a social discount rate. That is what is done in appendix D. If there is reason to believe the two are substantially different, then this assumption can be varied — either through adjusting the discount rate or adjusting the benefit and cost streams.[7]

[7] See Sjaastad and Wisecarver (1977, pp. 530-37) who present a more detailed analysis using a simple macro-economic model. They also consider other factors that may affect the discount rate, such as the direct effect of project execution (as opposed to financing) on private expenditures.

The traditional approach usually assumes tax financing reduces current consumption. But tax finance shifts the supply of savings to the left by λb, and so reduces investment. The more account is taken of future output (the greater λ), the greater the leftward shift in the supply of savings. If $\lambda = 1$, when the government imposes taxes now in order to provide benefits in the future people respond by decreasing their savings by the amount of the tax to restore their desired consumption path over time. That gives the same outcome as the traditional bond-financing case (the supply of savings to the private sector would fall by b, the full amount of public investment). If $\lambda = \varphi = 1$ then it would be the same as bond-financing (and the traditional bond-financing approach).

If the public does not fully account for future tax liabilities, $\varphi < 1$, then the effects of bond and tax financing differ, with bond–financing shifting the supply curve of savings to the private sector to the left by more than with tax–financing. Bond–financing is equivalent to tax financing plus a bond–financed tax cut. A bond–financed tax cut leaves the supply curve of savings to the private sector unchanged when full Ricardian equivalence holds ($\varphi = 1$) and shifts it to the left when it does not.[8] As a result, if $\varphi < 1$, bond-financing reduces investment more than tax-financing and so the appropriate discount rate is greater for bond-finance than for tax-financing (because bond-financing aggravates the capital tax distortion).

Harberger (1987) argues for a convention that the funds for the marginal outlay, the marginal project, the marginal program would always be 'sourced' in the capital market. In a functioning capital market, the reactions to changes in the supply of funds is likely to remain similar from one case to the next. The same basic mechanism would be at work all the time, with increased demand for funds displacing consumption and investment in proportions based on the relative elasticities of saving and investment to interest rates.

By contrast, each tax change is very different from the last, and the changes in consumption and investment may vary from case to case, depending on whether extra tax revenue comes from increasing income tax rates, extending the GST coverage or tightening income tax loopholes. There are a host of potential fiscal adjustments with little predictability about what the next one would be, and so we cannot specify the typical proportions that come from consumption or investment. The effects of tax financing are not stable or predictable.

Not only is capital market sourcing convenient, it is realistic, since on a day to day, month to month, and even year to year basis, governments get their needed marginal

8 See Sjaastad and Wisecarver (1977, pp. 528-537).

cash from the capital market, and typically allow any periodic cash surpluses to be reflected in the reduction of outstanding debt or invested in the capital market.[9]

Moreover, an alternative approach is to assume that the funds invested in the project could always be used to retire debt or be directly invested in the capital market (for example, through a future fund). If so, the opportunity cost of using those funds is determined in the capital market, as in figure B.3, whether tax-financed or bond-financed. In this case, using $100 million dollars for a project reduces the supply of funds to the private capital market by $100 million.[10]

When a regulatory proposal imposes current costs on firms and individuals, it increases the demand for capital. For example, firms may need to borrow, or forgo other investment, to cover increased compliance costs. That is akin to bond financing, with the borrowing repaid using revenue raised from consumers through higher prices (with similar effects to taxes used to repay bonds).

Regulations that impose costs and benefits directly on consumers may be akin to a tax financed project, and the discount rate would be closer to the consumption rate than with bond-financing.

The appropriate discount rate, therefore, depends on the project and how it is financed, and the distribution of benefits and costs. It is important to carefully specify the relevant project. For example, it may be that an (after-tax) consumption discount rate is appropriate to work out a consumer's value of statistical life.[11] But that does not mean it is relevant to determine the present value of costs and benefits from a government safety project that saves lives. The benefits, which would include the value of statistical lives saved, would be discounted by a rate that depended on project characteristics and how the project was financed. An analogy is that the discount rate used to calculate the value of an insurance company would be different from the rate its customers use to calculate the value they place on its insurance policies.

Further, because mortgage and consumption borrowing is not tax deductible, many consumers' discount rate is a before-tax rate. Capital market imperfections, such as credit rationing and borrowing constraints, may raise the relevant discount rate even further. As a result, even if a project was mainly financed from forgone consumption, the relevant discount rate may be closer to the before-tax rate, rather than the after-tax consumption rate. Direct estimates of the consumers' rate of time

9 See Harberger (1987, pp. 170-71).

10 See Sjaastad and Wisecarver (1977, p. 532).

11 The value of statistical life is defined in appendix C.

preference usually find it to be extremely high, even for determining the value of statistical life.[12]

Again, the equity effects of a government project are more complicated than the efficiency effects — depending not only on the parameters that determine the efficiency effects, but also how the burden of taxes and project benefits are split between generations.

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.

C Inter-generational comparisons: the social welfare function approach

C.1 Inter-generational comparisons: the representative consumer

Discounting costs and benefits received in the far future involves valuing the effects of policies on future generations, raising ethical issues. For example, the current generation can adopt policies which harm future generations – not only a different group of people, but one that is not around to defend its interests.

Discounting in the long-term involves comparing the welfare of different generations. The usual approach is to do so explicitly with a representative, infinitely–lived household who maximises the following utility function:[1]

$$U = u(c_0) + (1+p)^{1-\varepsilon}\gamma u(c_1) + (1+p)^{2(1-\varepsilon)}\gamma^2 u(c_2) + ... = \sum_{j=0}^{\infty}(1+p)^{j(1-\varepsilon)}\gamma^j u(c_j)$$

where c_t is the per person consumption of a typical household member during period t. Each period represents a different generation. The population (and so representative household) grows at the exogenous rate *p*. $\gamma = \frac{1}{(1+\theta)}$ with $\theta > 0$ is the pure rate of time preference, which is used to discount the utility of future generations..(1-ε) is the weighting the representative gives to the number of people in each generation.

There are two possible ways to proceed using this function: a prescriptive (normative) approach or a descriptive (positive) approach. The prescriptive approach interprets the utility function of the infinitely lived household as an inter-generational social welfare function. The social welfare function represents some ethical judgement about the appropriate distribution of welfare across generations. Just as a utility function shows how a person ranks different combinations of consumption goods, the social welfare function represents a value judgement of

1 See, for example, McCallum (1996).

how society should rank different distributions of utility across generations. It formalises the trade–off between efficiency and equity. A social welfare function gives the welfare of the whole society as a function of the utilities of individuals, just as the utility function gives the welfare of the individual as a function of the quantities of goods the individual consumes.

The descriptive approach uses this function as positive model of how the economy actually works and how the interest rate is determined. For example, a competitive economy with overlapping generations could operate in this manner (with c_t being the value of lifetime consumption of generation t).

Special cases of the utility function are often used under both approaches. For example, if ε=1, then $U=\sum_{j=0}^{\infty}\gamma^j u(c_j)$. Utility depends only on the average consumption of family members and the number of descendants does not affect the representative consumer's utility.[2] Alternatively, a constant population (p=0) would give the same result, and now models a constant population of identical individuals with the same utility function.

If ε=0, then $U=\sum_{j=0}^{\infty}(1+p)^j\gamma^j u(c_j)$. This has been called a 'Benthamite' formulation, where the household maximises the present value of the total utility of all current and future household members. Because the household cares about both the average utility and number of descendants, when the population is growing, it cares more about the future than a household that cares only about its descendants' average utility (an increase in future consumption per person has a larger effect on utility).

C.2 Deriving the discount rate

The production side of the economy is a standard Solow model, with exogenous labour–augmenting (Harrod neutral) technical progress at rate g. Output per household each period is a constant returns function of per household capital K and efficiency units of labour, N: $Y_j = Y(K_j, N_j)$ where $N_j=(1+g)^j(1+\pi)^j$. That is, exogenous labour–augmenting (Harrod neutral) technical progress increases the number of efficiency units of labour per head of population at rate g.

Now express the production function in per efficiency unit terms: let $y_j = Y_j/N_j$ and $k_j = K_j/N_j$ and $y_j=(1/N)Y(K/N,1)=y(k)$.

2 For example, Blanchard and Fischer (1989, pp. 38-39) use this case.

In the steady state, the effective labour supply, real income and the capital stock all grow at rate $g + p + gp$. Therefore Y/K is constant and so is the rate of return on capital, $Y_K = \partial F/\partial K = \partial f/\partial k$.[3] The share of income going to capital, $Y_K K/Y$ is constant. The share going to labour is constant. The distribution of income is steady, which requires Harrod neutral technical progress.[4]

In the steady state, output, consumption, the wage and capital per efficiency unit are constant. But it is people, not efficiency units, that receive income and consume. Population is growing at rate p. Consumption, capital, the wage and output per person grow at rate g, the rate of technical progress.

A steady state with c_t growing at rate g, is only possible if u′ has a constant elasticity with respect to c_j.[5] Assume, therefore, that the household has a power utility function: $u(c_j)=\frac{c_j^{1-\eta}}{1-\eta}$ with $\eta > 0$, which gives $u(c_j) = \ln (c_j)$ for $\eta = 1$. Power utility has a constant elasticity of marginal utility of $-\eta$ (which is also the co-efficient of relative risk aversion). It determines the rate at which an individual's marginal utility falls as income rises.

The household's budget constraint is:

$$y(k_j) = c_j + (1+p)k_{j+1} - (1-d)k_j$$

where capital depreciates at rate d.

The household chooses values of c_j and $k_1, k_2, \ldots$ to maximise utility:

$$u(c_0) + (1+p)^{1-\varepsilon}\beta u(c_1) + (1+p)^{2(1-\varepsilon)}\beta^2 u(c_2) + \ldots$$

subject to the budget constraint and given k_0.

The first order conditions give: $(1+\theta)(1+p)^{\varepsilon}(c_j / c_{j+1})^{-\eta} = y'(k_{j+1}) + 1 - d$

That is the interest rate is:

$$r = y'(k) - d = (1+\theta)(1+p)^{\varepsilon}(1+g)^{\eta} - 1$$

Or approximately (true when expressed as continuously compounded rates):

$$r = y'(k) - d = \theta + \varepsilon p + \eta g$$

3 ∂F/∂K = ∂(Nf)/∂K=N(∂k/∂K)(∂f/∂k)= ∂f/∂k

4 In the case of a Cobb-Douglas production function, Harrod neutral technical progress is also Hicks neutral. $Y = K^{\alpha}[a(t)L]^{1-\alpha} = [a(t)]^{1-\alpha} K^{\alpha} L^{1-\alpha}$

5 That is, $c_j u''(c_j)/u'(c_j)$ is constant. see McCallam (1996, footnote 11, p. 46).

In equilibrium the net return on capital equals the marginal rate of time preference equals the interest rate. As there are no taxes, the consumer equals the investment rate. If the prescriptive approach specifies a Ramsey rate that is below the marginal rate of return on investment, welfare would be improved by increasing the capital stock until the equilibrium condition holds.

We have derived the Ramsey formula for the social discount rate accounting for population growth: $r = \theta + \varepsilon p + \eta g$ where θ is the pure rate of time preference used to discount utility. ε depends on how changes in the population are valued and determines the effect of population growth, p on the discount rate. Consumption per person grows at rate g 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). Most authors ignore the effect of population growth on the discount rate — assuming either constant population (p=0) or a Benthamite social welfare function (ε=0), which gives a lower discount rate (when population growth is positive).

The interest rate r is the appropriate rate to discount consumption. The consumption discount rate is higher than the utility discount rate because the growth in per head consumption means the marginal utility of future consumption is less than current consumption (and η shows the rate at which it falls). A unit of consumption now gives more utility than a unit of consumption in the future.

C.3 The prescriptive approach

The prescriptive approach uses the Ramsey formula as a framework to guide the ethical choice of a social discount rate. It involves specifying the parameters to reflect some ethical beliefs, sometimes literally appealing to philosophers. There is much disagreement about these parameters. The prescriptive approach requires a subjective assessment of them and risks the analyst imposing his own judgement.

The parameter θ is now the pure social rate of time preference,[6] which determines the relative weights put on the welfare of different generations in the social welfare function, so so-called distributional weights approach. A positive parameter means less weight is put on the utility of future generations. One reason for a positive θ is the chance of some catastrophic event eliminating human life on earth. More consumption for future generations is worth less if there is some chance they will not exist. The probability of extinction is usually considered quite low. For

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

example, Stern sets it at 0.1 per cent (that is an annual probability of one in a thousand).

The other reason for a positive θ is if utility to future generations counts less than utility to the current generation. Some authors argue that θ should be set to zero (that is, $\gamma = 1$) so that all generations count equally. For example, Stern argues that it is ethical to give the same weight to the utility of different generations. Individuals often discount their own future utility because of impatience. Cowen argues that utility is not 'productive' over time as is invested capital and that impatience is not relevant in an inter-generational setting because future generations are not impatient to be born — and they do not experience a disutility of waiting to be born.

Others disagree and argue that if future generations are better off than current generations, their utility should be discounted. The appropriate weight to put on the interests of different generations involves comparing the welfare of different people (inter-personal justice). There is a well–developed literature on use of social welfare functions to compare people's welfare. As Brennan (2005) points out, Stern's preferred social welfare function has the Benthamite form:

$$U = u(c_0) + (1+\pi)u(c_1) + (1+\pi)^2 u(c_2) + ... = \sum_{j=0}^{\infty} (1+\pi)^j u(c_j)$$

where social welfare is simply the sum of the individual utilities.

This additive social welfare function, which maximises the sum of the utilities is often described as utilitarian, representing Bentham's philosophy that 'the greatest happiness of the greatest number is the foundation of morals and legislation'.[7] But it is an extreme case that implies social welfare increases when total utility is increased, no matter who gets it. The utilitarian social welfare function only considers total utility, and is not concerned with the distribution of utility between generations. It also sets $\varepsilon = 0$, meaning population growth does not increase the discount rate.

At the other extreme, the social welfare function could be min $(u(c_1), u(c_2), ...)$. That is, social well–being is judged by the welfare of the worst off member. This 'perfect egalitarian' case is sometimes known as the Rawlsian social welfare function, after philosopher John Rawls. It considers only the distribution of utility and pays no attention to total utility. If future generations were expected to be better off than the current generation, then the Rawlsian function suggests it should never transfer anything to future generations.

7 Introduction to the Principles of Morals and Legislation (1789).

A more general case would be a social welfare function analogous to power utility:

$W = \sum_{j=1}^{J} \frac{u_j^{1-\varphi}}{1-\varphi}$ where $\varphi \geq 0$ is the constant relative inequality aversion parameter (and setting population growth equal to zero for simplicity). It reflects a value judgement. The higher is φ, the greater the aversion to inequality in utilities. $\varphi = 0$ corresponds to the utilitarian assumption. As $\varphi \to \infty$, we approach the Rawls case. If $\varphi = 1$ then the social welfare function is log linear – or Cobb Douglas. This is the Nash bargaining outcome.

Brennan (2006) argues that the standard social welfare function approach implies when people in a generation are better off, they should receive a lower weight in the social welfare function on equity grounds (extra utility to the better off is not worth as much to society as extra utility to someone worse off). If ongoing economic growth is expected to make future generations better off and if the social welfare function values equality in utility, the social rate of time preference should be positive.

The social rate of time preference is zero with a utilitarian social welfare function, but this an extreme representation of social preferences that is unconcerned with equality between generations.

A positive social rate of time preference does not mean future lives are valued less than current lives. It means that future generations are expected to be better off than the current one, and so extra utility to the future is valued less at the margin than extra utility to the current generation.

But there is another reason for discounting consumption received by future generations. Even if θ, the social rate of time preference is set equal to zero, future consumption is still discounted (the ηg part of the discount rate). The growth in per head consumption means future generations have a lower marginal utility of consumption and so get less from a unit of consumption than earlier generations, so called concavity of the utility function. A dollar to a future generation is less valuable than a present dollar from a social point of view.

Again, this is controversial: for example Kaplow argues that technology may enable future generations to make more effective use of resources to get more out of a unit of consumption.[8]

Although people disagree about the importance of the two reasons for discounting future consumption, there is widespread agreement that future consumption should

8 Kaplow (2008, p. 386).

be discounted.[9] The best way to resolve the disagreements is to express the social welfare function in terms of the consumption of each generation as: $V(c) = \sum_{j=1}^{J} \frac{c_j^{1-\mu}}{(1-\mu)}$ where μ is the relative coefficient of aversion to inequality in consumption.[10] We ignore population issues (again setting population growth $p = 0$).

It is more intuitive to directly think about the effect of the distribution of consumption on social welfare when determining the discount rate for the consumption of future generations. For example, if $\mu = 1$ then we have a log social welfare function, which means 1 per cent of consumption always has the same social value. For example, \$200 to someone on \$20,000 a year has the same social value as \$2,000 to someone on \$200,000 a year.

More generally, the contribution to social welfare of a marginal increase in consumption of generation j is $c_j^{-\mu}$. If generation A has k times the consumption of generation B, then the social value of an extra unit of consumption to B is k^{μ} times the value to A. If consumption grows at annual rate g, a generation n years in the future would be $(1+g)^n$ times richer than the current generation and the social value of extra consumption to the current generation is $(1+g)^{n\mu}$ times greater. That is equivalent to discounting future consumption at the rate μg per year (approximately — exact if the variables are expressed in continuously compounded terms).

The coefficient of aversion to inequality in consumption, μ, determines the amount of inefficiency we are willing to bear to pursue redistribution. If generation A has k times the consumption of generation B, then a redistribution from A to B that wasted $1 - k^{-\mu}$ of the transfer would be marginal.[11]

---

9 For example, Arrow et al. (1996), all the contributors in Portney and Weyant (1999a).

10 Note that combining the standard consumer's power utility function and the constant inequality aversion social welfare function does not yield a simple constant elasticity $V(c)$ expression that combines the curvature measures η and φ in the expected way. Although Boadway (2006, p. 3) claims it does, with $\mu = \varphi + \eta + \varphi\eta$, he makes a simple arithmetic error. In fact, substituting the utility function into the social welfare function to express it in terms of consumption means the co-efficient on consumption in $V(c)$ is $(1-\mu) = (1-\varphi)(1-\eta)$, which gives $\mu = \varphi + \eta - \varphi\eta$. But if $\eta > 1$, then $u(c_j) < 0$ creating problems, as in the expression for $W(u)$ it is raised to a real exponent, which is undefined unless the exponent is an integer. Further, a higher φ would reduce rather than increase μ. Kaplow (2003) discusses these problems and the true relationship in more detail.

11 More generally, If $c_i > c_j$ then $SMU_i < SMU_j$ and we transfer from i to j, where SMU is social marginal utility. If making the transfer costs a proportion d of the transfer, we are willing to transfer until $SMU_i = (1 - d)SMU_j$ (ie until the loss to i equals the gain to j). that is, we are willing to put up with waste $d = 1 - SMU_i/SMU_j$. $SMU_i = c_i^{-\mu}$. So are willing to waste $1 - c_i^{-\mu}/c_j^{-\mu} = 1 - (c_j/c_i)^{\mu}$. Note that $(c_j/c_i) < 1$, so the higher μ, the more you are willing to waste, $(c_j/c_i)^{\mu}$ gets smaller.

For example, if μ =1 a redistribution that took $2,000 from someone on $200,000 a year, wasted 90 per cent of it and gave the remaining $200 to someone on $20,000 would keep social welfare constant. If μ =2, a transfer that wasted up to 99 per cent would be worthwhile (for example, a transfer that took $2,000 from A and gave $20 to B would be acceptable).

Harberger (1978) points out this, and other disquieting implications, of the distributional weights approach when considering transfers within a generation and argues that it results in unacceptable outcomes. He concludes that the distributional weights approach does not capture how most people think about distributional issues. It does not represent the value system of most citizens and risks economists' peculiar opinions on distributional issues to swamp all other considerations, something that is beyond the economist's professional role.[12]

In contrast, Stern (2008) and Dietz and Stern (2008) also argue that most people would consider such levels of waste in redistribution undesirable, but conclude that means a μ higher than 2 is implausible, whether arrived at through concavity of the utility function (as in Stern) or through discounting future utility (as in Kaplow and Brennan). They are concerned with action to mitigate global warming. They assume that mitigation would impose a cost on the current generation and benefit much richer future generations. If per capita consumption grows at 1.3 per cent per year (Stern's base case), those living in 100 years time would be 3.6 times richer than people today. Those living in 200 years would be 13.2 times richer.

Their logic is that if it makes no sense to take $2000 off a rich generation to make a poor generation $20 better off, taking $20 off a poor generation to give $2,000 to the rich generation is justified. They assume that the distributional weights approach captures the relevant ethical considerations (although they do agree it is 'a very narrow view of ethics').[13] If $\mu < 2$, then the rate used to discount the future consumption benefits should be no greater than $2g$ per year (g is the annual growth in consumption) to account for future generations being better off.

A co-efficient of relative inequality in consumption aversion, μ, in the range 1 to 2 would be explained by a coefficient of relative risk aversion η in the range 1 to 2 even with social inequality aversion φ equal to zero. That is, concavity of the utility function in line with standard empirical estimates gives substantial redistribution with all social welfare functions, even a utilitarian one.

Further, although a μ of 2 would be given by $\eta = 2$ or $\varphi = 2$ or some combination, the roles of η and φ in determining μ are not additive. Kaplow (2003) shows that

[12] See Harberger (1978, pp. S118-S119).

[13] Dietz and Stern (2008, p. 104).

the more concave the individual utility functions (the higher is η), the less the effect of the degree of concavity in social welfare (φ) on the co-efficient of relative consumption inequality aversion (μ), especially at high income levels. He concludes 'debates about whether the proper social welfare function is utilitarian or strictly concave (and, if so, how concave) may have diminished practical significance.'[14] The social welfare function may be quite concave in consumption, even if it is not in utility. For example, if η =1, then u_j = $ln(c_j)$ and $W=\sum_{j=1}^{J}\frac{\ln(c_j)^{1-\varphi}}{1-\varphi}=\sum_{j=1}^{J}\ln(c_j)=V(c)=\sum_{j=1}^{J}\frac{c_j^{1-\mu}}{(1-\mu)}$ with μ = 1, independent of φ.

Even with a utilitarian social welfare function, which gives a zero pure rate of time preference (θ = 0), the Ramsey formula is consistent with a wide range of rates to discount consumption (see table 1.3). Suggested consumption growth rates (g) range to 1 to 2 per cent and the co-efficient of relative risk aversion (η) from 1 to 4, which give a discount rate (ηg) of anywhere from 1 to 8 per cent, wide enough to encompass most views on the social discount rate. Not discounting future utility is consistent with substantial discount rates for future consumption.

C.4 Problems with the social welfare function approach

There are a number of problems with this social welfare function approach. First is the lack of agreement on the appropriate function and parameters to choose.[15] How future generations 'should' count is a value judgement and is inherently controversial. There is no consensus about how to value benefits to future generations.

Another problem with the social welfare function approach is that it is totally impractical. To use a social welfare function to do public policy analysis requires measuring utility on a cardinal scale and inter-personal utility comparisons[16] – often assuming that all individuals have the same utility function. If there is to be a trade–off between more happiness for one person, and less for another, we need to be able to measure in a comparable way the changes in happiness accruing to different people. But there is no consensus on whether utility can be measured.

[14] Kaplow (2003, p. 14).

[15] For example, the power utility form is homothetic. Buchanan and Hartley (2000, pp. 135-37) argues that is not appropriate — and the social welfare function should be positively skewed to reflect compassion rather than envy.

[16] Layard and Walters (1978, pp. 45-46).

Inter-generational comparisons are even more difficult. Even if we are happy to assume that the whole current generation can be represented by a single utility function, little is known about the preferences of future generations.

The social welfare function approach is a consequentialist moral theory. It says we should judge policies only in terms of their consequences and the only relevant consequences are individual's gratifications. There is the problem of what weight to attach to preferences that involve envy and maliciousness towards others. It ignores other social goals such as liberty, justice, order, community – goals which transcend individual wants. Further, social choice may be concerned with means.

The social welfare function approach is one particular view of social choice that may not capture how most people think about social welfare or account for equity. For example, the standard form of social welfare function focuses on equality. But most people would be unconcerned about a transfer of income from a very rich person to a rich person, yet a standard social welfare function would say it raises social welfare.

Harberger suggests that, judging by people's charitable giving, redistributions within their family and their gambling behaviour, most seem to care about alleviating poverty rather than equality.[17] He suggests a basic needs approach. Rather than rely on the differential weighting of the welfare of different individuals, this approach imputes external benefits connected with the improvement in the circumstances of others. Most people genuinely believe it is good for the sick to be healed, the homeless sheltered and so on. But it is not the recipient's utility that enters the donor's utility function but the consumption of particular goods and services (food, education, medical care, housing etc) or the attainment of certain states (better nourished, better housed etc) that are closely correlated with the adequate consumption of certain goods and services.

The efficiency approach separates equity and efficiency issues. An advantage is it allows us to explicitly consider the appropriate ethical obligations to future generations. The standard economist's social welfare function approach is only one way to do that. For example, it is not clear that the best way to account for equity effects on future generations is to lower the discount rate used in the social welfare function. Alternative ethical perspectives are possible, which may provide vastly different policy prescriptions.

For example, Nordhaus, suggests:

> Quite another ethical stance would be to hold that each generation should leave at least as much total societal capital (tangible, natural, human, and technological) as it

[17] See Harberger (1984).

> inherited. This would admit a wide array of time discount rates. A third alternative would be a Rawlsian perspective that societies should maximize the economic well–being of the poorest generation. The ethical implication of this policy would be that current consumption should *increase* sharply to reflect the projected future improvements in productivity.
>
> Yet another approach would be a precautionary (minimax) principle in which societies maximize the minimum consumption along the riskiest path; this might involve stockpiling vaccines, grain, oil, and water in contemplation of possible plagues and famines. Yet further perspectives would consider ecological values in addition to anthropocentric values.[18]

Some of these decision rules reflect the basic needs approach — people care about ensuring future generations' basic needs are met: not leaving them in dire circumstances and avoiding catastrophes that may threaten society.

C.5 The descriptive approach

The Ramsey formula approach can be interpreted as a positive model of how the economy works. The Ramsey formula then describes the equilibrium market interest rate – the marginal return to capital and the rate at which consumers trade consumption over time. A descriptive approach would then use real world observations to determine the parameters of the model, in contrast to basing them on ethical principles.

One approach is to interpret the utility function of the representative, infinitely–lived household:

$$U = u(c_0) + (1+\pi)^{1-\varepsilon}\gamma u(c_1) + (1+\pi)^{2(1-\varepsilon)}\gamma^2 u(c_2) + ... = \sum_{j=0}^{\infty}(1+\pi)^{j(1-\varepsilon)}\gamma^j u(c_j)$$

as a Barro and Becker dynastic utility function with overlapping generations.[19] In the model parents are altruistic toward their own children. The utility of parents depends on their own consumption, the number of children they have and the utility of each child. Altruistic parents choose family size, consumption and intergenerational transfers by maximising a dynastic utility function.

Altruism justifies the assumption that heads of dynastic families effectively have infinite lives. Because the dynastic head cares about his children's utility, and they care about their children, the head's utility depends on the consumption and number

[18] Nordhaus (2007, pp. 692-93).

[19] See Becker and Barro (1988) and Barro and Becker (1989).

of descendants in all generations. In effect, he acts as if he lives forever and can choose the entire time path of his descendants' consumption. Since the objective function is time consistent, the descendants face a problem of the same form and they have no incentive to deviate from the choices made initially. The dynastic head wants them to maximise their utility, and that is what they do.

The rate of time preference, γ, and the effect of population growth on utility, ε, are now parameters that depend on how the altruism of parents works. That is, the discount rate depends on the preferences of the current generation.

Parents allocate income across bequests, human capital investments in children and their own future consumption so that a marginal dollar gives the same value in all the uses (gives the same marginal contribution to dynastic utility).

Becker and Barro also make fertility and population endogenous, which can dramatically change the model's implications. For example, as the rate used by the current generation to discount the consumption of future generations depends on fertility, it too becomes endogenous.

In the simple infinitely lived consumer models, there are no distortions like public goods, externalities and taxes. The path of consumption determined by utility maximisation and competitive markets is Pareto optimal – one generation cannot be made better off without making some other generation worse off. The social discount rate would be the market rate determined within the model: an efficiency analysis would discount with the market rate of interest. In the presence of distortions, the social discount rate would need to be adjusted, but would still be based on market rates.

If a representative person in the current generation had the same preferences about population and consumption of future generations as a social planner (i.e. if the weights put on future generation's consumption in the social welfare function were the same as in the dynastic utility function), then both would make the same choices for fertility, consumption and investment and both would discount future consumption at the market interest rate.[20]

If the social planner who determines the social welfare function cares more for future generations than does the current generation, the prescriptive approach leads to a lower discount rate than the efficiency based descriptive approach. That is, the social planner places more weight on future numbers of people and consumption than the current generation and would discount the future with a lower rate to maximise social welfare (as judged by the planner's preferences).

[20] Barro and Becker (1989).

C.6 Value of life

Some argue that life should not be discounted, that a life saved in the future is no less valuable than a life saved today.[21] Lives cannot be invested and earn interest.

Certainly it is difficult to put a dollar value on a life. Willingness to pay does not work well. For most people, there is no amount that you could pay them to accept immediate, certain death (money is of little use to a corpse). Perhaps some people would accept a large sum knowing it would go to their family — but that puts a lower value on the lives of the most altruistic people. Moreover, the amount most people would pay to avoid immediate certain death would be only limited by their wealth.

Cost benefit analysis does not place a value on human life. Instead, it uses the value of statistical life, which values the reduction in statistical deaths arising from small risks, which is what most government policies affect. It is based on people's observed willingness to pay for small reductions in the risk of death rather than buying out the risk of certain death. For example, suppose a random event, affecting everybody equally, kills one person in a million and that each person is willing to pay $5 to eliminate this risk. A group of one million people is willing to pay $5 million to eliminate the risk of one statistical death to their group. In this example, the value of statistical life is $5 million, which may be greater than each person's wealth. It is $5 divided by the risk reduction of one chance in a million of death.[22] Most government policies are about small reductions in mortality risk and the value of statistical life is the correct way to value the benefits from risk reduction. That is, we are not discounting lives, but the money value of life saving measures. Money can be invested and so the money value of costs and benefits received in the future needs to be discounted.

If regulators valued future lives saved the same as current lives, then it would never be worth spending to save a life today. Money spent today to save lives could instead be invested to produce a larger lifesaving budget in the future, saving more lives. All the more so if technological progress makes the cost of saving lives fall over time. If the value of future lives saved is not discounted, then there is a higher marginal productivity in future spending on lifesaving and all lifesaving resources should be channeled towards the future. But the same argument applies each year. Life-saving expenditures would be delayed indefinitely.[23]

---

[21] For example, Cowen and Parfitt (1992).

[22] Example based on Viscusi (2006, p. 7).

[23] Schelling, (1987).

Choices need to be made between expenditures on reducing the risk to future lives and other goods, not the least of which is saving lives in the present. Putting a money value on the benefits from life saving expenditure just makes explicit what people implicitly do when making choices. People make trade-offs between safety and other uses of resources all the time, for themselves and on behalf of others. People are willing to trade risks to their own lives for quite minor pleasures. Life saving has a financial cost, discounting just allows assessment of the value of expenditures at different periods. 'If willingness to pay to reduce risk is the appropriate metric for allocating regulatory resources, discounting merely adjusts that metric to make expenditures comparable through time'.[24]

[24] Sunstein and Rowell (2007, p. 171).

D The shadow price of capital and the weighted average discount rate

The shadow price of an output or input measures all its social costs and benefits. Just as a businessman evaluates a project by its prospective profit, the government can evaluate a project by its increase in efficiency. Profit measures the extra consumption the project generates for its owner. The net present value of a government project measures the extra consumption it generates for society. Profit is calculated by assessing all the project inputs and outputs relevant to the project owner and converting these into costs and revenues using market prices. These are the prices paid by the business and determine profits. The 'social profit' (or efficiency effects) of a government project are calculated by assessing all of the project inputs and outputs for society and converting these into costs and benefits using shadow prices, which reflect the social value of project inputs and outputs.

Shadow prices differ from market prices when taxes, subsidies, externalities, monopoly, and price and quantity controls distort markets.[1] For example, when a government project uses a taxed input, there is a standard technique in cost benefit analysis for deriving its shadow price (that is, determining its social value).[2] When an input is taxed, the tax drives a wedge between its after-tax supply price and its before-tax demand price. When a government project demands extra units of the input, its price increases to decrease private demand and increase supply until the market satisfies the extra demand. The social cost of the extra input depends on whether it mainly decreases private consumption of the good or increases private supply. That part of government input use that comes from increased private supply is valued at the supply price, that which decreases consumption at the demand price. The shadow price, or social value, of the output is a weighted average of the demand and supply price — with the displacement of private demand and supply determining the weights.

Another way to look at it is that the supply price is the social cost of providing extra units of input the project uses. But when demand is reduced, that decreases tax

1 See Kanbur (2008).

2 See, for example, Department of Finance and Administration (2006, pp. 35-36). For more detail, see Harberger (1971).

revenue — a social cost, increasing the social cost of the extra input above the (after-tax) supply price.

Capital taxes mean that forgone private investment has a cost that is not reflected in project cost and benefit flows. \$1 of private investment has a social value greater than \$1. When determining the social value of a government project which gives costs and benefits over time, we need to shadow price the capital it uses. That is, we account for the consumption that forgone private investment would have produced, adding it to project costs to reflect the full impact of the project on consumption and discount with the consumption rate.

It is more complicated than the standard input case outlined above, as capital lasts for multiple periods and depreciates. The case considered here is a closed economy where all capital returns are taxed at a single rate.

In a closed economy, government investment must come at the expense of consumption or private investment. The shadow price approach calculates the social value of all consumption and investment impacts, converting them all to consumption equivalents. The first step is to determine whether private investment flows will be altered by a policy. Changes in investments are converted into equivalent units of consumption. All flows of consumption and consumption–equivalents are then discounted using the (after-tax) consumption rate of interest, the rate consumers would use in discounting future consumption benefits.

It turns out that discounting ordinary cost and benefits flows using a weighted average discount rate (usually) gives the same answer.

D.1 A perpetuity

Consider a one unit public investment that generates a perpetuity with a real rate of return δ. As the investment is a current cost borne for future benefits, it uses a unit of capital. The capital gives a flow of consumption benefits of δ per year forever. A fraction a of the investment takes place at the expense of private investment and a fraction $(1-a)$ at the expense of consumption (see appendix B for how these fractions are determined).

If private investment produced a perpetual real return i per year and the consumption is discounted at rate r, with a tax rate t on the real return so that $r = (1-t)i$. Then the present value of the consumption produced by \$1 of private capital is:

$$i/(1+r) + i/(1+r)^2 + i/(1+r)^3 + \ldots = i/r = 1/(1-t) .$$

which is greater than \$1. One dollar of the capital stock produces a stream of annual consumption benefits worth i cents per year, with $r = (1 - t)i$ cents going to the owner and $(i - r) = ti$ cents in tax payments to the government. The present value to the owner is the value of r cents forever discounted at r per cent, which is \$$r/r$ = \$1. The present value of the tax revenues is: $\$(i - r)/r = \$(i/r - 1)$. The present value of the benefits to society from an extra unit of private investment is the sum of these two: $\$(i/r - 1) + \$1$ = \$ i/r.

The project displaces private investment that would have produced ai dollars per year. The social cost of the project is the present value of forgone consumption (discounting the consumption flows with the consumption interest rate):

$$(1-a)+ai/(1+r)+ai/(1+r)^2+ ai/(1+r)^3+\ldots = (1-a) +ai/r.$$

This is the shadow price of capital used in the project, its opportunity cost. It is the value of the consumption that the capital used in the project would have produced. That is, the project produces output (a benefit), but uses resources, reducing private sector output (a cost).

The project improves efficiency if its benefits exceeds its costs. That is, if the present value of consumption the project produces is greater than its cost

$$\delta/(1+r)+\delta/(1+r)^2+ \delta/(1+r)^3+\ldots.= \delta/r \geq (1-a) +ai/r. \text{ Or}$$

$$\delta \geq (1-a)r +ai = w$$

The project is worthwhile if the return on investment is greater than the weighted average discount rate, where the weights depend on the amount of investment and consumption the project displaces. In other words, if the project has a positive net present value using a discount rate w:

$$\delta/(1+w)+\delta/(1+w)^2+ \delta/(1+w)^3+\ldots.= \delta/w \geq 1$$

It generates enough benefits to more than compensate private investors and consumers for their forgone consumption. With perpetuities, the shadow pricing approach and weighted average discount rate are two equivalent ways to work out the opportunity cost of capital used in a project.

D.2 In a two period world

The same is true in a two period world. Take a government project that costs 1 unit in period 1 and produces $1+\delta$ of benefits in period 2 and displaces $(1-a)$ units of consumption and a units of investment. The displaced private investment would have produced $a(1+i)$ of benefits in period 2.

The project is worthwhile if:

$$(1+\delta)/(1+r) \geq (1-a)+ a(1+i)/(1+r)$$

The left hand side is the present value of the period 2 consumption benefits the project produces. The right hand side is the shadow price of the capital used in the project, its opportunity cost – the forgone consumption in period 1 and the present value of the forgone consumption in period 2 from displaced private investment. The above project acceptance criterion reduces to:

$$\delta \geq (1-a)r + ai = w$$

The project is efficient if its rate of return exceeds the weighted average discount rate, w. Again the shadow pricing approach and weighted average approach are equivalent.

D.3 In a multi-period world

In a multi-period world, we need to worry about re-investment of project output. As the present value of consumption from \$1 of investment is greater than \$1, the amount of re-investment will affect the value of the project.

For example, consider the appropriate discount rate for a two period project in multi-period world. This is quite general as multi-period projects can be expressed as a sequence of two period projects. For example, an investment that pays off in two periods time is equivalent to an investment that pays off next period combined with a project which reinvests that output and pays off the following period.

One unit is invested in the current period, which produces output $(1+\delta)$ in the following period 2. The project is sourced a portion a from private investment and $(1-a)$ from consumption. Assume a unit of private investment produces a perpetuity paying i per period. Now if no project output is re-invested, the project is desirable if:

$$(1+\delta)/(1+r) \geq (1-a) + ai/r$$

The left-hand side is the present value of the project output. The right-hand side is the shadow price of capital, the present value of the consumption resources used in the project would have produced. This reduces to:

$$\delta \geq (1-a)r + ai +a(i-r)/r = w + a(i-r)/r$$

which substantially raises the rate of return the project must earn. For example, if $a = 0.75$, $i = 3$ per cent and $r = 1$ per cent (plausible risk-free numbers), then $w = 2.5$ per cent and $a(i–r)/r = 150$ per cent. The project must earn 152.5 per cent to be desirable – far above the weighted average rate – because it permanently reduces the private capital stock by 0.75 units, which reduces the present value of consumption by 0.75*0.03/0.01 = 2.25 units.

If some project output were re-invested, then that would increase the private capital stock, reduce the consumption cost of the project and decrease the required project return. If the private capital stock depreciates rather than produce a perpetuity, that would reduce the consumption cost from reducing private investment.[3]

In his influential and masterful survey of the social discount rate, Lind (1982a) recommends shadow pricing capital and discounting with the consumption rate. It is fair to say that very few cost–benefit guides or studies have followed the Lind recommendation because the informational requirements make it impractical. As Lind points out, estimating the shadow price of capital requires we have estimates of the marginal return on private capital, the consumption rate, the length of life of the typical private investment, the displacement and stimulative effects of public investment on private capital formation, the amount of re-investment of public and private project output.[4] The alternative of working out a threshold rate of return to use as a discount rate involves shadow pricing capital and so the information required is exactly the same.

But Lind (1982a) does not refer to the analysis of Sjaastad and Wisecarver (1977).[5] They independently show that re-investment of project output is a crucial determinant of the social discount rate. They also show that if society does not treat public–project depreciation as income subject to current consumption, but instead intend to save (and hence re-invest) all of that depreciation, the social discount rate is again the weighted average discount rate. Only the net return on capital is treated as current income to be consumed.

3 See Bradford (1975) for how the appropriate discount rate varies with the different assumptions about the re-investment of project output and life of the private capital stock. He finds that it can lie above *i* and below *r*.

4 Lind (1982a, (pp. 48-50, 77-78).

5 Lind does consider the re-investment issue on pp.48-50. Lind points out that Bradford (1975) assumes depreciation of private investment not re-invested but is treated as any other component of current income. Lind states it is ‘reasonable to assume’ (p.50) that private project depreciation is fully re-invested. Sjaastad and Wisecarver (1997) assume depreciation is saved in both the private and public sector. Recognising these different assumptions is needed to reconcile the equations in Bradford (1975) with those in Sjaastad and Wisecarver (1977).

Again consider a two–period project where 1 unit is invested and produces gross output $1+\delta$ one period later. Depreciation is 1, the investment fully depreciates in one period. If consumers intend to save depreciation, it shifts the savings supply curve to the right by 1 unit, having effects exactly symmetric with the extraction of 1 unit of capital in the first period. That is, it will reduce interest rates, increasing actual investment by a units and consumption by $(1–a)$ units. The intended re-investment of project depreciation increases actual re-investment by a. But that would, once public investment has fully depreciated, restore the total capital stock, and future income from that stock, to the paths that would have existed if the project had never been undertaken.[6] That is, the project only affects the private capital stock during the life of the project.

Whether the project increases efficiency involves comparing the increase in consumption in period 2 with the consumption forgone in period 1. The increase in consumption in period 2 is the net return from the project, δ, less the consumption forgone from the one-period reduction in private capital, ia, plus the extra consumption from the saving of project depreciation, $1–a$. That is, the project improves efficiency if:

$(\delta - ia + 1-a)/(1+r) \geq (1 - a)$ or $\delta \geq (1 - a)(1+r) + ia - (1-a)$ or

$\delta \geq (1 - a)r + ai.$

That is, the present value of consumption is higher with the project when the project return exceeds the weighted average discount rate, w, or when the project has a positive net present value using a discount rate w (that is, $(1+\delta)/(1+r) \geq 1)$. Again the weighted average and shadow pricing approach are equivalent.

Further, the weighted average approach gives the same project rankings as the shadow pricing approach.[7]

If society recognises the difference between depreciation and net benefits from a public investment project and attempts to save all the benefits which represent depreciation, then the weighted average discount rate is the social discount rate. That people try to save depreciation may appear to be a strong assumption. Consumers of government provided goods are unlikely to be aware what part of

6 Sjaastad and Wisecarver (1977, p. 522).

7 Project 1 has costs C_1 and benefits B_1. Project 2 has costs C_2 and benefits B_2. Project 1 is better under the shadow pricing approach if:

$[B_1 - C_1 - iaC_1 + (1-a)C_1]/(1+r) - (1-a)C_1 > [B_2 - C_2 - iaC_2 + (1-a)C_2]/(1+r) - 1-a)C_2$

where project depreciation is saved so the weighted average pricing approach applies. This reduces to: $B_1/(1+w) - C_1 > B_2/(1+w) - C_2$ which is the ranking when project flows are discounted with the weighted average discount rate, w.

project benefits represents net output and what part depreciation. On the other hand, if people do consume depreciation and run the capital stock down, it needs to be explained why they did not take this opportunity in the absence of the project. There are sound economic reasons for people to try to maintain the capital stock to keep consumption at desired levels.

It may be reasonable to assume re-investment of depreciation, especially with regulation. The government and firms can identify depreciation and reinvest it. For regulatory projects, firms often undertake the project investment, and recoup their costs from consumers. It is reasonable to assume that they can distinguish depreciation from net output and would save depreciation. If the government implicitly commits to provide the good or service beyond the life of the project, that requires reinvestment of public sector funds approximately equal to the depreciation of the project's capital.[8]

When depreciation is reinvested, the impact of the project on the economy lasts only for the duration of the project. That is appears more reasonable than the alternative — that a government capital project permanently alters the mix of consumption and investment in the economy.

Certainly the Sjaastad and Wisecarver approach is more practical. It only requires information the project flows, the investment and consumption rates, and the displacement of private investment. The shadow pricing approach needs all this plus the length of life of the typical private investment and the amount of re-investment of public and private project output. The rate of re-investment of project output would need to be calculated for each project, yet the information for the calculation is unlikely to be available.

In the absence of the necessary information, the most reasonable, convenient and practical approach is to assume society saves project depreciation and to discount with the weighted average discount rate, which makes best use of the information we do have. Thinking in terms of adjusting the discount rate is easier and more informative than the shadow pricing approach when we are ignorant.

If society does not save project depreciation, but consumes it, re-investment is less and the appropriate discount rate is greater than the weighted average rate, perhaps substantially greater. In this case, the project permanently reduces the private capital stock and the resulting reduction in consumption each period is a cost of the project. Further, the rate of discount must be higher the shorter is the life of the project, as

8 Jenkins (1981, p. 400).

short–term projects lead to capital consumption at an earlier date than long–lived projects.[9] Short term projects would face a high hurdle.

If people do save depreciation and also re-invest out of net project output (which increases α), that would reduce the social discount rate below *w*. Sjaastad and Wisecarver find this effect to be quantitatively small.[10]

It is common in the literature to extol the virtues of the shadow pricing approach, while simultaneously dismissing the weighted average approach as incorrect.[11] Sjaastad and Wisecarver show that if society saves depreciation, then the two approaches are identical. Whether that is a reasonable assumption can be criticised, but most critics of the weighted average approach do not. Instead they often make invalid criticisms.

To turn it around, the weighted average approach is only correct if society saves public project depreciation. When weighted average critics present counter-examples, they invariably assume (usually implicitly, and unwittingly) that depreciation is not saved. If the project output flows and corresponding changes in consumption when depreciation is saved are correctly specified, the weighted average and shadow pricing approach give the same answer. If critics specify project flows and changes in consumption inconsistent with depreciation being saved, then the weighted average approach does not give the correct answer.

For example, the United Stated Environmental Protection Agency (EPA) claims that the weighted average approach is 'acceptable for similarly timed cost and benefit flows' and

> it is technically incorrect and can produce net present value results substantially different from the correct result (where 'correct' is defined by the consumption rate of interest–shadow price of capital approach). The problem with the simple weighted average approach is that it seeks to accomplish two tasks using the social discount rate — pure time discounting and adjusting for the displacement of private investments that yield pre–tax social returns higher than the consumption rate of interest.[12]

The EPA claims the weighted average approach 'over discounts' long–lived projects, giving the wrong answer when benefits are far in the future, and that the problem is worse the farther in the future the benefits occur. The next section demonstrates that the claim is incorrect.

---

9 Sjaastad and Wisecarver (1977, p. 523).

10 See Sjaastad and Wisecarver (1977, p. 527).

11 For example, Abelson (2000, p. 129), Bureau of Transport and Regional Economics (1999, p .71, footnote 30) and United States Environmental Protection Agency (2000, p. 4).

12 See United States Environmental Protection Agency (2000, p. 42) and Abelson (2000, p. 129) makes the same point.

D.4 Does the weighted average approach give the correct answer for long–term discounting?

As Sjaastad and Wisecarver prove, in a two-period project, when depreciation is saved then the shadow pricing and weighted average approach give the same answer. Any project can be expressed as a sequence of two period projects. If depreciation is saved at each step in the sequence, discounting the project flows with the weighted average discount rate gives the same answer as discounting the consumption flows with the consumption rate. The EPA does not allow for re-investment of depreciation in its calculations, or the consequences for consumption of not re-investing, and incorrectly specifies the relevant investments.

Take a project in which 1 unit is invested in period 0, and the only effect on consumption is to reduce it by $(1-a)$ units in period 0 and then increase it by $(1-a)(1+r)^n$ units in period n. That project would just break–even (we evaluate consumption changes with the consumption rate, r).

The project is equivalent to following sequence of two period projects, each with a gross rate of return equal to δ:

1. Invest 1 in period 0, consumption falls by $(1-a)$, private investment falls by a.

The investment gives gross output of $1+\delta$ in period 1. 1 is re-invested. The re-investment increases consumption by $(1-a)$ and private investment by a, restoring the private capital stock to its pre–project path. The project therefore gives $\delta - ia + (1-a) = z$ in (potential) extra consumption in period 1. Note that $z \geq (1-a)(1+r)$ if $\delta \geq ia+(1-a)r = w$.

Now invest $z/(1-a)$ in period 1. This will reduce period 1 consumption by z units, keeping consumption in period 1 unchanged. Private investment falls by $az/(1-a)$.

2. The potential consumption and fall in private investment fund the increased investment: $z+az/(1-a)=z/(1-a)$.

3. The investment allows $z^2/(1-a)$ (potential) extra consumption in period 2.

4. Invest $z^2/(1-a)^2$ in period 2. This keeps period 2 consumption unchanged and gives a potential increase in period 3 consumption of $z^3/(1-a)^2$.

5. Invest $z^j/(1-a)^j$ in period j $(j < n)$, which keeps consumption at the initial level in each period j.

6. When we reach period n, the final investment allows consumption of $z^n/(1-a)^{n-1}$. The project is efficient if the present value of the consumption

allowed in period n exceeds the consumption forgone in period 0: $z^n/(1-a)^{n-1}(1+r)^n \geq (1-a)$

7. If $\delta \geq w = ia + (1-a)r$ then $z = \delta - ia + (1-a) \geq (1-a)(1+r)$ and the condition holds. The return on government projects needs to at least w for an investment in period 0 to raise consumption by $(1-a)(1+r)^n$ in n periods and keep consumption constant in period $1, .., n-1$.

That is, if we discount the project flows with the weighted average discount rate w, it gives the same answer as discounting the changes in consumption with the consumption rate, r.

Take a simple example (based on the EPA example, but reduced to 3 periods so it can be presented in a short table). Assume the investment rate is 5 per cent, the consumption rate 3 per cent, $1 of public investment reduces private investment by $0.75 and consumption by $0.25. that is, $a = 0.75$. the weighted average discount rate is $w = 0.75*5 + 0.25*3 = 4.5$ per cent.

Now take a 3 period project where $1 is invested in period 0 which increases consumption in period 2 only. This is equivalent to a sequence of two 2-period projects. The investments earn a return of 4.5 per cent. Depreciation is saved for each project. The period by period flows are set out below in table D.1, 1 unit is invested in period 0, reducing consumption by 0.25 and private investment by 0.75 which produces output of 1.045 in period 1. Depreciation of 1 unit is saved, increasing the private capital stock by 0.75. The potential increase in consumption in period 1 is $\delta - ai + (1 - a) = 0.2575$. In period 1 a new investment of 1.03 is made, which reduces the capital stock by a*1.03 = 0.7725 and crowds out 0.25*1.03 = 0.2575 of consumption in period 1. the combined effect of the two projects is to keep period 1 consumption constant. The new investment produces output of 1.045*1.03 = 1.076 in period 2. Depreciation of 1.03 is saved, which restores the private capital stock to its initial level. Period 2 consumption rises by 1.076 – 1.03 – 0.05*0.7725 + 0.25 = 0.265

Using the shadow pricing approach, the project breaks even. The present value (discounted with the consumption rate 3 per cent) of the change in consumption from the project is $-0.25 + 0.265/(1.03)^2 = 0$.

Table D.1 **Project cost and benefit flows**

A three period example

Period	*New investment in project*	*Gross project output*	*Re-investment*	*Potential increase in consumption*	*Change in private investment from new investment*	*Change in consumption*
0	1	0			– a = – 0.75	–(1–a)=–0.25
1	1.03 = 0.2575/0.25	(1+*w*)=1.045	a*1 = 0.75	0.2575= (1+w)–1–ia+(1–a)	–0.7725 = –(1–a)*1.03	0
2	0	1.076= 1.03*1.045	0.7725	0.265= 1.076–1.03–0.05*0.7725 +0.25*1.03	0	0.265

The EPA suggests the weighted average approach incorrectly rejects this marginal project because if we discount the period 2 gross output at 4.5 per cent, the net present value is $1.076/(1.045)^2$ = \$0.986, less than the \$1 cost of the project. They fail to carefully specify the investments needed to produce the above consumption pattern or account for re-investment. They do not account for the fall in consumption in each project period from the project induced decline in the private capital stock. Further, they have no re-investment in the final project period, which would permanently lower the private capital stock, yet the EPA fails to account for the resulting fall in consumption which would continue after the project finishes. The weighted average approach, properly applied, calculates the present value of the gross project output less the present value of investments made. The net present value of the project flows (columns 2 and 3 in table D.1) is:

$1.076/(1.045)^2 + 1.045/1.045 - 1.03/1.045 - 1 = 0$, consistent with the shadow pricing approach.

A project can be expressed as a sequence of two period projects. Each of the two period projects would increase the present value of consumption only if it earned the weighted average rate of return.

The EPA suggest that a project which involved investing 1 unit in period 0 with a payoff of $(1.03)^2 = 1.0609$ in period 2, with no further investments, would be marginal. The project would reduce the private capital stock by 0.75 units and consumption by 0.25 in period 0. That would reduce consumption by 0.75*0.05 = 0.0375 units in period 1. In period 2 consumption would increase by: $1.0609 - 0.0375 - 1 + 0.25 = 0.2734$ if the public intend to save project depreciation (if it did not, then the capital stock would be permanently reduced, increasing the required project return substantially). The present value of the change in consumption is:

$$0.2734/1.03^2 – 0.0375/1.03 – 0.25 = – 0.029$$

That is, the project return of 3 per cent is not enough to cover its costs. The EPA do not account for the fall in consumption in the intermediate periods from the fall in the private capital stock. Discounting with the consumption rate would only be correct if the government project did not crowd out private investment ($a = 0$), in which case the weighted average rate would be the consumption rate.

D.5 Discounting with the government bond rate

Many cost benefit guides, and practitioners, recommend the before–tax return on government bonds as the social discount rate. Sometimes using the government bond rate is justified as reflecting the cost of funds to the government.[13] But that does not account for the indirect cash flows when the government goes into debt. When we account for these flows, the weighted average discount rate gives the cost of borrowing. For example, if the parameters are as in the previous section, ($r = 3$ per cent, $i = 5$ per cent, $t = 40$ per cent, $a = 0.75$ and $w = 4.5$ per cent), when the government borrows $100 it:

- makes a $5 interest payment
- receives $2 in tax payments from those receiving the interest
- crowds out $75 of private investment, which reduces tax revenue by $75*0.05*0.4 = $1.50.

The total annual cost to the government of borrowing $100 is $5 – $2 + $1.50 = $4.50, or 4.5 per cent, the weighted average rate. Alternatively, extra tax revenue only comes from the newly stimulated savings, so is $25*0.05*0.4 = $0.50. So the total cost is $5 – 0.50 = $4.50, which is the weighted average rate times $100. The weighted average rate gives the cost of financing.

Really, whether the government bond rate is appropriate turns on how the cost benefit analysis accounts for the cost of risk, which is considered in the risk section.

13 See, for example, Department of Treasury and Finance Tasmania (1996, p. 16).

E Foreign Capital

So far we have analysed a closed economy,[1] but increased demand in the capital market can also be met from increased capital funds from abroad. The weighted average approach can be readily adapted to incorporate foreign capital. The weighted average discount rate would then be: $ai + (1–a–f)r + fMC_f$, where MC_f is the marginal cost of foreign funds, f is the proportion of investment sourced from foreign funds and a the proportion that comes at the expense of domestic investment.[2]

The option of borrowing from overseas makes the supply of capital more elastic, which reduces the rise in domestic interest rates and the amount of private investment crowded out in response to increased capital demand from a government project. If the supply of foreign funds were perfectly elastic (i.e. the supply curve horizontal), the weights on i and r would be zero and the social discount rate would be the marginal cost of foreign funds (as all project costs would come from foreign borrowing with no domestic consumption or investment being crowded out).

The effects of integration of the Australian capital market with international capital markets depends on complicated details of international tax law and how Australia's borrowing affects the interest rate that foreigners charge us.

E.1 Perfectly integrated markets

The supply curve is horizontal in the case of perfectly integrated markets (or perfect capital mobility), where Australia could borrow as much as it liked at a given world interest rate (that is, Australia is a price taker). That world rate would be the marginal cost of foreign funds and depends on how foreigners are taxed.[3]

For debt financing, international tax treatment of interest payments is on a residence-based income tax structure. That is, each country taxes its own residents on their world-wide interest income (from any source) and does not tax non-

[1] See appendix B.

[2] Edwards (1985), Sandmo and Dreze (1971).

[3] The following analysis of taxation of international capital flows is from Cohen et al. (1999, pp. 203-07).

residents. With perfect capital mobility, the supply of funds would be horizontal at the real foreign before-tax interest rate, which would be the marginal cost of foreign funds, independent of both domestic and foreign tax rates on interest income (see box E.1).

Box E.1 Cost of foreign debt in a perfectly integrated capital market with residence-based taxation

In a perfectly integrated capital market covered interest parity holds. For a marginal risk neutral investor, the nominal after-tax rate of return on Australian debt instruments equals the exogenous after-tax return on a foreign debt instrument plus the expected rate of depreciation of the Australian dollar. With residence based taxation the applicable tax rate for Australians is the Australian tax rate, while the applicable tax rate for foreigners is the foreign tax rate. This gives two parity conditions. For Australians: $\rho(1-\tau) = \rho^*(1-\tau) + \Delta s(1-\tau_g)$ where ρ^* is the exogenous foreign nominal interest rate and s is the expected log future spot dollar value of foreign exchange and so Δs is the expected rate of depreciation in the Australian dollar. τ_g is the tax rate on foreign exchange gains. Then Australians will be indifferent between investing home and abroad.

For foreigners the parity condition is: $\rho(1 - \tau^*) = \rho^*(1 - \tau^*) + \Delta s(1 - \tau_g^*)$ where τ^* is the foreign tax rate on income, τ_g^* the foreign tax rate on foreign exchange gains.

If foreign exchange gains are taxed at the same rate as other income, then before-tax interest rates will be equalised (adjusted for exchange rate changes): $\rho = \rho^* + \Delta s$

If the expected rate of depreciation equals the difference in inflation rates: $\Delta s = \pi - \pi^*$ then $\rho = \rho^* + \pi - \pi^*$ and so $\rho - \pi = \rho^* - \pi^*$ or $i = i^*$ In this case the real before-tax interest rate is determined by the real foreign before-tax interest rate and the traditional Fisher effect holds: $d\rho/d\pi = 1$. A 1 percentage point rise in the Australian inflation rate increases our nominal interest rates by 1 percentage point. The evidence on whether the Fisher effect holds is mixed (see appendix A.3). Some studies find it holds since deregulation, but others reject it. Changes in the relative price of traded and non-traded goods (including terms of trade changes) would cause the expected rate of depreciation to differ from the difference in inflation rates. Further, if large countries determine the foreign interest rate, then changes in their inflation rates would affect the before-tax world real interest rate (just as in a closed economy, see appendix A.3). If foreign and Australian inflation rates were correlated, then that it would appear as if the Fisher effect did not hold because Australian inflation rates and real interest rates would both change.

Source: Cohen et al. (1999, pp. 203-04).

For equity finance, source-based taxation is more applicable. Under pure source-based financing, income originating in a country is taxed in that country regardless of where the recipient lives. A country does not tax its residents' foreign sourced income. Real world tax laws are more complicated, with countries taxing their own residents foreign source income, but giving credit for taxes paid to foreign governments (if a tax agreement with that country is in place), which means residents pay taxes on their foreign sourced income at the higher of the two rates.

When foreigners supply equity funds, it is taxed at Australian tax rates and the revenue raised accrues to Australians. With perfect capital mobility, foreigners would supply equity funds so long as the after-tax return exceeded their after-tax return from investing at home. The cost to Australia of extra foreign equity investment would be the after-tax real foreign return (which would be below the Australian after-tax real return if Australian tax rates were lower than foreign ones).[4]

E.2 An upwards sloping supply from foreigners

If extra borrowing by Australians bids up the price that foreign lenders charge (that is, the supply of foreign funds is upwards sloping) then the marginal cost of funds is greater than the average cost. The marginal cost of foreign funds would be greater than the foreign interest rate. Further, a government project would then displace some domestic consumption and investment.[5]

The possibility of sovereign default means the supply of foreign capital to Australia slopes up, even if Australia is too small to affect the world interest rate.[6]

Sovereign default occurs when a country repudiates its debt or negotiates a partial default. For example, the government imposes exchange controls which control domestic private borrowers' repayment of foreign debt, and foreign lenders cannot take effective legal action to receive more than the government stipulated repayments.

---

4 One problem with this analysis is that if the Australian and foreign tax rates differ, then the equilibrium for foreign investors is not compatible with the equilibrium for Australian investors. If $\tau^* < \tau$, foreigners face a higher tax rate on their Australian investment than their home investments and foreign equity investment in Australia occurs until $(1-\tau^*)i^* = (1-\tau)i$. But Australians investing overseas would pay the same (Australian) tax rate on all investments, and would want to invest until $(1-\tau)i^* = (1-\tau)i$. The two are incompatible if $\tau \neq \tau^*$. If $\tau^* > \tau$, then Australians invest until $(1-\tau^*)i^* = (1-\tau)i$, which is incompatible with the foreigner's equilibrium, $(1-\tau^*)i^* = (1-\tau^*)i$ when tax rates differ. Further, when tax rates differ between countries, countries with high corporate tax rates may specialise in providing in supplying debt and those with lower rates supplying only equity. See Jones (2008, p. 233).

5 Edwards (1985, pp. 11-12) and Sandmo and Dreze (1971, pp. 403-04).

6 The following analysis is from Fane and Applegate (1995).

Lenders charge more to compensate for the risk of default. For example, say the world risk free rate is 3 per cent and that at current levels of Australian foreign borrowing, there is a 1 per cent annual chance of default (a once in a century event). The rate foreigners charge is 4% = 1.03/0.99 – 1.[7]

Increased borrowing by any of the country's residents increases the probability of national default and increases the cost of borrowing (the country risk premium). When one resident borrows more, it increases the interest rate other borrowers pay.

The chance of sovereign default creates other externalities from increased foreign borrowing. When a resident borrows more from overseas, it raises the probability that other domestic borrowers will be able to avoid making contracted repayments. Further, when individuals borrow more, they ignore the costs to others of the penalties imposed in the event of default. These two externalities cancel out if the sovereign borrower chooses to default when the penalty for default is less than the amount owed, leaving the increase in the cost of borrowing as borrowing increases as the only relevant externality. That is, with sovereign default, the borrowing country's risk premium is equal to the net negative externality from foreign borrowing.

The optimal tax on foreign borrowing to take account of this externality, under perfect capital mobility, is exactly the same as the optimal tax to take advantage of monopsony power when a country faces an upwards sloping supply curve of foreign debt and there is no sovereign risk. That is, sovereign risk results in an effective upwards sloping supply curve under perfect capital mobility, just as monopsony power would. The marginal social cost of extra foreign borrowing is greater than the average cost. That is, the rate of return required to cover the cost of increased foreign borrowing is higher than the rate of interest required by foreign lenders, as it must account for the extra borrowing bidding up that rate. The marginal cost equals $i_f(1+1/e)$ where e is the elasticity of the supply of foreign capital to Australia.[8]

7 More generally, if lenders require a return of i and the probability of default on the whole debt is λ, then the interest rate charged on the loan, i', would be $(1+i) = (1+i')(1-\lambda)$. Therefore $i= i' - \lambda - i'\lambda$ and so $i' = (i + \lambda)/(1- \lambda)$. Further, if default is more likely to occur in a recession, then there would be a further premium to account for the market risk that the chance of default introduces to the payoffs.

8 Interest payments to foreigners are i_fD, where i_f is the interest rate charged by foreigners and D is the total amount of foreign debt. i_f increases with D. Extra borrowing increases total interest payments by $\partial i_fD/\partial D = i_f + D\partial i_f/dD = i_f(1+1/e)$ where $e = (\partial D/\partial i_f)(i_f/D)$, the elasticity of the supply of foreign capital to Australia. With linearity, the marginal cost line should have twice the slope of the supply curve. The premium in the marginal cost of funds is double the country risk premium. $i_f = aD+b$. $\partial i_f/\partial D = a$. $TC = i_fD = aD^2+bD$. $MC = dTC/dD = 2aD+b$. The effect of taxes imposed on foreigners depend on the elasticity of supply and demand for foreign capital, as in figure B.1. If foreign capital is only a small portion of total capital in Australia, then the

As illustrated in figure E.1, if the world risk free rate is 3 per cent and the lenders charge 4 per cent at current borrowing levels, assuming linearity, the elasticity of supply is 4.[9] The marginal cost of foreign borrowing is 4%(1+1/4) = 5%. Even a small probability of default can make a large difference to the marginal cost of funds.

Figure E.1 **The cost of foreign funds to Australia with a 1 per cent probability of default**

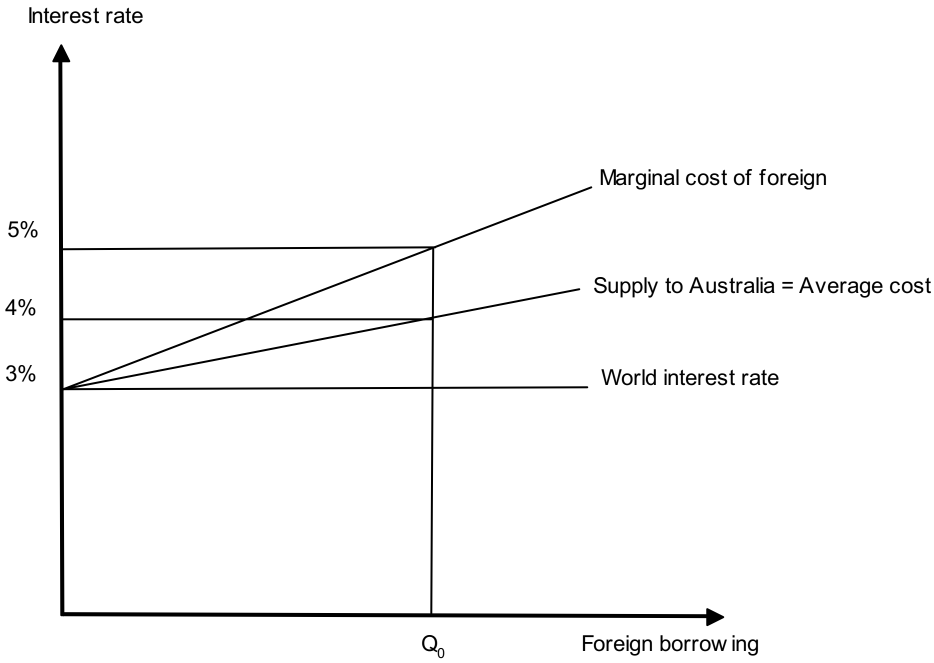

Alternatively, even a small open economy with perfect capital mobility will faced an upwards sloping supply curve of foreign funds if lenders and borrowers have different perceptions of default risk.[10]

E.3 Conclusions for the social discount rate

In a closed economy, it was concluded that the weighted average discount rate was likely to be close to the before-tax investment rate because the supply of savings was relatively inelastic compared with the demand for investment. Foreign capital investment makes the supply of capital to Australia more elastic and the weighted

demand for it is likely to be quite elastic and the burden of the taxes on it is likely to be borne by the foreign lender.

9 If the current level of foreign borrowing is Q_0, the slope of the supply curve is Q_0/1% and so the elasticity of supply is (Q_0/1%)*4%/Q_0 = 4.

10 Edwards (1985, p. 11).

average discount rate is: $ai + (1-a-f)r + fMC_f$, where MC_f is the marginal cost of foreign funds, f is the proportion of investment sourced from foreign funds and a the proportion that comes at the expense of domestic investment.

But the foreigners are likely to supply funds for debt at the before-tax interest rate. Further, the supply of foreign borrowing is likely to slope upwards, which makes the marginal cost of foreign borrowing greater than the average cost. The net effect is to leave the conclusion that the weighted average rate is close to the before-tax return unchanged. For example, if the risk free before-tax rate were 4 per cent, the after-tax risk free rate 1 per cent and 30 per cent of funds came from sources from which society earned the after-tax return and 70 per cent from sources yielding the before-tax return, then the weighted average is above 3 per cent: 0.3*1% + 0.7*4% = 3.1%. Further, because foreign equity financing has an marginal cost greater than the before-tax return and foreign debt financing greater than the before-tax return, the true weighted average rate is higher still.

F Deriving the Ramsey equation[1]

As set out in Appendix H, $p = E(mx)$ is the basic consumption based pricing equation, where p is the price of the asset, x is its payoff next period and $m = \gamma \frac{u'(c_{j+1})}{u'(c_j)}$ is the stochastic discount factor. This is an equation that holds for each individual, for each asset to which he has access and for any two periods of a multi–period model. If the equation does not hold, then the investor would buy more or less of the asset (and change current and future consumption) until it does hold. It comes from the first order conditions for a consumer with the utility function $U(c_j, c_{j+1}) = u(c_j) + \gamma E_j[u(c_{j+1})]$. The assumption that the consumer maximises time and date separable expected utility could be dropped. The basic consumption pricing equation would hold for any general non-separable utility function. The relation between the discount factor and real variables would, however, be more complicated. It would be a function of many variables, not just the discount on future utility and present and future consumption.

For a risk free asset the basic pricing equation becomes $1 = E(mR^f) = R^f E(m)$ where R^f is the gross risk-free return. This gives $R^f = 1/E(m)$

If the consumer has a power utility function $u(c_j) = \frac{\left(c_j^{1-\eta} - 1\right)}{(1-\eta)}$ with $\eta > 0$, which gives $u(c_j) = ln(c_j)$ for $\eta = 1$. Power utility has $u'(c) = c^{-\eta}$ and so a constant elasticity of marginal utility (or co-efficient of relative risk aversion) of $-\eta$. Then the basic pricing equation for a risk free asset is: $R^f = \frac{1}{\gamma}\left(\frac{c_{j+1}}{c_j}\right)^{\eta}$

Taking logs gives the Ramsey equation: $r_t^f = \theta + \eta g$, where $r_t^f = \ln R_t^f$ is the continuously compounded risk free rate (see Appendix A for the relationship between continuously compounded rates and their discrete per period counterparts; they are approximately equal), $\gamma = e^{-\theta}$ and $g = ln(c_{j+1}/c_j) = lnc_{j+1} - lnc_j$ is the continuously compounded growth rate of consumption.

1 Cochrane (2005, pp. 10-11, 35-36).

The power utility assumption is a convenient simplification that makes the elasticity of marginal utility independent of consumption. It gives the simple form of the Ramsey equation, where the risk free rate depends on consumption growth. It is a first approximation to the more general case.

G Rates of return in Australia

G.1 The risk free rate of return

Asset holders have a return they expect to make over the period that they will hold the asset. The actual return is uncertain and may differ from expected returns – creating risk. Risk is the variance in actual returns around the expected return. An asset is risk free if the actual always equals the expected return.[1]

A riskless asset should be free from all risk. Possible risks include:

- default risk, the risk that the seller of the asset will not make the required repayments;
- inflation risk, the risk that inflation will differ from its expected level, causing the real returns realised on assets which pay specified nominal amount to vary;
- reinvestment risk, the risk of interest rate changes when asset life differs from the investor's desired investment horizon, and;
- liquidity risk, the risk of being unable to sell the instrument for cash at short notice without significant costs.

The government bond rate is usually considered the riskless rate because the government can compel taxpayers to finance repayments, making default risk low. But many types of government bonds exist, with different returns. All government bonds are subject to some risk, and the risks vary across different types of government bonds. For example, most bonds offer to pay some nominal amount at a future date, subjecting bond holders to the risk of unexpected changes in inflation and making the real return from holding bonds risky. Unexpected higher inflation reduces their real return.

In the United States, economists tend to use the return on short term (90 day) Treasury Bills as the riskless interest rate. The inflation risk is small because of the short time period. US Treasury bills have averaged (over long time periods) an

1 Damodaran (2008, p. 3).

annual realised real return of around 1 per cent.[2] It could be argued that the real return on short term bonds are negatively correlated with the market (the market does well in the short–run when inflation is unexpectedly high and poorly when inflation is unexpectedly low). If so, the short term bond return may understate the riskless interest rate.[3]

Bonds with a longer maturity have more inflation risk, because inflation further in the future is more difficult to predict, and inflation can vary much more over long periods. Consequently, the real yields on long–term bonds are uncertain. If they are positively correlated with the market (the market does poorly when inflation is unexpectedly high over long periods) then the expected real yield on long term bonds must be higher than the risk free rate to compensate.

On the other hand, long–term bonds contain less reinvestment risk than short–term term bonds. As Campbell points out:

> For long–term investors, money market investments are *not* riskless because they must be rolled over at uncertain future interest rates. Just as borrowers have come to appreciate that short–term debt carries a risk of having to refinance at high rates during a financial crisis, so long–term investors must appreciate that short–term investments carry the risk of having to reinvest at low real rates in the future. For long–term investors, an inflation–indexed long–term bond is actually less risky than cash. A long–term bond does not have a stable market value in the short term, but it delivers a predictable stream of real income and thus supports a stable standard of living in the long term.[4]

To minimise reinvestment risk, the risk free rate used should match the duration of the cash flows being discounted. Most government projects have a long duration, and so short term bond rates are not appropriate.

Since 1986, the Australian government has issued inflation–indexed bonds, which offer a guaranteed real before-tax return, free from inflation risk. They tend to be long–term bonds with 10–year terms, an appropriate time period for long–term government projects.

Bonds may have different amounts of liquidity risk. For example, assets that sell in thin markets (when sales volumes are low and may easily dry up) have higher liquidity risk. The higher liquidity risk, the higher the premium investors require to hold the asset. Differences in liquidity risk could explain the low return to US

2 See Goetzmann and Ibbotson (2005, p. 18); Mehra (2006, table 2.1, p. 5); Mehra (2003, table1, p. 55) and Cochrane (2005, p. 21).

3 Sieper (1981, p. 23) and Summers (1990, p. 121).

4 Campbell (1999, p. 2).

Treasury Bills compared with Australian bonds (the US Treasury Bill market is the largest and most liquid market in the world).

Indexed bonds are less liquid than other government bonds, because its market is much smaller. It must pay a higher yield to compensate investors for this relative illiquidity. For example, in the United States, it is recognised that part of the return on Treasury Inflation Protected Securities compensates for lower liquidity, but that 'does not appear to be constant and can change on a weekly basis'.[5] The liquidity premium can become extremely high in times of financial distress. The inflation risk premium, however, is estimated to increase the returns on long-term nominal bonds, compared with indexed bonds, by about 1 percentage point

The expected real return on bonds is the nominal return less the expected rate of inflation over the period. Unfortunately, it is difficult to observe expected inflation. Some analysts use an estimate of expected inflation to estimate the expected real interest rate from nominal returns. For example, they take the nominal 10–year bond return (or some average of recent rates), currently around 6 per cent, and subtract the mid–point of Reserve Bank's target inflation rate (2.5 per cent) to get a real return of 3.5 per cent.[6] But the Reserve Bank's target is often breached.

Another approach is to infer the expected real return from realised real returns. If inflationary expectations are unbiased then realised inflation would equal expected inflation plus an error term. The average of realised inflation rates over long periods would give the expected inflation rate.

Table G.1 sets out realised annual real returns on government bonds since 1971. Clearly inflationary expectations can be wrong for long periods of time. For example, realised real returns on short–term 90 day bonds were negative in the 1970s, a sign that inflation was higher than anticipated.

What is the relevant risk free rate for a project that gives a return over 10 years? A 90 day Treasury bill rate would not be risk free – because there is the reinvestment risk of being subject to interest rate changes when rolling over to new bills every 90 days and inflation risk. It may overstate or understate the true risk free rate. The expected return on a ten–year zero coupon bond would give the nominal risk–free interest rate, but the real return is subject to inflation risk. It overstates the risk free rate. The returns on indexed bonds may overstate the riskless real interest rate because of the liquidity premium.

5 From Federal Reserve Bank of Cleveland (2009).

6 For example, Access Economics (2006, p. 47).

It turns out that the returns on short– and long–term government bonds in Australia have been very similar, earning identical returns over the whole period and with average nominal returns each decade differing by less than 1 percentage point.

Table G.1 **Government bond average annual realised returns 1971–2008 Australia**

arithmetic means

Time period[a]	*10 year govt bonds (nominal)*	*90 day bank accepted (nominal)*	*90 day bank accepted (real)*	*Indexed (from July 1986)*	*Inflation*[b]
	%	%	%	%	%
1971–2008	9.2	9.2	2.9		6.2
1970s	8.4	8.7	–1.5		10.4
1980s	13.3	14.1	5.2	4.9	8.4
1990s	8.9	7.9	5.2	4.6	2.6
2000s	5.8	5.7	2.2	3.0	3.2
1987–2008	8.2	7.9	4.1	4.0	3.6
Standard deviation	3.2	4.0	3.5	1.0	4.2

[a] Years are financial years. For example, 2007 is the year to 30 June 2007. 1970s is 1 July 1970 to 30 June 1979. [b] Inflation measured by CPI.

Source: Derived from tables at www.rba.gov.au.

The real return to short–term government bonds has averaged 2.9 per cent since July 1971. If we exclude the 1970s, the return is higher (averaging over 4 per cent for July 1986 to 2007). Further, measured inflation (such as the change in the Consumer Price index, CPI) overstates the true level of inflation, favouring an upwards adjustment in the riskless real interest rate over the realised real return on government bonds (including indexed bonds because their inflation adjustment payments over–compensate consumers for price changes). Researchers estimate that the US CPI overstates inflation by about 1 percentage point per year.[7]

These results are consistent with other studies on the real returns on government bonds in Australia. Real returns on 90 day bills averaged over 2 per cent real (arithmetic mean, geometric mean only slightly lower) from 1883-2005, 1958-2005, 1883-1957. But from the 1980s on, the real return rose to just over 4 percent (with a 95 per cent confidence interval of about plus or minus one percentage point).[8] Real ten year bond returns are about half a percent higher – 2.55 per cent for the long term and 4.5 per cent since the 1980s. Their confidence intervals are about the same.

7 See box 1.1.

8 The 95 per cent confidence interval is the mean plus or minus $1.96\sigma/\sqrt{n}$. For 90 day bank accepted bonds σ = 3.5 per cent. For a 27 year horizon, the confidence interval is plus or minus 1.3 per cent.

The long–term average real return on indexed bonds (from 1987) is very close to the realised real return on 90 day government bills (4.0 per cent v 4.1 per cent). Their return was lower than the return on bills in the 1980s and 1990s, but higher in the 2000s.

The realised real return on bonds has not been stable. For example, real rates were negative in the 1970s, averaged around 5 per cent in the 1980s and 1990s, and 3 per cent or less in the 2000s, changes unlikely to have happened by chance.[9] A reasonable estimate of the before–tax riskless return is Australia is around 3–4 per cent. All government bonds likely pay a premium above the risk free rate to compensate for the risks they impose. On the other hand, because the CPI overstates inflation, realised bond returns may understate the true real risk-free return.

G.2 The market risk premium

The risk premium is usually measured by the difference in the return investors receive on a broad share market index and the government bond rate (a measure of the risk free return). This measure probably understates the share market premium over the risk-free rate because government bond nominal returns include a premium to compensate for the risks they impose (although the premium could be negative for short-term bonds).

Table G.2 summarises the realised real annual returns from July 1970 to June 2008 on a broad Australian share market index. Over that period, the real return averaged 6.7 per cent, giving a 3.8 per cent premium over 90 day treasury notes (which earn a return close to 10 year bonds, see table G.1). As both returns are deflated with the CPI, its overstatement of the inflation rate does not bias the estimate of the premium.[10] These estimates do not account for imputation credits, and so understate the before-tax returns to individuals since imputation was introduced (on July 1, 1987). The Sharpe ratio, a measure of the ratio of reward to risk in asset markets (see appendix H), is the risk premium divided by the standard deviation in market returns (3.8 per cent/19.3 per cent) is 0.2.

---

9 With a standard deviation of 3.5 per cent, the 95 per cent confidence interval with 10 years data is plus or minus 1.1 per cent. For the inflation risk-free indexed bonds, the real return is more stable and the confidence interval with 10 years data is less than 1 per cent.

10 But the premium estimated from real returns is slightly lower from the one estimated from nominal returns because the real premium is $r_e - r_f$. The nominal premium is $(1+\rho_e) - (1+\rho_f) = (1+r_e)(1+\pi) - (1+r_f)(1+\pi) = r_e - r_f + \pi(r_e - r_f)$ where π is the rate of inflation, r_e is the return to equity and r_f the risk free return. Further, inflation may affect the different asset classes differently.

Table G.2 uses the arithmetic average of annual returns. When expected returns are constant (independently and identically distributed — i.i.d.), the arithmetic mean $\bar{r}_a$ is an unbiased estimate of the expected rate of return, E(r). If returns are uncertain and the time period is longer than a year, the arithmetic mean lies above the geometric mean, the annualized measure of the actual compounded change in the index over the period (as if there had been no volatility in return).[11] The difference will be greater the longer the time period and the greater the volatility.[12] The mean geometric return on shares from 1971 to 2008 was 5.5 per cent per year.

Table G.2 **Realised real annual returns, Australia.**[a]

arithmetic average

Time period[b]	*Share market index*	*Equity premium over 90 day bank accepted bills*	*Return to capital in the national accounts*
	%	%	%
1971–2008	6.7	3.8	8.5
Standard deviation	19.3	18.8	3.0
95 % confidence interval[c]	6.2	6.0	1.0
1970s	–5.6	– 4.1	8.2
1980s	14.0	8.8	6.9
1990s	9.0	3.8	7.7
2000s (to 2008)	8.3	6.0	11.7
1980–2008	10.5	6.2	8.6

a Nominal returns are deflated by the Consumer Price Index.· **b** Years are to 30 June. 1970s is 1 July 1970 to 30 June 1979. **c** The confidence interval is the mean plus or minus this number (which is 1.96σ/√n where σ is the standard deviation and n is the number of observations.

Source: Vanguard http://www.vanguard.com.au/Personal_Investors/Tools_and_education/Calculators_and_tools/Index_chart/index.aspx,;Dolman (2007) and www.rba.gov.au. The share market index prior to 1980 is the MSCI Australia gross total return Index. After 1980 it is the S&P/ASX All ordinaries Accumulation Index. The author converted nominal returns into real returns and calculated the equity premium.

Using the arithmetic mean to discount future payments over-estimates the correct discount factor and so the arithmetic mean understates the discount rate that should be used to discount cash flows. The discount factor for a payment received in n periods time is $[1/(1+E(r)]^n < E[1/(1+\bar{r}_a)^n]$ as $\bar{r}_a$ contains estimation error that is compounded when transformed into a discount factor for payments received more than one period in the future.[13] As the geometric mean is below the arithmetic mean, it is even more biased.

[11] If the index rises by x per cent over n years, the average annual geometric return is $(1+x)^{1/n}$.

[12] McCulloch (2003).

[13] See Cooper (1996) p.160. Cooper also derives unbiased estimators of $[1/(1+E(r)]^n$.

The evidence is that returns are not i.i.d.[14] Nevertheless, the arithmetic mean is likely to be closer to the true discount rate than the geometric mean if transient disequilibrium causes the serial correlation in returns. If the risk premium changes over time, required rates of return at any date would have to be estimated conditional on the set of variables that predict the risk-premium using the appropriate model.[15]

The geometric mean may be relevant if there is mean reversion in stock returns, where periods of relatively high returns tend to be followed by periods of relatively low returns and vice versa. There have been claims that long–horizon (two– to ten–year) stock index returns tend to follow a mean–reverting pattern.[16] Cochrane argues that although aggregate stock returns show predictability over long periods, the evidence for mean reversion is weak.[17]

The risk premium is the extra return needed to compensate investors for the expected risk of holding shares and cannot be observed, as expected returns cannot be measured. We only observe realised returns – and negative realised premiums demonstrate that realised returns can differ greatly from expectations. Average realised returns estimate the true risk premium if expectations are correct on average.

The variance in market returns is so great, that the market return cannot be estimated precisely, even with 38 years of data. Over the past 38 years, the Australian share market real annual return has varied from 59 per cent to minus 36 per cent. The standard deviation in share market returns is 19.3 per cent, implying a 95 per cent confidence interval of plus or minus 6.2 percentage points. That is, we can only be 95 per cent certain that the market return lies somewhere between 0.5 and 12.9 per cent. Certainly, substantial estimation error is possible. The return is very sensitive to the starting and finishing dates of the measurement period. For example, measuring from 1980, the share market real return averages 10.5 per cent, a premium of 6.2 per cent over short term bonds.

Over any one decade, the realised equity premium has great variability. Although from 1970 to 2008 it averaged 3.8 per cent, it was negative in the 1970s and over 8 per cent in the 1980s.[18] With only 10 years data, the 95 per cent confidence interval would be plus or minus 12 percentage points.

---

14 See the next section.

15 See Cooper (1996, pp. 163-65).

16 See Jones and Netter (2008).

17 Cochrane (2005, pp. 415-22).

18 Richards (1991, p. 11) finds a negative risk premium of -1.9 per cent for 1970-90.

More years of data can help pin down the estimated risk premium more precisely, but the standard deviation is still large, giving a 95 per cent confidence interval of at least plus or minus 3 percentage points even with 100 years of data.

Studies of the risk premium (over 10 year government bond returns) over long periods in Australia tend to find that it averages at least 6 per cent, often the estimates are over 7 per cent.[19] Relative to Treasury Bills, the estimates are higher (usually about half a percentage point). Geometric returns are about 1.5 to 2.5 percentage points lower. The studies use data that goes back to 1883. But we cannot be certain that the average risk premium has been constant over long time periods. Does anyone believe the risk premium in the 1890s tells us much about today's market? 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 the riskless rate.[20]

To complicate matters, a decline in the risk premium that people require to hold equities would boost share prices. That is, realised returns jump, increasing the measured risk premium, when the true premium falls. Conversely, when risk premiums increase (as they have in response to the sub–prime mortgage crisis), realised returns fall. In 2008-09 (to February 28) the Australian share market fell by 33 per cent in nominal terms (more in real terms). That reduces the measured average risk premium since 1970 by more than 1 percentage point, just when actual risk premiums have risen. The effect on 100 year estimates of the risk premium would be small.

Dimson et al (2002) estimate global equity premiums from 1900 to 2001. They find Australia's equity risk premium to be 7.9 per cent over ten year bonds (8.5 per cent over bills), above the world average of 5.4 per cent (5.9 per cent over bills). Australia was unusual in having higher risk premiums (over bills) in the first 50 years of the 20th century than in the next 52 (about 9.25 per cent v 7.75 per cent, whereas for the world it was 4.1 per cent v 7.7 per cent arithmetic average).[21] In

19 Brailsford et al. (2007) provide a comprehensive study of equity, bond and bill rates for varying time periods dating back to 1883. See also Gray and Officer (2005); Lally (2005); Dimson et al. (2002) and Bishop (2007).

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

21 Dimson et al. (2002).

contrast, Brailsford et al (2007) find a premium over bills of 6.8 per cent from 1900 to 2001, and from 1883 to 1957 of 6.4 per cent and 1958 to 2005 (chosen because the data quality is better after 1958) of 6.8 per cent.

An alternative approach to estimating the market risk premium is to use the equation of yield. At any point in time the expected annual rate of return on an asset is given by: $r = D/P + \Delta P/P$ where D is the expected cash flow (dividend payment) over the year, P is the price of the asset and ΔP is the expected capital gain over the year. D/P is the dividend yield and ΔP the expected rate of capital gain. All the terms should be consistent (e.g. in nominal terms for the nominal return, in real terms for the real return).

The equation of yield comes from the present value equation for an asset's price:

$P_t = D_t/(1+r_t) + D_{t+1}/(1+r_t)(1+r_{t+1}) + D_{t+2}/((1+r_t)(1+r_{t+1})(1+r_{t+2}) + \ldots$ is the asset price at the beginning of time t (with cash flows received at the end of the period) where r_i is the expected rate of return in period i.[22] The price of an asset is the present value of its expected future cash flows. The equation holds for expected returns (where each term is the expected value) and also defines realised returns.

For the share market, the real expected return is $r_f + e = D/P + \Delta P/P$ where r_f is the expected risk free return and e is the risk premium.

Some experts 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 over the past 60 years, and the true equity premium is closer to the 4 per cent of recent decades. That is, people have learnt the true distribution of returns is better than originally thought. Share prices would increase as people revised downwards their expectation of these events occurring, and the average of past capital gains would over-estimate expected future capital gains. That is, past realised returns overstate expected future returns and so the average equity risk premium in the past is an over-estimate of expected future risk premiums.

Instead of estimating the expected capital gain with past capital gains, authors use the past growth in dividends or the expected growth in GDP (usually based on past growth rates) to estimate the expected capital gain. Then plugging in values for the expected dividend yield and risk free rate gives the expected risk premium. Alternatively, plugging in the realised average dividend yield and risk free rate over

22 Likewise, $P_{t+1} = D_{t+1}/(1+r_{t+1}) + D_{t+1}/(1+r_{t+1})(1+r_{t+2}) + D_{t+2}/((1+r_{t+1})(1+r_{t+2})(1+r_{t+3}) + \ldots$. Substituting this into the expression for P_t gives: $P_t = D_t/(1+r_t) + P_{t+1}/(1+r_t)$ which can be written as: $r_t = D_t/P_t + (P_{t+1} - P_t)/P_t$ which is the equation of yield.

long time periods gives the average past expected risk premium. Most people find it to be in the range of around 2.5 to 5 per cent.[23]

The trailing dividend yield measures the past 12 months dividends divided by the current share price (and so is a real return). It averaged between 3 and 4 per cent between 2002-03 and 2007-08. In 2007-08 it rose steadily to its current level (as of February 2009) of above 7 per cent (over 9 per cent when grossed up to account for franking credits), close to all time highs in Australia.[24] Although the expected dividend yield could be lower – especially because 40 per cent of companies announced a cut in dividends in the previous month – the high dividend yield (and low bond yield) indicate that risk premiums are high (which is typical in a recession). With a dividend yield of 7 per cent, risk free rate of 2 per cent and expected capital gain of 2 per cent, the risk premium is 7 per cent.

The long-term average dividend yield is around 4 per cent, the long-term risk free rate around 3-4 per cent and real growth in dividends (or, alternatively real growth in GDP or earnings) around 3 to 4 per cent. The equation of yield implies the expected risk premium has averaged around 3 to 5 per cent.

The equity premium is large. At around 3 to 8 per cent it is larger than the risk free rate, and so constitutes more than half of the market return. The quantity and price of an asset's risk is a major determinant of its value. The size of risk premium to include in a project discount rate makes a huge difference.

The growth in dividends and risk free rate are both fairly stable. If the expected return were constant, that would imply that the expected capital gain would equal the expected growth in dividends and the expected dividend yield would be constant. Realised returns would fluctuate randomly around the expected return.

But the dividend yield changes over time, which means returns are not i.i.d. Price/dividend ratios can only move if they forecast future returns, if they forecast future dividend growth or if there is a bubble (where the price/dividend ratio would

---

23 For example, in the United States, the realised equity premium averages around 6-8 per cent for long periods. Using the expected capital gains approach, DeLong and Magin (2009) p. estimate 2.5 per cent, Fama and French (2002) between 2.5 and 4.3 for 1951-2000, Jagannathan, McGrattan and Scherbina (2001) estimate from 3.4 per cent to 5.9 per cent from 1926 to 1999. Siegal (2004) estimates the future risk premium to be 3 per cent. McCulloch and Leonova (2005) provide a nice survey of the evidence. They conclude the equity risk premium is in the range 3 to 5 per cent.

24 See the Reserve Bank share market data (which includes dividend yield) at http://www.rba.gov.au/Statistics/Bulletin/F07.pdf and Oliver (2009).

grow explosively). Historically, most variation in price/dividend ratios comes from varying expected returns, reflecting a changing risk premium.[25]

That is, the risk premium changes over time and returns are predictable. If price/dividend ratio forecasts of dividend growth are not sufficient to explain the variance of price/dividend ratios, then the price/dividend forecasts of returns must fill the gap. Return predictability turns out to be exactly the same as excess price volatility (relative to constant discount rate present-value models). If share prices are 'too high', so that the price/dividend ratio is high, subsequent returns will be low as prices adjust to their correct levels. A high price/dividend ratio implies either investors expect dividends to rise in the future or expect returns to be low or expect prices to rise forever (a bubble).[26]

It is controversial whether the residual discount rate movement is variation in real investment opportunities or represents 'irrational' behaviour and market inefficiency. Volatility tests are only tests of discount rate models, and are not tests of market efficiency. In fact, some discount rate process can rationalise any forecastability of returns (or price volatility). The issue is whether macroeconomic events, which change investment opportunities, justify the implied time variation in discount rates. Those who think excess volatility is evidence of market inefficiency must believe the expected returns on assets vary too much over time.[27]

As the implied variation in discount rates gives rise to small utility costs, corresponds to reasonable measures of time varying investment opportunities, and appear related to business cycles, the evidence is at least consistent with efficient markets.

Reasons why the evidence suggests imperfections in the current discount rate models, rather than irrationality, include that the discount rate errors are small and are strongly suggestive of economic explanation. With a constant dividend growth g and discount rate r, $P/D = 1/(r–g)$. A small change in the discount rate has a large effect on price. For example, if $r–g$ changes from 5 per cent to 3 per cent, price rises by 66 per cent. This overstates the effect of changes in expected returns on price because expected returns are not constant, but they are stable enough for small market imperfections in expected returns to translate into a large effect on price. Large differences between the prices of similar assets are consistent with small differences in expected returns.

[25] Cochrane (2005, pp. 396-97).

[26] See Cochrane (2005, pp. 396-409).

[27] See Cochrane (1991).

Return predictability may suggest ‘second-order corrections to a basically correct efficient-markets view, such as better measures of fundamental movements in discount rates or incorporation of frictions like taxes, transactions costs and market microstructure’[28] rather than irrational investor behaviour. The financial economics literature interprets return predictability as efficiently reflecting time variation in risk premiums, induced by time variation in real investment opportunities available in the real economy. Time variation in discount rates appear to be related to business cycles. For example, risk premiums rise during recessions.[29]

G.3 The marginal return to capital

The share market returns set out in table 2.2 are before-personal income tax (except for imputation credits) and after company tax and measure the return to levered equity. Working out the before-tax return to investment from share market returns is difficult and requires accounting for the effects of company tax and gearing levels.[30] Even then it is only the return to listed companies.

The marginal return to private capital is a weighted average of the returns in the various sectors, with the weights depending on the responsiveness of investment in each sector to changes in the interest rate.[31] The best 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.[32] The average return to investors over a long time period should provide a reasonable estimate of the return to capital.

Hall and Scobie (2005) find the annual return to capital from 1978 to 2002 consistently exceeds 11 per cent real.[33] Mohun finds the real return to capital fell from a peak of 16.2 per cent in 1969 to a trough of 6.5 per cent in 1983, and then recovered to reach 13.4 per cent by 2001.[34] Makin (2006) finds the real net

---

28 Cochrane (1991, pp. 463-64).

29 Cochrane (undated; 2005, pp. 410, 442; 1991, pp. 478-82). He suggests other reasons for rejecting the irrationality view.

30 See New Zealand Treasury (2008) for an example of the adjustments that need to be made.

31 See Harberger (1969) for details.

32 As suggested by Harberger (1987, pp. 177-78; 1992, p. 29) and by Poterba (1997, p. 10–1).

33 See Hall and Scobie (2005, figure 14, p. 18). It is operating surplus divided by the net capital stock. It is not clear whether the operating surplus is net or gross, so they could be gross returns.

34 These are net returns, GDP less wages, depreciation, and imputations divided by the capital stock. See Mohun (2004, pp. 104-09).

marginal return to capital to be in the range 8.2 per cent to 9.2 per cent from 1995-96 to 2004-05, averaging 8.8 per cent with a standard deviation of 0.3 per cent.[35] In 1997 the ABS produced estimates of the net return to capital from 1983-84 to 1995-96. They range from 15.3 to 18.5 per cent, averaging 16.8 per cent (with a standard deviation of 1.1 per cent).[36]

Some studies of international returns to capital estimate the returns in Australia. Pyo and Nam (1999) find the real return to capital in Australia averaged 10.2 per cent from 1970 to 1995. Caselli and Feyrer (2007) estimate the return in Australia in 1996. Their estimates vary from 7 to 13 per cent real (they prefer 8 per cent).[37]

This is consistent with the US evidence. The real pre–tax return on U.S. non-financial corporations over the last four decades has averaged about 6.6 per cent per year, while the returns to U.S. non-financial industries over the 1997–2006 period averaged 8.9 per cent per year.[38]

Poterba (1997) finds over the 1959 to 1996 period, the real pre-tax return on capital in the non-financial corporate sector averaged 8.5 per cent. He estimates corporate capital income (net of depreciation) divided by a measure of the corporate capital stock. They follow a similar pattern to Dolman's estimates (falling from the mid-1960s to the 1980s, before recovering in the 1990s).

He also looks at the return to capital (business assets) in G7 nations from 1965 to 1996. They follow the same pattern and the average for all G7 was 14.3 per cent, slightly less than the US average of 15.4 per cent (these are gross figures).

Summers (1990) finds the national accounts estimate of the average return in the non-financial corporate sector from 1952-1987 to be 9.6 per cent, and 7.4 per cent for the private sector as a whole. His measure of the capital stock includes land and inventories. After accounting for various biases, Summers concludes the return to capital is in the order of 7 per cent. He too finds a fall from the mid-1960s to the 1980s.[39]

---

[35] See Makin (2006, table 1, p. 230).

[36] See Australian Bureau of Statistics (1997, table 26, p. 55). The estimates are the ratio of net operating surplus to the net capital stock. The capital stock is non-dwelling construction and equipment. It is not clear whether these are real or nominal rates of return.

[37] All these estimates appear to be gross returns (before depreciation), but do not include any capital gains from changes in capital good prices. For example, without any relative price changes, the value of the capital stock would be expected to increase at the rate of inflation.

[38] Nordhaus (2007, pp. 689-90).

[39] See Summers (1990, pp. 118-120).

Mulligan (2003) estimates the capital income net of depreciation per dollar of capital in place at the beginning of the year. Capital is fixed assets valued at current cost, and is aggregated for private residential and non-residential sectors. The capital measured is employed domestically (including foreign-owned capital), includes owner-occupied housing, excludes consumer durables, and excludes the government sector. The capital income measure is before all taxes and ignores capital gains (other than the effects of depreciation). From 1947-96 the average real before-tax rental rate, net of depreciation averaged 9.5 per cent, varying from 7.3 to 12.1 per cent with a standard deviation of 1.1 per cent.[40] In Mulligan (2002) he estimates the pre-direct-tax capital rental measure from 1900 to 1996. It ranges from 3.7 to 12.5 percent. Mulligan then establishes that his estimated after-tax marginal return to capital is a much better predictor of consumption growth than interest rates on financial assets.

All the national accounts based estimates have a number of biases. For example, income of unincorporated enterprises needs to be divided into income accruing to labour and income accruing to capital.

Estimates of capital income often include income accruing to land, whereas the estimates of the capital stock (usually derived using the perpetual inventory method from investment flows) usually exclude land. For example, Caselli and Feyrer (2007) find that accounting for land and other natural resources reduces their estimated marginal product of capital in Australia from 13 per cent to 7 or 8 per cent.[41]

On the other hand, estimates of capital income usually exclude indirect taxes and property taxes. Harberger suggests the social yield on capital investment is the private return (profits, interest and rent) and the taxes paid on it – including:

> property taxes on the land, building and other assets used in the process of production; capital's share of value added taxes paid on the value added of the entity (enterprise, industry, sector, total private sector) in question; capital's share of other indirect taxes paid as consequence of transforming raw materials and purchased components into the final products of the enterprise.[42]

Poterba (1997) adds in property taxes to calculate the before-tax return. The property tax adjustment raises returns a full percentage point.[43] The Mulligan (2003) estimates also add property taxes.

---

40 Mulligan (2003, p. 26).

41 See Caselli and Feyrer (2007, table II, p. 549).

42 See Harberger (1992, appendix 1).

43 See Poterba (1997, p. 8–9).

The rate of return is calculated on the basis of those fixed assets that are measured in the national accounts, but if there are other, unmeasured assets that provide capital services, the measured rate of return on capital overstates the true return. For example, there are good reasons to argue that at least part of capital income comes from intangible assets. On the other hand, investment in intangible capital is currently measured as expenditure on intermediate inputs, reducing gross operating surplus. The net bias in the measured rate of return is not clear, and depends on the size of the measured return relative to the accumulation rate of intangible capital.[44] Further, although intangible assets have become more important over time, there is no upwards trend in the measured return to capital. Rates of return over the past decade are comparable with the returns in the 1960s.

An alternative way to proceed is to estimate the before-tax cost of capital to companies, which would equal the marginal return to investment in equilibrium. The cost of capital is the minimum rate of return that an investment project must earn in order to cover its funding costs and any tax liabilities. Although the supply of funding depends on expected returns, the cost of capital is usually estimated as a weighted average of realised debt and equity returns. Working out the before-tax return to investment from share market and debt returns is difficult. Gearing levels, dividend policy and effective tax rates are crucial, yet are difficult to measure, change over time, may change with changes in expected inflation and may interact with each other. Real world complications, such as incomplete loss offset and expected inflation varying over time, make the effects of taxes difficult to model and hard to measure.

Observed share-market returns (presented in section G.3) are after company and property taxes. Using financial market data to estimate the before-tax return is difficult. It requires accounting for the complicated effects of company tax, which depends on its interaction with inflation (e.g. through historical cost depreciation provisions and the deductibility of nominal interest payments), with personal taxes (which was affected by the introduction of imputation) and the effects of international capital flows (and their tax treatment).[45] Even then we have only measured the return to listed companies. Further, rates of return in financial markets are extremely volatile and subject to swings in expectations. Further, mistaken

[44] Let the earnings on capital be E, investment in intangibles X, the measured capital stock K and the intangible capital stock I. The measured rate of return on capital is $(E - X)/K$, whereas the true return is $E/(K + I)$. The measured return overstates the true return when:

$(E - X)/K > E/(K+I)$ or $(E - X)(K+I) > KE$ or $EI - XK - XI > 0$ which can be written as

$(E - X)/K > X/I$ *or as* $E/(K + I) > X/I$. That is, only if the rate of accumulation of intangible capital is less than the return on capital.

[45] See, for example, Cohen et al. (1999) on the effects of inflation on the cost of capital.

expectations can give negative realised real returns for long time periods. High volatility makes it difficult to pin down an estimate. For example, the standard deviation of realised annual share market returns is around 20 per cent, giving a 95 per cent confidence interval of plus or minus 4 percentage points even with 100 years of data. The standard deviation of national accounts based measures is usually around 1 per cent. National accounts measures are broader, steadier and more reliable. Moreover, Mulligan (2002) finds that the return on capital is correlated with growth but financial market returns are not.[46]

But estimates of the cost of capital in Australia have been made, and they are consistent with the returns estimated from the national accounts. Lally (2000) found the consensus cost of capital for Australia to be 10 per cent real. Dews et al (1992) estimate the annual real pre-tax cost of capital from 1963-64 to 1983-84. The estimates fluctuate widely over time and also vary greatly depending on the weighting assumed for debt and equity (whether the weights are average or marginal shares or use stock or flow data). But the preferred measure is above 10 percent except in the mid-1970s.[47] The authors also survey previous studies, which have produced a wide range of estimates using various methodologies. The average cost of capital estimated varied from 18 per cent (Brunker 1984), 4 per cent from 1962-1986 and 8 per cent (1975-1986) (Carmichael and Stebbing 1981; Dews 1988; 1989), 10 per cent (Australian Manufacturing Council 1990), 11 per cent (long-run ex post) (Johnston, Parkinson and McCray 1984), 12 per cent ex ante using data from 1976-77 to 1985-86 (Department of Finance 1987), 6 per cent (ex-post) from 1967-68 to 1982-83 (Swan 1988).[48] The authors warn 'This paper demonstrates that a wide range of outcomes can result from often arbitrary assumptions used in constructing measures of the cost of funds.'[49] Other estimates range from 4.5 per cent (in 1999) to 10 per cent nominal (in 1998) to 10 to 11 per cent real.[50]

Rough calculations confirm that the cost of capital estimated from share market returns are consistent with the national accounts estimates.[51] Expressing the rates of return as continuously compounded, the before-tax real rate of return $i = \rho - \pi$ where ρ is the nominal return and π the expected rate of inflation. If the company tax rate is τ then $\rho = [(i_f + \pi)(1 - \tau) + \beta e]/(1 - \tau_e)$ where i_f is the risk free real rate of return, β is

[46] Harberger (1998, pp. 24-26) finds a positive relationship between the return to capital and rates of total factor productivity improvement.

[47] See Dews et al. (1992, p. 180).

[48] See Dews et al. (1992, appendix 1, pp. 24-26) for details.

[49] Dews et al. (1992, p. i).

[50] Bureau of Transport Economics (1999, p. 67), Commonwealth Competitive Neutrality Complaints Office (1998, pp. 9–11), and Department of Finance (1991, p. 57).

[51] These calculations are based in New Zealand Treasury (2008).

the asset beta, e is the equity risk premium and τ_e is the effective rate of company tax. Nominal returns are taxed at the rate τ_e and the interest paid on bonds is tax deductible at the statutory rate. The effective rate of taxation may differ from the statutory rate. For example, because depreciation provisions that differ from economic depreciation. The effective rate would be higher than the statutory rate when only the historical or original cost of a capital asset may be written off even if the cost of replacing the asset is rising over time. Provisions such as accelerated depreciation would lower it. Further, inability to deduct losses would raise the effective rate further. For simplicity, I assume that $\tau = \tau_e$ and so $i = i_f + \beta e/(1-\tau)$. Table G.3 shows how the before-tax real rate of return varies with the market risk premium and risk free rate of return, assuming a company tax rate of 30 per cent and an asset beta of 0.67. The market beta is one, which gives an average asset beta of 0.67 if the average level of gearing is 33 per cent.[52]

Table G.3 **The before-tax real rate of return**

With a company tax rate of 30 per cent and asset beta of 0.67

	Risk free rate of return (r_f)	
Market risk premium (e)	*3 per cent*	*4 per cent*
3 per cent	5.9%	6.9%
4 per cent	6.8%	7.8%
6 per cent	8.7%	9.7%

The results are close to the national account measures of the marginal return to capital, which average 8.9 per cent since 1965. The numbers are indicative only, and are presented just to demonstrate that a before-tax real return of around 9 per cent is reasonable.

Further, Poterba (1997) finds the weighted average of equity and bond returns to be (on average) close to the average after-corporate tax national accounts return (for the US from 1959 to 1996), suggesting the before-tax national accounts returns are accurate. Summers (1990) also finds the two types of estimates to be consistent.

All the variables that determine the cost of capital vary over time. For example, the company tax rate was greater than 30 per cent for most of that time and imputation was not introduced until 1986. If the realised return to capital varies with the cost of

[52] The gearing ratio is $g = D/(D+E)$ where D is the amount of debt and E the value of equity. But $g/(1-g) = D/E$ and so a gearing ratio of 33 per cent corresponds to a debt-equity ratio of 0.5. In 2007, the debt-equity ratio for the Australian share market was 0.37 (Jones (2008, p. 208), but the value of equity has fallen substantially since then. The relationship between market beta β_M and asset beta β_a is $\beta_M = \beta_a(1+D/E)$. See New Zealand Treasury (2008, p. 4-5). The lower the debt-equity ratio, the higher the asset beta (the closer to 1).

capital, then the historical average depends on the historical pattern of the risk free rate, gearing ratio, market risk premium and company tax rate.

An alternative approach would be to forecast the likely future cost of capital by estimating the future values of the variables that determine it. But that is difficult, all the more so if we took account of real world complications such as the effect of capital gains taxes. How to do so is often controversial.[53] A pragmatic approach is to use historical averages over long time periods as an indicator of the future. A disadvantage is that it encourages assuming that the cost of capital is constant. Explicitly estimating the future cost of capital (if possible) can account for predictable changes in the risk premium. In fact, if credit market conditions are temporarily unusual (or likely to change), it is desirable to explicitly allow the discount to rate to change over time, towards its normal (or predicted) level.

G.4 The effect of tax on the risk premium

The risk premium is a substantial portion of the market return. It would appear that taxation of risky returns would mean a smaller after-tax premium. For example, if the before-tax market rate of return was 8 per cent and the risk-free rate 4 per cent, the risk premium is 4 per cent. With a 50 per cent tax rate, the after-tax market rate of return is 4 per cent and the after-tax risk free rate is 2 per cent, the after-tax risk premium is 2 per cent. The tax reduces the risk premium proportionately.

But with proportional taxation and full loss offset, the after–tax price of per unit of risk is the same as the before-tax price. Although part of the risk premium is taxed, the variance of the asset is reduced proportionately. These two offset, leaving the ratio of reward to risk (the Sharpe ratio) unchanged.

A tax on the whole real return is equivalent to taxing only the riskless part of the return for market traded assets. The appropriate risk premium is unaffected. The pre–tax risk premium should be used to judge the risk coming from the government project, it is the social price of risk.

To see this, consider a world where only the riskless return, r is taxed at rate t. An investor invests X in a risky asset that pays a gross return R_i in state of the world i. After the risk is realised and tax is paid on the riskless return, a tax payment of trX, the gross after-tax return in state i is $R_i - trX$.

[53] For example, the standard approach is to account for the burden capital gains taxes impose on the cost of equity with the accrual equivalent tax rate of capital gains (as done, for example, by Cohen etc al. (1999, p. 205). Jones (2009) argues this is incorrect because personal taxes are only paid when capital gains are realised and not when they accrue.

Now assume instead the tax (at rate *t*) is imposed on all investment returns, including the risky component, and there is full deductibility of losses. If the consumer increases his investment in the risky asset from *X* to *X/(1–t)*, borrowing the additional funds at the riskless rate *r*, then in state *i* the investor will receive (after paying tax and repaying the loan):

$$\frac{R_i}{1-t} - \frac{tX}{1-t}(1+r) - t\left[\frac{R_i - X}{1-t} - \frac{tX}{1-t}r\right] = (1-t)\frac{R_i}{1-t} - \frac{tX}{1-t}(1+r-1-tr)$$
$$= R_i - \frac{tX}{(1-t)}r(1-t) = R_i - trX$$

The first term is the gross return on the investment of *X/(1–t),* the second term is the repayment of the loan *tX/(1–t)* with interest (or the forgone income from reducing investment in the risk free asset), and the third term is the tax owed (including the deduction allowed for interest payments on the loan).

The investor receives the same after-tax return when all returns are taxed as when only the riskless return is taxed for every possible outcome (for every state of nature). The same result would apply if the investor funded the increased investment in the risky asset by reduced investment in the risk free asset rather than by borrowing. It would apply if nominal interest payments were taxed.

Taxes on the risky component of returns have no effect on individuals, who simply gross up their investments in risky assets to offset the effects of the tax. Portfolio adjustments mean the taxation system only taxes the riskless part of the return and not the risk premium. A tax on all returns is equivalent to a tax only on the riskless return.

Taxing risky returns does not, in the end, provide insurance since taxpayers' portfolio adjustments unwind this effect of taxation. Taxpayers wanting less risk could have invested in less risky portfolios in the absence of taxation; if subject to a system that automatically absorbs a portion of any risk they take, they simply gross up the extent of their pre-tax risk exposure in an offsetting manner.[54]

Figure 6 illustrates the effects of taxation with the standard capital market diagram, which shows the portfolio frontier and capital market line. The slope of the capital market line, $\frac{r^M - r^f}{\sigma_M}$, is the Sharpe ratio. A tax on the risk free part of the return shifts the portfolio frontier and capital market line down in a parallel fashion. A tax on all returns shrinks in the portfolio frontier (reducing variance and average returns

[54] This proof and explanation is from Kaplow (2008, pp. 239-241).

by a proportion *t*) and gives the same capital market line, and so the same opportunities.

This is a first approximation, and is true for proportional taxation of the real return with full loss offset when the investor can borrow and lend at the risk-free rate. In reality marginal rate progression, certain restrictions on the use of tax losses, limits on the ability to gross up certain types of investments and borrowing constraints will affect the risk premium. Borrowing constraints are not obviously important for those holding most assets, assets obtained when grossing up might themselves be used as collateral and constrained individuals can purchase assets with greater risk rather than personally borrowing to buy more assets of a given risk. But the lack of full loss offset would raise the effective rate of taxation on risky assets. Grossing up is difficult for untraded and illiquid assets. Forced diversification through the tax system would be inefficient if the imperfection that caused the illiquidity was moral hazard. If may be efficient if adverse selection restricted liquidity.[55]

Figure G.1 **Taxing all returns is equivalent to taxing the risk free return**

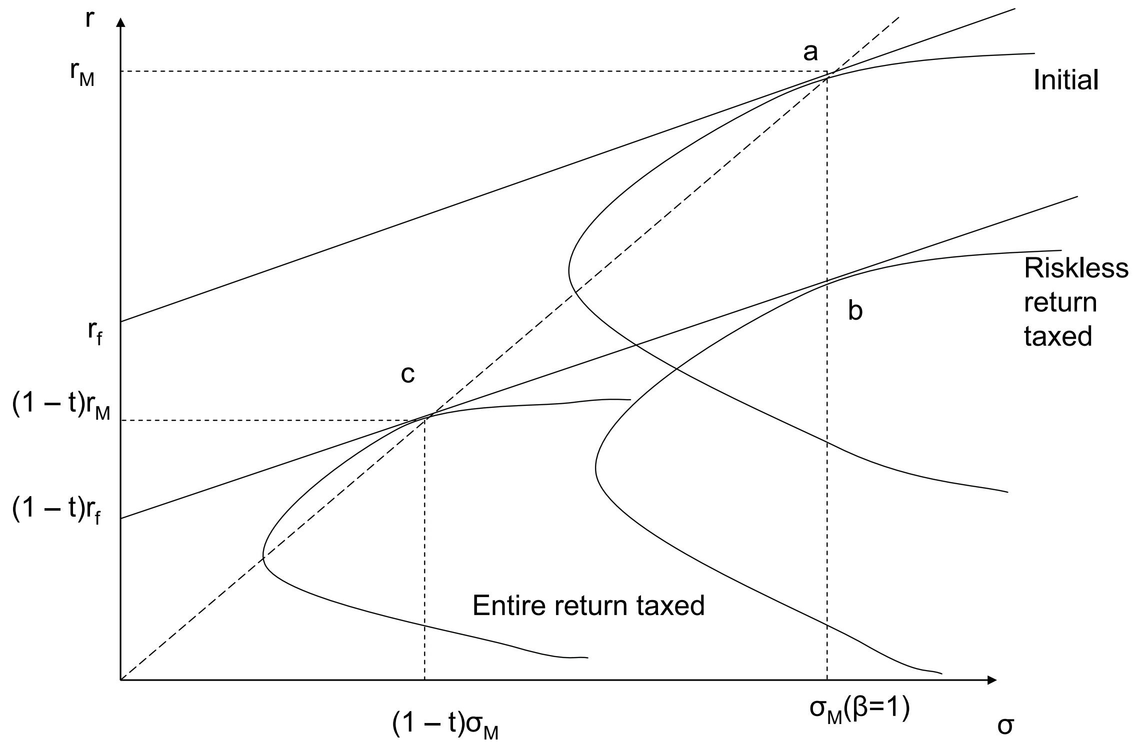

55 Kaplow (2008, p. 241, footnote 35).

H Asset pricing

'The centrality of the trade–off between risk and expected return infuses all investment decisions'. Bernstein (2007, p. xvii).

H.1 The consumption based model

The starting point for any theory of asset pricing must be the investor's first–order condition for savings and portfolio choice – the basic consumption based model.[1] The consumer has a utility function defined over current and future consumption:

$$U(c_j, c_{j+1}) = u(c_j) + \gamma E_j[u(c_{j+1})]$$

An investor must decide how much to save and how much to consume. Let an asset have an uncertain payoff of x_{j+1} in time $j+1$. The investor can buy and sell as much of the payoff x_{j+1} as he wishes, at price p_j. The marginal utility loss of consuming a little less today and buying a little more of the asset should equal the marginal utility gain of consuming a little more of the asset's payoff in the future. That is, the consumer's first order condition is:

$$p_j = E(m_{j+1} x_{j+1})$$

$$m_{j+1} = \gamma \frac{u'(c_{j+1})}{u'(c_j)}$$

where m_{j+1} is the stochastic discount factor or marginal rate of substitution and E is the expected value, based on information at time j. If the equation does not hold, then the investor would buy more or less of the asset (and change current and future consumption) until it does hold.

The equation can be rewritten as

$$p_j = \frac{E(x_{j+1})}{R^f_{j+1}} + \operatorname{cov}(m_{j+1}, x_{j+1}) = \frac{E(x_{j+1})}{R^f_{j+1}} + \frac{\operatorname{cov}\left[\gamma u'(c_{j+1}), x_{j+1}\right]}{u'(c_j)}$$

1 This section is based on Cochrane (2005, pp. 3-17, 35-37).

where the risk free gross return is $1+r^f_{j+1} = R^f_{j+1} = 1/E(m_{j+1})$. That is, expected returns are proportional to the covariance of returns with the discount factor. The first term on the right hand side is the asset's price is a risk neutral world. The second term is a risk adjustment. An asset whose payoff covaries negatively with the discount factor has its price decreased. Marginal utility declines as consumption increases, so the asset's price decreases if its payoff covaries positively with consumption. Such an asset makes the consumption stream more volatile. If the asset covaries negatively with consumption, it has a high price.

That is, it is the covariance of a payoff with consumption rather than its variance that determines its riskiness to an investor. What matters to the investor is the volatility of his consumption.

If the asset price is set equal to 1, then the payoff is a gross return – say R^k for asset k. Then the equation can be written as:

$$E(R^k) - R^f = -R^f \operatorname{cov}\left(m, R^k\right) = -\frac{\operatorname{cov}[u(c_{j+1}), R^k_{j+1}]}{E[u'(c_{j+1})]}$$

That is, all assets have an expected return equal to the risk free rate plus a risk adjustment. Assets whose returns covary positively with consumption make consumption more volatile and must offer a higher expected return to induce investors to hold them. Idiosyncratic risk, uncorrelated with the discount factor, does not affect asset prices or generate a higher return. Only systematic, or market, risk generates a risk correction in the price, or a return different from the riskless return (a risk premium in the discount rate).

The basic pricing equation can be written as:

$$E(R^k) = R^f + \left(\frac{\operatorname{cov}(R^k, m)}{\operatorname{var}(m)}\right)\left(-\frac{\operatorname{var}(m)}{E(m)}\right) = R^f + \beta_{k,m}\lambda_m$$

where $\beta_{j,m}$ is the regression coefficient of the return R^k on m (which varies from asset to asset). λ_m is often interpreted as the price of risk and $\beta_{k,m}$ as the quantity of risk in each asset. The price of risk depends on the volatility of the discount factor. This is a beta pricing model, where each expected return is proportional to $\beta_{k,m}$.

So far the only assumption made is that investors can make small marginal investments in the asset. The basic pricing equation does not assume complete markets, a representative investor, normally distributed asset returns, two periods, quadratic utility, no labour income or market equilibrium. The equation applies to each individual, for each asset to which he has access and for any two periods of a multi–period model. Further, the assumption that the consumer maximises time and date separable expected utility could be dropped. The basic consumption pricing equation would hold for any general non-separable utility function. The relation between the discount factor and real variables would, however, be more complicated. It would be a function of many variables, not just the discount on future utility and present and future consumption.[2]

H.2 The Capital Asset Pricing Model

Further assumptions can be added to derive other asset pricing models, such as the widely used Capital Asset Pricing Model (CAPM). The CAPM is an example of a factor pricing model, where the discount factor m_{j+1} is a linear function of factors. The factors are plausible proxies for marginal utility. That is, factor models are special cases of the consumption based model – special cases that allow marginal utility growth to be proxied from some other variable.[3]

In the CAPM the market portfolio is the only source of the consumer's income so that $m_{j+1} = a + bR^{W}_{j+1}$, where R^W is the return on a claim to total wealth, often proxied by a broad share market portfolio. a and b are parameters. Future consumption is funded solely from returns to portfolios of securities. For example, there is no labour income.

The CAPM's simplifying assumptions make asset prices linear in the single risk factor, the market return. For example, that security returns are jointly normally distributed. Alternatively that investors have quadratic utility. In a multi–period model the assumptions that returns are independent and identically distributed (no shifts in the investment opportunity set), and the interest rate and relative commodity prices are constant, need to be added. If consumption risk changes over time, the CAPM fails to hold.[4]

Further, investors have homogenous expectations and there are no transactions costs. A risk free security exists and there are no borrowing constraints.[5] Alternatively, in the absence of a riskless asset, that there are no short–selling constraints so that the investor can form a zero beta security.

2 Cochrane (2005, pp. 35-36).

3 Cohrane (2005, pp. 149-52).

4 See Cochrane (2005, pp. 152-55); Jones (2008, p. 129).

5 See Jones (2008, pp. 123-29). These assumptions mean all consumers face the same capital market line, which relates portfolio variance and return.

The CAPM prices assets relative to a market portfolio. The CAPM uses the return on the market portfolio as the single factor. It can be expressed as a beta pricing model:

$$E(R^k) = R^f + \beta_{kM}[E(R^M) - R^f] = \beta_{kM}E(R^M) + (1-\beta_{kM})R^f$$

where $\beta_{kM} = \sigma_{kM} / \sigma_M^2 = COV(R^M, R^k) / VAR(R^M) = \rho_{kM}\sigma_k / \sigma_M$ measures the amount of market risk in security *j,* the risk in the asset that cannot be diversified away and $\rho_{kM} = \sigma_{kM} / \sigma_k \sigma_M$ is the coefficient of correlation of the return on security *k* and the market return, R_M. This CAPM asset pricing equation is known as the security market line.

In the CAPM, all investors measure and price risk identically. Asset prices depend solely on the expected cash payoff, the quantity of market risk provided by the asset, the risk free rate and the price of risk — which are market–determined variables. Individuals who perceive the same distribution of payoffs for a risky asset will price it exactly the same way, independent of their individual utility functions and attitudes to risk.[6] The CAPM approach does not require everyone has the same utility function — just common expectations

In the CAPM, what matters is not the variability of an asset's income flows themselves, but the contribution they make to overall portfolio risk (and, ultimately, to consumption risk). For example, holding a diversified portfolio can eliminate diversifiable risk, unexpected losses on some assets are offset by gains on others.

Market risk, by contrast, is associated with economy–wide increases or decreases in asset values (resulting in changes in aggregate consumption), so it cannot be eliminated through portfolio diversification. Market risk arises from the volatility of the economy and from associated changes in the value of aggregate wealth. Because those changes create undesired uncertainty, they are costly to investors and are reflected in market prices.

Only non-diversifiable, or market, risk is costly. The level of market risk associated with an asset determines the return demanded by the market above the risk free return, that is the risk premium. Assets that only have diversifiable risk earn the riskless rate of return (a zero risk premium). Assets which are positively correlated with market return (contain market risk) must earn a higher rate of return. The higher the correlation with market return, the higher the required rate of return. Assets that are negatively correlated with the market reduce portfolio risk (they provide insurance) and earn less than the riskless return.

6 Copeland and Weston (1983, p. 196).

The risk premium is about how much market risk is contained in the stream of payments, it is a not about downside risk, which should be accounted for in the expected value estimates.

H.3 Multifactor models

The CAPM has the advantage of convenience and 'seductive simplicity'.[7] Although the CAPM is widely used and has proved to be a powerful tool for assessing risks, it fails many empirical tests: much variation in expected returns is unrelated to market beta. That is not surprising, as the CAPM is based on many unrealistic assumptions, such as that investors only care about the mean and variance of one–period portfolio returns, which miss important dimensions of risk. For example, investors also care about how their portfolio return covaries with labour income and future investment opportunities (such as changing expectations of future returns) — which affect marginal utility. Market beta, then, is not a complete description of an asset's risk and so does not explain all differences in expected asset returns.[8] There are assets, portfolios and strategies whose average returns cannot be explained by their market beta. Additional sources of priced risk beyond movements in the market portfolio are need to explain why some average returns are higher than others.[9]

An extension has been multifactor models, such as the intertemporal capital asset pricing model (ICAPM), where investors care how their wealth varies with future state variables that affect their opportunities to invest or consume the asset's payoff, such as labour income, the prices of consumption goods, investment opportunities and financial distress factors. Multifactor models add additional beta, along with a market beta, to explain expected returns.

Multifactor models measure risk by covariances with several common factors — usually based on empirical fit rather than theory (which gives rise to the data–snooping critique), arguing that the factors reflect unidentified state variables that produce undiversifiable risks (covariances) in returns that are not captured by the market return and are priced separately from market betas.[10]

Multifactor models dominate the empirical description, performance attribution and explanation of average returns, but give little guidance about choosing factors or answering what forces determine factor risk prices. What are the sources of

7 Fama and French (2004, p. 44).

8 See Fama and French (2004).

9 Cochrane (2005, pp. 390-91).

10 See for example, Fama and French (2004).

macroeconomic risk that drive asset prices? It seems a recession or financial distress factor lies behind many asset prices. People forgo large return premiums to avoid assets that fall in recessions. For example, it takes a risk higher premium to get people to hold stocks at the bottom of a recession. Asset returns depend on standard market beta and four liquidity premiums: expected returns are higher for stocks that are illiquid on average, for stocks whose illiquidity gets worse when the market goes down, for stocks whose price goes down when the market gets more illiquid, and for stocks which become more illiquid when the market gets more illiquid.[11]

H.4 Asset return puzzles

The equity premium puzzle

The equity premium puzzle is that the standard consumption-based model of asset pricing can only explain the observed equity premium by assuming unrealistic levels of risk aversion.[12]

The Sharpe ratio, a measure of the ratio of reward to risk in asset markets, is:

$\frac{E(R^k)-R^f}{\sigma_k}$ for asset k (it is proportional to the price of risk derived in H.1 and is sometimes described as the market price per unit of risk). It is invariant to leveraging.

From H.1, the basic asset pricing equation in terms of an asset's return is:

$$1 = E(mR^k) = E(m)E(R^k) + \rho_{m,R^k}\sigma\sigma_k$$

Where $\sigma = \sqrt{\text{var}(m)}$ is the standard deviation of the stochastic discount factor, σ_k is the standard deviation of R^k and ρ_{m,R^k} is a correlation co-efficient. This can be written as: $E(R^k) = R^f - \rho_{m,R^k}\frac{\sigma}{E(m)}\sigma_k$. As the absolute value of correlation co-efficients cannot be greater than 1, this gives: $\left|\frac{E(R^k)-R^f}{\sigma_k}\right| \le \frac{\sigma}{E(m)}$

[11] Cochrane (2004, p. 9).

[12] This section summarises Cochrane (2005, pp. 17-22, 455-59).

Now for power utility, $m = \gamma\left(\frac{c_{t+1}}{c_t}\right)^{-\eta}$ (see appendix F). If consumption growth is normally distributed, then we get the Hansen-Jagannathan bound:

$\left|\frac{E(R^k) - R^f}{\sigma_k}\right| \leq \frac{\sigma}{E(m)} = R^f\sigma \approx \eta\sigma_c$ where σ_c is the standard deviation of consumption growth.

In the United States, historic measures of the equity premium for the broad share market are around 6 to 8 per cent and the standard deviation of the market return around 16 to 20 per cent, giving a Sharpe ratio of 0.3 to 0.5.[13]

Complete markets mean that individuals can diversify any idiosyncratic risk in consumption, so that only fluctuations in aggregate consumption are relevant for consumers. The standard deviation of aggregate consumption, σ_c, is about 1 to 2 percent. The Hansen-Jagannathan bound implies that the co-efficient of relative risk aversion (η) must be at least 15 (for σ_c = 2 and a Sharpe ration of 0.3) and could be as high as 50 (for σ_c = 1 and a Sharpe ration of 0.5), an incredibly high level of risk aversion.

Many believe that that the historically high risk premium reflects luck and the expected future risk premium is in the range 3 to 5 per cent, giving a Sharpe ratio of 0.15 to 0.3 (the ratio in Australia for the past 38 years was 0.2).[14] That reduces the required co-efficient of risk aversion to as low as 7.5. But even that is difficult to explain. Empirical estimates of η range from 0.2 to 4 (with most between 1 and 2).[15] A η between 1 and 2 and a standard deviation of consumption growth of between 1 and 2 per cent imply a Sharpe ratio between 0.01 and 0.04 and, at most, a risk premium of 0.008 (less than 1 per cent). Observed levels of these variables are greater by a factor of 10.

Ways of generalising the calculation make matters worse. Even broad share market indexes are only a portion of the true market portfolio. But that means they likely lie inside the mean-variance frontier and so their Sharpe ratio understates the true Sharpe ratio. The above calculations assume that asset returns and consumption growth are perfectly correlated. In the United States aggregate consumption has a 0.2 correlation with the market return, raising the required level of risk aversion fivefold (to 75–250).[16]

---

[13] Cochrane (2005, p. 456); Mehra and Prescott (2003, p. 908).

[14] See appendix G.

[15] Zhuang et al. (2007), Weitzman (2007, p. 709).

[16] Cochrane (2005, p. 21).

The risk free rate puzzle

If the level of risk aversion was really as high as the simple consumption based model suggests, then a risk free rate puzzle would arise. For example, individuals should strongly want to smooth consumption over time, because consumption shortfalls have a much larger effect on utility than consumption increases. As the economy is growing, individuals should be keen to borrow to increase current consumption. Everyone trying to borrow should result in high interest rates, but the observed riskless rate is low. In a certain world, the riskless rate is $r^f = \theta + \eta g$.[17] g (average consumption growth) is around 2 percent, and so an η equal to 15 or 50 implies a riskless interest rate of over 30 or 100 percent, so long as θ (the pure rate of time preference) is not negative (that is, so long as there is a preference for current over future utility). These rates are vastly above observed rates.

The Ramsey formula for the risk free rate can be extended to account for uncertainty. If consumption growth is independently and identically distributed normal with a known mean g and variance σ_c^2, then the formula for the risk free interest rate is:[18]

$$r^f = \theta + \eta g - \frac{\eta^2 \sigma_c^2}{2}$$

In a world of uncertainty, the real interest rate is also negatively related to the degree of risk aversion (for example, there is a precautionary motive for saving, which decreases interest rates). For example, extremely high levels of risk aversion can produce low risk free rates, even puzzlingly low. For example, if we substitute g = 2 percent, θ = 2 percent and η = 100, 200 into the formula for the risk free rate under uncertainty, we get a risk free rate of 2 percent and – 398 percent.

Although high levels of risk aversion can (almost by co-incidence) give reasonable levels of the risk free rate, small changes in parameters would result in dramatic changes in the risk free rate. That is why such high levels of risk aversion are inconsistent with real world observations. A high η means that consumers are unwilling to substitute consumption over time, so it would take large variations in the interest rate to get them to adapt to small variations in consumption growth. For example, a one percentage point increase in consumption growth would raise interest rates by η percentage points, totally implausible if η is high.[19] Observed real

[17] This formula is derived in appendix F.

[18] This model is set out in Weitzman (2007, p. 710). The equations are derived in Cochrane (2005, p. 12).

[19] See Cochrane (2005, pp. 457-458).

interest rates are not only low, but relatively stable (for example, annual realised real returns on 90 day Treasury Bills since 1971 had a standard deviation of 3.5 per cent).

Reasonable parameter values, such as θ = 0 to 2 per cent, η = 1 to 2, g = 1 to 2 per cent and σ_c = 1 to 2 per cent, give a negligible risk premium and a risk-free rate of anywhere from 1 to 6 per cent. In the US the return on equities has averaged around 7 per cent real and on short term Treasury bonds around 1 per cent, and the puzzle is the gap between them. The Ramsey equation does not resolve the puzzle of which rate to choose.

Explanations of the equity premium puzzle

A high equity premium must come from a lot of risk or a lot of risk aversion. The explanations of the equity premium puzzle fall into three main categories: preference modifications, individual consumption risk and fear of disasters.[20] No one explanation has been universally accepted, but each explains part of the premium, and together, perhaps all.

Preference modifications

Modifying preferences can make consumers extremely adverse to share market losses and yet also imply a low risk free rate. For example, in the standard approach the one parameter determines both risk aversion and time preference (η is the coefficient of risk aversion and the reciprocal of the elasticity of intertemporal substitution). Some authors have suggested utility functions (Epstein-Zin preferences) which break this link, and so allow high risk aversion without the risk free rate puzzle.

Standard preferences assume time separability and state independence. Changes to these assumptions that make consumers effectively highly risk averse include making utility depend on past own consumption (habits), past consumption of others or the current consumption of others, so that changes in aggregate growth have a large effect on well-being.[21]

Another approach are ad hoc behavioural explanations. For example, Siegal and Thaler (1997) put the equity premium puzzle down to loss aversion on the part of investors combined with annual portfolio evaluations. Others claim investors

[20] Survey of the equity premium puzzle include DeLong and Magin (2009); Mehra and Prescott (2003); Mehra (2006, 2003); Cochrane (1997) and Kotcherlakota (1996).

[21] See Jones (2008, pp. 147-48); Cochrane (2005, pp. 465-81; 1997, pp. 19-22).

irrationally over–estimate the riskiness of equity because of behavioural biases, such as focusing on the short-term.

These preferences can explain at least some of the equity premium and risk free rate puzzle. But they still require high levels of risk aversion to account for the equity premium.

To the extent that preference based explanations of the equity premium are correct, the resulting premium should be included in the social discount rate. It represents the real cost of risk to consumers based on their own evaluations.

Individual consumption risk

In the consumption based pricing model, consumers have the same (aggregate) consumption risk because they costly eliminate diversifiable risk, Incomplete markets, borrowing constraints and transactions costs can stop consumers from diversifying risks and make individual consumption risk higher than aggregate consumption risk.[22] For example, most people have a job that could be under threat in a recession. Uninsurable human capital risk can make people unwilling to hold the additional aggregate risk associated with shares. It can also encourage saving, driving down the risk free rate.

Other reasons for increased individual consumption risk include consumption commitments and concentrated share ownership.[23] Only a small minority of population that holds significant amounts of shares. Their consumption is more volatile and more highly correlated with the stock market return. Consumption commitments amplify risk aversion.

Uninsured individual risk is not an obvious solution to the equity premium puzzle. Individual consumption is not volatile enough to satisfy the Hansen-Jagannathan bound, and is less correlated with stock returns than aggregate consumption – and the more volatile it is, the less correlated. Idiosyncratic shocks uncorrelated with asset returns have no effect on asset pricing. Shocks correlated with asset returns are quickly traded away.[24] For example, by using borrowing to create ‘home made consumption smoothing’. Although borrowing constraints may limit consumers’ ability to smooth, in that case their discount rates would be high.

Consantinides and Duffie do present an idiosyncratic risk model, but prevent people self–insuring by making the shocks permanent. The variance of idiosyncratic risk

[22] See, for example, Mankiw (1986).

[23] See Mankiw and Zeldes (1991) and Chetty and Szeidl (2007).

[24] Cochrane (2005, p. 459-60; 1997, pp. 23-24).

rises when the market declines. The model shows how high consumption volatility is consistent with low aggregate volatility. But the distribution of consumption does not spread out as much as the model requires to explain all the equity premium. With any plausible characterization of the actual labour income process, the model requires high consumption volatility or high risk aversion to explain the observed risk premium.

To the extent that high individual risk causes the equity premium, it represents a real cost of risk to individuals and should be included in the social discount rate unless the government has some advantage that allows it to overcome the imperfections and diversify risk that individuals cannot. If a government project imposes risk on individuals, the market price of risk indicates the cost of that risk .

Fear of disasters

Weitzman (2007) points out that the standard finance models assume agents know the economy's stable structural parameters: such as the mean and variance of the (assumed) normal distribution of the economy's growth rate. But if instead these structural parameters are uncertain, evolving and have to be discovered, then the major asset return puzzles (including the equity premium puzzle) become explicable because of investor aversion to structural uncertainty and fear of a negative growth 'disaster'. The puzzles involve

> skittish investors nervously reacting with unsure expectations to deeper forces of shifting structure … to an outside observer it looks like inside investors are being rewarded by an inexplicably-high empirical asset return, while actually they are bearing the extra risk of rare disasters in the left tail of the distribution that happen not to have materialized within the limited sample.[25]

The result 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.[26] To complicate matters, the expected probability of disasters may change over time.

Weitzman concludes that makes sense to use the market premium rather than the one derived from observed variance in aggregate consumption

> there is some theoretical justification for treating the subjective variability of the future growth rate as if it were equivalent in welfare to the observed variability of a

[25] Weitzman (2007a, p. 1104).

[26] Weitzman (2007, p. 1104).

> comprehensive economy-wide index of equity-wealth returns. For this as-if- REE high-growth-variability interpretation, the simple standard model of asset pricing may have the potential to be a decent shortcut conceptualization of what is actually happening in a complicated ever-changing world where unforeseen bad events—scary, disruptive, and without precedent— may evolve at any future time.[27]

That is, the share market is much more volatile than consumption, but it more accurately reflects the true risks to consumption.

Weitzman's analysis dovetails with the Barro approach. Barro has a series of papers explaining the premium puzzle by a small probability of a disaster.[28] It has to be a disaster that reduces the returns to equities by more than bonds — such as the Great Depression.

Barro and Ursua (2008) find that the probability of a 'consumption disasters' (a greater than 10 percent decline in aggregate consumption), estimated from historical data, justifies the equity premium. They assume a known proportional reduction in consumption occurs with a known probability. Weitzman emphasises that asset puzzles dissolve when we also account for the process of learning about these parameters.

As the benefits from government investments are likely to vary with aggregate consumption and fall during consumption catastrophes, a reasonable assumption is that government investments face the same risk as private investments and the cost of capital should include the market risk premium caused by fear of disasters.

Other explanations

Transactions costs can explain the equity premium when equity is more costly to trade than debt. The equity premium may reflect the lower liquidity of equity compared with bonds. The risk free asset should be free from liquidity risk, there may be a premium for the risk from holding less liquid assets.[29]

If the high return on equities is to compensate for the costs that arise because of their lack of liquidity, the bond return still does not correctly measure the costs from government projects. If anything the liquidity premium for government assets should be higher because from the point of view of the taxpayer, a government

[27] Weitzman (2007a, p. 1128).

[28] See Barro (2006; 2007; 2009), Barro and Ursúa (2008). See also Wachter (2008), who examines a time varying probability of disaster.

[29] See Swan (2006) and Levy and Swan (2008).

investment is even less liquid than equity. Unlike equity it cannot be traded or disposed of by an individual taxpayer.

If government does have an advantage in raising money, then that is an argument for expanding government financing role. But it should still invest in the projects with the highest return, an objective promoted by using a market-based discount rate.

Further, the equity premium puzzle is the high Sharpe ratio for stocks. Cochrane points out that

> There are large Sharpe ratios *between* stocks (as in the value-growth premium studied by Fama and French) ignoring bonds all together. High sample Sharpe ratios are pervasive in finance and not limited to the difference between stocks and bonds.[30]

Differences in the liquidity of stocks and bonds cannot cause the entire equity premium.

McGrattan and Prescott (2003) argue that tax and regulatory changes boosted the demand for shares — which raised realised returns but lowered expected returns, accounting for the equity premium puzzle.[31]

H.5 Two approaches to accounting for risk: certainty equivalent v risk premium.

For a one period valuation, using a risk–adjusted discount rate on expected values or the risk free rate on certainty equivalents are alternative ways to account for risk.[32] Both should give the same answer. For example, if the CAPM holds, then

$E(R^k) = R^f + \beta_{kM}[E(R^M) - R^f] = R^f + \psi COV(R^M, R^k)$ where

$\beta_{kM} = COV(R^M, R^k)/VAR(R^M)$ and $\psi = [E(R^M) - R^f]/VAR(R^M)$.

Now if P_0 is the price of the asset and P_E is the risky payoff received at the end of the period then $E(R^k) = E(P_E)/P_0 = R^f + \psi COV(R^M, R^k)$.

Rearranging this gives the risk adjusted rate of return valuation formula:

$$P_0 = E(P_E)/[R^f + \psi COV(R^M, R^k)].$$

30 Cochrane (2008, p. 259).

31 McGrattan and Prescott (2003).

32 The following is from Copeland and Weston (1983, pp. 195-96).

The value of the asset is the expected value of the risky payment discounted with a risk adjusted discount rate. If the asset has positive market risk, a risk premium is added to the risk free rate. The higher the covariance with the market return, the more market risk it contains and the higher the required rate of return.

An equivalent way to value the asset is to deduct a risk charge from the expected value of the risky payoff to convert it into a certainty equivalent and discount with the risk free return. The certainty equivalent of a risky payment is the smallest certain payment which the household would accept in exchange for the risky one.

The covariance between the risky asset and the market can be written as:

$\mathrm{cov}(R^M, R^k) = \mathrm{cov}\left[R^M, \frac{P_e}{P_0}\right] = \frac{1}{P_0} COV(P_e, R^M)$. Substituting this into the risk-adjusted rate of return equation gives:

$P_0 = \frac{E(P_e)}{R^f + \psi(1/P_0)COV(P_e, R^M)}$. Rearranging this gives the certainty equivalent valuation formula:

$$P_0 = \frac{E(P_e) - \psi COV(P_e, R^M)}{R^f}$$

The risk charge depends on the level of market risk in the benefit flow. When a benefit co-varies with the market return, it contains market risk and the certainty equivalent is less than the expected value. The more market risk, the lower the certainty equivalent.

The CAPM pricing equation can be expressed as expected values discounted by a risk adjusted discount rate or as their certainty equivalents discounted by the risk free rate. The two approaches are equivalent for one-period valuation models and require exactly the same information. They are just different ways of expressing the same calculation.

A major advantage of the standard CAPM model is that the value of an asset depends solely on the expected cash pay offs, the quantity of market risk in the payoffs, the risk free rate and the price of risk — which are all market determined variables. Individuals who perceive the same distribution of payoffs for a risky asset will price it exactly the same way, independent of their preferences and attitudes to risk.[33] The CAPM approach does not require everyone has the same utility function — just common expectations.

33 Copeland and Weston (1983, p. 196).

If a benefit is uncorrelated with the market (it only has diversifiable risk), its certainty equivalent is equal to its expected value and the appropriate discount rate is the riskless return. If it is negatively correlated with the market, the risk–adjusted discount rate is below the riskless rate, and the certainty equivalent would be greater than the payment's expected value (it provides insurance and reduces portfolio risk).

H.6 Discounting in a multi-period world

When the net benefit flows from a regulation extend over many periods, the conditions under which adding a constant risk premium to the discount rate is a correct way to account for risk are more stringent. It requires the certainty equivalent relative to the expected value of the net benefits to decline steadily over time, which implies the risks of the benefit stream increase over time. That is, $CE_n/(1+r_f)^n = EV_n/(1+r^*)^n$ where CE_n is the certainty equivalent of a payment received in period *n, EVn* is its expected value and r^* is the risk adjusted discount rate. For this to be true, *CE/EV* must fall systematically with *n*. It would be correct if the riskiness of income flows is greater the further in the future they are received. That is, if the project's cumulative risk increases at a constant rate so that the certainty equivalent approach requires a larger deduction for risk the further into the future the cash flow is received.

Using the market risk premium is justified when the net benefits from the proposal follow a Weiner process.[34] A Weiner process is the same as Brownian motion. The variance of the change in a Weiner process grows linearly with the time horizon. Over the long run, its variance goes to infinity. It is a common assumption in financial economics. The proportional risk associated with holding such an asset grows linearly over time.

The two main reasons why market risk of net cash flows can change over time is investor reassessment of the project risk and changes in aggregate consumption risk.

The standard CAPM can be used in a multi-period model only if the project's beta each period is known with certainty. Intermediate uncertainty is only admissible if it does not affect the beta coefficients.[35]

[34] For example, if $NB_j = E_{j-1}(NB_j)(1+\varepsilon_j)$ where NB_j is the net benefit at time *j*, $E_{j-1}(NB_j)$ is its expected value at time $j-1$ and ε_j is a random variable with zero mean and constant covariance with the return on the market portfolio. Then net benefits at time *j* evolves as a martingale and the risk-adjusted discount rate can be used to compute its present value. See Copeland and Weston (1983, pp. 361-63).

[35] See Brearley and Myers (1991, p. 201) and Jones (2008, pp. 153-156).

When the risk contained in net benefits flows does not increase linearly through the life of the project, using one constant discount rate leads to incorrect decisions, and the certainty equivalent approach is preferable. It is more flexible and allows separate risk adjustment for net benefits received in different periods.

To calculate certainty equivalents is difficult. It is equally difficult to calculate the corresponding discount rate structure, which requires the same information. A common practical solution is to use a constant risk–adjusted discount rate, which uses market information on the current price of market risk.

Myers and Turnbull (1977) find that this traditional CAPM approach gives close to correct answers in a multi–period world providing the right asset beta is used to calculate the discount rate. But that depends on project life, growth trend in expected cash flows and other variables, which are not usually considered important in assessing business risk.[36]

A constant risk adjusted discount rate gives the correct answer (that is, the same as discounting the certainty equivalents at the riskless rate) if cash flows follow a pure random walk. Otherwise, conventional capital budgeting is wrong, but it often gives close to the right answer. Myers and Turnbull (1997. p. 330) conclude:

> no serious errors are introduced by discounting cash flow streams at one–period expected rates of return inferred from observed beta. … conventional valuation formulas based on discounting expected cash flows give a good approximation to asset values derived from rigorous analysis of equilibrium market values.

But determining beta is more complicated. Beta depends on the link between cash flow forecast errors and forecast errors on the market return. It depends on asset life, the growth trend in cash flows and on the pattern of expected cash flows over time (which affects the relative weights on the components of cash flows).

Further, the discount rate in a multi–period world may include extra risk premiums for other risks — such as a changing riskless rate of return and changes in the economy.

[36] Appendix I, section I.5 sets out how beta in the CAPM falls with asset life.

I Discounting the distant future

I.1 Uncertainty about the appropriate rate — averaging discount factors

When there is uncertainty about what discount rate to use, Weitzman argues that the appropriate discount rate to calculate present value is lower the further in the future the payments are received.[1]

Assume that the relevant discount rate is constant for the whole period, but we simply are not sure what it is. It is uncertainty about the rate to use, not an uncertain interest rate that fluctuates over time. To put risk aside, assume risk neutrality. Weitzman shows that as the time horizon of the project increases, the appropriate discount rate falls, because it is discount factors, not discount rates, that should be averaged.

Although it may seem strange that the discount rate is persistent, yet unknown, there is genuine uncertainty about the correct discount rate, even when interest rates are known.

For example, say we are uncertain whether the discount rate should be 3 or 7 per cent per year and each is equally likely. The expected present value of a dollar received in one year's time is:

$$0.5*\$1/1.03 + 0.5*\$1/1.07 = 0.5*(\$0.97 + \$0.93) = \$0.953 = 1/1.0496,$$

which implies a discount rate of 4.96 per cent, close to the average of 3 and 7 per cent.

The expected present value of a dollar received in 100 years time is:

$$0.5*\$1/(1.03)^{100} + 0.5*\$1/(1.07)^{100} = 0.5*(\$1/19.2 + 1/867.7) = \$1/37.6 = 1/(1.037)^{100}$$

The implied discount rate is 3.7 per cent, much closer to 3 than to 7 per cent.

1 So called gamma discounting. See Weitzman (2001).Many other authors present this argument, including Pindyck (2006, p. 5); Pearce et al. (2003, pp. 127-130) and Cowen (2008, pp. 12-13).

The expected present value of a dollar received in 200 years time is:

$$0.5*\$1/(1.03)^{200} + 0.5*\$1/(1.07)^{200} = 0.5*(\$1/369.36 + 1/752{,}931.6) = \$1/738.35 = =1/(1.034)^{100}$$

The implied discount rate is 3.4 per cent, even closer to 3.

That is, if we are uncertain about which discount rate to use, the appropriate 'certainty equivalent' discount rate declines over time. The further in the future the payment is received, the lower the discount rate we should use to calculate its present value. For payments in the distant future, we should be using discount rates from the lower end of the spectrum of possible values. The lower rates have a greater relative weight the further we look into the future.

For example, often a proposal's risk characteristics are not clear, nor is the rate of return for that risk class or the weights to put on different sources. Averaging of discount factors favours using the lower possible rates, such as a lower risk premium.

When the discount rate declines the further into the future a payment is received, it is called hyperbolic discounting. There is some empirical evidence that people discount in a hyperbolic fashion – but this evidence also shows very high short–term discount rates (17–30 per cent).[2]

But Gollier (2003) turns the Weitzman logic on its head. Consider a riskless public investment project of one unit, which generates a single payoff of $(1.04)^n$ in *n* years time. Column 2 in table I.1 lists the payoffs for payments received in 1100 and 200 years time.

The public project comes at the expense of private investment. We do not know what the return to private capital will be. Suppose there is a 50 per cent chance it will be 7 per cent and 50 per cent chance it will be 3 per cent. The expected payoffs that a unit of private investment would produce is set out in column 3 of table I.1.

According to the Weitzman logic, the 1 year public project should be rejected. Its return of 4 per cent is less than the certainty equivalent discount rate (4.96 per cent), so it has a negative net present value. But we should accept the 100 and 200 year projects — as the rate of return is 4 per cent, above the appropriate discount rates (which are near 3 per cent, at 3.7 and 3.4 per cent).

2 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, 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.

Table I.1 **Payoffs from a one unit investment**

Time horizon	*Public project with a 4 per cent return*	*Expected payoff with equi-probable 3 or 7 per cent return*
1	1.04	0.5*1.03+ 0.5*1.07 = 1.05
100	1.04^{100} = 50.5	$0.5*1.03^{100}+0.5*1.07^{100}$= 9.6+433.9= 443.5 = 1.063^{100}
200	1.04^{100}= 2 550.75	$0.5*1.03^{200}+0.5*1.07^{200}$= 184.7+376 465.8 = 376 650.5 = 1.066^{200}

But with risk neutrality, if we maximise future value it is clearly better to invest in the market rather than the project. For the 100 year public project, the future value of the private investment is almost 9 times greater than the project and the expected return is 6.3 per cent (closer to 7 than 3 per cent, and well above the project's 4 per cent return). For the 200 year project the expected future value of the private investment is almost 150 times greater and the expected return 6.6 per cent. The certainty equivalent discount rate rises if we evaluate the project further in the future.

Gollier argues that different investment projects should be ranked according to their expected net future value and we should take a larger interest rate to discount long–term cash–flows with respect to short–term ones. As the time horizon of the project increases, we should be using interest rates from the upper end of the spectrum of possible values — the opposite of the Weitzman result.

Both results are correct about the effects of uncertainty about the discount rate. The further into the future a payment is received, the lower the discount rate used to calculate expected net present value (Weitzman). The further into the future we evaluate a project, the higher the discount rate used to calculate net future value (Gollier 2003). The expected net present value and expected net future value criteria can recommend different courses of action.

The paradox disappears if the investment problem is carefully specified. For example, Gollier and Weitzman (2009) show that in a more rigorous formulation of the problem, the two approaches give the same discount rate and the puzzle is resolved. They examine a model where a project decision must be made, then the interest rate is revealed and the consumer determines his optimal consumption path. Whether the cash flows are converted into present consumption or future consumption makes no difference to the discount rate so long as the values are adjusted by marginal utility at the relevant time period, which adjusts for the risk associated with financing the project.

In this case the risk-adjusted discount rate has properties that resemble the Weitzman recommendations. The rate is lower the further in the future the payment is received, and approaches the bottom of the interest rate distribution. In the case of log utility, the original Weitzman rule is correct.

I.2 The benefits of delay

Hyperbolic discounting is extremely favourable for projects, like global warming abatement, with immediate costs but benefits that are received in the distant future. The immediate costs would be discounted at a high discount rate. The future benefits are discounted at a low rate, raising the net present value of the project compared with a constant rate.

But we need to account for the benefits of delay. Even if a project has a positive net present value, delaying it for a year may increase the net present value.[3] If the project is undertaken immediately (in period *0*), the cost C_0 is borne in period *0* and a benefit B_0 is received in period n, far in the future. C_0 could be the present value in period *0* of a stream of costs and B_0 could be the present value in period n of a stream of benefits. If we delay the project one period (undertake it in period 1), then C_1 is the cost borne in period 1 and B_1 is the benefit received in period n. An alternative assumption is that the benefits will also be delayed by one period, but with hyperbolic discounting, that would make little difference to the period *0* present value of the benefits (the marginal discount rate from period n to period $n+1$ is almost 0). If we use i_s to discount short-term costs and i_L to discount long term benefits, then it is efficient to delay the project for a year if the present value of starting the project now is less than the present value of waiting a year:

$$-C_1/(1+i_s)+B_1/(1+i_L)^n > -C_0+B_0/(1+i_L)^n \text{ or}$$

$$-C_1/(1+i_s) < C_0+(B_1-B_0)/(1+i_L)^n$$

But when the benefits are received far in the future (n is large), then the last term is negligible and the project should be delayed if $C_1/(1+i_s)<C_0$. That is, so long as costs grow at a slower rate than the short term discount rate. In the context of global warming abatement, say the policy is to reduce atmospheric concentration of CO_2 to some target level by 2100. The present value of the benefits from that are likely to be fairly constant. If the present value of the costs if we start today are more than the present value of the costs if we delay starting the project for one period, then delaying the delaying the project increases its value. The relevant discount rate is i_s, the short term discount rate.

3 See Layard (1994, pp. 43-44).

Although hyperbolic discounting increases the net present value of long term projects, it may also make delaying them attractive. Delay may have a further benefit if more information becomes available and some uncertainty about costs and benefits is resolved.

To find the best starting date for a project, its net present value for different starting dates should be compared. At minimum, the option of delaying for a year should always be considered.

I.3 Uncertainty about the path of rates

A different source of uncertainty is the lack of knowledge of the future paths of interest rates. There is a risk that interest rates may change.

Gollier models the interest rate process to determine the term structure of interest rates.[4] The result is time consistent. He shows the answer mainly depends upon the time horizon under scrutiny, the degree of relative prudence and the degree of resistance to intertemporal substitution. For commonly accepted levels of these indexes, the effect of uncertainty on the socially optimal discount rate may vary from 2 to 8 per cent. He finds that the discount rate to be used for long–lasting investments should be a decreasing function of their duration, because of the negative effect of accumulating the per period growth risk in the long run.

That is, uncertainty about future interest rates means we should use a lower than current real interest rate to discount costs and benefits received further in the future (the yield curve for real interest rates slopes down).[5]

Newell and Pizer (2003) model possible future time paths of interest rates, past on historical interest rate behaviour. Future rates can vary widely. They construct expected discount factors to calculate present value, which gives a lower discount rate for benefits received further in the future — again because the lower rates have a relative weight in the discount factor the further we look into the future.

4 See Gollier (2002; 2002a; 2004).

5 Note the usual yield curve slopes up – but that is for the yield on nominal bonds. The longer the term of a nominal bond, the greater is inflation risk and the higher the yield needed to compensate.

I.4 Real options

Long term projects should be treated as a problem of sequential decision making under uncertainty rather than a 'one-shot' benefit cost analysis.[6]

A project may have a positive net present value but still not be accepted right away. The firm may gain by waiting and accepting the project in a future period, for the same reasons that investors do not always exercise an option just because it is in the money. This is more likely to happen if the firm has the rights to the project for a long time, and the variance in project inflows is high.

When a project has a positive net present value, it is necessary, but not sufficient condition for an efficiency improvement. For example, say a project costs \$20 and gives \$1000 benefit in 101 years. Today the interest rate is 5 per cent, but next year it will change. There is a 50 per cent chance it will rise to 7 per cent and a 50 per cent chance it will fall to 3 per cent. It will then remain at this new level. The expected present value of the payment next year (as worked out in 'Uncertainty about the appropriate rate — averaging discount factors') is:

$$\$1000/(37.6) = \$1000/(1.037)^{100} = \$26.60.$$

The project has a positive net present value. If our choice was, invest today or never invest, we would invest today. Further, the uncertainty about the future rate (holding the expected value of the interest rate constant) increases the expected value of the project. If the interest rate stayed at 5 per cent, the project would be worth $\$1000/1.05^{100}= \$1000/131.5 = \$7.60$ next period and it would not be worth undertaking.

But what if we wait one year before deciding whether to invest. If the interest rate rises to 7 per cent, then the value of the benefit is: $\$1000/1.07^{100} = \$1000/867.7 = \$1.15$, less than the cost of the investment.

If it fell to 3 per cent, the value of the benefit next year is: $\$1000/1.03^{100} = \$1000/19.2 = \$52.03$.

We would only invest if the interest rate fell. Further, the value of the project is double compared with not waiting. The interest rate uncertainty creates a value to waiting for new information and an incentive to wait rather than invest now. If interest rates are uncertain, it may be worth waiting to see whether they rise or fall.[7]

---

6 Weyant (2008, p. 78) .See also Lind (1995).

7 In the Weitzman example, the uncertainty about the interest rate never resolves itself. We go on being uncertain about the appropriate interest rate for the whole period.

With any irreversible investment under uncertainty, the timing of the investment is crucial. The decision is not only whether to invest, but when to invest, there can be benefits from delay. Even if the net present value of the project is positive, it still may be worth delaying it. When there is uncertainty over net benefits, flexibility is more valuable.

For most government projects, there is an option to delay and they are difficult to reverse. When the project involves making public policy or providing a public good, the government is usually a monopoly, and can choose when to start, unlike a competitive firm, which has to worry whether a competitor will jump in and make the investment.

The ability to delay an irreversible investment profoundly affects the decision to invest. This has been called real option theory: a firm with an opportunity to invest is holding an option – it has the right, but not the obligation, to buy an asset at some future time of its choosing. When it invests, it exercises its option. It 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.[8]

As a result a positive net present value is a necessary, but not sufficient condition for a project to be efficient. The net present value must exceed the value of keeping the option alive. This option value can be large, explaining why firms often set hurdle rates three or four times the cost of capital. The ‘trigger’ value is not when the expected benefit becomes positive, but when it is sufficiently high. The trigger is higher when there is greater uncertainty, especially if future information would help resolve it. Ironically, a lower interest rate increases the value of waiting and the investment hurdle.

The option value of delay increases the appropriate discount rate to apply to government projects. Even a high discount rate is conservative when the project is an irreversible investment with uncertainty. Even if the net value of the project is positive, it still may be worth paying to delay it.[9] The option value of waiting is valuable when there is uncertainty and you may acquire some information. There may also be a benefit from delay simply by waiting to bear costs.[10]

8 Dixit and Pindyck (1995).

9 See Dixit and Pindyck (1994, chapter 1) and Dixit (1992).

10 See appendix I.2.

The value of delaying the project increases with:

- The sunk cost of the investment
- The degree of uncertainty over benefits.

The higher the degree of uncertainty, the greater the benefit from waiting. Uncertainty about net benefits increases the hurdle rate for the project. It is not enough that the net present value is positive, it must be larger than the option value of delay.

I.5 The risk premium for the distant future

There are reasons for using a lower risk premium for assets that pay off in the distant future. For example, Myers and Turnbull (1977) find that beta in the CAPM falls with asset life. Projects that pay off in the distant future have a lower beta, and so a lower risk premium.

The cash flow is decomposed into components, each with its own risk. They assume that the CAPM model holds each period and each component's beta is independent of time. Each component is like a distinct asset. The composite asset's beta is a weighted average of each component's own beta. The weights are the proportional contribution of the component to the composite asset's price. For example, a net benefit flow could be decomposed into the cost flow and benefit flow.[11]

Myers and Turnbull show that when the component's have different betas. The beta for the composite asset must be a declining function of asset life because asset life affects each component's weight. The expected annual cash flow generated by each component is held constant. But as the horizon is expanded, the present value of each stream increases, but not at the same rate. The cash flows of the stream having the higher beta are discounted at a higher rate and so the weight put on the high beta stream declines as asset life increases. The longer the horizon, the greater the proportion of the asset's price generated by the safer stream and the lower the asset's beta and so the lower the relevant discount rate for evaluating the project.

[11] See Jones (2008, pp. 151-53) for details.

References

Abelson, P. 2000, 'Public Economics principles and practice', *Applied Economics*.

Abusah, S. and de Bruyn, C. 2007, 'Getting Auckland on Track: Public Transport and New Zealand's Economic Transformation' *Ministry of Economic Development Working Paper*, August.

Access Economics 2006, 'Opal Cost Benefit Analysis' Report by Access Economics Pty Limited for The Opal Alliance, February.

Ackerman, F. and Heinzerman, L. 2002, 'Pricing the Priceless: Cost-Benefit Analysis of Environmental Protection', *University of Pennsylvania Law Review*, vol. 150, pp. 1553-84.

Allsop, V. 1995, *Understanding Economics,* Routledge.

Ariel, R. 1998, 'Risk adjusted discount rates and the present value of risky costs' *Financial Review*, vol. 33, pp. 17-30.

Arrow, K. 1966, 'Discounting and Public Investment Criteria,' in Kneese, A.V. and Smith, S.C., (eds), *Water Research*, Baltimore

______ 2007, 'Global Climate Change: A Challenge to Policy' *Economists' Voice*, June, pp. 1-5, www.bepress.com/ev.

Arrow, K. and Lind, R. 1970, 'Uncertainty and the Evaluation of Public Investment Decisions' *American Economic Review*, vol. 60, no. 3, pp. 364–78. Reprinted in Layard and Glaister (eds) 1994, pp. 160–78.

——, Cropper, M., Eads, G., Hahn, R., Lave, L., Noll, R., Portney, P., Russell, M, Schmalensee, R., Smith, V. and Stavins, R. 1996, *Benefit-Cost Analysis in Environmental, Health, and Safety Regulation*, American Enterprise Institute, The Annapolis Center, and Resources for the Future.

Atkins, F. 1989, 'Co-integration, error correction and the Fisher effect' *Applied Economics*, vol. 21, pp. 1611-20.

Australian Government 2007, *Best Practice Regulation Handbook*, Canberra.

Australian Office of Financial Management. 2002, 'Terms and Conditions Of Issue Treasury Indexed Bonds' Canberra, 18 February, http://www.rba.gov.au/FinancialServices/CGSTenderInfo/terms_and_conditions_ib.pdf

Bailey, M., and Jensen, M. 1972, ‘Risk and the Discount Rate for Public Investment’, in Michael C. Jensen (ed)., *Studies in the Theory of Capital Markets,* New York, Praeger, pp. 269–93. Page references are to the mimeographed copy at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=390110

Baker, R., Barker, A., Johnston, A. and Kohlhaas, M. 2008, *The Stern Review: an assessment of its methodology*, Productivity Commission Staff Working Paper, Melbourne, January.

Barro, R. 2009, ‘Rare Disasters, Asset Prices, and Welfare Costs’ *American Economic Review*, vol. 99, no. 1, March, pp. 243-64.

—— 2007, ‘On the Welfare Costs of Consumption Uncertainty’ mimeo, Harvard University, January.

—— 2006, ‘Rare Disasters and Asset Markets in the Twentieth Century’ *Quarterly Journal of Economics*, vol. 121, August, pp. 823-66.

Barro, R. and Becker, G. 1989, ‘Fertility Choice in a Model of Economic Growth’ *Econometrica*, vol. 57, no. 2, March, pp. 481-501.

Barro, R. and Ursúa, J. 2008, ‘Macroeconomic Crises since 1870’ Harvard University, March.

Bazelon, C. and Smetters, K. 1999. ‘Discounting inside the Washington DC Beltway’ *Journal of Economic Perspectives*, vol. 13, pp. 213-28.

Becker, G. and Barro, R. 1988. ‘A Reformulation of the Economic Theory of Fertility’ *Quarterly Journal of Economics*, vol. 103, no. 1, February, pp. 1-25.

Bernstein, P. 2005, *Capital Ideas* John Wiley and Sons, (originally published in 1993).

—— 2007, *Capital Ideas Evolving,* John Wiley and Sons.

Blanchard, O., and S. Fischer. 1989, *Lectures on Macroeconomics*. Cambridge, Mass.: MIT Press.

Boadway, R. 2006, ‘Principles of cost-benefit analysis’ *Public Policy Review*, vol.2, no.1, pp.1-44.

Boardman, A. and Greenberg, D. 1998, ‘Discounting and the Social Discount Rate’ chapter 8, pp. 269-318 in Thompson and Decker (eds) 1998.

Boardman, E.A., Greenberg, D.H., Vining, A.R. and Weimer, D.L. 2006, *Cost–Benefit Analysis: Concepts and Practice*, 3rd edition, Pearson Prentice Hall, New Jersey.

Boskin, M., E. Dulberger, R. Gordon, Z. Griliches and D. Jorgenson. 1998, 'Consumer Prices, the Consumer Price Index, and the Cost of Living' *Journal of Economic Perspectives,* v. 12, no.1, Winter pp.3–26.

Bradford, D. 1975, 'Constraints on Government Investment Opportunities and the Choice of Discount Rate' *American Economic Review*, vol. 65, no. 5, December, pp. 887-899.

Brailsford, T., Handley, J. and Maheswaran, K. 2007, 'A Re-Examination of the Historical Equity Risk Premium In Australia'. April, http://ssrn.com/abstract=982043.

Brearly, R. and Myers, S. 1991, *Principles of Corporate Finance*, 4th edition, McGraw-Hill.

——, Cooper, I. and Habib, M. 1997, 'Investment Appraisal in the Public Sector' *Oxford Review of Economic Policy,* vol. 13, no. 4, pp. 12-28.

Brennan, G. 2006, 'Discounting the future, yet again' seminar paper for the Research Triangle Ethics Circle, http://www.unc.edu/~gsmunc/EthicsGroup/Brennan%20--%20Discounting%20the%20Future.pdf.

Browning, E. and J. Browning. 1994, *Public Finance and the Price system* 4th edition, Prentice-Hall.

Buchanan, C. and P. Hartley. *Equity as a social goal.* New Zealand Business Roundtable, Wellington, New Zealand.

Buchanan, J. 2008, 'Public debt' *The New Palgrave Dictionary of Economics*. 2nd edition. Steven N. Durlauf and Lawrence E. Blume, (Eds) Palgrave Macmillan.

Bureau of Transport and Regional Economics. 1999, *Facts and furphies in benefit cost analysis: transport*. Report 100. Commonwealth of Australia, Canberra.

—— 2005, 'Risk in cost benefit analysis' Report 110, Department of Transport and Regional Services, Commonwealth of Australia.

Byatt, I. 2008, 'Weighing the present against the future: the choice, and use, of rates of discount in the analysis of climate change' pp.92-113 in *Climate Change Policy: Challenging the Activists*, Institute of Economic Affairs.

Campbell, J. 1999, 'Who Should Buy Long-Term Bonds?' mimeo, Harvard University.

Caselli, F. and Feyrer, J. 2007, 'The Marginal Product of Capital' *Quarterly Journal of Economics*, May, pp. 535-68.

Cawley, J. and Philipson, T. 1999, 'An Empirical Examination of Information Barriers to Trade in Insurance', *American Economic Review*, vol. 89, no. 4, pp. 827–846, September.

Chetty, R. and Szeidl, A. 2007, 'Consumption Commitments and Risk Preferences' *Quarterly Journal of Economics,* May, pp. 831-77.

Cline, W.R. 1993, 'Give greenhouse abatement a fair chance' *Finance and Development,* vol. 30, no. 1, pp. 3-5.

Cochrane, J. 1991, 'Volatility tests and efficient markets: a review essay' *Journal of Monetary Economics*, vol. 27, pp. 463-485, North-Holland

—— 1997, 'Where is the market going? Uncertain facts and novel theories' *Economic Perspectives*, Federal Reserve Bank of Chicago, November/December, vol. XXI, no. 6, pp.3-37.

—— 2004, 'Asset Pricing: Liquidity, Trading, and Asset Prices' *NBER Reporter*, Winter 2004/5, pp.1-12.

—— 2005, *Asset Pricing*. Revised edition. Princeton University Press.

—— 2008, Financial Markets and the Real Economy' *Handbook of the Equity Risk Premium*, chapter 7, pp. 237-330, Elsevier.

—— Undated. 'Efficient Markets Today' http://faculty.chicagogsb.edu/ john.cochrane/research/Papers/Cochrane_efficient_markets.doc

Cohen, D., Hasset, K. and Hubbard. R.G. 1999, 'Inflation and the user cost of capital: does inflation still matter' in M. Feldstein (ed) *The Costs and Benefits of Price Stability*, University of Chicago Press.

Cole, D. 2007, 'The *Stern Review* and Its Critics: Implications for the Theory and Practice of Benefit-Cost Analysis' October.

Commonwealth Competitive Neutrality Complaints Office (CCNCO). 1998, *Rate of Return Issues,* CCNCO Research Paper, December, Productivity Commission, Canberra.

Congressional Budget Office, 2004, 'Estimating the Value of Subsidies for Federal Loans and Loan Guarantees' *CBO Study.* August. Congress of the United States.

Copeland, T. and Weston, J. 1983, *Financial Theory and Corporate Policy*, 2nd edition, Addison-Wesley.

Cooper. I. 1996, 'Arithmetic versus geometric mean estimators: setting discount rates for capital budgeting' *European Financial Management,* vol. 2, no.2, pp. 157-67.

Cooray, A. (undated) 'The Fisher effect: a review of the literature'.

Cowen, T. 2001, 'What is the Correct Intergenerational Discount Rate?' Mimeo. Department of Economics, George Mason University.

—— 2008, 'Caring about the Distant Future: Why It Matters and What It Means' *University of Chicago Law Review*, vol. 74, no. 5, pp. 5-40.

—— and Parfit, D. 1992, ‘Against the social discount rate’ in Peter Laslett and James S. Fishkin (eds.) *Justice between age groups and generations*, Yale University Press: New Haven, pp. 144-161.

Cropper, M. and Laibson, D. 1999, ‘The Implications of Hyperbolic Discounting for Project evaluation’, chapter 16, pp. 163-72 in Portney and Weyant, *Discounting and Intergenerational Equity*, Resources for the Future Press,

Currie, D. 2000, ‘Funding the London Underground’ *London Business School.* http://facultyresearch.london.edu/docs/paper35.pdf

Damodaran, A. 2008, ‘What is the riskfree rate? A Search for the Basic Building Block’ mimeo, December, Stern School of Business, New York.

Dasgupta, P. 2006, ‘Comments on the Stern Review's Economics of Climate Change’ December 12. Seminar organised by the Foundation for Science and Technology at the Royal Society, London,

DeLong, B. and Magin, K. 2009, ‘The U.S. Equity Return Premium: Past, Present, and Future’ *Journal of Economic Perspectives*, vol. 23, no.1, Winter, pp. 193-208.

Department of Finance, 1991, *Handbook of Cost–Benefit Analysis,* Australian Government Publishing Service, Canberra.

Department of Finance and Administration 2006, *Handbook of Cost–Benefit Analysis*, Financial Management Reference Material No. 6, Australian Government, Canberra, January.

Department of Health and Ageing and enHealth Council. 2003, *Guidelines for Economic Evaluation of Environmental Health Planning and Assessment. Volume 1 — The Guidelines.* Commonwealth of Australia, January.

Department of Treasury and Finance Tasmania. 1996, *Project Initiation Process*. December.

Devlin, W. and Brummitt, B. 2007, ‘A Few Sovereigns More: The Rise of Sovereign Wealth Funds’, *Economic Round-up*, Spring, pp. 119-36.

Dews, N., Hawkins, J. and Horton, T. 1992, ‘Measuring the cost of capital in Australia’ *Reserve Bank Research Discussion Paper,* no. 9205, June, Economic Analysis Department, Reserve Bank of Australia

Dietz, S. and Stern, N. 2008, ‘Why Economic Analysis Supports Strong Action on Climate Change: A Response to the Stern Review’s Critics’ *Review of Environmental Economics and Policy*, vol. 2, no. 1, Winter, pp. 94-113.

Dimson, E., Marsh, P. and Staunton, M. 2002, ‘Global evidence on the equity risk premium’ London Business School, September.

Dixit, A. 1992, ‘Investment and Hysteresis’, *Journal of Economic Perspectives*, vol. 6, no. 1, Winter, pp. 107-32.

Dixit, A. and A. Williamson. 1989, ‘Risk-Adjusted Rates of Return for Project Appraisal’, Policy, Planning, and Research Working Papers, Agricultural Policies, Agriculture and Rural Development Department, World Bank, November, WPS 290

—— and Pindyck, R. 1994, *Investment and Uncertainty*, Princeton University Press.

—— and —— 1995, ‘The options approach to capital investment’, *Harvard Business Review*, May, pp. 105-15.

Dolman, B. 2007, ‘The distribution of recent economic gains: some early observations’, presentation to the 3rd Productivity Perspectives Conference, December, Canberra.

Dreze, J. 1974, ‘Discount Rates and Public Investment: A Post-Scriptum’ *Economica*, vol. 41, no. 161, February, pp. 52-61.

Economic Planning Advisory Commission. 1995, *Private Infrastructure Task Force Report*. September, AGPS, Canberra.

Edwards, S. 1985, ‘Country risk, foreign borrowing and the social discount rate in an open developing economy’ *NBER Working Paper Series*, Working Paper no. 1651.

Ergas, H. 2008, ‘Shelter from the Storm: Reflections on the “Risk Society”’, *Quadrant*, January, vol. LII, no. 1–2.

Fama, E., and French, K. 2002, ‘The Equity Premium’, *Journal of Finance*, April, vol. 57, no. 2, pp. 637–59.

—— and —— 2004, ‘The Capital Asset Pricing Model: Theory and Evidence’ *Journal of Economic Perspectives,* vol. 18, no. 3, Summer, pp. 25–46.

Fane, G. and Applegate, C. 1995, ‘Sovereign default risk and the social cost of foreign debt’ chapter 9, pp. 161-70 in J. Pitchford (ed), *The Current Account and Foreign Debt*, Routledge, London.

Federal Reserve Bank of Cleveland, 2009, ‘Calculating Inflation Expectations Using Two Kinds of Treasury Securities’ http://www.clevelandfed.org/research/data/TIPS/bg.cfm accessed August 14, 2009

Feldstein, M. 1972, ‘The Inadequacy of Weighted Discount Rates’ in *Cost Benefit Analysis,* Layard. R. (ed), pp. 311-31, Harmondsworth UK, Penguin Books.

Foldes, L and Rees, R. 1977, ‘A Note on the Arrow-Lind Theorem’ *The American Economic Review*, vol. 67, no. 2, March, pp. 188-193.

Frederick, S., Loewenstein, G. and O'Donoghue, T. 2002, 'Time discounting and time preference: a critical review', *Journal of Economic Literature*, vol. 40, June, pp. 351–401.

Flyvbjerg, B., Bruzelius, N. and Rothengatter, W. 2003, *Megaprojects and Risk: An Anatomy of Ambition,* Cambridge University Press.

Garnaut, R. 2008, *The Garnaut Climate Change Review*, final report, Cambridge University Press.

Goetzmann, W. and Ibbotson, R. 2005, 'History and the Equity Risk Premium' *Yale School of Management*, mimeo, 18 October.

Gollier, C. 2002a, 'Discounting an uncertain future' *Journal of Public Economics*, vol. 85, pp.149–166.

—— 2002, 'Time Horizon and the Discount Rate', *Journal of Economic Theory*, vol. 107, pp. 463–473.

—— 2003, 'Maximizing the expected net future value as an alternative strategy to gamma discounting', University of Toulouse, September 1.

—— 2004, 'The consumption-based determinants of the term structure of discount rates' University of Toulouse, July 24.

—— 2006, 'An evaluation of the Stern's Report on the economics of climate change', http://idei.fr/doc/wp/2006/stern.pdf

Gollier, C. and Weitzman, M. 2009, 'How Should the Distant Future be Discounted When Discount Rates are Uncertain?' November.

Gordon, R. 2006, 'The Boskin Commission Report: A Retrospective One Decade Later' *National Bureau of Economics Research Working Paper no.12311*, National Bureau of Economic Research.

Gordon, M. and Gould, L. 1984, 'The Nominal Yield and Risk Premium on the TSE-300, 1956-1982', *Canadian Journal of Administrative Sciences*, vol.1, no.1, pp.50-60.

Grant, S. and J. Quiggin. 2003, 'Public Investment and the Risk Premium for Equity' Economica, vol. 70, pp.1–18

—— 2006, 'The Risk Premium for Equity: Implications for Resource Allocation, Welfare and Policy' *Australian Economic Papers*, September, pp.253–268.

Gray, S. and R. Officer. 2005, 'A Review Of The Market Risk Premium and Commentary On Two Recent Papers' A Report Prepared for the Energy Networks Association, 15 August.

Grout, P. 'Public and Private Sector Discount Rates in Public-Private Sector Partnerships' *Economic Journal,* v.113*,* March) pp. C62-68.

Harberger, A. 1997, ‘Economic Project Evaluation: Some Lessons for the 1990s’ *Canadian Journal of Program Evaluation*, Special issue, pp.5-48.

—— 1992, ‘Towards an operational approach to social cost-benefit analysis’ in S.B. Dahiya (ed.), Project Evaluation, New Delhi, Concept Publishing Co.

—— 1987, ‘Reflections on Social Project Evaluation’ pp.153–188 in G. Meier (ed) *Pioneers in Development Second Series*. World Bank Oxford University Press.

—— 1984, ‘Basic needs versus distributional weights in social cost‑benefit analysis’, *Economic Development and Cultural Change*, vol. 32, no. 3, pp. 55-474.

—— 1978, ‘On the Use of Distributional Weights in Social Cost-Benefit Analysis’, *Journal of Political Economy*, vol. 86, no. 2, part 2, pp. S87-S120

—— 1971, ‘Three Basic Postulates for Applied Welfare Economics: An Interpretive Essay’, *Journal of Economic Literature*, vol. 9, no. 3, pp. 785-97.

—— 1969, ‘On measuring the social opportunity cost of public funds’ in *The Discount Rate in Public Investment Evaluation,* western Agricultural Economics Research Council, Report no.17, Denver, Colorado. Reprinted as chapter 4, pp. 94-122 in A. Harberger *Project Evaluation: Collected Papers*.

Hathaway, N. 1997, ‘Privatisation and the Government Cost of Capital’, *Agenda*. vol. 4, no. 2, pp. 155–164.

—— 1997a, ‘Rejoinder’, *Agenda*, vol. 4, no. 4, pp. 486–488.

Hawtrey, K. 1997, ‘The Fisher effect and Australian interest rates’, *Applied Financial Economics,* vol. 7, pp. 337-346.

Hepburn. C. and Groom, B. 2007, ‘Gamma Discounting and Expected Net Future Value’, *Journal of Environmental Economics and Management*, Elsevier, vol. 53, no. 1, pp. 99–109, January.

Heyne, P. 2002, ‘Interest’, http://www.econlib.org/library/Enc1/Interest.html.

Hirshleifer, J. 1964, ‘Efficient Allocation of Capital in an Uncertain World’, *American Economic Review*, vol. 54, no. 3, pp. 77-85, May

HM Treasury. 2003, *The Green Book: Appraisal and Evaluation in Central Government*, Treasury Guidance. London.

Inder, B. and Silvapulle, P. 1993, ‘Does the Fisher effect apply in Australia’, *Applied Economics*, vol. 25, pp.839-843.

Infrastructure Australia 2008, *Outline of Infrastructure Australia's Prioritisation Methodology*, September.

—— 2008a ‘National PPP Guidelines Discount Rate Methodology Guidance’ vol. 5 December.

Jagannathan, R., McGrattan, E. andScherbina, A. 2001, 'The Declining US Equity Premium', *Federal Reserve Bank of Minneapolis Quarterly Review*, vol. 24, no. 4, pp. 3–19.

Jenkins, G. 1981, 'The Public-Sector Discount Rate for Canada: Some Further Observations', *Canadian Public Policy*, vol. 7, pp. 399-407.

Jensen, M. and Bailey, M. 1972, 'Risk and the Discount Rate for Public Investment', *Studies in the Theory of Capital Markets,* M. Jensen, (Ed), Praeger Publishers, http://papers.ssrn.com/abstract=390110.

Jones, C. 2008, *Financial Economics,* Routledge.

—— 2008a, 'Realisation vs. Accrual Based Taxes on Capital Gains' mimeo, Collage of Business and Economics, Australian National University, July.

—— 2005, *Applied Welfare Economics*, Oxford University Press.

Jones, S. and Netter, J. 2008, 'Efficient Capital Markets' from *The Concise Encyclopedia of Economics*, D. Henderson (ed), 2nd edition, http://www.econlib.org/library/Enc/EfficientCapitalMarkets.html

Kaplow, L. 1994, 'Taxation and Risk Taking: A General Equilibrium Perspective', *National Tax Journal*, vol. XLVII, no. 4, pp. 789–98.

—— 2008, *The Theory of Taxation and Public Economics*, Princeton University Press.

—— 2006a, 'Taxation', John M. Olin Center For Law, Economics, and Business Discussion Paper No. 542, Harvard Law School, Cambridge, USA. Forthcoming in *Handbook of Law and Economics,* Polinsky and Shavell, (eds.), Elsevier.

—— 2006b. 'Discounting Dollars, Discounting Lives: Intergenerational Distributive Justice and Efficiency', p*repared for the University of Chicago Conference on Intergenerational Equity and Discounting*, April 10.

—— 2003 'Concavity of Utility, Concavity of Welfare, and Redistribution of Income', *Harvard Law School John M. Olin Center for Law, Economics and Business Discussion Paper Series,* Paper 437.

Kanbur, R. 2008. 'Shadow pricing' *The New Palgrave Dictionary of Economics*. 2nd Edition. Steven N. Durlauf and Lawrence E. Blume (Eds), Palgrave Macmillan.

Klein, M. 1996, 'Risk, Taxpayers, and the Role of Government in Project Finance' Policy Research Working Paper 1688, The World Bank, December.

—— 1997, "The risk premium for evaluating public projects' *Oxford Review of Economic Policy*, vol. 13, no. 4, pp. 29-42.

Kocherlakota, N. 1996, 'The Equity Premium: It's Still a Puzzle', *Journal of Economic Literature*, vol. 34, no. 1, March, pp. 42-71.

Lally, M. 2000, *The Real Cost of Capital in New Zealand: Is It Too High?*, New Zealand Business Roundtable, Wellington.

—— 2002, 'The cost of capital under dividend imputation', Prepared for the Australian Consumer and Competition Commission, Victoria University of Wellington, June.

—— 2005, 'Structural Breaks and the Estimation of the Market Risk Premium in Australia', Report prepared for the Queensland Competition Authority

Layard, R. 1994, 'Introduction', pp. 1-56, in Layard and Glaister (eds).

Layard, R. and Glaister. S. (eds), 1994, *Cost Benefit Analysis*. 2nd edition, Cambridge University Press.

Layard, R. and Walters, A. 1978, *Microeconomic Theory*, McGraw-Hill.

Levy, A. and Swan, P.2008, 'The Liquidity Component of the Equity Premium', mimeo, Australian School of Business, University of New South Wales, 21 February.

Lind, R. (ed).1982, *Discounting for Time and Risk in Energy Policy,* Baltimore, Johns Hopkins University Press for Resources for the Future.

—— 1982a, 'A Primer on the Major Issues Relating to the Discount Rate for evaluating National Energy options' chapter 2, pp. 21–94.

—— 1982b, 'The Rate of Discount and the Application of Social Benefit-Cost Analysis in the Context of Energy Policy Decisions', chapter 12, pp. 443-457.

—— 1995 'Intergenerational equity, discounting, and the role of cost–benefit analysis in evaluating global climate policy', *Energy Policy*, vol. 23, no. 4–5, pp. 379-89.

Lo, A. 2008, 'Efficient markets hypothesis', *The New Palgrave Dictionary of Economics*, 2nd edition, S. Durlauf and L. Blume (eds.)

Makin, T. 2006, 'Has Foreign Capital Made Us Richer?', *Agenda*, vol. 13, no. 3, pp. 225–37.

Malkiel, B. 2003, 'The Efficient Market Hypothesis and Its Critics', *Journal of Economic Perspectives*, vol. 17, no. 1, Winter, pp. 59–82

Mankiw, G. 1986, 'The Equity Premium and the Concentration of Aggregate Shocks', *Journal of Financial Economics*, vol. 17, pp. 211-219.

—— and Zeldes, S. 1991, 'The consumption of stockholders and nonstockholders', *Journal of Financial Economics*, vol. 29, pp. 97-111.

McCallum, B. 1996, 'Neoclassical vs. Endogenous Growth Analysis: An Overview' *Federal Reserve Bank of Richmond Economic Quarterly*, vol. 82, no. 4, Fall, pp. 41–71.

McCulloch, B. 2003, ‘Geometric Return and Portfolio Analysis’, *New Zealand Treasury Working Paper 03/28*, December.

McCulloch, B. and Leonova, D. 2005, ‘The Market Equity Risk Premium’, *Treasury Paper*, New Zealand Treasury, May.

McGrattan, E. and E. Prescott. 2003, ‘Average Debt and Equity Returns: Puzzling?’, *American Economic Review*, vol. 93, no. 2, May, pp. 392-397.

Mehra, R. 2006, ‘The Equity Premium Puzzle: A Review’, *Foundations and Trends in Finance,* vol. 2, no. 1, pp. 1–81

—— 2003, ‘The Equity Premium: Why Is It a Puzzle?’, *Financial Analysts Journal,* vol. 59*,* January/February, pp.54-69.

Mehra, R. and Prescott, E. 2003, ‘The Equity Premium in Retrospect’ in *Handbook of the Economics of Finance,* G. Constantinides, M. Harris and R. Stulz, (eds) Elsevier, pp. 887-936

Merton, R. 1980, ‘On Estimating the Expected Return on the Market’, *Journal of Financial Economics*, vol. 8, pp.323-61

Mishan, E. 1972, ‘Uncertainty and the Evaluation of Public Investment Decisions: Comment’ *American Economic Review*, vol. 62, no. 1-2, pp. 161–164.

Mohun, S. 2004, ‘The Australian Rate of Profit 1965-2001’, *Journal of Australian Political Economy*, no. 52, pp. 83-112.

Mulligan, C. 2003, ‘Capital Tax Incidence: Fisherian Impressions from the Time Series’, *National Bureau of Economics Research Working Paper no.9916,* August.

—— 2002, ‘Capital, interest, and aggregate intertemporal substitution’ *National Bureau of Economics Research Working Paper no.9373*.

Murphy, K. 2008, ‘Some simple economics of climate change’, presentation to Mont Pelerin Society 60 the Anniversary Conference, Tokoyo, September 7.

Myers, S. and Turnbull, S. 1977, ‘Capital budgeting and the capital asset pricing model: good news and bad news’, *Journal of Finance*, vol. XXXII, no. 2, May, pp. 321–333.

Newell, R and Pizer, W. 2003, ‘Discounting the distant future: How much do uncertain rates increase valuations?’, *Journal of Environmental Economics and Management*, vol. 46, no. 1, pp. 52–71.

New South Wales Treasury 1997, *Guidelines for Economic Appraisal*, Office of Financial Management, Treasury Policy and Guidelines Paper 97–2.

New Zealand Treasury 2008, ‘Public Sector Discount Rates for Cost Benefit Analysis’, July

Nordhaus, W. 2007, ‘A Review of the Stern Review on the Economics of Climate Change’ *Journal of Economic Literature*, vol. XLV, September, pp. 686–702

Olekahns, N. 1996, ‘Further evidence on the Fisher effect’, *Applied Economics*, vol. 28, pp. 851-856.

Oliver, S. 2009, *Oliver’s Insights*, AMP Capital Investors. 37th edition, March.

Partnerships Victoria 2003, *Use of Discount Rates in the Partnerships Victoria Process*, Technical Note, July.

Pearce, D., Groom, B., Hepburn, C. and Koundouri, P. 2003, ‘Valuing the Future: Recent advances in social discounting’, *World Economics*, vol. 4, no. 2, April–June, pp. 121–141.

Perkins, F. 1994, *Practical Cost Benefit Analysis*, MacMillan, Melbourne.

Pindyck, R. 2006, ‘Uncertainty in Environmental Economics’, AEI-Brookings Joint Center For Regulatory Studies, Related Publication 06-39, December.

Portney, P. and Weyant, J. 1999, ‘Introduction’, *Discounting and Intergenerational Equity*, Resources for the Future Press, pp. 1-11.

Poterba, J. 1997, ‘The rate of return to corporate capital and factor shares: new estimates using revised national income accounts and capital stock data, *National Bureau of Economics Research Working Paper no.6263*, November.

PC (Productivity Commission) 2005, ‘The Private Cost Effectiveness of Improving Energy Efficiency’, Report no. 36, Canberra.

Queensland Treasury 2006, *Cost-Benefit Analysis Guidelines Achieving Value for Money In Public Infrastructure and Service Delivery.* Project Assurance Framework Supplementary Guidance Material, Version 0.3, July.

Quiggin, J. 1996, ‘The partial privatisation of Telstra’, supplementary submission’ Submission to Senate Environment, Recreation, Communications and the Arts Reference Committee inquiry into the Telstra, Dilution of Public Ownership Bill

—— 2006, ‘Stern and the critics on discounting’, http://johnquiggin.com/wp-content/uploads/2006/12/sternreviewed06121.pdf.

—— 1997, ‘The Equity premium and the Government Cost of Capital: A Response to Neville Hathaway’, *Agenda*, vol. 4, no. 4, pp. 475–85.

Richards, A. 1991, ‘The Cost of Equity Capital In Australia: What Can We Learn From International Equity Returns?’, *Reserve Bank Research Discussion Paper,* no. 9107, September, Economic Analysis Department, Reserve Bank of Australia.

Rosen, H. 1999, *Public Finance*. 5th edition, Irwin McGraw-Hill.

Sandmo, A. 1972, 'Discount Rates for Public Investment Under Uncertainty', *International Economic Review*, vol. 13, no. 2, June, pp. 287–302.

Sandmo, A. and Dreze, J. 1971, 'Discount Rates for Public Investment in Closed and Open Economies', *Economica,* vol. 38, pp. 395-412.

Schelling, T. 1987, 'Value of Life', Eatwelkl, J., Milgate, M. and Newman, P. (eds), *The New Palgrave World of Economics,* Macmillan Press, pp. 706-12

Siegel, Jeremy J. 2005, Perspectives on the Equity Risk Premium,*Financial Analysts Journal*, vol. 61, no. 6, pp. 61-71, November/December, http://ssrn.com/abstract=872783

Siegel, J. and Thaler, R. 1997, 'Anomalies The Equity Premium Puzzle' *Journal of Economic Perspectives*, vol. 11, no. 1, Winter, pp. 191–200.

Sieper, E. 1981, 'The Choice of Discount Rate in the Assessment of Damages', Mimeo, Departments of Economics, Australian National University, July.

Sjaastad, L. and Wisecarver, D. 1977, 'The Social Cost of Public Finance', *Journal of Political Economy*, vol. 85, no. 3, June, pp. 513–547.

South African Department of Environmental Affairs and Tourism. 2004, C*ost Benefit Analysis*, Integrated Environmental Management Information Series 8, Pretoria, http://www.deat.gov.za

South Australian Department of Treasury and Finance. 2007, *'Partnerships SA' Private Sector Participation in the Provision of Public Services Guidelines for the Public Sector,* Projects and Government Enterprises, July.

Spackman, M. 2001, 'Risk and the Cost of Risk in the Comparison of Public and Private Financing of Public Services', *Topics,* no. 23, March, National Economic Research Associates.

—— 2004, 'Time Discounting and of the Cost of Capital in Government', *Fiscal Studies;* vol. 25, no. 4, December, pp.467–518.

—— 2006, 'Social Discount Rates For The European Union: An Overview', Università degli Studi di Milano, Dipartimento di Scienze Economiche, Aziendali e Statistiche, Working Paper no. 2006-33, October.

Stapleton, R. and Subrahmanyam, M. 1978, Capital market equilibrium in a mixed economy, optimal public sector investment decision rules, and the social rate of discount', *The Quarterly Journal of Economics,* August, pp. 399-411.

Stern, N.. 2007, *The Economics of Climate Change: The Stern Review*, Cabinet Office – HM Treasury, Cambridge University Press, UK.

—— 2008, 'The Economics of Climate Change', *American Economic Review: Papers & Proceedings,* vol. 98, no. 2, May, pp. 1–37.

Stiglitz, J. 1982, ‘The Rate of Discount for Benefit–Cost Analysis and the Theory of the Second Best’, in Lind (ed), pp. 151–204. Reprinted in Layard and Glaister (eds), pp.116–159.

Summers, L. 1990 ‘What is the social return to capital investment’, in *Growth, Productivity, Unemployment*, Peter A. Diamond (ed.), Cambridge, MA, MIT Press, pp. 113-141.

Sunstein, C. and Rowell, A.. 2007, ‘On Discounting Regulatory Benefits: Risk, Money, and Intergenerational Equity’, *University of Chicago Law Review*, vol. 74, pp. 171-208.

Swan, P. 2006, ‘Optimal Portfolio Balancing Under Conventional Preferences and Transaction Costs Explains the Equity Premium Puzzle’, mimeo, University of New South Wales, March 1.

Thompson, F. and Green, M. (eds) 1998, *Handbook of Public Finance,* Marcel Dekker, New York.

Treasury Board of Canada 1998, *Benefit-Cost Analysis Guide*. July, Ottawa.

—— 2007, *Canadian Cost-Benefit Analysis Guide Regulatory Proposals*, http://www.tbs-sct.gc.ca.

United States Environmental Protection Agency 2000, *Guidelines for Preparing Economic Analyses*, Office of the Administrator, Environmental Protection Agency, 240-R-00-003, September.

United States Office of Management and Budget 2003, *Regulatory Analysis*, Circular A–4, September.

Van Ewijk, C. and Tang, P. 2003, ‘How to Price the Risk of Public Investment?’, *De Economist,* vol. 151, no. 3, pp. 317-28.

Victorian Competition and Efficiency Commission. 2007, *Guidance Note on Discounting*, Melbourne

Victorian Department of Treasury and Finance. 2007, *Victorian Guide to Regulation,* 2nd edition, April.

Viscusi, K. 2007, ‘’Rational discounting for regulatory analysis’, *University of Chicago Law Review*, vol. 74, pp. 209-246.

____ 2006, ‘Monetizing the Benefits of Risk and Environmental Regulation’, A*EI-Brookings Joint Center For Regulatory Studies Working Paper* 06–09, April. http://aei-brookings.org/admin/authorpdfs/redirect-safely.php?fname=../pdffiles/phpzl.pdf.

Warner, J. and Pleeter, S. 2001, 'The Personal Discount Rate: Evidence from Military Downsizing Programs', *American Economic Review*, vol. 91, no. 1, March, pp. 33-53.

Warr. P. and Wright, B. 1981, 'The Isolation Paradox and the Discount Rate for Benefit–Cost Analysis', *The Quarterly Journal of Economics*, vol. 96, no. 1, pp. 129–145.

Watcher, J. 2008, 'Can time-varying risk of rare disasters explain aggregate stock market volatility?' mimeo, Department of Finance, The Wharton School, University of Pennsylvania, September 24.

Weimer, D. 2008 'Cost-Benefit Analysis', *The New Palgrave Dictionary of Economics*, 2nd edition, S. Durlauf and L. Blume (eds).

Weisbach, D. and Sunstein, C. 2008, 'Climate Change and Discounting the Future: A guide for the Perplexed', *Reg-Markets Center Working Paper no.08-19, University of Harvard Public Law Working Paper no.08-20.*

Weitzman, M. 2001, 'Gamma Discounting', *American Economic Review*, vol. 91, no. 1, March, pp. 260–271.

—— 2007, 'A Review of The Stern Review on the Economics of Climate Change', *Journal of Economic Literature*, vol. XLV, September, pp. 703–724.

—— 2007a, 'Subjective Expectations and Asset-Return Puzzles', *American Economic Review,* vol. 97, no. 4, pp. 1102-1130.

Weyant, J. 2008, 'A Critique of the *Stern Review's* Mitigation Cost Analyses and Integrated Assessment', *Review of Environmental Economics and Policy,* vol. 2, no. 1, Winter, pp. 77-93.

Wright, S., Mason, R. and Miles, D.2003, 'A Study into Certain Aspects of the Cost of Capital for the Regulated Utilities' February 13. Report commissioned by the United Kingdom economic regulators and the Office of Fair Trading. http://www.ofcom.org.uk/static/archive/Oftel/publications/pricing/2003/capt0203.pdf

Young, L. 2002, 'Determining the Discount Rate for Government Projects', *New Zealand Treasury Working Paper* 02/21.

Zhuang, J., Liang, Z. Lin, T. and De Guzman, F. 2007, 'Theory and Practice in the Choice of Social Discount Rate for Cost–Benefit Analysis: A Survey', *ERD Working Paper No. 94*, Asia Development Bank, May.