

Creedy, John; Passi, Hemant

Working Paper

Public Sector Discount Rates: A Comparison of Alternative Approaches

New Zealand Treasury Working Paper, No. 17/02

Provided in Cooperation with:

The Treasury, New Zealand Government

Suggested Citation: Creedy, John; Passi, Hemant (2017) : Public Sector Discount Rates: A Comparison of Alternative Approaches, New Zealand Treasury Working Paper, No. 17/02, ISBN 978-1-988534-14-5, New Zealand Government, The Treasury, Wellington

This Version is available at:

<https://hdl.handle.net/10419/205705>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by/4.0/>

Public Sector Discount Rates: A Comparison of Alternative Approaches

John Creedy and Hemant Passi

New Zealand Treasury Working Paper 17/02

2017



THE TREASURY

Kaitohutohu Kaupapa Rawa

New Zealand Government

DISCLAIMER

The views, opinions, findings, and conclusions or recommendations expressed in this Working Paper are strictly those of the author(s). They do not necessarily reflect the views of the New Zealand Treasury or the New Zealand Government. The New Zealand Treasury and the New Zealand Government take no responsibility for any errors or omissions in, or for the correctness of, the information contained in these working papers. The paper is presented not as policy, but with a view to inform and stimulate wider debate.

**NZ TREASURY WORKING
PAPER 17/02**

Public Sector Discount Rates: A Comparison of Alternative Approaches

MONTH/YEAR

June 2017

AUTHORS

John Creedy
Victoria University of Wellington
23 Lambton Street
Wellington 6011
New Zealand
Email John.creedy@vuw.ac.nz
Telephone [+64 4 4637422]

Hemant Passi
New Zealand Treasury
1 The Terrace
Wellington 6011
New Zealand
Email hemant.passi@treasury.govt.nz
Telephone +64 4 917 6305
Fax +64 4 890 6686

ISBN (ONLINE)

978-1-98-853414-5

URL

Treasury website at June 2017:
<http://www.treasury.govt.nz/publications/research-policy/wp/2017/17-02>
Persistent URL: <http://purl.oclc.org/nzt/1962.txt>

ACKNOWLEDGEMENTS

We should like to thank Girol Karacaoglu for supporting this work. We also benefited from helpful discussions with many colleagues in the New Zealand Treasury and a number of other government departments. In particular we should like to thank Kirsten Jensen, Dieter Katz, Joanne Leung, Chris Parker and John Yeabsley for their helpful comments on an earlier draft of this paper.

NZ TREASURY

New Zealand Treasury
PO Box 3724
Wellington 6008
NEW ZEALAND
Email information@treasury.govt.nz
Telephone 64-4-472 2733
Website www.treasury.govt.nz

Abstract

This paper sets out the alternative approaches to the public sector discount rate and explains the assumptions involved. There are two main ways of thinking about the discount rate. First, the social opportunity cost of capital approach (SOC) defines the discount rate as the rate of return that a decision-maker could earn on a hypothetical ‘next best alternative’ to a public investment. Second, the social rate of time preference approach (SRTP) defines the discount rate as the rate of return that a decision-maker requires in order to divert resources from use in the present, to a public investment. In an ‘ideal’ market, these two rates are brought into alignment in equilibrium. However, as there are no markets for public investments, there are no market signals to equate preferences for investing in such projects with rates of return. There is no completely objective way of determining public sector discount rates. Essentially the discount rate reflects how the government values the future when making decisions on behalf of society: value judgements and assumptions are necessary. The paper aims to clarify these judgements. Elements of both approaches may be relevant to many policy and operational decisions that require discounting, in which case different approaches may be relevant for different contexts. The paper also briefly considers the use of hyperbolic discounting.

JEL CLASSIFICATION

H43; H50

KEYWORDS

Public sector discount rate; social opportunity cost; social time preference rate.

Executive Summary

The public sector discount rate reflects how the government values outcomes that occur in the future relative to those that occur in the present. It is used across central and local government to ‘weight’ future costs and benefits when agencies carry out cost-benefit analysis (CBA), and to estimate the cost to the Crown of investing in public assets (the capital charge calculation). Despite many years of debate, there is no consensus on this topic, either in academic research or in policy guidance. This paper sets out the alternative approaches to discounting and explains the assumptions involved.

The discount rate can be interpreted as the minimum rate of return that the government expects from its investments. This gives two ways of thinking about the discount rate.

The social opportunity cost of capital approach (SOC) defines the discount rate as the rate of return that a decision-maker could earn on a hypothetical ‘next best alternative’ to a public investment

The social rate of time preference approach (SRTP) defines the discount rate as the rate of return that a decision-maker requires in order to divert resources from use in the present, to a public investment.

In theory, in an ‘ideal’ market, these two rates are brought into alignment in equilibrium. Rational decision-makers continue to invest as long as the rate of return that can be earned on public investments exceeds the rate demanded. However, as there are no markets for public investments, there are no market signals to equate preferences for investing in such projects with rates of return. This leaves the two alternative ways of approaching public sector discount rates introduced above.

SOC is typically measured by reference to the rate of return on private-sector investments with similar risk characteristics to the public project under consideration. Under this approach, the discount rate is composed of a risk-free rate of return plus a risk-based premium which varies according to the riskiness of the project. This is the basis for the NZ Treasury’s current advice.

SRTP is a direct statement of the decision-maker’s preferences for valuing the future. Loosely speaking, it states that the discount rate depends on any intrinsic preference for trading-off the future relative to the present (so-called ‘pure time preference’), and the extent to which decision-makers may wish to prioritise the present, if there is an expectation that living standards are expected to grow in the future.

There is no completely objective way of determining public sector discount rates. Essentially the discount rate reflects how the government values the future when making decisions on behalf of society. It is therefore natural to expect that value judgements and assumptions are necessary. First, a market-based SOC assumes that political decision-makers *should* trade-off the future for public investments in the same way that individuals and businesses do when making decisions about their own personal consumption and investment. However, it is possible that many individuals in their political roles as citizens might be more concerned about future social outcomes than is reflected in their decisions about their own personal consumption and investment.

Second, a market-based SOC assumes that the government cares about the same types of risk, and demands the same risk-premiums, as private investors. Risk-premiums in financial markets are determined by the volatility of individual asset returns, and how those returns are expected to vary relative to the investor's overall portfolio. This approach to measuring and pricing risk may not be meaningful in a public sector context. Furthermore, governments' ability to manage risk is likely to be different to that of private sector firms. These considerations mean that the risk-premium component of public sector discount rates could, depending on how the government evaluates and prices risk, be different from those implied by private sector rates of return.

Elements of both approaches may be relevant to many policy and operational decisions that require discounting, in which case a hybrid approach might be appropriate. Different approaches may be relevant for different contexts.

The standard model of discounting involves discounting future amounts at a constant proportional rate, however long the time horizon under consideration. This is known as constant exponential discounting. One concern that has been raised with this approach is that, even with a low discount rate, the weights that a decision-maker attaches to future time periods eventually converge toward zero. An alternative approach is to apply a progressively lower rate to net benefits in more distant time periods. This is known as hyperbolic discounting, and has the effect of scaling-up the weight attached to the more distant future relative to exponential discounting.

Table of Contents

Abstract	i
Executive Summary	ii
1 Introduction	1
2 Discounting and Present Values	3
2.1 The need to discount	3
2.2 The Present Value of a Benefit Stream	3
2.3 Discounting over long periods	4
3 Thinking About Discount Rates	5
3.1 Time preference.....	5
3.2 The optimal level of consumption and investment.....	6
3.3 Market rates of return.....	6
3.4 Public sector discount rates.....	7
4 The social opportunity cost of capital	8
4.1 The rationale for a SOC approach.....	8
4.2 Determination of SOC rate	9
4.3 The definition of risk.....	10
4.4 Current public sector discount rates in New Zealand	11
5 The social rate of time preference	12
5.1 Determination of the SRTP.....	12
5.2 The social rate of time preference	13
5.3 Specifying the parameters	14
6 Comparison of alternative approaches	15
6.1 Reflecting opportunity cost	16
6.2 Reflecting time preference.....	17
6.3 Public sector risk.....	18
6.4 Measurement and parameter values	19
6.5 International approaches	22
7 Time-varying discounting	23
7.1 Discounting over the long term	23
7.2 Hyperbolic discounting.....	23
7.3 Time inconsistency	24
8 Conclusions	25
References	26

Public Sector Discount Rates: A Comparison of Alternative Approaches

1 Introduction

The public sector discount rate reflects how the government values outcomes that occur in the future relative to those that occur in the present. It is used to value the future costs and benefits of public projects and to estimate the cost of investing in public assets (the capital charge calculation).¹ As a result, the discount rate plays an essential role in guiding public spending and investment decisions.

The public sector discount rate can be interpreted as the minimum rate of return that the government expects from its investments. As a result, it represents the benchmark against which the returns from public sector investments should be assessed. Too high or too low a discount rate could cause the government to make the 'wrong' investments. In particular, setting the discount rate too high could lead to the rejection of public spending initiatives that would otherwise be regarded as valuable.

Furthermore, a high discount rate, by lowering the weight that is attached to future outcomes, tilts decisions towards projects that deliver net benefits in the near term. Conversely, a low discount rate tilts decisions towards projects which deliver net benefits over the longer term.

The recommended discount rate is, not surprisingly, highly controversial. Its choice involves both technical issues and a range of value judgements. This paper provides a comparison of alternative approaches. The aim is to clarify alternative viewpoints and so provide the basis for a balanced debate about advice regarding public sector discount rates. There are likely to be many competing opinions, many of which may ultimately boil down to value judgements which cannot be reconciled objectively. The paper aims to distinguish where objective considerations can be used to make recommendations and where value judgements are involved.

Emphasis is given to the two main approaches to the public sector discount rate. These are the social rate of time preference approach and the social opportunity cost of capital approach. It also covers some other key considerations relevant to how public projects might be discounted, such as whether discount rates should be constant or whether they should gradually fall over the length of life of a project.

It is of course recognised that the question of the public sector discount rate and the total amount of government expenditure are in principle intimately related. For example, a lower

¹ In New Zealand, other discount rates are used across government, most notably, the risk-free rate for accounting valuations: see <http://www.treasury.govt.nz/publications/guidance/reporting/accounting/discount-rates>. These are used as the basis for valuing a number of public sector assets and liabilities; for example, insurance liabilities (particularly ACC's insurance claims liability), employee benefits (such as pensions), and the student loan book. These discount rates are outside the scope of this project.

discount rate may be expected to be associated with higher planned levels of current expenditure. In practice however, current and planned spending levels are determined by a complex political process which involves many considerations other than the discount rate. These broader considerations are covered by New Zealand's Fiscal Management Approach and are therefore not discussed here; see Lomax, McLoughlin and Udy (2016). The focus here is purely on the approaches taken to thinking about a public sector discount rate in the context of comparing cost-benefit profiles for the appraisal and ranking of public sector projects.

Similarly the need to attach a monetary value to costs and benefits in each period during the life of a public project, where market prices are not available to provide a guide or markets are subject to distortions, presents a huge challenge that has also given rise to an extensive literature. The measurement of costs and benefits and the choice of public sector discount rate are often conflated, in particular in dealing with risk and uncertainty. Also, it is sometimes argued that, where a discount rate is used that is lower than the market rate, the cost of a project should include an associated opportunity cost. Furthermore, it is often suggested that the 'excess burden' of taxation needed to finance the project should be taken into account. Sometimes the problems of measuring future benefits are simply assumed away when discussing discounting. Again, while recognising these issues, given the aims of this paper, focus is largely on the discount rate alone, although the questions of risk and opportunity costs are necessarily discussed.

Finally, it is important to stress that the purpose of this paper is to clarify issues and provide a framework for discussion. It is not intended to propose any particular recommendations, or even reach unambiguous conclusions. As mentioned earlier, rational policy analysis requires clarification of the separate roles of value judgements and economic technicalities.

Section 2 provides a brief introduction to the need for discounting, along with the basic mechanics of discounting and computing present values. Section 3 outlines the basic intuition behind approaches to discounting. Sections 4 and 5 go on to outline the two main approaches to thinking about public sector discount rates. These are, as mentioned above, the social opportunity cost of capital, and the social rate of time preference approaches. Section 6 compares these two approaches and so 'sets the scene' for debate. Section 7 introduces the issue of time-varying discount rates. Section 8 concludes.

The public sector discount rate has been the subject of an extensive and often highly technical and controversial literature over a very long period.² This paper makes no pretence to be comprehensive or to provide a review of the vast literature: references to this literature are therefore highly selective. As mentioned earlier, its aim is to set out the issues in a way that can stimulate rational debate. While it is not possible to avoid some technicalities, every attempt has been made to make the discussion clear and widely accessible.

² A flavour of the controversy can be obtained from the debate following the Stern Report (2006) on climate change: see for example, Carter *et al.* (2006), Dasgupta (2006), Nordhaus (2006) and Varian (2006). For broader literature reviews, see Harrison (2010) and OECD (2015).

2 Discounting and Present Values

2.1 The need to discount

When appraising public projects a decision-maker is typically faced with decisions that deliver different profiles of costs and benefits over time. In such cases, the decision-maker typically cares not just about the absolute value of the costs and benefits, but also about *when* those costs and benefits materialise. For instance, a project that delivers a series of benefits in the near term is not usually considered to be the same as a project which has exactly the same pattern of costs, but delivers its benefits ten years later. Benefits that materialise sooner are usually considered to be more valuable than those that occur later.

How should alternative cost-benefit profiles be compared? One way of answering this question is to ask, 'What is the maximum amount that a decision-maker would pay now to secure the stream of net benefits available from a given project?' In other words, what is the *present value* of a given stream of net benefits? Discounting provides a way of answering this question by formally specifying the value that a decision-maker assigns *now* to outcomes that occur in the future. In doing so, discounting allows projects with different net benefit time profiles to be converted into single values with a common valuation date. These values can then be compared in order to reject proposals which do not yield a positive net benefit, and rank those that do.³

2.2 The Present Value of a Benefit Stream

It is therefore necessary to consider how the present valuation of any future cost or benefit depends on its timing. For example, how would the present valuation of \$100 received T years in the future vary with T ? As mentioned above, benefits that materialise sooner are usually considered to be more valuable than those that occur later (for reasons explored shortly).

At its simplest level, discounting amounts to little more than scaling-down the value of costs and benefits that occur in the future by a factor that increases with the length of the delay. The standard way of doing this is to scale the value of any cost or benefit down by a *constant* factor for each additional year into the future it is expected to occur. This constant factor is determined by the *discount rate*, and this approach is known as *constant exponential discounting*.

Suppose \$1 is invested, earning an interest rate of r per period. After one period it is therefore worth $\$(1+r)$. Hence it is only necessary to invest $\$1/(1+r)$ in order to receive \$1 after one period. The present value of \$1 to be received in one period's time is thus $\$1/(1+r)$. The same argument leads to the result that the present value of \$1 to be received in two period's time is $\$1/(1+r)^2$. Hence the present value, PV , of the net benefit stream, $x_0, x_1, x_2, x_3 \dots x_T$, is thus:

$$PV = x_0 + \frac{x_1}{1+r} + \frac{x_2}{(1+r)^2} + \dots + \frac{x_T}{(1+r)^T} \quad (1)$$

This can be interpreted as the maximum amount that the decision-maker would be willing to pay now to secure the net benefit stream under consideration. In other words, resolving

³ An axiomatic approach to discounting was first proposed by Koopmans (1960); see Creedy and Guest (2008) for a simplified exposition. Without discounting, it is also not possible to compare infinitely lived alternative projects: however, on partial orderings using an 'overtaking' criterion, see von Weizacker (1965) and Heal (1998).

net benefit streams into their PVs can be used to compare and rank alternative cost/benefit profiles.


2.3 Discounting over long periods

For any given discount rate, r , there is thus a *discount factor* for any period, t , equal to $1/(1+r)^t$. This obviously declines as t increases. Figure 1 demonstrates how discount factors decline, under constant exponential discounting for a range of discount rates, as the time period increases. The rate of decline is clearly much more rapid for higher discount rates. For example, under a 5% discount rate, \$100m arising 50 years from now has a present value of roughly \$9m.

A concern that has been raised with constant exponential discounting is that, even with a low discount rate, the weights that the decision-maker attaches to future time periods eventually become very small. In practical terms, this means that decision-maker effectively ‘stops caring’ about outcomes in the future, provided these outcomes are distant enough.⁴

In response to this concern, some economists have proposed using time-varying or *hyperbolic* discount rates, where the discount rate is progressively reduced as the time when net benefits are received increases. This topic is considered in section 7.

Figure 1 Discount Factors for Alternative Discount Rates and Time Periods



⁴ It is convenient to think in terms of the ‘half life’, defined as the time taken for the value in that year to be reduced by half. Where r is the discount rate, the half life is equal to $\log 2 / \log(1+r)$. For example, annual rates of interest of 2, 5 and 10% have half lives of respectively 35, 14.2 and 7.3 years.

3 Thinking About Discount Rates

It has been seen that the discount rate captures how a decision-maker trades-off future versus present benefits. This may seem like a difficult concept to observe or measure, but in fact people make this type of trade-off all the time. When deciding whether to spend now or to invest, people have to balance their own individual preferences for consumption in the present with the fact that investment is productive, and can therefore generate more consumption for the future.

By making this type of decision, people implicitly reveal how they discount the future. This section shows that, in equilibrium, the discount rate can be thought of as the rate which just balances individuals' preferences for investment with the rate of return on investment. Although it involves strong assumptions, this is a key insight as it provides the basis for the two main ways of thinking about the discount rate. These are:

- opportunity cost of capital approaches – based on observing market rates of return as a measure of how people *can* trade-off consumption in the future versus the present
- time preference approaches – based on peoples' preferences for how they *wish* to trade-off future versus present consumption

It will be seen that in a large competitive market with no distortions, these two approaches are equated in equilibrium. That is, the rate at which people wish to trade-off their own present and future consumption will equal the rate at which markets allow them to shift consumption through time via investment. As long as these conditions hold, both approaches should therefore give the same discount rate. This provides the basis for setting private sector discount rates by reference to market rates of return (that is, on an opportunity cost of capital basis).

However, it is unlikely that these conditions hold for public sector projects. As a result, the equivalence between time preference and opportunity cost approaches can break down. It is then necessary to choose which of the two approaches is most appropriate for determining *public sector* discount rates, or whether they can be combined in some way.

3.1 Time preference

When considering whether to spend money now or invest for the future, individuals typically postpone a dollar's worth of consumption now only if they are compensated with more than a dollar of consumption in the future. The rate at which an individual needs to be compensated for postponing consumption is known as the *rate of time preference*.⁵

⁵ The term *consumption* is used frequently in the following sections. It is important to stress that consumption should be interpreted more broadly than the everyday material sense of the word. Specifically, consumption should be taken to mean the use of currently available resources to deliver a benefit, *of any kind*, in the *present*. This includes any immediate social or non-material benefit that people value. For example, consumption could refer to the use of current resources to deliver social services which generate benefits in the present. Consumption of any durable good occurs over the lifetime of that good, not simply when it is purchased.

This is usually specified in percentage terms, in the same way as an interest rate.⁶ It varies:

- for a given individual according to how much that individual is consuming/investing at any given time, and
- from individual to individual according to personal preferences.

As a result, time preference captures how individuals are willing to trade-off present and future consumption. The fact that individuals typically require a positive rate of return to postpone consumption amounts to saying that individuals value a dollar now more than a dollar in the future. In other words, people behave as if they ‘discounted’ the value of amounts received in the future.

3.2 The optimal level of consumption and investment

Spending one’s income now means sacrificing the opportunity to invest and yield more consumption potential for the future. As a result, individuals must balance the preference to bring consumption forward against the fact that investment is productive. A rational decision-maker manages this trade-off by continuing to invest an additional dollar of income as long as the rate of return from that investment more than compensates for the sacrifice of present consumption.

It is typically assumed that the more an individual invests now, the more unwilling that person is to sacrifice further, and therefore the greater the required increase in future consumption that is needed as compensation. In other words, the rate of compensation that an individual requires to make an investment is assumed to increase the more that is invested now.

This means that an individual continues to invest until the rate of return obtained from the last unit of investment is just sufficient to compensate for the corresponding loss in present consumption. Or, stated differently, rational decision-makers continue to invest until their willingness to trade present for future *consumption* is equated to the rate of return that can be earned on their *investment*.

3.3 Market rates of return

For many individuals interacting in a competitive market economy, they can be thought of as borrowing and lending to plan their own consumption over time such that in equilibrium *each* individual’s willingness to trade consumption over time is equated to the rate of return available in the market.

Moreover, in the economy as a whole, market rates of return themselves vary according to the aggregate amount of capital supplied by individuals. It is usually assumed that there are decreasing returns to capital, so that the rate of return on investment declines the more that is invested. Therefore in a well-functioning market, the rate at which individuals are willing to trade consumption over time is jointly determined along with the market rate of return.

This last statement provides an important insight into how to think about the discount rate. Specifically, it states that (in a well-functioning market), the rate at which individuals wish to

⁶ For example, a rate of time preference of 5% means that in order to convince an individual to delay \$100 of consumption by a year, that individual would need to be offered (at least) \$105 in a year’s time. It should also be noted that time preference need not always be positive. If people expect little or no income in the future, they may be willing to save at zero or even negative rates of return. However, this discussion is concerned with the general case.

trade-off consumption over time is equal to the rate at which the market allows them to make this trade-off. Stated differently, in equilibrium, market rates of return indicate:

- the rate of return on productive capital that is available to individuals in the market (that is, the opportunity cost of capital), and
- the minimum rate of return that individuals require to sacrifice a unit of present consumption (that is, time preference).

If presented with the option to finance a new investment project, this reasoning indicates that the rate of return on that project should be no less than the market rate of return. Otherwise the rate of return on the new project would be inferior to the rate of return on existing investment possibilities and not exceed the cost to investors of sacrificing present consumption.

This reasoning therefore provides the basis for discounting at market rates of return. It states that discount rates should be set by reference to *both* the rate of return that is demanded by individuals to postpone consumption *and* the rate of return that can be earned by individuals on a comparable investment project. In a well-functioning market with no distortions, these two rates are expected to be equal in equilibrium.

This provides a convincing rationale for why market rates of return are used to discount *private* decisions about how individuals should invest. Indeed, it provides the basis for how companies discount their investments (see section 4). However, it is not clear that this rationale necessarily extends to investments by the public sector.

3.4 Public sector discount rates

By definition, there are no large competitive markets for public sector investments. Public sector projects are typically very different from private projects. They are likely to give rise to externalities and have distributional implications. Therefore there is no mechanism to elicit peoples' time preferences for public sector projects. Nor are there reliable ways to estimate the rates of return on a broad range of public sector investments (financial and social-sector) for the purposes of setting opportunity costs.⁷ In addition, public projects must be financed by either taxation or debt. In both cases a present transfer of real resources is made from private to public sectors. In addition, taxation (and possibly debt – to the extent that it is simply delayed taxation) involves deadweight losses that should be included as a cost of undertaking public projects.⁸

Despite the fact that there are no well-functioning markets for public sector projects, the considerations for thinking about the discount rate are the same. Namely:

- What is the rate of return that a decision-maker could earn on a hypothetical next best available alternative? That is, what is the opportunity cost of a public project?
- What rate of return would a decision-maker require to sacrifice a unit of present consumption in order to invest in a public project?

⁷ Furthermore, some public sector projects may be sufficiently large for general equilibrium effects to arise.

⁸ That is a tax, by altering the prices at which economic transactions would otherwise be traded, typically causes some consumers to be less willing to buy, and some producers to be less willing to sell, than in the absence of the tax. This distortion gives rise to a 'deadweight loss' or 'excess burden' of taxation, and should be regarded as a cost of public financing. As this concerns the measurement of costs, rather than problem of setting the discount rate, it is not discussed further here.

In principle, these two rates of return would also be equated in equilibrium, as a public sector decision-maker would continue to invest until their willingness to trade present for future consumption is equated to the rate of return that can be earned on their investment. However, in the absence of markets for public projects, it is not possible to rely on the usual market signals. Furthermore, even in 'perfect' markets, equilibrium prices have efficiency properties but are not necessarily associated with 'optimal' outcomes, particularly as the latter depend on distributional considerations.

As a result, there are two different ways of thinking about public sector discount rates. These are:

- The social opportunity cost of capital approach, and
- The social rate of time preference approach

These two approaches are considered in turn in the following two sections.

4 The social opportunity cost of capital

The social opportunity cost of capital (SOC) approach takes the view that public projects should be discounted by reference to the rate of return that could be earned from the *next best alternative* use of public funds.

This next best alternative is usually taken to be a private sector project with similar risk characteristics to the public project under consideration. In other words, private sector rates of return are usually considered to be the relevant measure of *opportunity cost*. This is the approach that the New Zealand Treasury currently uses to recommend discount rates.

4.1 The rationale for a SOC approach

The rationale for the SOC approach is based on the idea that the return on public projects must at least meet the 'hurdle' of the next best rate of return available to the public. Otherwise in principle an improvement could be made by investing public funds at the higher rate of return available in the private sector (or letting individuals do this themselves), and then distributing the proceeds from that investment to society. In New Zealand, one way of achieving this, for example, would be to increase the government's investment in the Superannuation Fund. As a result, the SOC plays an important role in disciplining public sector investment.

Another important property of the SOC is that market rates of return reflect not only the return that can be earned on market investments, but also the return that is demanded by individuals for such investments.⁹ That is, as mentioned above, in large well-functioning markets, the SOC reflects how the market as a whole balances individuals' preferences for trading-off the future with the productive potential of investment.

However, this also means that taking a market-based SOC as the basis for public sector discount rates takes the view that the government *should* trade-off the future for public sector investments in the same way that individuals and businesses do when making decisions about their own personal consumption and investment. This might be appropriate for investments that are expected to earn a return for present generations, but some people might argue that it is not appropriate for projects which have inter- and intra-generational

⁹ These two concepts are equated in equilibrium, as discussed in section 3.

impacts. The benefits arising from public projects may also be, by their nature, very difficult to value in money terms compared with private investments. The assumption that they can all easily be incorporated in the calculation of a project's benefits is a strong one.

In addition, the next best use of the funds might, in the absence of a public project, be to take on less risk and a lower return than a comparable private sector investment. Therefore, taking the traditional SOC view that the next best use of the funds would be to invest in a private project with a given risk profile imposes a key assumption about individuals' preferences for taking on risk in the absence of a government investment.

A SOC-based view also assumes that the government is concerned about the same types of risk, and prices risk in the same way as markets. These assumptions may not be appropriate for all government projects. For example, the risks associated with public projects may not be strongly correlated with those in the market.

As a result, a SOC-based approach is not, as is often assumed, a completely objective way of determining public sector discount rates. Rather, the decision to use a market-based SOC involves several implicit assumptions that need to be explicitly recognised. These issues are explored in further detail in section 6.

4.2 Determination of SOC rate

At a high level, SOC-based approaches use asset-pricing models to estimate the expected rate of return from a public sector project. This is carried out by benchmarking or comparing the public project against private sector projects or companies considered to have similar risk characteristics.

Given the choice to use a SOC-based approach, determining discount rates is then largely a technical exercise using well-established methods in the finance literature. A choice can be made between several different asset pricing models. These include the capital asset pricing model, arbitrage pricing theory, and multi-factor models.

This is essentially the approach that private companies use to estimate their discount rates and cost of capital. This is also the basis for the NZ Treasury's current approach to recommending public sector discount rates, using the Capital Asset Pricing Model (CAPM). The central insight of the CAPM is that, in a competitive market, the expected rate of return on *any* asset is equal to the risk-free rate of return, plus an equity premium that varies in direct proportion to the riskiness of that asset. This insight is captured by the basic CAPM formula:

$$r = r_f + \beta(r_m - r_f)$$

Where:

- r is the rate of return required, that is, the opportunity cost of investing in a public project. In other words, it is the required discount rate.
- r_f is the *risk-free rate of return*. This is estimated by reference to the yields on long-term (10 year) government bonds.
- $r_m - r_f$ is the *equity-risk premium*. This is the difference between the average rate of return available in the stock market (r_m) and the risk-free rate (r_f).

- β is a measure of the riskiness of the private sector asset/company against which the government project is being benchmarked. It measures how sensitive the returns on the asset are (on average) to overall market returns.

Therefore the CAPM formula can be used to estimate the discount rate as follows:

1. Select a set of private sector projects/companies that are considered to have similar risk characteristics to the public sector project under consideration
2. Estimate the average beta of these companies/projects (discussed further below)
3. Input this estimated beta, along with estimates of the risk-free rate and the equity premium, into the CAPM formula.¹⁰

The result can be interpreted as the expected rate of return on a market-based portfolio with similar risk characteristics to the public project under consideration. In other words, it is the 'next best alternative' rate of return available to the public for a given risk appetite.

4.3 The definition of risk

It is important to stress that β measures risk in a precise sense. Usually, risk is understood to mean the variability of asset returns. However, if an investor chooses assets carefully, the portfolio can be balanced so as to offset at least some of the variability of individual assets while still yielding a positive expected rate of return on the portfolio overall. In other words, investors can diversify.

If diversification is approximately costless, all investors seek to balance their portfolios so as to eliminate all diversifiable risk. In other words, investors hold well-diversified portfolios in equilibrium.

This means that the only risk that investors bear is that part of an asset's variability that cannot be diversified away by adding it to a well-balanced portfolio. This is referred to as 'non-diversifiable' risk. This is what is measured by β , and it is estimated by regressing the returns of an individual asset on market returns. It can be interpreted as the sensitivity of the returns of an individual asset to overall market movements.¹¹

The key insight of the CAPM is that beta is the only notion of risk that is relevant to determining an asset's risk premium. In particular, the CAPM predicts that this risk premium takes a particularly simple form: it is proportional to beta.

However, this measure of risk does not automatically carry over well to public sector portfolios. It may not even be meaningful in a public sector context. A further issue is the way in which risks of public projects are correlated with the private market. This type of question is considered further in section 6.

¹⁰ In reality, some technical adjustments are made to the basic CAPM formula to account for the effects of taxation, inflation and the capital structure of the companies/projects being used for benchmarking. Further detail on the precise formula currently used to estimate public sector discount rates is available at:
<http://www.treasury.govt.nz/publications/guidance/planning/costbenefitanalysis/discount rates>

¹¹ For example, an asset with beta of 1.3 means that when the market rises or falls by an extra 1 per cent, *on average* the asset price will change by 1.3 per cent.

4.4 Current public sector discount rates in New Zealand

In New Zealand, the Treasury is responsible for advising central and local government agencies on the discount rate to be used for appraising public spending initiatives financed by general taxation. In addition to the default public sector discount rate (which is used to discount most government initiatives as well as for calculating departmental capital charge) Treasury also currently recommends specific discount rates for buildings, infrastructure and technology. These discount rates are set out in Table 1, and are specified in real pre-tax terms.

Table 1 Recommended Public Sector Discount Rates in New Zealand¹²

Category	Annual rate
Default rate	6.0
General purpose office and accommodation buildings	4.0
Infrastructure and special purpose (single-use) buildings:	6.0
<ul style="list-style-type: none">• Water and energy• Prisons• Hospitals• Hospital energy plants• Road and other transport projects	
Telecommunications, media and technology , IT and equipment, Knowledge economy (R&D)	7.0

They are set by reference to the rates of return that the government could hypothetically earn by investing public funds in a private sector project with similar risk characteristics to public investments. Treasury also allows for agencies to use project-specific discount rates, if these can be determined on clearly rationalised grounds for the case in hand.

The 6 per cent default rate is the rate used to appraise initiatives in CBAX – Treasury’s cost-benefit analysis tool. However, agencies can also conduct sensitivity testing using a 3 per cent discount rate, which reflects a risk-free rate of return set by reference to government bond yields.¹³

¹² New Zealand Treasury (2016).

¹³ For further details of how Treasury determines public sector discount rates, see NZ Treasury (2008) and NZ Treasury (2016).

5 The social rate of time preference

While the SOC attempts to specify the discount rate by trying to identify the rate of return on the *next best alternative* to a public project, the social rate of time preference (SRTP) takes the alternative route of attempting directly to specify preferences for trading-off the future when considering public projects.¹⁴

As mentioned above, there is no market for public projects, so there is no way of observing individual preferences for how to value the future in such cases. Nor is there any reliable way to survey and aggregate the preferences of individuals with regard to how time should be valued for public projects. As a result, the SRTP should not, despite the conventional language used, be interpreted as the preferences of 'society' in any aggregate sense. Rather, it is more appropriate to interpret the SRTP as reflecting the preferences of a decision-maker acting on behalf of society. The framework set out below can then be thought of as an attempt to formalise what considerations a rational decision-maker should take into account when setting the discount rate.

5.1 Determination of the SRTP

The most common way of approaching the SRTP is to consider how a rational decision-maker might optimise the consumption and investment path of the economy, supposing it could do so. Carrying out this thought experiment sheds light on what considerations such a decision-maker would take into account in setting a 'socially optimal' discount rate, and how those considerations depend on the preferences, or values, of the decision-maker.

This approach can be formalised by supposing that the decision-maker can choose a consumption and investment path for the economy so as to maximise some function that evaluates different consumption paths; that is, it attaches a numerical value or 'score' to each path. Specifically, suppose that the decision-maker takes the following steps.

- Assign a 'welfare score' to the level of consumption in each time period. This score reflects the decision-maker's preferences for consumption *within* that period. In particular, it is generally assumed that the decision-maker has some preference for smoothing consumption over time.¹⁵
- Discount these welfare scores (as described in section 2) according to the decision-maker's own specific rate of time preference. This reflects the decision-maker's own values about how future welfare should be valued relative to the present
- Sum these discounted scores together to formulate an aggregate measure of 'social welfare' (that is, an abstract measure of the overall score that the decision-maker assigns to a given consumption path). This is often referred to as the decision-maker's 'social welfare function'.

¹⁴ For further discussion and references, see Creedy and Guest (2008).

¹⁵ More accurately, it is assumed that the decision-maker has a *decreasing marginal valuation* for consumption in any given period. That is, the greater is consumption, the less that additional increments to consumption add to the decision-maker's score. This implies consumption smoothing, as the more a decision-maker reallocates from one year to another, the more unwilling they are to do so further. Thus the decision-maker is somewhat averse to unequal consumption streams over time. The degree of this aversion depends on the specification of preferences.

This evaluation or social welfare function is therefore considered to be additive (the marginal benefit from increasing consumption in one period does not depend on consumption in other periods). It can be written formally as follows:

$$W(C) = U(c_0) + \frac{U(c_1)}{1 + \rho} + \frac{U(c_2)}{(1 + \rho)^2} + \dots + \frac{U(c_T)}{(1 + \rho)^T}$$

Where:

- $W(C)$ represents the aggregate ‘social welfare’ that the decision-maker seeks to maximise by choosing the consumption profile $C = (c_0, c_1, \dots, c_T)$.
- $U(c_t)$ is the function that the decision-maker uses to assign a ‘welfare score’ to the level of consumption (c_t) in any given year t .¹⁶
- ρ represents the decision-maker’s *pure rate of time preference*. It is important to stress that this parameter is different from the discount rate that we are seeking to specify. Namely this parameter ρ is used to discount *welfare scores*, whereas the discount rate that we are seeking to specify is used to discount *cash flows*. Nonetheless, ρ is a key *determinant* of the overall discount rate.
- T is the length of the decision-maker’s evaluation horizon.

Importantly, the above expression differs from the present value formula given above. The term $W(C)$ is the present value of the stream of $U(c)$ values, discounted at the pure time preference rate.

The term, ρ , in particular attracts a great deal of controversy, as it implies that the decision-maker values future generations less than the present. As this is a crucial value judgement, it is not surprising that views differ regarding the value of ρ that ‘should’ be imposed for public projects, and that these views are often expressed in strong terms. This parameter is discussed further in section 6 (particularly box 1).

5.2 The social rate of time preference

Given the form of objective – the welfare function given above - the decision-maker is then assumed to choose a consumption and investment path for the economy so as to maximise the value of $W(C)$ subject to the constraints of available resources, productivity, and so on. A further important assumption is made about the weighting function $U(c_t)$: this is considered to be ‘iso-elastic’. That is, the elasticity of U with respect to c is assumed to take the constant value, say θ . The maximisation problem can be solved to show that money values, that is the consumption values, c , rather than the ‘welfare’ values, U , are discounted at the rate, r , where:

$$r = \rho + \theta g$$

and:

- ρ is the decision-maker’s pure rate of time preference (as defined above)

¹⁶ In the context of individual multi-period optimisation, U represents a utility function. In the present context it is instead a cardinal weighting function. In the individual context, the term, ρ , is referred to as a ‘utility discount rate’.

- g is the annual growth rate of per capita consumption, generally assumed to be constant.
- θ is a term that reflects the nature of the decision maker's preferences with respect to consumption smoothing. As mentioned above it represents the (assumed to be) constant elasticity of U with respect to variations in c . It is a parameter of the welfare scoring function, U . Loosely speaking, it captures how averse the decision maker is to unequal consumption streams across time, or its 'aversion to intertemporal inequality'.

Then r is the *social rate of time preference*. In cost-benefit analyses, the present value (in money terms) of a stream of annual net benefits (interpreted here as the stream, c_t) is evaluated using this social time preference rate, r .¹⁷

The equation above is known as the 'Ramsey equation', following Ramsey (1928).¹⁸ In the last two decades, a number of countries have moved to using this equation as the basis for setting public sector discount rates according to a social rate of time preference approach: see section 6 for a summary of international approaches.

The equation states that the socially efficient consumption discount rate is equal to:

- the decision-maker's rate of pure time preference, ρ ,
- plus a 'wealth effect'. This reflects the reasoning that if the decision-maker cares about equalising consumption over time, and per capita consumption is expected to grow over time, future outcomes should be discounted due to the fact that people in the future enjoy higher living standards. The size of this wealth discount effect depends on the decision-maker's aversion to intertemporal inequality, as captured by θ .

This implies that even if the decision-maker chooses to value the *welfare* of all time-periods and generations equally (that is, chooses a rate of pure time preference equal to zero) there are nevertheless other reasons to discount cash flows of consumption in the future because of the wealth effect.¹⁹

5.3 Specifying the parameters

The SRTP approach involves a number of assumptions and value judgements, particularly in relation to the choice of variables in the Ramsey equation. These value judgements need to be made as explicit as possible. Many commentators attempt to set the values of ρ and θ by reference to observed individual behaviour. However, these approaches involve a number of strong assumptions, and also make the implicit assumption that what *is*, is an appropriate signal of what the relevant values *should be*.

Therefore it is more appropriate to understand the implications of different choices, and treat these decisions as value judgements that have to be made openly. How this might be possible is considered in section 6. Table 2 illustrates some of the parameter choices that have been used internationally.

¹⁷ It is tempting to think that discounting U using ρ is equivalent to discounting c using r . However, the two are not necessarily equivalent.

¹⁸ It plays an important part in optimal growth models involving a 'representative agent': see, for example, Blanchard and Fischer (1989) and Barro and Sala-i-Martin (1995).

¹⁹ Some economists argue for the inclusion of an additional term in the Ramsey equation to account for the chance that there will be some catastrophe event so devastating that the outcomes of the policy become irrelevant. In reality, this term is difficult to quantify, and is likely to be small enough to ignore.

Table 2 Percentage Discount Rates used with the Ramsey Equation

Country	ρ	θ	g	SRTP
UK baseline public sector discount rate	1.5	1	2	3.5
France baseline public sector discount rate ²⁰	1	2	1.5	4
Stern climate change review	0.1	1	1.3	1.4
Nordhaus critique of the Stern review	1.5	2	2	5.5
Harmonised European Approaches for Transport Costing	1.5	1	1.5	3

Sources: HM Treasury (2003), Cropper et al (2014), Stern (2007), Nordhaus (2007), HEATCO (2006)

6 Comparison of alternative approaches

Sections 4 and 5 have looked at two seemingly very different approaches to determining the discount rate. Under the social opportunity cost of capital (SOC) approach, the discount rate is interpreted as the rate of return that the government foregoes when it invests public funds on behalf of society. Under the social rate of time preference (SRTP) approach, the discount rate is interpreted as the rate of return required by a ‘socially-minded’ decision-maker in order to defer a unit of consumption from the present to the future.

This section attempts to weigh the advantages and disadvantages of these two approaches with reference to choices faced by governments. A crucial point to stress is that neither approach offers a completely objective way of determining public sector discount rates. Value judgements are inevitable.

SOC-based approaches can seem appealing, as basing discount rates on observable market returns appears to be more objective than approaches based on SRTP. However, the very decision to select a SOC-based approach carries a number of implicit assumptions and value judgements. SRTP-based approaches require more transparent statements of the decision-maker’s value judgements.

The following discussion is organised into four subsections. The first two subsections seek to assess how SOC and SRTP compare in relation to the way they reflect the government’s opportunity cost and the government’s time preference rate. Accounting for public sector risk and the overall ease and practicality of measurement are discussed in the next two subsections. This is followed by a brief summary and a selection of international approaches.

²⁰ The French guidance does not explicitly state the values used to derive its public sector discount rate: the 4 per cent baseline discount rate was chosen as a central value.

6.1 Reflecting opportunity cost

The key rationale for the SOC approach is that if the return on public projects does not at least meet the hurdle of the next best rate of return available to the government, then government investment displaces, or crowds-out, an investment that would have generated more overall value.

Under the New Zealand Treasury's current approach, share market returns are considered to be the most appropriate measure of the next best alternative to the government. This is because the companies that constitute the market face incentives to carry out the most productive investments in the economy (both locally and overseas), and these returns are available to the public. Moreover, the government does in fact invest in the share market through the New Zealand Superannuation Fund. It could choose to substitute between this investment and other public projects.

However, this approach to defining the SOC assumes that, in the absence of undertaking a public sector project, the next-best use of the funds would be to invest in a private sector project of equal magnitude and risk. In other words, it assumes that a public sector project fully crowds-out or displaces a private sector investment of the same cost and risk profile. This is a key judgement. It is clearly the appropriate counterfactual for *private sector* investment decisions. However, in most cases the appropriate counterfactual for setting the government's opportunity cost is likely to be one of the following two other possibilities.

1. *The funds for the public project would never have been raised.* In this case the next best use of the funds would have involved a mixture of consumption and investment by private individuals and businesses, most likely with a lower risk-return profile. Therefore, to the extent that private consumption (as opposed to investment only) has been displaced as a result of the public project, the relevant measure of opportunity cost should also reflect the rate at which private individuals are willing to trade-off current and future consumption by saving. This rate is far from clear, but it would generally be expected to be lower than the opportunity cost of a risky private investment.²¹ Therefore, allowing for the fact that consumption (and not just investment) is displaced by the financing of a public project would tend to lower the opportunity cost of capital relative to the current approach.
2. *The funds would have been invested in an alternative public project.* In the case of core public service delivery, it is unlikely that the government would consider not raising the funds as a feasible counterfactual. In most cases it is also unlikely that the government would be willing freely to substitute between public service delivery and investment in the Superannuation Fund (or any other private investments). In this case, the appropriate measure of opportunity cost would be defined by reference to an alternative method of delivering a comparable service. This may be observable for some projects for which there are well-defined private alternatives (for instance communications networks). However, this is unlikely to be observable for most social-sector projects. For such cases, setting the SOC by reference to share market returns is a strong assumption.

As a result, a SOC approach based on share market returns does not provide an objective measure of the opportunity cost of public projects for all cases.

However, an SRTP approach does not account for the fact that raising public funds crowds out at least some private investment. For instance, supposing a public sector project costs \$100m and that half of this money would have been allocated towards private investment

²¹ It is often set by reference to some measure of government bond yields as a proxy for a risk-free rate of saving.

in the absence of the public project, so that the remaining half displaces private consumption. The relevant opportunity cost of the public project is:

- \$50m of present consumption which is displaced, *and*
- The present value, in consumption terms, of the remaining \$50m that would otherwise have been invested at a private rate of return. That is, the present value of the \$50m invested at a private rate of return, but discounted at the SRTP.

Therefore, an SRTP approach should take account of the fact that a public project displaces not only present consumption, but also private investment that would have generated a *stream of future consumption*. The present value of such consumption streams that are displaced should be recognised as a cost of the public project when conducting the cost-benefit analysis (CBA). This is known as the *shadow price of capital*. Estimating the shadow price of capital raises additional challenges, but methods have been proposed in the CBA literature.²²

In summary, a SOC based on share market returns tends to overestimate social opportunity cost, as it assumes full crowding-out. On the other hand, a SRTP approach is likely to underestimate this cost, as it assumes no crowding-out of private investment projects. A pragmatic solution might involve either of the following approaches.

- Take a weighted average of a market-based SOC and SRTP. The weights can depend on the proportion of investment and consumption in the counterfactual case where resources are left in private hands. However the assignment of these weights is not a simple matter, and such an approach may yield an intermediate number which does not cope well with addressing the variability of individual projects.
- Apply a SOC-based discount rate in cases where the main effect of a proposal is to displace or alter the use of capital in the private sector, and apply an SRTP-based discount rate when a proposal primarily and directly displaces private consumption.

6.2 Reflecting time preference

The SRTP approach to determining discount rates is a direct way of trying to think about the considerations that are relevant to determining a decision-maker's time preference. However, actually specifying these considerations requires subjective decisions about somewhat abstract concepts. These issues are discussed below when looking at the measurability of the two approaches. Furthermore, although the underlying theory used to formulate the version of SRTP presented above is widely used and cited, it involves a number of assumptions which need to be scrutinised.

On the other hand, using a market-based SOC to determine public sector discount rates (whether this is based on share market returns or risk-free rates, or a combination of the two) assumes that market rates of return contain all relevant preferences about how the future should be traded-off versus the present in all cases. That is, a market-based SOC imposes the value judgement that political decision makers *should* trade-off the future for public investments in the same way that individuals or businesses do when making decisions about their own personal consumption and investment.

This might be appropriate for some government investments whose principal purpose is to earn a financial rate of return for society, for example the New Zealand Superannuation

²² See for example Boardman *et al.* (2006), Boardman *et al.* (2008), Lind (1990) and Parker (2011)

Fund, and some investments by state-owned enterprises. However, the relevant question to ask is whether individuals would trade-off the future for, say, publicly-funded health, education, arts and community projects in the same way that they do for decisions about their own personal consumption and investment. That is, would individuals demand the same rate of return in order to sacrifice a unit of present consumption for a social investment, as that which they demand for a private investment?

It is possible that individuals in their political roles as citizens might be more concerned about future social outcomes than is reflected in their day-to-day decisions about their own personal consumption and investment. Specifically, people may care about the wellbeing of current and future generations from both a financial and social justice perspective (involving intra-generational and inter-generational equity). These preferences are unlikely to be reflected in market-based instruments, the primary purpose of which is to shift private consumption through time. If this were the case, it would mean that individuals would be more willing to invest in public sector projects than is implied by market rates of return. That is, individuals would be willing to invest in public projects up to a point where the rates of return on public investments are lower than those observed in financial markets. This would mean that the discount rate implied by a market-based SOC would be too high for certain public projects, particularly in the social sector, as it would overlook preferences that would tend to lower the discount rate. An SRTP approach is a direct way of trying to capture these preferences.

A separate issue is whether government agencies are able to value all relevant costs and benefits, including social costs and benefits, when carrying out cost-benefit analyses (CBA) of policy initiatives. The current New Zealand Treasury approach assumes that this is the case.²³ As mentioned earlier, the question of whether all costs and benefits can be measured in money terms is outside the main focus of this project. However it is recognised that this is likely to be a strong assumption, and that if some of these benefits cannot be valued, or are otherwise systematically undervalued in CBA, then a dollar yield on a public sector project is not equivalent to a dollar yield on a private sector investment. Under these circumstances it is sometimes argued that it is appropriate to trade-off the future for social sector projects differently than for private sector projects. In other words, it is argued that the required rate of return for some public investments might be lower than that implied by a market-based SOC.

It is certainly important to know whether some types of cost and benefit might be systematically under- or over-estimated in CBA. However, it is not usually advisable to try to correct for these biases through the discount rate. The wrong assessment of costs and/or benefits could go in either direction, and they are likely to vary widely from project to project.

6.3 Public sector risk

Section 4 highlighted the point that private sector measures of risk may not automatically carry over to the public sector. In particular, the CAPM (and other market-based models of risk-premiums) assume that the decision-maker is willing to freely substitute between different assets in order to balance the risk and return of the overall portfolio in accordance with the decision-maker's appetite for risk. For example, this definition of risk means that governments are comfortable offsetting the risk of say a health intervention 'failing' with an

²³ The New Zealand Treasury's CBA guidance states, 'assuming all benefits have been valued correctly, we should be indifferent between one kind of benefit and another if their value is the same'.

improvement to communications infrastructure ‘succeeding’, and vice versa, as long as the expected net effect on the risk and return of the overall portfolio is favourable.

However, some people may argue that governments select public sector investments on the basis of their ability to deliver specific economic and social goals. Under this view, the relevant concept of portfolio could be defined more narrowly as the set of possible investments that contribute to the delivery of a particular economic or social goal. If this view is considered reasonable, governments might be willing to make investments so as to offset risks within any given portfolio, whilst being much less willing to substitute freely between portfolios.

Furthermore, there are reasons to expect that the public sector is better placed to manage risk than the private sector. For example, public sector risk is spread much more thinly than most private sector risk in the sense that, for any given project considered in isolation, the success or failure of that project does not affect the government’s ability to deliver its other economic and social commitments in the same way as it would for most businesses in the private sector. The government also has an implicit taxpayer guarantee, given its constitutional power to tax, and can issue highly liquid low-risk debt. These options are not usually available to private companies.

These points raise the question of whether taking a market-based SOC as the basis for recommending public sector discount rates is appropriate. Specifically, this approach assumes that market-based models for pricing risk are also applicable to how the government prices risk. Again, this amounts to imposing a potentially significant assumption, at least for some projects, which ought to be scrutinised.

The SRTP approach does not explicitly account for risk through the discount rate. Rather, it usually relies on separate risk analyses such as scenario analysis, sensitivity analysis and/or Monte Carlo simulations to take account of project-specific risk relating to particular cost items and benefits.²⁴ Whether this is the best way to account for the overall risk that is faced by the government is a question for debate.

6.4 Measurement and parameter values

The SRTP offers a formula expressed in terms of easily interpreted parameters. However, the question arises of how transparently this can be made operational? Doing so requires decision-makers to specify the values of their pure rate of time preference, ρ , and their aversion to intertemporal inequality, θ . Both of these values represent abstract constructs that cannot be easily gauged intuitively.

Consider the variable, θ , which reflects the decision-maker’s degree of aversion to unequal consumption streams across time, or the ‘aversion to intertemporal inequality’. It is a measure of the sacrifice that a decision-maker is willing to make for the purposes of redistributing consumption over time. It is related to the concavity of the welfare evaluation, or scoring function, $U(c)$.

Supposing consumption is expected to grow over time, people in the future will enjoy higher material living standards than those in the present. If the decision-maker has some preference for smoothing the living standards of society over time, one can carry out the

²⁴ That is, specify probabilities for uncertain outcomes and cash flows, then statistically sample from these distributions to obtain probability distributions and confidence intervals over the range of all possible outcomes.

thought experiment of what loss the decision-maker would be willing to tolerate when making a redistribution across time.

An analogy can be made with using a 'leaky bucket' to transfer water from where it is plentiful to where it is scarce. How leaky a bucket would the decision-maker be willing to tolerate before judging the waste to imply that the benefit of the transfer exceeds its benefit? In the same way, suppose that a decision-maker can redistribute consumption over time, but only through a leaky budget envelope. That is, transferring consumption from one time period to another would involve some irretrievable loss or leak. How much of a leak would a decision-maker be willing to tolerate when making a transfer?²⁵

Attempting to answer this question can be used to gauge the meaning of different values of θ . For example, assuming annual per capita consumption growth of 2 per cent, setting $\theta=1$ (as assumed in the UK in the specification of SRTP) implies that the decision-maker would be willing to smooth consumption by taking \$1 from ten years in the future in return for \$0.82 for the present. In this way, examining the implications of different values can be used to gauge what choices most closely match the decision-maker's values.

Rather than trying to specify these variables directly, following the kind of thought experiment discussed above, some economists have tried to infer peoples' underlying preferences for trading-off time and inequality by reference to how people make actual decisions that involve making these trade-offs (that is, by using 'revealed preference' approaches). Examples include looking at savings and investment decisions and studying the degree of redistribution implied by income tax systems.

Whilst these studies can be used to provide some information about how individuals make the relevant trade-offs, they have very strong limitations. In particular, much of this research involves estimating individuals' preferences in contexts that are not always directly comparable with how individuals might be expected to behave with regards to public sector investment decisions.²⁶

This is essentially the same criticism discussed above in relation to SOC-based approaches. Namely that relying on the results of such studies assumes that decision-makers should trade-off the future for public decisions in the same way that individuals do for their own private decisions (or decisions made in other contexts that may not be strictly comparable).

As a result, this research may be used only to help guide a broader discussion about what parameter values are reasonable. A more appropriate way of doing this is to point out the implications of different parameter choices, and choose a parameter combination that most closely matches the decision-maker's social values.

The preceding sub-sections have suggested that that SOC appears to be easier to measure than SRTP. This discussion need not be repeated in detail here, except to restate the general conclusions that a first-best measure of SOC is not clearly observable in all cases and that the risk-premiums on private projects may not always be a good indicator of the risk-premiums required by the government.

²⁵ The 'leaky bucket' experiment was introduced by Atkinson (1970) in discussing the concept of inequality aversion. A somewhat different approach was used by Okun (1975) who discussed transfers between *groups* of individuals. For its use in the present context, see Creedy (2007).

²⁶ The alternative approaches to 'measurement' are critically examined by Creedy (2007).

The longest maturity on New Zealand government bonds is approximately 20 years, and even in international markets, securities with horizons of more than 30 to 40 years are unlikely to be traded with sufficient volume and liquidity to provide a reliable ‘noise-free’ signal of how markets discount the long term, and a representative signal of how the long term is valued.

Therefore, rather than SOC being an entirely objective or technical exercise, taking the next-best alternative use of public funds to be a private sector investment with similar risk characteristics amounts to imposing a potentially significant assumption. As a result, there may be no clear benchmark from financial markets to help determine either the risk-free or the risk-based component of SOC. If this is the case, SRTP may be the only feasible approach for valuing outcomes in the very long term. The issue of discounting over long time horizons is considered further in section 7. A brief summary of the comparisons is given in Table 3.

Table 3 Comparison of the Two Approaches

1. Capturing opportunity cost	2. Capturing time preference
<ul style="list-style-type: none"> • The current market-based social opportunity cost (SOC) approach is likely to overestimate the public sector discount rate for many projects, as it assumes full crowding out. • A social rate of time preference (SRTP) approach that does not account for the shadow price of capital is likely to underestimate the public sector discount rate. 	<ul style="list-style-type: none"> • The current market-based SOC may overestimate the public sector discount rate if CBA fails to account for socially-motivated preferences. • SRTP is difficult to quantify in a transparent way.
3. Accounting for risk	4. Overall measurability
<ul style="list-style-type: none"> • The market-based SOC may overestimate the public sector discount rate, as governments may be better placed to handle risk. • A pure SRTP does not account for risk in the discount rate: risk is dealt with separately in CBA. 	<ul style="list-style-type: none"> • The SOC is observable for public projects that have a clear market alternative, but not all public projects • The SRTP requires an explicit statement of value judgements, and requires dealing with abstract theoretical constructs.

6.5 International approaches

Table 4 provides some information about approaches used in a number of other countries.

Table 4 International Approaches

Country	Method	Percentage real discount rate
Australia	No central recommendations. Recommendations vary widely by agency and state	
Canada	SRTP and SOC	8 for SOC 3 for SRTP
Denmark	SOC	4 for years 0-35 3 for years 36-70 2 for years 71+
France	SRTP	4 for projects 0-30 years 2 for years 31+
Germany	Unknown	3
Japan	SOC	4
Netherlands	SOC	5.5
Norway	SOC	4 for years 0-40 3 for years 40-75 2 for years 75+
Sweden	No cross-sector national guidance, but SRTP for transport sector	3.5 for transport sector
United States	SRTP and SOC	7 for projects whose main effect is to displace private capital (to approximate SOC) 3 for projects whose main effect is to displace private consumption (to approximate SRTP) Sensitivity testing at 1-3 encouraged for projects which have a significant intergenerational impact
United Kingdom	SRTP	3.5 for years 0-30 3 for years 31-75 Reducing to 1 over years 75-300+
New Zealand	SOC	6 default rate, plus variations for general purpose office and accommodation buildings (4), infrastructure (6) and technology (7)
Sources: Harrison (2010), Gollier and Hammit (2014), HM Treasury (2003), Johansson and Kristrom (2007), OECD (2015), Office of Management and Budget (2003), Mouter (2015), Treasury Board of Canada (2007)		

7 Time-varying discounting

The discussion has so far assumed that future costs and benefits are discounted at a constant proportionate rate, however long the time horizon under consideration. As explained in Section 2 above, this is referred to as *constant exponential discounting*.

One concern that has been raised with constant exponential discounting is that, even with a low discount rate, the compounding effect of applying a constant rate can cause decision-makers to attach a very low value to more distant future outcomes. Some people are concerned that this builds ‘short-termism’ into political decision-making.

This section outlines some of the characteristics of an alternative model of discounting, namely time-varying or hyperbolic discounting, which involves applying a progressively lower discount rate for more distant outcomes.

7.1 Discounting over the long term

Even with relatively low discount rates, constant exponential discounting causes the discount schedule – that is, the weights that the decision-maker attaches to future time periods – to converge towards zero for more distant net benefits: this is illustrated above in Figure 1. In practical terms, this means that decision-makers effectively ‘stop caring’ about outcomes in the future, provided these outcomes are distant enough.

As is evident from the Figure 1, discount *factors* of less than 10% are reached within timeframes that are reasonable for some public sector projects, even with relatively low discount rates.²⁷ In particular, constant exponential discounting can imply the following.

- It is not worth present generations incurring a small cost now to avoid potential costly events in the distant future.
- The outcomes of distant generations are discounted *relative to each other* in the same way that more proximate generations are discounted. Some people may argue that, beyond some time horizon, T , they have no particular preference for discounting the wellbeing of generation $T + 1$ much more than the wellbeing of generation T .²⁸

As mentioned in Section 6, there may not be any market instruments that can be used reliably to gauge rates of return in the distant future. That is, an opportunity cost of capital approach is likely to break down in the context of very long term projects.

7.2 Hyperbolic discounting

In response to these concerns, some people have proposed using time-varying discount rates whereby progressively lower rates are used as the time period at which net benefits are received becomes more distant. This approach is known as *hyperbolic discounting*. It has the effect of scaling-up the weight attached to the more distant future relative to exponential discounting. Clearly, this is distinct from the idea of using a lower exponential rate for longer-term projects.

²⁷ Specifically, discount factors of less than 10 per cent are reached within 80 years at a discount rate of 3 per cent; and within 50, 35 and 25 years for discount rates of 5, 7 and 10 per cent respectively.

²⁸ Although this view does not take into account the wealth effect in the Ramsey equation.

In the last decade or so, some countries – most notably the UK, France, Denmark and Norway – have adopted hyperbolic discount rates as their default guidance for public sector CBA.

Two main rationalisations have been suggested for hyperbolic discounting in the case of individual decisions. Both are based on uncertainty about the future.

First, it can be shown that if the discount rate is expected to be constant over time, but there is uncertainty over its precise value, then the term structure of discount rates declines over time.²⁹

Second, under a Ramsey SRTP formulation of the discount rate, suppose the following assumptions are introduced. There is uncertainty over the future growth rate of per capita consumption, g_t , and this uncertainty is positively correlated, so that there is some tendency for positive or negative consumption shocks to accumulate. In this case, it can be shown that the decision-maker's discount rate declines over time. The reason for this result is that, given a preference for consumption smoothing, the decision-maker cares more about a negative shock than an equal and opposite positive shock. This asymmetry introduces a precautionary effect which causes the decision-maker to place a greater weight on the future, and hence to lower the discount rate, relative to a situation of complete certainty. The longer the time horizon, the more time there is for positive and negative shocks to accumulate, and so the variance of potential outcomes increases with time. This increasing uncertainty raises the strength of the precautionary effect, causing decision makers to apply a declining discount rate schedule.

Other arguments relate to statements about basic value judgements, involving a compromise between the situation under which there may be said to be a 'dictatorship' of the present (constant exponential discounting) and a 'dictatorship' of the future (zero discounting).³⁰

7.3 Time inconsistency

One issue that arises under hyperbolic discounting is that it results in time-inconsistent decisions. Specifically, with hyperbolic discounting it is possible that a decision-maker may wish to change an investment decision taken in the past, even though the forecasts of the costs and benefits made at the date of the decision turn out to be correct.

This problem arises because when making a decision now that has impacts at some future time t , the decision-maker uses a long-term (low) discount rate to evaluate the impact between periods t and $t + 1$. However, when time t actually arrives, he or she will apply a short-term (high) discount rate to period $t + 1$. In other words, when the future arrives the decision-maker's relative valuation of benefits in subsequent periods is lower than when the decision was originally taken. This could cause the decision maker to want to change the decision from that point onwards.

In such cases, the decision-maker's wish to make a different decision is due solely to the passage of time, because hyperbolic discounting applies different discount rates according to when in time things are evaluated. Constant exponential discounting does not encounter this problem as it treats time consistently in the sense that the *relative* valuation of any two benefits separated by a given time interval does not depend on when they are evaluated.

²⁹ See, for example, Weitzman (2007).

³⁰ See Chichilnisky (1997) and Li and Lofgren (2000) for alternative specifications.

Time inconsistency raises an interesting question, as it effectively introduces two ‘versions’ of the decision maker. One is a forward-looking version who discounts more distant outcomes at a low rate, and the second is a more present-focussed person who discounts the immediate future at a high rate. This comes into conflict when the more distant future becomes the immediate future.

8 Conclusions

The purpose of this paper has been to explain the alternative ways of thinking about public sector discount rates, and thereby provide the basis for a rational debate. It has been shown that there are two main ways of thinking about public sector discount rates. These depend largely on the assumption regarding the way the resources would otherwise be used in the absence of the public project being considered. They are as follows.

- *The social opportunity cost of capital approach* – the rate of return that a decision-maker could earn on a hypothetical next best available alternative (or the opportunity cost of a public project)
- *The social rate of time preference* – the rate of return that a decision-maker requires to sacrifice a unit of present consumption in order to invest in a public project.

It has been stressed that neither approach offers an objective way of determining public sector discount rates. Value judgements are inevitable.

SOC-based approaches are traditionally regarded as being more objective, based as they are on observable market returns. However, the very decision to select a SOC-based approach carries a number of implicit assumptions. These include assumptions that market-based counterfactuals provide an appropriate counterfactual for public projects, public projects fully crowd-out private sector projects of equal magnitude, political decision-makers should trade-off the future in the same way that individuals and businesses do when making decisions about their own personal consumption and investment – or all relevant preferences are measurable and accounted for in CBA cash flows. Also it is assumed that the way markets evaluate and price risk are also how governments evaluate and price risk.

Furthermore, SRTP-based approaches require explicit statements of the decision-maker’s value judgements. Indeed, the necessity to specify these aspects explicitly, rather than concealing them, may be considered a strength of the approach. Only very limited information can be drawn from financial markets to help calibrate SRTP. Nevertheless, the determination of the appropriate discount rate involves complex concepts that are not easy to specify precisely.

It has been seen that there is no unambiguously clear answer to the question of what public sector discount rate to use. The choice depends on a complex range of technical judgements as well as value judgements, so it is not surprising that widespread agreement about the rate is very difficult, if not impossible, to achieve. After first making a decision regarding the general approach to be taken, many further questions remain involving orders of magnitude of crucial variables. It is hoped that the present paper can contribute to rational policy analysis by clarifying the essential features of alternative approaches and highlighting the central role of value judgements.

References

- Atkinson, A.B. (1970) On the measurement of inequality. *Journal of Economic Theory*, **2**, 244–63.
- Barro, R.J. and Sala-i-Martin, X. (1995) *Economic Growth*. New York: McGraw-Hill.
- Blanchard, O.J. and Fischer, S. (1989) *Lectures on Macroeconomics*. MIT Press, Cambridge, MA.
- Boardman, A., Greenberg, D., Vining, A. and Weimer, D. (2006) *Cost-benefit analysis: concepts and practice*. Prentice-Hall, New Jersey.
- Boardman, A., Moore, M. and Vining, A. (2008) *Social Discount Rates for Canada*. Available from Jdi-
legacy.econ.queensu.ca/Files/Conferences/PPPpapers/Moore%20conference%20paper.pdf
- Caplin, A. and Leahy, J.V. (2000) The social discount rate. *National Bureau of Economic Research Working Paper*, No. W7983. Cambridge, MA.
- Carter, R.M., de Freitas, C.R., Goklany, I.M., Holland, D., Lindzen, R.S., Byatt, I., Castles, I., Henderson, D., Lawson, N., McKittrick, R., Morris, J., Peacock, A., Robinson, C. and Skidelsky, R. (2006) The Stern Review: A dual critique. *World Economics*, **7**, 165–232.
- Chichilnisky, G. (1997) What is sustainable development? *Land Economics*, **73**, 467–91.
- Creedy, J. (2007) Policy evaluation, welfare weights and value judgements: a reminder. *Australian Journal of Economics*, **10**, 1-15.
- Creedy, J. and Guest, R. (2008) Discounting and the time preference rate, *Economic Record*, **84**, 109-127.
- Cropper, M.L., Freeman, M.C., Groom, B. and Pizer, W.A. (2014) Declining discount rates. *American Economic Review*, **104**, 538-43
- Dasgupta, P. (2006) *Comments on the Stern Review of Climate Change*. University of Cambridge, Cambridge, UK.
- Gollier, C. and Hammitt, J.K. (2014) The long-run discount rate controversy. *Annual Review of Resource Economics*, **6**, 273-295.
- Harrison, M. (2010) Valuing the Future: the social discount rate in cost-benefit analysis. Available from <http://www.pc.gov.au/research/supporting/cost-benefit-discount>
- Heal, G. (1998), *Valuing the Future: Economic Theory and Sustainability*. Columbia University Press, New York.
- HEATCO (2006) Developing Harmonised Approaches for Transport Costing and Project Assessment – Proposal for Harmonised Guidelines. Available from <http://www.transport-research.info/project/developing-harmonised-european-approaches-transport-costing-and-project-assessment>

- H.M. Treasury (2003) *The Green Book*. Available from <https://www.gov.uk/government/publications/the-green-book-appraisal-and-evaluation-in-central-government>
- Johansson, P.O. and Kristrom, B. (2015) *Cost-Benefit Analysis for Project Appraisal*. Cambridge: Cambridge University Press.
- Koopmans, T.C. (1960) Stationary ordinal utility and impatience. *Econometrica*, **28**, 287–309.
- Laibson, D. (1996) Hyperbolic discount functions, undersaving, and savings policy. *National Bureau of Economic Research Working Paper* No. 5635. NBER, Cambridge, MA.
- Mouter, N. (2015) Why do discount rates differ? Analyzing the differences between discounting policies for transport Cost-Benefit Analysis in five countries. MKBA working paper. Available from <http://www.mkba-informatie.nl/mkba-voorgevorderden/working-papers/mouter-why-do-discount-rates-differ-analyzing-differences-fi/>
- Li, C.-Z. and Lofgren, K.-G. (2000) Renewable resources and economic sustainability: a dynamic analysis with heterogenous time preferences. *Journal of Environmental Economics and Management*, **40**, 236–50.
- Lind, R. (1990) Reassessing the government's discount rate policy in light of new theory and data in a world economy with a high degree of capital mobility. *Journal of Environmental Economics and Management*, **18**, S8-S28.
- Lomax, N., McLoughlin, S. and Udy, B. (2016) The New Zealand Fiscal Management Approach. Available from http://www.nzae.org.nz/wp-content/uploads/2016/10/3228152_FINAL-Conference-Paper-Guide-to-NZs-FMA.pdf
- Nordhaus, W.D. (2006) The stern review on the economics of climate change. *National Bureau of Economic Research Working Paper* No. 12741. Cambridge, MA.
- Nordhaus, W.D. (2007) A review of the Stern Review on the Economics of Climate Change. *Journal of Economic Literature*, **45**, 686-702.
- NZ Treasury (2008) Public Sector Discount Rates for Cost Benefit Analysis. Available from <http://www.treasury.govt.nz/publications/guidance/planning/costbenefitanalysis/discountrates>
- NZ Treasury (2016) Current Discount Rates. Available from <http://www.treasury.govt.nz/publications/guidance/planning/costbenefitanalysis/currentdiscountrates>
- OECD (2015) Adapting Transport Policy to Climate Change: Carbon Valuation, Risk and Uncertainty. Available from <http://www.oecd.org/publications/adapting-transport-policy-to-climate-change-9789282107928-en.htm>
- Office of Management and Budget (2003) Circular A-4. Available from <https://www.transportation.gov/regulations/omb-circular-no-4>
- Okun, A.M. (1975) *Equality and Efficiency: The Big Trade-Off*. The Brookings Institution, Washington, DC.

- Parker, C. (2011) Valuing the future appropriately: Case for using the shadow price of capital in social discount rate policy. NZIER report for Auckland Council. Available from:
<http://www.aucklandcouncil.govt.nz/SiteCollectionDocuments/aboutcouncil/committees/economicdevelopmentforum/meetings/economicfmag20111011.pdf>
- Ramsey, F.P. (1928) A mathematical theory of saving. *Economic Journal*, **38**, 543–59.
- Stern, N. (2006) *The Economics of Climate Change*. Available from http://www.hm-treasury.gov.uk/independent_reviews/stern_review_economics_climate_change/sternreview_index.cfm.
- Stern, N. (2007) *The Economics of Climate Change: The Stern Review*.. Cambridge University Press, Cambridge and New York.
- Treasury Board of Canada (2007) Canadian Cost-Benefit Analysis Guide: Regulatory Proposals. Available from <http://publications.gc.ca/site/eng/456648/publication.html>
- Varian, H.L. (2006) Recalculating the costs of global climate change. *New York Times*, 14 December.
- Weitzman, M.L. (2007) Subjective expectations and asset-return puzzles. *American Economic Review*, **97**, 1102-1130.