How Should the New Zealand Government Discount Future Payoffs?

Arthur Grimes

Motu Economic and Public Policy Research

A discussion paper prepared for The Treasury

5 June 2023

Motu economic & public policy research

I thank Kirsten Jensen, Chris Parker, Andrew Coleman and colleagues at Motu for very helpful comments while preparing this discussion paper.

Executive Summary

This report responds to a Treasury request "to provide an independent expert perspective on New Zealand public sector discounting, specifically hyperbolic discounting. In this context and based on the provider's existing expert knowledge of the literature and international experiences, from the perspectives of theory, methodologies and applied policy:

(a) What are the advantages and disadvantage of implementing hyperbolic discounting, using different variants, compared with the exponential discounting currently used for applying discounting across time?

(b) Are there interactions between (i) the method for applying discounting across time (hyperbolic vs exponential) and (ii) the method for setting the discount rate (SRTP vs SOC) relevant to a preferred discounting regime?

(c) What are other advantages and disadvantages (such as SRTP vs SOC), and international trends and lessons for New Zealand, for the Treasury to consider in retaining or changing the current discounting regime for the New Zealand public sector to reflect long-term and intergenerational impacts?"

In providing a perspective, this report considers the rationales for using different formulations of the social discount rate (SDR), including reviewing and critiquing the literature related to the setting of the SDR and, in one instance, extending that literature. Key points in the paper are as follows:

- Treasury currently adopts a default real discount rate of 5% p.a. when evaluating public sector projects, with 6% being used for some sectors; robustness checks can incorporate a lower rate of 2%. The default discount rate incorporates both the risk-free discount rate plus a default risk premium. The discount rate is used in conjunction with an exponential discounting formula.
- There are two common approaches to determining the discount rate: the 'social opportunity cost of capital' (SOC) approach and the 'social rate of time preference' (SRTP) approach. The former bases the public sector discount rate on the observed return on the next best alternative investment that has the same degree of risk. The latter is based on a social welfare function defined over present and future utilities.
- With exponential discounting, use of a 5% discount rate when evaluating a public sector project translates into an implicit decision to treat $1 of benefit today as worth more than $100 of benefit that occurs in a century's time. Similarly, $9 of benefit today is worth more than a $100 of benefit that occurs in 50 years' time.
- For projects in which all payoffs (costs and benefits) are valued in the market, government should use the same exponential discount rate (and hence the same discount factor) as does the private sector when undertaking competing activities (including the same risk premium for a similar project); otherwise, the choice of (a lower) discount rate would artificially bias production away from private towards public producers.
- It follows that the SRTP approach applies only to situations for which there are non-market payoffs. Returns to future generations, who are not represented in today's market, can be considered as one form of non-market payoff.
- Hyperbolic (and quasi-hyperbolic) discounting applications often employ a lower discount factor (higher discount rate) for near-term payoffs and then have a higher discount factor for longer term payoffs than is the case with exponential discounting. The effect of hyperbolic discounting is to have a declining discount rate (DDR) over future years.
- Hyperbolic discounting in general leads to time inconsistent decisions. If hyperbolic discounting were to be adopted, it would need to be accompanied by an institutional framework that makes it difficult to reverse policy decisions, so building in time consistency for policymaking.

- Ramsey showed that under certain restrictive assumptions the SOC and SRTP discount rates are identical. The SRTP reflects three factors: the rate of pure time preference in the utility function (i.e. the degree to which the utility function exponentially discounts future consumption), the elasticity of utility with respect to consumption, and the per capita steady state growth rate of the economy. If future growth is higher, the SRTP discounts future consumption more highly so that some future consumption is brought forward to the present.
- The standard derivation of Ramsey's result is based on *intra-generational* optimisation rather than *inter-generational* considerations, with the pure rate of time preference within and between generations assumed to be the same. Ethically, however, there is no basis to say that future generations should be discounted by a person today in the same way as that person might discount their own consumption in their old age. Hence the standard exponential discounting approach to the social discount rate does not have an ethical basis to underpin it.
- One formulation used to reflect inter-generational concerns is to adopt a DDR that converges to zero, implicitly increasing the weight placed on future generations' utility as time evolves, so avoiding (or mitigating) one generation playing a dictatorial role over another.
- Private rates of return do not incorporate externalities, so in cases of negative externalities such as environmental degradation, the social rate of return will be less than the private return. In addition, the economy may not already be on the optimal consumption path. Each of these considerations imply that a discount rate determined by the SOC approach is not valid.
- Macroeconomic uncertainties relating to the future economic growth rate, and to the appropriate future discount rate, optimally result in the use of a DDR, similar to (but not the same as) use of hyperbolic discounting.
- Uncertainty about future consumption is especially relevant when dealing with tipping points, for example potential loss of biodiversity for a habitat. More generally, degradation of an environmental resource (and potentially also of a 'resource' such as social capital) should result in the payoff to that resource increasing as the resource becomes scarce. While, in some cases, this mathematically may be equivalent to adopting a stable payoff coupled with a DDR (an option referred to as 'dual discounting'), it conceptually makes more sense to model this situation as one with a stable discount rate coupled with an increasing payoff as the resource (or amenity) becomes scarce, since that payoff represents the correct shadow price at the time.
- Another form of dual discounting mooted in this report for further investigation, is to drop the common assumption that all components of utility share the same rate of pure time discount and instead to assume that different components of utility are discounted by the individual (or society) at different rates. This approach is consistent with the treatment of goods in Debreu's *Theory of Value*. The effects of including a lower discount rate on certain non-market outcomes (such as the environment) is to increase resources spent now to maintain resource stocks, so increasing the service flow from those stocks in future years. The approach does not necessarily lead to DDRs, so has a different rationale than does hyperbolic discounting, and the effects of the two approaches are different.
- Treasury's default social discount rate includes the market average risk premium, based on the capital asset pricing model (CAPM). An alternative risk framework is the consumption CAPM (CCAPM) in which a project's risk premium is determined by the correlation of its returns with the performance of the economy (rather than with the market investment portfolio). The CCAPM corresponds to the models underlying the SRTP approach within a risky environment.
- Use of the average risk premium (under either the CAPM or CCAPM approaches) is inappropriate for many government projects. Government projects are often chosen to address quite different features of society (e.g. public goods, combatting externalities and merit goods) than is the case for

private sector projects. Many of these projects have low, zero, or even negative correlation with economic outcomes (or with market average returns). They may also be chosen to reflect the interests of future generations whose preferences are not represented in current market prices, and whose welfare is discounted by private sector decision-makers.

- The market return incorporates a liquidity component plus a component relating to project risk. The liquidity component should not be incorporated into the SDR as government incorporates such impacts already through its fiscal envelope for expenditure to protect its balance sheet.
- The component related to project risk comprises both diversifiable risk (represented by its covariance with the market return) and undiversifiable risk. The latter is relevant for irreversible projects in which a decision to invest foregoes the real option of waiting for more information prior to investment. This cost occurs only once, at the time the investment is begun, so should be reflected in a once-only discount rather than as a component of an ongoing discount rate.
- The preceding observations mean that for many government projects, it is inappropriate to apply the average market risk premium as incorporated in Treasury's default discount rate.
- Treasury's current assumed real risk-free rate is, however, set too low based on historical observation. Based on the decade preceding the pandemic, a real risk-free rate of between 0% and 1% is more appropriate than the current negative rate assumed by Treasury.
- Approaches used to set the SDR vary considerably across jurisdictions around the world, even amongst developed countries. Some countries use SOC, others use SRTP; some have DDRs, others do not; some have different discount rates for different sectors, others do not; amongst developed countries, default rates vary from a low of 1% p.a. (Germany) to 8% in Canada. Hence there is no consensus on the setting of the SDR across countries.
- One factor that is not clear in comparing cross-country practices is the extent to which government investment in each country is constrained by a broader fiscal envelope. Further investigation into the role that the SDR plays in determining the quantum, rather than the ranking, of investments across countries is warranted.
- On balance, the conceptual arguments indicate that the use in New Zealand of a 5% default discount rate for (non-commercial) public sector projects, coupled with exponential discounting, leads to an inappropriately high discounting of future returns for many projects.
- One avenue to explore is the use of a DDR. The use of DDRs may, however, lead to issues of time consistency; longer term reductions in the SDR, as used for instance in several countries including the UK, are unlikely to have major time consistency issues but are also unlikely to substantively change estimated present discounted values for projects with long-run payoffs.
- An alternative avenue that is recommended here is to attribute (potentially non-linearly) increasing returns to a resource as it becomes scarce. This approach may be supplemented by exploration of different rates of pure time preference for different goods reflecting societal preferences that value the preservation of certain resources more highly than the preservation of other resources. This approach more explicitly deals with both preferences and scarcity than does hyperbolic discounting or other DDR frameworks.
- The most important immediate action that Treasury can undertake to adopt an appropriate set of discount rates is to ascribe project-specific risk premia that reflect the distribution of a project's returns relative to the state of the economy.

1. Introduction

When conducting project evaluation, the New Zealand Treasury uses a social discount rate (SDR) based on exponential discounting to convert future payoffs (benefits and costs) into current values.[1] Dasgupta (2008) describes the intuition behind using a SDR when evaluating multi-period projects:

> *Society now conducts a thought experiment … by asking how much additional consumption it would demand on behalf of tomorrow's people in payment for a reduction in today's consumption by one unit. We say that the "social rate of discount" between today's and tomorrow's consumptions is that additional consumption demanded, less unity."*

Treasury adopts a default real discount rate of 5% p.a.[2] when evaluating public sector projects, with 6% being used for the telecommunications, media and technology, IT and equipment, and knowledge economy (R&D) sectors; robustness checks can incorporate a lower rate of 2% (Treasury, nd; Treasury, 2022). The default discount rate incorporates both the risk-free discount rate plus a default risk premium (Treasury, 2008). Appendix 1 sets out Treasury's formula for calculating its default discount rate, highlighting the importance of the risk premium component.

To focus on key concepts, this report concentrates initially on the setting of the risk-free discount rate or, equivalently, the rate applied to a project that has zero correlation with aggregate consumption shocks (Breeden, 1979). The paper will, however, discuss how global uncertainty about real macroeconomic parameters such as growth rates, and the risk-free discount rate itself, may affect the risk-free rate used for project analysis. Discussion of the appropriate risk premium to add to the risk-free discount rate is included subsequently in the report.

In 2021, the Parliamentary Commissioner for the Environment (PCE) made the recommendation: "Modify the social discount rate used to evaluate initiatives and replace it with one that better reflects the longer-term, intergenerational costs and benefits that pertain to the environment." (PCE, 2021, p.120.) One option for consideration that the PCE presented was (proportional) hyperbolic discounting which, as shown later in this report, discounts future payoffs at a lesser rate than is commonly the case with exponential discounting.

Previous papers have discussed the appropriate discount rate to use for public sector projects in New Zealand. A comprehensive discussion is provided by Creedy and Passi (2017). Their survey made the common distinction between two approaches to determining the discount rate: the 'social opportunity cost of capital' (SOC) approach and the 'social rate of time preference' (SRTP) approach. The former bases the public sector discount rate on the observed return on the next best alternative investment that has the same risk, which may be a private sector investment. The latter is based on a social welfare function defined over present and future utilities. As shown later in this report, the SRTP typically incorporates three influences: the rate of 'pure time preference' (i.e. how society values a future generation relative to the present generation), the growth in aggregate consumption across time, and the effect of an increase of consumption on individuals' utility. Under extreme assumptions, SOC=SRTP but this equality cannot be expected to hold in most real-world circumstances. One reason that the two rates may not be equal is that future generations play no part in determining current market rates (so current rates may discount future generations' welfare severely) whereas a decision-maker acting on behalf of future generations may decide that future generations are as deserving of resources as are the current generation. Creedy and Passi conclude

[1] Note that all values throughout this report refer to inflation-adjusted (i.e. 'real') values, both in terms of payoffs and the discount rate.
[2] Henceforth in this report, any discount rate that is not accompanied by a timeframe refers to an annual rate.

"that neither approach offers an objective way of determining public sector discount rates. Value judgements are inevitable."

Grimes (2010) also contrasted the SOC and SRTP approaches. He concluded that a public sector project with returns that mirror those in the private sector (e.g. a government-owned commercial financial institution) should use the same discount rate (including the same risk premium) as used in the private sector (i.e. the SOC approach) whereas a project that yields intangible (non-market) returns should be framed in terms of a standard welfare analysis, reflecting a SRTP approach. This report returns to this idea of setting different discount rates for different types of public sector project in later sections.

Before outlining approaches that have been formulated to set the public sector discount rate, it is important to set out why the setting of the public sector discount rate is important. As shown in section 2 of this report, which sets out the formal basis of exponential discounting, use of a 5% discount rate when evaluating a public sector project translates into an implicit decision to treat $1 of benefit today as worth more than $100 of benefit that occurs in a century's time. Similarly, $9 of benefit today is worth more than a $100 of benefit that occurs in 50 years' time. Such implicit trade-offs make it difficult to implement policies that favour future generations (as may be envisaged, for instance, in a social investment or a social wellbeing approach to public policy). It is because of these apparently incongruous results from the use of exponential discounting with a 5% discount rate, that other options should be investigated, both conceptually and empirically.

The paper proceeds as follows: Section 2 outlines the process of exponential discounting (which can incorporate a discount rate chosen either using the SOC or SRTP approach). Section 3 outlines an alternative discounting approach – hyperbolic discounting – which has achieved prominence especially as a result of experiments within the behavioural economics literature. Section 4 outlines the core approach to determining the SRTP, based on the pioneering work of Frank Ramsey. Section 5 explores extensions to the Ramsey framework which can result in a declining discount rate being applied to future years (i.e. with a similar effect to the use of hyperbolic discounting). Section 6 explores 'dual discounting' (or, more generally, 'multiple discounting') drawing on insights from Debreu's (1959) *Theory of Value*. Section 7 discusses the issue of how the public sector should handle risky returns to its projects when setting a project-specific discount rate. Section 8 briefly describes other countries' setting of the SDR and relates their practices to the foregoing considerations. Section 9 concludes and provides three main recommendations.

2. Exponential discounting

To set the scene for the remainder of the report, consider how a government decision-maker converts net benefits (i.e. benefits less costs) for a multi-period project that covers periods $t=0$ to $t=T$ within a standard cost benefit analysis (CBA) context. The discounted sum of net benefits (*NetBen*) is given by:

$$pdv(NetBen) = \sum_{t=0}^{T} NetBen_t (1+d)^{-t} \quad (1)$$

The exponential discount factor, $D(t)$, applying to net benefits that occur at time *t* is then given by:

$$D(t) = (1+d)^{-t} \quad (2)$$

where d is the chosen discount rate. Column (1) of Table 1 presents values for the discount factor, $D(t)$, for a selection of future years (t) using exponential discounting with a discount rate (d) of 5%.[3] Given the values in this example, we multiply benefits that occur in 1 year's time using the discount factor $D(1) = 0.952$ [= $(1.05)^{-1}$]. Similarly, using the same 5% discount rate, we multiply benefits that occur in 100 year's time using the discount factor $D(100) = 0.008$ [= $(1.05)^{-100}$].

As noted in the Introduction, use of exponential discounting with a 5% discount rate implies that a benefit which accrues in a century's time is considered to be worth less than 1% of the same sized benefit that accrues today. Hence benefits in a century's time of mitigating climate change or of protecting biodiversity are today considered trivial and not worth the costs of mitigation if we were to naively implement this approach with these values.

Nevertheless, corporates commonly use an exponential discount rate when evaluating long-term projects (i.e. in discounting future cashflows). The use by corporates of this approach (with discount rates that may exceed 5%) creates a practical and conceptual problem for government. For instance, imagine a government which considers undertaking a commercial operation that competes directly with the private sector. It may create a state-owned enterprise (SOE) that has equal productivity as do private sector firms. If the private sector firm uses a discount rate of 5% to evaluate its investment in a project with the default degree of risk while the government uses a lower discount rate (i.e. a higher discount factor), the effect will be that government crowds out the private sector investment. Indeed, the SOE could be less productive than the private sector firm but the government's discounted cashflows could still be larger than is the case for the private firm undertaking the same project (with higher productivity). The result could be a government sector that crowds out the private sector with the likely effect of lower overall productivity across the economy.

This simple example provides a reason underpinning why government commercial projects should use the same discount rate (i.e. the same discount factor) as does the private sector when undertaking competing commercial activities.[4] This point is emphasised by Arrow et al. (2014) who argue that the rate of return to capital should be used as the discount rate when a public sector project displaces a private project. They note that this reasoning underpins the US Office of Management and Budget (OMB) adoption of "an estimate of the average pretax rate of return on private capital in the U.S. economy" when the effect of a project "is to displace or alter the use of capital in the private sector" (OMB, 2003).[5]

This reasoning can be extended beyond government provision of commercial projects to any project that may be funded by government in which only market payoffs (costs and benefits) exist. In such circumstances, a private firm will make an investment decision based on all payoffs relevant to society (since, by assumption, all payoffs are reflected in the market). It follows that a discount rate based on SRTP (if SRTP≠SOC) should apply only to projects that have non-market payoffs. Non-market payoffs may be related to the standard reasons that underpin government involvement in a project such as dealing with externalities and provision of public goods or merit goods. In an inter-generational context, government could also be involved to protect the rights of future generations whose preferences may not be reflected in current market prices.

[3] All discount factors are rounded to three decimal places.

[4] Similarly, in a risky environment, government commercial entities should use the same risk premium as a competing private sector enterprise to avoid artificially creating a playing field tilted in favour of SOEs.

[5] Note, however, that this reasoning falls short when the intention is to alter the private sector outcome as a result of the private sector causing an externality in the economy. In such a case, the private sector rate of return differs from the societal rate of return.

Table 1: Discount factors with three different discounting approaches

	Discounting approach		
Year	**Exponential**	**Hyperbolic**	**Quasi-hyperbolic**
0	1.000	1.000	1.000
1	0.952	0.952	0.952
2	0.907	0.909	0.943
3	0.864	0.870	0.933
4	0.823	0.833	0.924
5	0.784	0.800	0.915
6	0.746	0.769	0.906
7	0.711	0.741	0.897
8	0.677	0.714	0.888
9	0.645	0.690	0.879
10	0.614	0.667	0.870
11	0.585	0.645	0.861
12	0.557	0.625	0.853
13	0.530	0.606	0.844
14	0.505	0.588	0.836
15	0.481	0.571	0.827
16	0.458	0.556	0.819
17	0.436	0.541	0.811
18	0.416	0.526	0.803
19	0.396	0.513	0.795
20	0.377	0.500	0.787
50	0.087	0.286	0.582
100	0.008	0.167	0.352
200	0.000	0.091	0.129
500	0.000	0.038	0.006
1000	0.000	0.020	0.000

Notes: All discount factors are presented to 3 decimal places.
The exponential discount factors are calculated using a discount rate (d) of 0.05 based on equation (2).
The hyperbolic discount factors are calculated using a discount rate (d) of 0.05 based on equation (3).
The quasi-hyperbolic discount factors are calculated using values of β=0.962 and δ =0.99 based on equation (4).

3. Hyperbolic and quasi-hyperbolic discounting

The concept of hyperbolic discounting arose principally from observations in the behavioural economics literature (Frederick et al., 2002). In many experiments, subjects are observed to apply a high discount rate to near-term benefits with the per-period discount rate thereafter declining substantially; i.e. they have declining discount rates (DDRs). The effect is to apply a lower discount factor to near-term benefits than is usual with exponential discounting but then for the discount factor to decline more slowly than it does under exponential discounting. Using previous terminology, proportional hyperbolic discounting is formulated as follows:[6]

$$D(t) = \frac{1}{(1+td)} \quad (3)$$

Column (2) of Table 1 presents values for the discount factor, $D(t)$, based on hyperbolic discounting for a selection of future years with a discount rate of 5%. It is readily seen that returns in years beyond $t = 1$ are discounted using a higher discount factor when using hyperbolic discounting relative to the use of exponential discounting.

The idea behind hyperbolic discounting (that current payoffs are favoured over future payoffs, but agents do not discount the future as aggressively as under exponential discounting) has been used to formulate an expression for 'quasi-hyperbolic discounting', as follows (Phelps and Pollak, 1968; Laibson, 1997):

$$D(0) = 1, \quad (4a)$$

$$D(t) = \beta\delta^t \text{ for } t>0 \quad (4b)$$

where $\beta < \delta < 1$. It is also normal to assume that $(1 - \delta) < d$ where d is the discount rate typically used in exponential discounting. Column (3) of Table 1 illustrates the discount factor using quasi-hyperbolic discounting with $\beta = 0.962$ and $\delta = 0.99$ (which together produce the same discount factor in $t = 1$ as in the previous two columns). Figure 1 depicts the discount factors corresponding to the three columns in Table 1 (noting that the intervals beyond year 20 do not represent annual increments in time).

Figure 1: Discount factors with three different discounting approaches

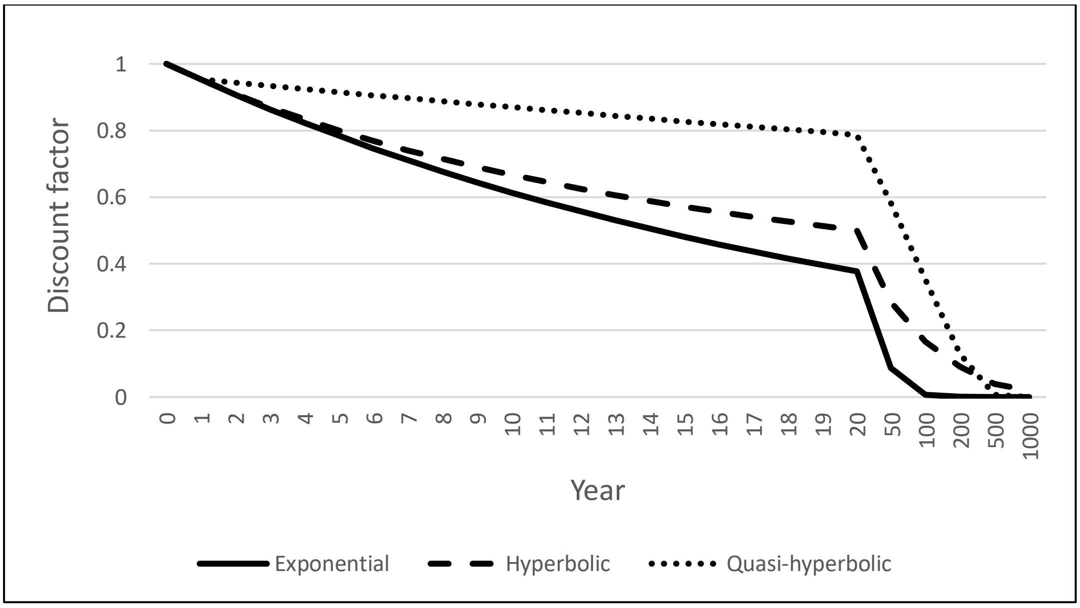

Notes: The three lines in Figure 1 correspond to the three columns in Table 1. The rates chosen for each discounting approach yield a discount factor of 0.952 in year 1.

[6] A more general form of hyperbolic discounting is given by the expression: $D(t) = \frac{1}{(1+\alpha t)^{\gamma/\alpha}}$ where $\alpha \geq 1$ and $\gamma \geq 1$. However, this form is less commonly used in the economic discounting literature.

For public policy purposes, a question arises as to whether hyperbolic discounting should be applied on the grounds that (some) people in experiments personally discount their future payoffs hyperbolically rather than exponentially. At the personal level, hyperbolic discounting gives rise to a familiar time consistency problem that can lead to procrastination. For instance, using the quasi-hyperbolic formulation for simplicity, $(1-\beta)$ is frequently observed in experiments to be much larger than any value of d that is used by policy agencies, so people put off an action with costs that are borne today that might reap benefits tomorrow (e.g. starting a diet), saying to themselves that they will be prepared to undergo those costs in a future period. However, when it comes to that future period, they will again be faced with a high value of $(1-\beta)$ so will again put off the action that will produce those costs. The procrastination occurs even though there is no new information arriving or any other change in circumstances for the individual. This lack of time consistency (and rationality) provides a poor basis to guide long-term public policy decisions.

If hyperbolic discounting were to be adopted, it would need to be accompanied by an institutional framework that incorporates difficulties to reversing policy decisions, so building in time consistency for policy. For instance, if government were to state that further development of a wetland were to cease completely in 5 years' time – with the delay reflecting a low value for β in (4b) – the policy environment would need to be such that this decision could not be reversed. If it could be reversed, then the initial decision is not credible (i.e. it is not time consistent) and so the application of hyperbolic discounting would be ineffective. Arguably, the slow pace of action to mitigate greenhouse gas emissions in New Zealand is an example of time inconsistent policy-making in which decisions to implement mitigation policies are signalled but then subsequently diluted as the near-term costs of the policies cause further delay and/or commitment to implementation.

A key lesson here is that unless the institutional structure produces time consistent decisions, it is difficult to justify the use of hyperbolic discounting for policy purposes. Strulik (2021) sets out a time-consistent approach to hyperbolic discounting whereby the hyperbolic discount rate depends on actual time but not on decision time (i.e. is independent of when a decision is taken). However, for many situations, it is unclear what type of realistic policy environment would result in this setup other than one in which irreversible policies are legislated.[7] In many (realpolitik) applications, time consistent decision-making is not likely. In these cases, just because we may observe individuals (inconsistently) discounting their own payoffs hyperbolically, it does not follow that hyperbolic discounting is appropriate for rational assessments of public policy interventions.

[7] One example, discussed in Grimes (2010), where such an approach might work is when a decision must be taken about the choice of quality (and hence longevity) of an infrastructure investment. For instance, a concrete based road may last with little maintenance for 50 years, whereas a bitumen road may require substantial maintenance to still be operating in peak condition in 20 years' time. If there is concern that a future decision-maker may not allocate resources to adequately maintain the infrastructure in future (because of the procrastination incentive given by hyperbolic discounting), then the current decision-maker may choose to invest in "future-proofed" infrastructure that mitigates the potential for infrastructure decline over the next 50 years. In this situation, a concrete-based road (which builds in time consistent decision-making) may be preferred to a bitumen-based road (with the potential for time inconsistent decisions) even if a conventionally-calculated BCR for the latter is higher than for the concrete road.

4. Ramsey discounting

A fundamental result in the discounting literature concerns the equivalence, under certain conditions, of the SOC and SRTP approaches to discounting. Specifically, Frank Ramsey (1928) showed that under certain assumptions:

$$r = \rho + \gamma g \quad (5)$$

where r is the risk-free market return (i.e. an investment of 1 unit made in year $t=0$ returns $1+r$ in $t=1$, with compounding thereafter), ρ is the rate of pure time preference in the utility function (i.e. the degree to which the person exponentially discounts future utility), γ is the elasticity of marginal utility with respect to consumption, and g is the per capita steady state growth rate of the economy (and the growth rate of per capita consumption). A simple derivation of (5) is given in Appendix 2.

In an economy with no market imperfections and no uncertainty we could therefore use either the (SOC) rate of return to a risk-free investment to discount the future or use the SRTP, often termed the 'consumption discount rate', i.e. the right hand side of (5). In each case, an appropriate risk premium would be added to the risk-free rate.

The consumption discount rate comprises three elements. The rate of pure time preference is included within the discount rate (as intuitively expected) but it is supplemented by the product of the growth rate of consumption and the elasticity of utility with respect to consumption. The intuitive reasoning behind inclusion of this product term is that if growth in consumption is present (e.g. due to technical progress) then we should spread that consumption over both the present and the future, which is equivalent to discounting future consumption so consuming more now as g increases. The degree to which we increase the discount on the future is determined by how much utility increases as consumption increases, which is determined by γ, the elasticity of marginal utility with respect to consumption.

The standard derivation of the Ramsey equation (including that in Appendix 2) is based on a representative individual. Consequently (as well as assuming zero market imperfections such as the absence of borrowing constraints), it is based on *intra-generational* optimisation rather than *inter-generational* considerations. This distinction is important because, in this specification, ρ is treated both as the *intra-generational* and the *inter-generational* pure rate of time preference. Ethically, however, there is no basis to say that future generations should be discounted by a person today in the same way as that person might discount their own consumption in their old age (Creedy, 2007; Cowen, 2007 and 2018; Dasgupta, 2008).[8] More generally, Chichilnisky (1996) argues that, in ethical terms, neither the future generation(s) nor the present generation should play a dictatorial role over the other.[9] For this reason, Pearce et al. (2003) consider that not discounting future generations' utility at all (i.e. using a discount rate of 0%) is unethical since it implies that the current generation

[8] Applying an inter-generational discount rate (i.e. valuing a future generation's welfare relative to our own) involves an inter-personal welfare comparison that is of the same type as trading off the welfare of one member of the current generation with that of another. Economists are typically reluctant to make inter-personal welfare comparisons within a generation, leaving these decisions to political decision-makers. CBA, however, typically treats the marginal dollar received by each individual as having equal impact on overall welfare whether it accrues to a starving individual or to a millionaire. In the absence of lump-sum taxes, this approach accords with no known philosophical school of thought (Blackorby and Donaldson, 1990; Sen, 2009).

[9] Helm (2019) is similarly concerned for the rights of future generations, arguing that the capabilities approach (Sen, 1999) implies future citizens are entitled to a set of renewable and non-renewable assets which gives them the opportunities to make their choices about how they want to live their lives.

.

should make large sacrifices for future generations. Chichlinisky argues that an outcome in which no generation tyrannises another generation represents an outcome with 'sustainable preferences'.[10] She proves that a welfare criterion based on the sum of discounted utilities (using any discount factor) is not consistent with such sustainable preferences.[11]

Dasgupta (2008) emphasises that γ performs three functions in the standard iso-elastic consumption function: (i) it is the elasticity of utility with respect to consumption, (ii) it is an index of risk aversion, and (iii) it is an index of inequality aversion. A higher value of γ increases the consumption discount rate for a given value of g since the weight placed on equity of consumption across generations is increased (i.e. more of the future richer generations' consumption is used for consumption today). Note that if there is uncertainty concerning the rate of growth (e.g. uncertainty about the rate of technical progress) then people will wish to save more (i.e. consume less) today than they would in a certain world to protect themselves against the downside risks to future utility arising from lower future consumption (in the presence of declining marginal utility from consumption).

A further criticism by Dasgupta of the basic Ramsey approach is that typically it is assumed that the elasticity of marginal utility with respect to consumption (i.e. γ) is constant over time and is independent of the level of consumption. He argues that this assumption "has only tractability to commend it". Pearce et al. (2003) make a similar point regarding the assumed constancy of γ. The possibility that tractability has caused other hidden drawbacks in current discounting approaches is developed further in section 6 of this report.

Finally, Dasgupta makes an important observation about the use of a SOC approach (based on private sector returns) to set the discount rate, as Nordhaus does in his work on climate change (Nordhaus, 2007). Private rates of return do not incorporate externalities, so in cases of negative externalities such as environmental degradation, the social rate of return will be less than the private return; it is the social rate of return that is relevant when considering public sector decisions. Furthermore, as Arrow et al. (2014) note, the equating of the SOC and SRTP rates in the Ramsey formula rests on the assumption that the (one-sector) economy is already on the optimal consumption path. If this is not the case, then the appropriate discount rate is not the SOC rate (even in the absence of externalities) but is instead the consumption discount rate derived from the social welfare function.

---

[10] If humanity faces a risk of extinction, it is rational to place some discount on the future (Yaari, 1965). Stern et al. (2006) assumed an extinction rate of 0.1% p.a. which he equated to the pure rate of time preference in his review of the economics of climate change.

[11] Caplin and Leahy (2004) run the thought experiment of testing whether inter-generational trade-offs are time consistent in the sense that the next generation will agree with the decisions made by the current generation. For instance, from today's viewpoint are we happy that the infrastructure deficit built up since the early 1980s was optimal, or do we instead consider that it would have been preferable for the last generation to provide more infrastructure investment than actually occurred? If, *ex ante*, the deferred investment was an optimal strategy but, *ex post*, that strategy is viewed as sub-optimal, then there is disagreement between the two generations about the rate of time discount. When this situation occurs, there is no single objective rate of time discount; we wish that the last generation had paid more heed to our generation when it was taking its expenditure decisions than it chose to. Cowen (2007) notes that substantial economic reforms may involve a significant loss of production during the reform period but they are implemented in anticipation of higher subsequent incomes once dislocation of the economy has run its course. Cowen considers that use of a high intertemporal discount rate rarely makes such economic reform worthwhile. Similarly, Grimes (2010) shows that, with a high discount rate, low growth policies that are already in place may be preferred to an economic policy that favours greater future growth to the benefit of future generations.

5. Extensions to Ramsey discounting

Building on the weaknesses identified above that are inherent in the simple Ramsey approach, many scholars have suggested additional considerations which should be considered when setting a discount rate, especially when incorporating long-term payoffs.

As noted by Dasgupta (2008), uncertainty about the rate of future consumption growth is one reason to discount the future more lightly (Gollier, 2012). Since uncertainty about future consumption is likely to grow as the time period to that consumption lengthens (especially if growth is non-stationary or exhibits substantial autocorrelation), the degree of discount (between future periods) should reduce based on this form of aggregate uncertainty.[12] (Note that we are dealing here with aggregate uncertainty, rather than to the risks surrounding any specific project.) The implication of this result is that a DDR should be incorporated into project analysis for public sector projects (that have non-market payoffs).[13]

Pindyck and Wang (2013) and Arrow et al. (2014) emphasise that uncertainty about consumption is especially relevant when dealing with tipping points, for example massive economic effects due to climate change or when addressing substantial loss of biodiversity for a habitat. In the presence of tipping points, not only is the scale of change inherently non-linear but it is also likely to be irreversible so making the negative outcome permanent. One approach to dealing with this issue suggested by Arrow et al. is to adopt a schedule of discount rates that fall as the likelihood of reaching a tipping point rises. However, this approach appears to confuse the payoffs in the period in question (i.e. around the time of the tipping point) with how we value those payoffs today. A conceptually superior approach is not to alter the discount rate but instead to ascribe non-linearly increasing payoffs to the avoidance of catastrophic effects as tipping points are approached.

A second source of aggregate uncertainty that affects the discount rate, emphasised by Weitzman (1998 and 2001), is uncertainty about what the appropriate (risk-free) discount rate should be. For instance, we may be certain that the most appropriate (constant risk-free real) discount rate is in the range of 1% to 5% but not certain about where in that range it should be. Weitzman shows that in this case, the discount rate applying to the first year should be the certainty equivalent (i.e. the weighted average using probability weights) of the possible rates; however as t increases, the certainty equivalent discount rate (applying to the period t to $t+1$) falls towards its lowest possible value. The reason is that the weights that we apply to future payoffs are related to the discount factor (rather than to the discount rate). Because of the exponential discounting formula, averaging the (uncertain) discount factors does not produce the same result as averaging the (uncertain) discount rates. A high discount rate reduces the discount factor asymptotically towards zero so the contribution of the lower discount rates when averaging the discount factor increases as t lengthens.

Appendix 3 provides a numerical example (based on an example in Pearce et al., 2003) where the policymaker has equal weights over five discount rates (1%, 2%, 3%, 4%, 5%). The appropriate discount rate starts at 3% then falls to 2.8% at 20 years, 2.5% at 50 years, 2.2% at 100 years and 1.2% at 1,000 years (where the latter may be relevant for a policy such as climate change). This example shows that even in the presence of considerable uncertainty about the discount rate (e.g. equal probability of 1% or 5% being the appropriate rate) the implied discount rate schedule experiences

[12] Dasgupta and Maskin (2005) similarly show that uncertainty about *when* payoffs are received leads to a DDR.
[13] Henceforth, the important caveat "that have non-market payoffs" applies to any references to public sector projects, unless indicated otherwise. The reason for the inclusion of this caveat is, as discussed earlier, that there would otherwise be a presumption that all or most long-term commercial projects should be delivered by government even where the private sector is the most efficient producer.

only a slow rate of decline. Therefore, except for very long-term projects, this 'gamma discounting' approach of Weitzman's is unlikely to imply major changes in project decisions.[14]

An approach that has some similarities with that of Weitzman, but resting on different fundamentals, is that of Li and Löfgren (2000). They consider two representative individuals: a utilitarian (who discounts the future in a standard exponential way with a discount rate incorporating the current generation's rate of pure time preference) and a conservationist representing future generations (who has zero rate of time discount). Reflecting Chichlinisky's framework, society's discount rate is a weighted average of the two rates. As time increases, the weight placed on the conservationist's (undiscounted) utility increases relative to that of the utilitarian's (discounted) utility so society's weighted discount rate declines towards zero over time.

Hoel and Sterner (2007) highlight that in an economy with two (or more) goods, e.g. a general consumption good and an environmental amenity, the relative price of the goods will change over time depending on scarcity. They assume a single pure rate of time preference over an individual's utility which is defined over both goods (with a constant elasticity of substitution utility function each period). They note that if there is a decline in the services from the environmental good then the marginal utility of that good will increase so optimally increasing the resources devoted to preserving it. The increase in marginal utility associated with the environmental good acts as a counterweight to the effect of the discount rate. Hence exponential discounting does not necessarily lead to a decline in the environmental resource.

Groom et al. (2022) note that an equivalent approach to that of Hoel and Sterner (in which a single social discount rate is used coupled with changing relative prices) is to use dual discount rates. The dual rates for a market consumption good, C, and a non-market environmental good, E, are again influenced by scarcity of the environmental good relative to the consumption good. The two rates are set as follows:

$$r_{Ct} = \rho + n_{CCt}g_{Ct} + n_{CEt}g_{Et} \qquad (6a)$$

$$r_{Et} = \rho + n_{EEt}g_{Et} + n_{CEt}g_{Ct} \qquad (6b)$$

where ρ is the (single) rate of pure time preference; n_{CCt} is the elasticity of marginal utility of market consumption with respect to itself, n_{EEt} is the elasticity of marginal utility of non-market consumption with respect to itself, n_{CEt} is the elasticity of marginal utility of consumption of one good with respect to the other, g_{Ct} is the growth rate of the market consumption good and g_{Et} is the growth rate of the non-market good.

The changing relative prices of the market and non-market goods inherent in the Hoel and Sterner (2007) and Groom (2022) specifications are important in practical applications in cost benefit analysis. If utility depends on a resource such as an environmental amenity, then the analysis should not assume a constant flow of utility per period from that resource; instead, the utility flowing from the amenity should reflect the quantity (and/or quality) of that flow at each period in the future. This analysis indicates that decisions taken using exponential discounting need not necessarily lead to depletion of a resource; it is the combined effect of the discount rate and the relative price effect that determines the path for that resource.

[14] Anchugina et al. (2017) derive a related result for the appropriate discount rate to use in the far-distant future for the case of hyperbolic discounting. They show that rather than converging to the lowest admissible rate, with hyperbolic discounting the far distant discount rate is given by the harmonic mean of the admissible rates.

6. Dual (and multiple) discounting

The Hoel and Sterner (2007) and Groom et al. (2022) contributions, amongst others, are important in recognising that utility may flow both from a general consumption good available in the market and a non-market amenity such as the services from an environmental resource. Like almost all analyses in this field, however, they assume a single rate of pure time preference to cover utility derived from all goods in each future period. This assumption appears to be made for the sake of tractability, mirroring Dasgupta's comment about the assumption of a constant elasticity of substitution of consumption. Kula and Evans (2011) make this point in advocating for a dual discounting approach (i.e. using different rates of pure time preference for different goods), but they do not present a formal model to show how dual discounting with multiple rates might be formulated and analysed.

There is nothing in economic theory to justify the common restrictive assumption that there is a single rate of pure time preference. To see that this is the case, consider the treatment of goods by Debreu (1959) in his *Theory of Value*. Debreu states that a commodity is characterised by its physical properties, its location, and the date at which it will be available. Thus a good with identical properties that is available at two different dates is treated as two separate goods. Furthermore, each good (defined by physical characteristics, location and date) has its own price.

For simplicity, consider a separable utility function, V, defined over four such goods (c_1, c_2, c_3, c_4) with four potentially different functional forms (v_1, v_2, v_3, v_4) for converting consumption into utility:[15]

$$V = v_1(c_1) + v_2(c_2) + v_3(c_3) + v_4(c_4) \tag{7}$$

We have no *a priori* reason to restrict the v_i functions to bear any specific relationship to one another. Now compare this general formulation with the way in which time discounting problems are commonly formulated. For instance, assume that c_1 and c_3 are market consumption goods at times *t=0* and *t=1* respectively (henceforth denoted as c_{m0} and c_{m1}) while c_2 and c_4 are non-market environmental amenities at times *t=0* and *t=1* respectively (henceforth denoted as c_{n0} and c_{n1}). Assuming separability of utility between all four goods (i.e. in each period and across time), the typical discount rate problem specifies:

$$V = [v_m(c_{m0}) + v_n(c_{n0})] + (1+\rho)^{-1}[v_m(c_{m1}) + v_n(c_{n1})] \tag{8}$$

Two restrictions are embodied in equation (8): the utility function for each good is assumed constant over time, and – most importantly for the idea of dual discounting – the same rate of pure time preference is assumed to hold for utility derived from both market and non-market goods. Nothing in Debreu's framework states that either of these restrictions needs to hold.

As an example of the importance of the second restriction, an individual (with known lifetime income and initial wealth) may be willing to trade off some market consumption in the future for greater consumption today, but may be insistent that the river that runs next to them remain clean throughout their lifetime. There is no reason for them to be forced to accept a reduction in river quality just because they are willing to accept a reduction in their consumption stream.[16]

---

[15] The mathematics would be more complex if we adopted non-separable utility, so the simpler assumption of separability is adopted to introduce the concepts.

[16] An example of how the first restriction may fail to hold is that an individual may rationally expect that their enjoyment of "head-banging" music (of whatever genre) will be less during their retirement than it is while still

To illustrate how the relaxation of a single rate of pure time preference may alter analysis of a project, consider the following problem. The individual maximises their lifetime (two-period) utility defined over market consumption (subscript m) and non-market consumption of an environmental amenity (subscript n) with the utility function given by (9) in which ρ_m is the rate of pure time preference for market consumption and ρ_n is the rate of pure time preference for the non-market environmental amenity. Utility is maximised subject to the budget constraint, (10), in which y is the individual's lifetime endowment, and subject also to a production technology for the environmental amenity, (11).

$$V = u(c_{m0}) + v(c_{n0}) + (1+\rho_m)^{-1}u(c_{m1}) + (1+\rho_n)^{-1}v(c_{n1}) \qquad (9)$$

$$c_{m1} = (y - c_{m0} - a_0)(1+r) \qquad (10)$$

$$c_{n1} = N + a_0 \qquad (11)$$

To ease the algebra, while fixing ideas, the utility functions for market goods in (9) are assumed to be constant across periods, as are the utility functions for the non-market good. Notably, in (9), utility from market consumption and utility from the non-market amenity may be discounted using different rates of pure time preference, reflecting the observations above.

In the budget constraint, r is the exogenous rate of return in *t=1* for one unit of the endowment invested at *t=0* (which is exogenously given, in a small open economy, by the world real interest rate), and a_0 is expenditure from the endowment in *t=0* used to enhance the non-market amenity in *t=1*; (11) shows that the flow from the non-market amenity, c_{n1}, is given by some exogenous amount N supplemented by the expenditure of a_0 in *t=0*. Henceforth, c_{n0} is assumed to be exogenously given by nature (and/or by past investments) and so is dropped from the analysis. Hence V is maximised through choices of c_{m0}, c_{m1}, and a_0 (or of c_{n1}).

An important feature of the system in (9)-(11) is that society must endure a resource cost in *t=0* to enhance the non-market amenity flow in *t=1*; that resource could instead have been used in the initial period for consumption or invested (returning r) to increase consumption in period *t=1*; hence a standard opportunity cost is incorporated into the analysis. A second important feature of the analysis is the presence of potentially different rates of pure time preference for non-market versus market goods in the utility function.

As shown in Appendix 4, the solution to the maximisation problem yields several interesting features. First, the standard Ramsey expression – in this case using the growth rate of *market* consumption and the rate of pure time preference for the *market* good – is retained. Second, as ρ_n declines relative to ρ_m, consumption of the environmental amenity in *t=1* increases and consumption of the market good in each period declines. Hence more of the endowment is spent in the initial period (via a_0) to protect the environmental amenity and less is spent on market consumption in each period.

This result is relevant for cost benefit analysis applied to public policy decisions. Typically, the same discount rate is used for a multitude of projects, differing perhaps only by a margin to allow for differing risk profiles. But this approach is entirely arbitrary. It assumes (even ignoring risk) that people are as relaxed about the state of the climate or a river in future relative to today as they are about their market consumption or the state of a work of art in the future relative to today. There is no reason to make such an assumption. The use of dual (or, more generally, multiple) discount rates

---

in their youth. For our purposes, it is not essential to relax this first restriction to demonstrate the potential importance of dual discounting so, for tractability (!), it is retained in what follows.

embodying differing rates of pure time preference for different goods can reflect differing societal weights placed on future streams of services from differing goods.

The adoption of dual discounting in the form suggested above is consistent with the idea of inter-generational 'existence value' in the environmental economics field (Krutilla, 1967). There are certain goods for which people may not necessarily gain use value but which they consider should exist for the foreseeable future; other goods are not accorded the same status. For instance, New Zealanders may consider that original copies of Te Tiriti o Waitangi should be preserved for the very long term but may be less concerned to preserve other documents (even if they consider that such documents would still be valued in future). Similarly, New Zealanders may consider that the Kiwi should continue to exist indefinitely whereas they may be less concerned to preserve a less iconic species. In these examples, people are effectively using a lower (or zero) rate of pure time preference for the Treaty or the Kiwi relative to other items that enter the utility function.

7. Risk premium[17]

As outlined in section 1 and in Appendix 1, Treasury's default discount rate adopts a risk premium based on the capital asset pricing model (CAPM; Sharpe, 1964 and Lintner, 1965). The CAPM determines a project's risk premium (relative to the risk-free rate) by virtue of its 'beta' which in turn reflects the sensitivity of the project's returns to those of the overall market (the market portfolio). A project that is expected to record high returns at the same time as the market portfolio is performing poorly (and vice versa) will have a low risk premium, since that project provides useful diversification properties within an investment portfolio. Likewise, a project with strongly pro-cyclical returns will have a high risk premium (no matter what its individual variance of returns). Treasury's default discount rate assumes a market equity premium of 7% and the market average beta (which converts to an asset beta of 0.67 assuming leverage of 33%). The risk component therefore comprises the bulk of the Treasury's default discount rate. Indeed, the current Treasury calculation assumes a *negative* real risk-free rate at -1.3%,[18] so the risk premium comprises more than the full discount rate. As discussed earlier, use of a risk premium that reflects the premium used by the private sector for an investment with a similar risk profile is appropriate for the public sector when it is considering an investment project for which all payoffs are observed in the market.

As an aside, the public sector already adopts different hurdle rates of return incorporating different risk premia for financial investments with differing risks when making financial portfolio investments. The Crown has substantial financial investments (e.g. via ACC and the NZ Superannuation Fund) which are spread across assets with different risks. Investments in bonds are expected to have a low rate of return relative to investments in market-based equities which in turn have lower expected returns than investments in private equity. Hence the principle of adopting an asset-specific risk premium is already well-established in Government investment practice.

More generally, how should the risk premium be calculated for a project with non-market returns that is undertaken by government? The main model to determine risk premia in such situations is the consumption capital asset pricing model (CCAPM; Breeden 1979) which corresponds to the type of optimisation model set out in Appendix 2 in an environment with risky payoffs. A project that is expected to record high payoffs at the same time as the broader economy (and consumption) is performing poorly will be accorded a low risk premium, since that project provides useful

[17] The discussion in this section draws on that in Grimes (2010).

[18] I.e. the real risk-free rate = (1.0065/1.02)-1; see values in Appendix 1.

diversification properties for the wider economy. A project with strongly pro-cyclical returns relative to the broader economy will be accorded a high risk premium. A project with returns that have a zero correlation with the state of the economy will have a zero risk premium.[19]

A major irrigation scheme is an example where the CCAPM may imply a low risk premium. Treasury research (Buckle et al., 2007) indicates that rainfall is a material determinant of New Zealand's cyclical economic outcomes, with drought negatively affecting each of national production (GDP), domestic demand (GNE) and exports. An irrigation scheme has its highest payoffs precisely when the economy is at a low ebb due to drought; thus the CCAPM implies that the irrigation scheme should be accorded a low risk premium. Indeed, if the returns to an irrigation scheme are counter-cyclical (i.e. a negative consumption beta) the appropriate societal discount rate would be below the risk free (government borrowing) rate.

From a government's narrow fiscal perspective, why might such an approach also make sense? If government is concerned with stability in its revenue stream, it will value highly a scheme that produces tax revenues at times when other tax revenues are declining. An irrigation scheme that meets this requirement acts to offset other risks to the government's revenue stream, thus having a risk-reducing role for government's overall revenues. This risk-reducing role again implies that the risk premium is negative for the scheme.

Now consider the nature of many public sector investments. The reason that the public sector undertakes investments (either through direct public provision or through funding of private provision) is often because: (i) these investments are considered to increase some societal welfare criterion, and (ii) these investments are not commercially viable for the private sector to undertake due to the presence of non-market payoffs. An example is investment in public goods (such as trapping of pests in a national park or remediation of a river); another is investment in merit goods (such as investment in mental health facilities).

The following question then arises when setting the risk premium on such a project: What is the covariance of its return relative to the state of the economy? For instance, do the returns to pest trapping in a national park vary when the economy is strong or weak? If the answer to this question is no, then the appropriate risk premium for the project is zero. As another instance, do the returns to a project that invests in mental health facilities vary with the state of the economy? Hypothetically, a case might be made that mental health deteriorates when the economy is at its trough and improves when the economy is strong. If this were the case, then the return to the mental health facility would be negatively correlated with the state of the economy (i.e. the service is most used, so has its highest return, when the economy is weak) and so a negative risk premium should be assigned to the project. In this second case, the appropriate discount rate would therefore be below the risk-free rate (rather than the current situation in which the opposite is the case and the market risk premium is added to the risk-free rate).

Several further points are warranted in considering matters related to risk. In setting the social discount rate, Coleman (2019) draws on modern finance literature to demonstrate that the market premium for a private sector firm comprises two components: (i) project earnings risk, and (ii) liquidity risk. The latter depends on the effect that undertaking the project may have on the firm's access to capital. It will tend to be higher for less liquid projects (e.g. for projects that cannot easily

[19] In considering the covariance of payoffs with the state of the economy, it may be most appropriate to consider payoffs over long timeframes (e.g. economic cycle lengths rather than day-to-day returns) since government's time horizon is such that it is longer term covariances that matter.

be on-sold) which may be especially relevant for long-lived indivisible projects. Households have a preference in their own portfolios to hold liquid securities, so the risk premium is heightened for illiquid projects. In the government sphere, liquidity concerns are taken care of through an overarching 'fiscal envelope' for expenditures (to ensure a sound public balance sheet) so, at the project level, this liquidity aspect should not be included in the project's risk premium. This consideration means that an illiquid long-term project (with market or non-market returns) may be better financed by government than the private sector. Examples might include investments to mitigate, or adapt to, climate change, set aside land from productive uses for a national park, or to create a long-lived educational institute or health facility. Note that this consideration can work alongside either an exponential discount rate or a DDR and can work with dual discounting options.

The component related to project risk comprises both diversifiable risk (represented by its covariance with the market return) and undiversifiable risk. The latter is relevant for irreversible projects in which a decision to invest foregoes the real option of waiting for more information prior to investment. This cost occurs only once, at the time the investment is begun, so should be reflected in a once-only discount rather than as a component of an ongoing discount rate.

The cost relating to the foregone real option in the case of an irreversible project implies a discount rate structure akin to that of quasi-hyperbolic discounting. The β parameter in equation (4b) formalises the real option cost of proceeding with the project, leaving δ to reflect the SRTP. This reasoning, based on rational decision-making, provides a much stronger case for adopting quasi-hyperbolic discounting than is provided by the experimental evidence on personal discounting behaviour. (Note that in practical applications β will vary across projects reflecting the real option cost involved for a specific project.)

While the long-term average risk-free market rate of return is itself uncertain (see the discussion in section 5), it is prone to change from one year to the next. Treasury's current best estimate of the real risk-free rate is, as shown earlier, -1.3%. Since 2010, an estimate of the real one year government bond rate[20] has varied from -6.05% to 3.07%. If we ignore the pandemic period, the real one year government bond rate over 2010-2019 varied between -1.30% and 3.07%, averaging 0.24%. For applications to (commercial) government projects that require a market rate of return, a real risk-free rate of 0% to 1% appears more appropriate than is the current rate of -1.3%. This consideration, however, does not affect the SRTP which is independent of the market interest rate.

The use of different risk premia for different projects (and for different financial investments), as is warranted either with a CAPM or a CCAPM approach, implies that public sector decision-makers already use multiple discount rates when assessing project benefits. Thus, while the rationale for using multiple discount rates when projects vary according to their risk is different to that discussed in section 6, the feasibility of the public sector adopting different discount rates for different projects is already well established.

8. Other countries' practices

Several studies have reviewed the practices adopted by authorities across multiple countries in setting their SDR. The most recent example of such a survey is that by Groom et al. (2022) which is used as the source of information for all international rates discussed in this section. Groom's Table 1

[20] Calculated as the one year swap rate less the average of the annual inflation rate to the current quarter and the annual inflation rate four quarters hence.

sets out the SDR for 18 countries (including EU as a country for this purpose)[21] plus two international bodies (which are ignored in what follows as their rationales for setting a particular SDR may well differ from a country's perspective).

Practices in setting the SDR differ widely across the 18 countries with 11 of the 18 countries claiming to use a SOC based approach, 5 using a SRTP approach and 2 using both approaches in setting the discount rate.

Of the 18 countries, 7 have chosen to adopt a DDR. For instance, the UK has a default rate of 3.5% for years 0-30, 3.0% for years 31-75, and 2.5% for years 76-125.[22] A further 2 countries adopt relative price adjustments for environmental goods which, as shown in section 5, is equivalent to instituting a DDR but with no price adjustment accorded to increasing environmental scarcity.

Dual (or multiple) discounting is used by 6 countries, with the main sectors having a lower SDR being health and environment. The USA uses a lower SDR for projects that have important intergenerational effects.

The resulting default SDR rates vary considerably, from a low of 1% in Germany to 10% in Mexico. Within countries, rates can vary considerably; for instance Canada has a default rate of 8% but an SDR of 3% may be applied to projects in the health and environment fields. Similarly, Australia has a default rate of 7% with 4% used for projects that have non-market benefits, while the USA has a default rate of 7% with 3% used for intergenerational benefits. The UK has a default SDR rate of 3.5% (based on a rate of pure time preference of 0.5%, a 1% catastrophic risk premium and using a consumption elasticity of 1), while using an SDR of 1.5% for health. Germany is the stand-out in terms of a very low SDR with its 1% rate being based on a 0% rate of pure time preference and a consumption elasticity of 1. Middle income countries (including Brazil, Columbia, Mexico and Peru) have SDRs (each based on a SOC approach) of between 8% and 10%, perhaps reflecting high market returns in countries that are in a catch-up phase in terms of development opportunities.

This quick tour of practice indicates that the appropriate method of setting the SDR is far from reaching a consensus, even across developed countries (viz Canada's 8% default rate versus Germany's default rate of 1%). What is not clear, however, is the extent to which government investment in each country is constrained by a broader fiscal envelope. For instance, it could be that in each of Canada and Germany, total public investment is set by some balance sheet constraint, and the role of the SDR is to rank projects within the resulting envelope, rather than determining the size of the envelope. This issue – i.e. the role of the SDR in determining the quantum, or just determining the ranking, of investments – is one that needs further exploration at the cross-country level.

---

[21] The 18 countries are: Australia, Brazil, Canada, Chile, Columbia, Denmark, EU, France, Germany, Ireland, Mexico, Netherlands, New Zealand, Norway, Peru, Sweden, UK and USA.

[22] HM Treasury, *Green Book 2022*, https://www.gov.uk/government/publications/the-green-book-appraisal-and-evaluation-in-central-governent/the-green-book-2020#a6-discounting. In practice, however, the UK's DDR schedule makes little difference to the pdv of a project with long-term payoffs. For instance, if a project yields a constant flow of benefits =100p.a. for years 0 to 125, then its pdv with the DDR schedule is 3,080 whereas the pdv with a constant 3.5% rate would be 2,918; a difference of only 5.6% in the pdv. The reason for the small difference is that, with exponential discounting, the first 30 years are the crucial years in terms of setting the discount factors for future years.

9. Conclusions

Treasury currently adopts a default real discount rate of 5% p.a., together with exponential discounting, when evaluating public sector projects. The default discount rate incorporates both the risk-free discount rate plus a default risk premium. For government decisions concerning (provision and/or funding) of a project that has solely market payoffs it is appropriate to use the same discount rate as does the private sector when undertaking a project that has an equivalent risk profile. If government did not do so, the choice of (a lower) discount rate would artificially bias decisions away from projects that the private sector would undertake based on all societal returns.

For projects with non-market returns, it is important to recall that the effect of using Treasury's current discounting practice is to treat $1 of benefit today as worth more than $100 of benefit that occurs in a century's time. Several scholars (e.g. Dasgupta, 2008; Arrow et al. 2014; Cowen, 2018) argue that the severe inter-generational discounting (for instance of non-market environmental outcomes) that results from this practice has no basis amongst standard ethical theories. Furthermore, private rates of return do not incorporate externalities, so in cases of negative externalities such as environmental degradation, the social rate of return will be less than the private return. In addition, the economy may not be on the optimal consumption path which is assumed to be the case when discounting non-market returns by a market rate. Each of these considerations implies that a discount rate determined by the SOC approach is not valid for many government projects.

Several rationales have been afforded for using declining discount rates (DDRs) over time in place of exponential discounting. One rationale is based on experimental evidence that subjects tend to use hyperbolic discounting when making intertemporal decisions, i.e. they employ a lower discount factor (higher discount rate) for near-term payoffs and then have a higher discount factor for longer term payoffs than is the case with exponential discounting. The effect of hyperbolic discounting is to have a declining discount rate (DDR) over future years. Other rationales for DDRs include the impact of uncertainty about future macroeconomic conditions and balancing the interests of future generations with the current generation. However, hyperbolic discounting in general leads to time inconsistent decisions. If hyperbolic discounting were to be adopted, it would need to be accompanied by an institutional framework that makes it difficult to reverse policy decisions, so building in time consistency for policymaking. It has been suggested that DDRs may be useful when dealing with issues of tipping points (for example potential loss of biodiversity for a habitat) with projects that seek to combat such risks being subject to discount rates that fall as the likelihood of reaching a tipping point rises. Conceptually, however, it makes more sense to model the payoff to the project as increasing sharply as the tipping point is approached, coupled with a constant discount rate; this latter approach is consistent with placing an appropriate value on the utility of the generation that is affected when it is approaching the tipping point.

A form of dual discounting, mooted in this report, drops the common assumption that all components of utility share the same rate of pure time preference and instead adopts the assumption that different components of utility may be discounted by the individual (and society) at different rates. This approach is consistent with the treatment of goods in Debreu's *Theory of Value*. The effects of including a lower discount rate on certain non-market outcomes (such as preservation of the Kiwi or of physical copies of Te Tiriti o Waitangi) is to increase resources spent now to maintain an existing stock, so increasing the service flow from that stock in future years (relative to the outcome with a higher rate of pure time preference). This approach does not necessarily lead to DDRs, so has a different rationale than does hyperbolic discounting, and the effects of the two approaches are different. The dual discounting approach is embedded in the utility function and so should not give rise to a time consistency issue unlike hyperbolic discounting.[23]

Treasury's default social discount rate includes the market average risk premium which should not be used when discounting most government non-commercial projects. Many such projects have low, zero, or even negative correlation with economic outcomes (or with market average returns). These correlations reflect the choice of government projects that are designed to address quite different features of society (e.g. provision of public goods and merit goods, or combatting externalities) than is the case for private sector projects. Different liquidity constraints faced by the government relative to private sector firms also mean that the market risk premium is inappropriate to apply to public sector non-commercial projects. Hence the current practice of incorporating the market risk premium means that many government projects face an inappropriately high risk premium under current policy settings. Note, however, that the Treasury's current assumed real risk-free rate is set too low; a real risk-free rate of 0-1% is more appropriate than the currently assumed negative rate. In addition, a once-only real option discount should apply to an irreversible project in the first year of its implementation.

Approaches used to set the SDR vary considerably across jurisdictions around the world, reflecting a lack of consensus on the setting of the SDR across countries. Hence the international experience is of little guidance for determining how our government should set its SDR.

On balance, the conceptual arguments indicate that the use in New Zealand of a 5% default discount rate for non-commercial public sector projects, coupled with exponential discounting, leads to an inappropriately high discounting of future returns for many public sector projects. While adoption of DDRs has been recommended elsewhere (on experimental and conceptual grounds), their use internationally – for instance in the UK – shows that values currently adopted by other countries in setting their DDR make little difference to long-term calculations of net present value for realistic cases.

An alternative avenue that is mooted here is to adopt different rates of pure time preference for different goods, coupled with increasing returns attributed to an increasingly scarce resource (in an evaluation). This approach more explicitly deals with both preferences and scarcity than does hyperbolic discounting or other DDR frameworks.

Recommendations

In summarising the materials discussed above, three recommendations are made for Treasury to consider. Incorporation of these recommendations into discount rate practice within the New Zealand government sector would go a substantial way to addressing the dictatorship of the current generation over future generations that arises from the existing approach to discounting.

1. The most immediate substantive change that Treasury could make to its discounting practice is for the SDR to incorporate appropriate risk premia for projects with different risk profiles. The risk premium for a project that is adopted because of the existence of non-market payoffs should be related to the covariance of the project's returns with the state of the economy and should not reflect the average market risk premium. The risk premium which reflects this covariance should then be added to a low real risk-free rate (probably in the

[23] This conjecture has still to be checked in future work. The intuition behind it reflects the Treaty example. If the current generation deems it worthwhile to maintain existing physical copies, and if preferences on this matter do not change, it is difficult to see why future generations would choose to deviate from the decision.

order of 0-1%). Irreversible projects (which can potentially be delayed and which could benefit from the gathering of further information over time) should have an added once-only discount placed on them to reflect the loss of a real option when the project is implemented. Further work is required to determine the appropriate size of this real option discount, which is likely to be project-specific.

2. Evaluations of long-term projects that may affect, or be affected by, tipping points should, where warranted, incorporate a strongly increasing payoff as an (uncertain) tipping point is approached, rather than using a declining discount rate.
3. Long-term payoffs to projects for which the populace is likely to have a lower rate of pure time discount compared with that for generalised consumption could have a lower SDR than the default rate (after allowance for risk). This proposal (for a dual discount rate based on differing rates of pure time preference) is one which warrants further investigation.

Appendix 1: Treasury's Calculation of the Default Public Sector Discount Rate

Treasury's current approach to setting the discount rate is based on a calculation of the (real) weighted average cost of capital (*WACC*) faced by a project with a given level of risk (reflected in its 'asset beta'). The formula is as follows (Treasury, 2008; based on Lally, 1998):

$$WACC = \left[\frac{1+WACCn}{1+i}\right] - 1 \qquad \text{(A1.1)}$$

with:

$$WACCn = [RFR(1 - Tc) + (Ep * \beta_a)]/(1 - Te) \qquad \text{(A1.2)}$$

Variables definitions, with current values in parentheses (Treasury, 2022), are:

WACC	real weighted average cost of capital, i.e. the default discount rate	(0.047);
WACCn	nominal weighted average cost of capital	(0.068);
Tc	corporate tax rate	(0.28);
Te	effective tax rate	(0.24);
Ep	equity risk premium	(0.07);
RFR	risk free rate	(0.0065);
i	inflation rate	(0.02);
β_a	asset beta[24]	(0.67).

This formulation demonstrates that the bulk of the discount rate comprises the risk premium. Indeed if $\beta_a = 0$ then *WACC* = -1.3% based on the values set out by Treasury. This negative discount rate reflects the negative real interest rate that prevailed through the covid pandemic. A more realistic inflation figure for the period (e.g. 7%) would give an even more negative (current) number. As discussed in section 7, the one-year real rate on government bonds averaged 0.24% over 2010-2019.

The 6 percentage point difference between the default beta discount rate (~ 5%) and the zero beta (risk free) discount rate (~ -1%), based on Treasury's own parameters, indicates that a material assumption implicit when using the default discount rate is that the project under consideration has an asset beta of 0.67. The importance of this point is discussed in section 7 of the report.

[24] Treasury (2008) defines the asset beta as "a measure of the sensitivity of an asset's return to that of the market portolio. A beta of one means that the expected return of the investment always moves with the market as a whole; a beta of zero means that the expected return of the investment is independent of the market. A beta of zero implies the risk premium is also zero. The weighted average beta of all stocks in a market equals one (same as the market beta); … the market average asset beta is 0.67 (assuming leverage of 33%)."

Appendix 2: Derivation of the Ramsey equation

The following provides a derivation of the Ramsey equation for discounting, equation (5), in the simplest possible two-period context.[25] There is no uncertainty or market imperfections and there exists a representative individual. The individual has time separable utility function defined over present and future consumption (c_0 and c_1 respectively) with rate of pure time preference, ρ:

$$U = u(c_0) + (1+\rho)^{-1}u(c_1) \quad \text{(A2.1)}$$

Utility is maximised by choices of c_0 and c_1 subject to the budget constraint:

$$c_1 = (y - c_0)(1 + r) \quad \text{(A2.2)}$$

where *y* is the individual's (or society's) resource endowment and *r* is the risk-free return gained from investing one unit of that resource from t=0 to t=1.

Maximisation of (A2.1) subject to (A2.2) yields:

$$\frac{u\prime(c_1)}{u\prime(c_0)} = \frac{1+\rho}{1+r} \cong 1 + \rho - r \text{ (for small } \rho \text{ and } r\text{)} \quad \text{(A2.3)}$$

where $u'(c_t)$ denotes marginal utility.

Hence:

$$r = 1 + \rho - u'(c_1)/u'(c_0) \quad \text{(A2.4)}$$

If the utility function is iso-elastic, we have:

$$u(c_t) = c_t^{1-\gamma}/(1-\gamma) \quad \text{(A2.5)}$$

and:

$$u'(c_t) = c_t^{-\gamma} \quad \text{(A2.6)}$$

For small g, (A2.3) then becomes:

$$r \cong \rho + \gamma g \quad \text{(A2.7)}$$

which is the Ramsey equation.

Note that utility here is defined over the absolute level of consumption. An interesting (and potentially important) extension to this standard approach would be to incorporate relative effects in which utility is defined over consumption of individuals in the country relative to consumption of individuals in comparator countries, reflecting an international version of the Easterlin Paradox (Easterlin, 1974; Becchetti et al., 2013; Grimes and Reinhardt, 2019).

[25] More general derivations from continuous time infinite horizon models produce the same underlying result.

Appendix 3: Example of gamma discounting

Consider an example in which the policymaker is unsure about the appropriate value of the (constant risk-free real) discount rate. She considers that it could be 1%, 2%, 3%, 4% or 5% with equal probability of each. Table A3.1 presents the resulting discount factor (calculated using exponential discounting) for each discount rate and year. The certainty equivalent discount factor is then given by the penultimate column. Pearce et al. (2003) show that for a certainty equivalent discount factor, *D*, the certainty equivalent discount rate, *s*, is given (in the case of equal probabilities) by:

$$(1 + s)^{-t} = D$$

Hence:

$$s = D^{-1/t} - 1$$

The final column presents the certainty equivalent discount rate (to 3 d.p.).

Table A3.1. Certainty equivalent discount factors and discount rates given a range of potential discount rates.

	Discount rate					Certainty	**Certainty**
	0.01	0.02	0.03	0.04	0.05	equivalent	**equivalent**
Year	Discount factor					discount factor	**discount rate**
0	1.000	1.000	1.000	1.000	1.000	1.000	**0.030**
1	0.990	0.980	0.971	0.962	0.952	0.971	**0.030**
2	0.980	0.961	0.943	0.925	0.907	0.943	**0.030**
3	0.971	0.942	0.915	0.889	0.864	0.916	**0.030**
4	0.961	0.924	0.888	0.855	0.823	0.890	**0.030**
5	0.951	0.906	0.863	0.822	0.784	0.865	**0.029**
6	0.942	0.888	0.837	0.790	0.746	0.841	**0.029**
7	0.933	0.871	0.813	0.760	0.711	0.817	**0.029**
8	0.923	0.853	0.789	0.731	0.677	0.795	**0.029**
9	0.914	0.837	0.766	0.703	0.645	0.773	**0.029**
10	0.905	0.820	0.744	0.676	0.614	0.752	**0.029**
11	0.896	0.804	0.722	0.650	0.585	0.731	**0.029**
12	0.887	0.788	0.701	0.625	0.557	0.712	**0.029**
13	0.879	0.773	0.681	0.601	0.530	0.693	**0.029**
14	0.870	0.758	0.661	0.577	0.505	0.674	**0.029**
15	0.861	0.743	0.642	0.555	0.481	0.657	**0.028**
16	0.853	0.728	0.623	0.534	0.458	0.639	**0.028**
17	0.844	0.714	0.605	0.513	0.436	0.623	**0.028**
18	0.836	0.700	0.587	0.494	0.416	0.607	**0.028**
19	0.828	0.686	0.570	0.475	0.396	0.591	**0.028**
20	0.820	0.673	0.554	0.456	0.377	0.576	**0.028**
50	0.608	0.372	0.228	0.141	0.087	0.287	**0.025**
100	0.370	0.138	0.052	0.020	0.008	0.117	**0.022**
200	0.137	0.019	0.003	0.000	0.000	0.032	**0.017**
500	0.007	0.000	0.000	0.000	0.000	0.001	**0.013**
1000	0.000	0.000	0.000	0.000	0.000	0.000	**0.012**

Appendix 4: Dual discounting example

In the model set out in section 6, the individual maximises their lifetime (two-period) utility defined over market consumption (with subscript m) and non-market consumption of an environmental amenity (with subscript n). The utility function is given by (A4.1) in which ρ_m is the rate of pure time preference for market consumption and ρ_n is the rate of pure time preference for the non-market environmental amenity. Utility is maximised subject to the budget constraint (A4.2) in which y is the individual's lifetime endowment, and subject also to a production technology for the environmental amenity, (A4.3).

$$V = u(c_{m0}) + v(c_{n0}) + (1+\rho_m)^{-1}u(c_{m1}) + (1+\rho_n)^{-1}v(c_{n1}) \quad \text{(A4.1)}$$

$$c_{m1} = (y - c_{m0} - a_0)(1+r) \quad \text{(A4.2)}$$

$$c_{n1} = N + a_0 \quad \text{(A4.3)}$$

The utility functions for market goods in (A4.1) are assumed to be constant across periods, as are the utility functions for the non-market good. In the budget constraint, r is the exogenous rate of return in *t=1* for one unit of the endowment invested at *t=0* (which is exogenously given, in a small open economy, by the world real interest rate), and a_0 is expenditure from the endowment in *t=0* which is used to enhance the non-market amenity in *t=1*; (A4.3) shows that the flow from the non-market amenity, c_{n1}, is given by some exogenous amount N supplemented by the expenditure of a_0 in *t=0*. Henceforth, c_{n0} is assumed to be exogenously given by nature (and/or by past investments) and so is dropped from the analysis. Hence V is maximised through choices of c_{m0}, c_{m1}, and a_0 (or of c_{n1}).

Maximising V subject to (A4.2) and (A4.3) yields expressions (A4.4) and (A4.5) which, together with the budget constraint, provides three equations for the three unknowns c_{m0}, c_{m1}, and c_{n1}.

$$\frac{u'(c_{m1})}{u'(c_{m0})} = \frac{1+\rho_m}{1+r} \cong 1 + \rho_m - r \quad \text{(A4.4)}$$

$$\frac{u'(c_{m1})}{v'(c_{n1})} = \frac{1+\rho_m}{(1+r)(1+\rho_n)} \cong 1 + \rho_m - r - \rho_n \quad \text{(A4.5)}$$

If we assume $u(c_{mt}) = c_{mt}^{1-\gamma}/(1-\gamma)$, then (A4.4) gives us a slight variant of the standard Ramsey equation as shown in (A4.6) in which g_m is the growth rate of market consumption:

$$r \cong \rho_m + \gamma g_m \quad \text{(A4.6)}$$

To ease complexity while retaining the core ideas, we henceforth assume that each of $u(c_{mt})$ and $v(c_{nt})$ is logarithmic (i.e. iso-elastic with $\gamma = 1$).[26] With this simplification, we solve for each of c_{m0}, c_{m1}, and c_{n1} as follows, in which we have used approximations of the type shown in (A4.4) above.

$$c_{m0} = (1+\rho_m)(y+N)\theta \quad \text{(A4.7)}$$

$$c_{m1} = (1+r)(y+N)\theta \quad \text{(A4.8)}$$

$$c_{n1} = (1+\rho_m-\rho_n)(y+N)\theta \quad \text{(A4.9)}$$

where: $\theta = (3 + 2\rho_m - \rho_n)^{-1}$

[26] Dasgupta (2008) notes that both Stern and Nordhaus adopt logarithmic utility in their analyses of climate change impacts; Dasgupta argues for a much higher value of γ, based on ethical reasoning.

Note that (y+N), which appears in the solution for each variable, is effectively the full resource endowment since it contributes directly to c_{n1}.

To illustrate the effect of changing the value of ρ_n relative to ρ_m, assume: y=2, N=1, r=0.05, ρ_m=0.03=ρ_n, which together yield the vector: (c_{m0}, c_{m1}, c_{n1}) = (1.0198, 1.0396, 0.9901).

Keeping other values constant, if ρ_n=0.00, the vector becomes (1.0098, 1.0294, 1.0098); i.e. as the rate of pure time preference for the non-market good falls, services from that good rise and consumption of the market good in both periods declines.

Thus, a lower rate of pure time preference for the non-market good has the effect that we intuitively expect: more resources are used to maintain this amenity while consumption of market goods declines relative to the case of assuming a single rate of pure time preference. Given the exponential nature of discounting, the effect of this dual discounting on the allocation will increase as time evolves in a model with more than two periods.

References

Anchugina N, Ryan M, Slinko A. Hyperbolic discounting of the far-distant future. *Economics Letters*, 155, 140-143.

Arrow K, Cropper M, Gollier C, Groom B, Heal G, Newell R, Nordhaus W, Pindyck R, Pizer W, Portney P, Sterner T, Tol R, Weitzman M. 2014. Should governments use a declining discount rate in project analysis? *Review of Environmental Economics and Policy*, 8(2), 145-163.

Becchetti L, Castriota S, Corrado L, Ricca E. 2013. Beyond the Joneses: Inter-country income comparisons and happiness. *The Journal of Socio-Economics, 45*, 187-195.

Blackorby C, Donaldson D. 1990. A review article: The case against the use of the sum of compensating variations in cost-benefit analysis. *The Canadian Journal of Economics,* 23(3), 471-494.

Breeden D. 1979. An intertemporal asset pricing model with stochastic consumption and investment opportunities. *Journal of Financial Economics*, 7, 265-296.

Buckle R, Kim K, Kirkham H, McLellan N, Sharma J. 2007. A structural VAR business cycle model for a volatile small open economy. *Economic Modelling*, 24(6), 990-1017.

Caplin A, Leahy J. 2004. The social discount rate. *Journal of Political Economy,* 112(6), 1257-1268.

Chichilnisky G. 1996. An axiomatic approach to sustainable development. *Social Choice and Welfare*, 13(2), 231-257.

Coleman A. 2019. *Liquidity, the Government Balance Sheet, and the Public Sector Discount Rate*. Working Paper 19-13. Wellington: Motu Economic and Public Policy Research.

Cowen T. 2007. Caring about the distant future: Why it matters and what it means. *The University of Chicago Law Review,* 74(1), 5-40.

Cowen T. 2018. *Stubborn Attachments*. San Francisco, Stripe Press.

Creedy J. 2007. Policy evaluation, welfare weights and value judgements: A reminder. *Australian Journal of Labour Economics*, 10(1), 1-15

Creedy J, Passi H. 2017. *Public Sector Discount Rates: A Comparison of Alternative Approaches*. New Zealand Treasury Working Paper 17/02. Wellington: The Treasury.

Dasgupta P. 2008. Discounting climate change. *Journal of Risk and Uncertainty*, 37, 141-169.

Dasgupta P, Maskin E. 2005. Uncertainty and hyperbolic discounting. *American Economic Review*, 95(4), 1290-1299.

Debreu G. 1959. *Theory of Value: An Axiomatic Analysis of Economic Equilibrium*. New Haven and London, Yale University Press.

Easterlin R. 1974. Does economic growth improve the human lot? In M Abramovitz, P David & M Reder (Eds.), *Nations and households in economic growth: Essays in Honor of Moses Abramovitz*. New York, Academic Press.

Frederick S, Loewenstein G, O'Donoghue T. 2002. Time discounting and time preference: A critical review. *Journal of Economic Literature*, 40(2), 351-401.

Gollier C. 2012. *Pricing the Planet's Future: The Economics of Discounting in an Uncertain World*. Princeton, NJ, Princeton University Press.

Grimes A. 2010. *The Economics of Infrastructure Investment: Beyond Simple Cost Benefit Analysis*. Working Paper 10-05. Wellington: Motu Economic and Public Policy Research.

Grimes A, Reinhardt M. 2019. Relative income, subjective wellbeing and the Easterlin Paradox: Intra- and inter-national comparisons. In: Rojas M. (ed.) *The Economics of Happiness: How the Easterlin Paradox Transformed our Understanding of Well-Being and Progress*, Switzerland, Springer, 85-106.

Helm D. 2019. Natural capital: Assets, systems, and policies. *Oxford Review of Economic Policy*, 35(1), 1–13.

Hoel M, Sterner T. 2007. Discounting and relative prices. *Climatic Change*, 84, 265-280.

Krutilla J. 1967. Conservation reconsidered. *American Economic Review*, 57(4), 777-786.

Kula E, Evans D. 2011. Dual discounting in cost-benefit analysis for environmental impacts. *Environmental Impact Assessment Review*, 31, 180-186.

Laibson D. 1997. Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics*, 112(2), 443-478.

Lally M. 1998. Public sector cost of capital: A comparison of two models. *New Zealand Economic Papers*, 32(2), 195-212.

Li C, Löfgren K-G. 2000. Renewable resources and economic sustainability: A dynamic analysis with heterogeneous time preferences. *Journal of Environmental Economics and Management*, 40, 236-250.

Lintner J. 1965. The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *Review of Economics and Statistics*, 47, 13-37.

Nordhaus W. 2007. The Stern Review on the economics of climate change. *Journal of Economic Literature*, 45(3), 686-702.

Office of Management and Budget (OMB). 2003. *Circular A-4: Regulatory Analysis*. Washington DC, Executive Office of the President.

PCE. 2021. *Wellbeing Budgets and the Environment: A Promised Land?* Wellington: Parliamentary Commissioner for the Environment.

Pearce D, Groom B, Hepburn C, Koundouri P. 2003. Valuing the future: Recent advances in social discounting. *World Economics*, 4(2), 121-141.

Phelps E, Pollak R. 1968. On second-best national saving game-equilibrium growth. *Review of Economic Studies*, 35(2), 185-199.

Pindyck R, Wang N. 2013. The economic and policy consequences of catastrophes. *American Economic Journal: Economic Policy*, 5(4), 306-339.

Ramsey F. 1928. A mathematical theory of saving. *Economic Journal*, 38(4), 543-549.

Sen A. 1999. *Development as Freedom*. Oxford: Oxford University Press.

Sen A. 2009. *The Idea of Justice*. Cambridge MA, Belknap Press of Harvard University Press.

Sharpe W. 1964. Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance*, 19, 425-442.

Stern N et al. 2006. *The Stern Review of the Economics of Climate Change*. Cambridge, Cambridge University Press.

Strulik H. 2021. Hyperbolic discounting and the time-consistent solution of three canonical environmental problems. *Journal of Public Economic Theory*, 23, 462-486.

Treasury. 2008. *Public Sector Discount Rates for Cost Benefit Analysis*. Wellington: The Treasury.

Treasury. nd. *Discount Rates*. Downloaded from: https://www.treasury.govt.nz/information-and-services/state-sector-leadership/guidance/financial-reporting-policies-and-guidance/discount-rates

Treasury. 2022. *CBAx Tool User Guidance*. Wellington: The Treasury.

Weitzman M. 1998. Why the far-distant future should be discounted at its lowest possible rate. *Journal of Environmental Economics and Management*, 36, 201-208.

Weitzman M. 2001. Gamma discounting. *American Economic Review*, 91, 260-271.

Yaari M. 1965. Uncertain lifetime, life insurance, and the theory of the consumer. *Review of Economic Studies*, 32(2), 137-150.