Green Book Discount Rate Review: Technical Annexes

Mark C. Freeman and Ben Groom*

June 30, 2026

* Mark Freeman, School for Business & Society, University of York, United Kingdom, mark.freeman@york.ac.uk. Ben Groom, Dragon Capital Chair of Biodiversity Economics, LEEP Institute, University of Exeter Business School, United Kingdom and Visiting Professor at the Grantham Research Institute for Climate Change and the Environment, London School of Economics and Political Science, b.d.groom@exeter.ac.uk.
This is a fully independent report commissioned by HM Treasury. Any views expressed are solely those of the authors and do not represent those of HM Treasury, nor do they constitute HM Treasury policy. This report should therefore not be reported as representing the views of, or being endorsed by, HM Treasury. The report also does not represent the views of the University of York, the University of Exeter, the Financial Conduct Authority or its independent Cost Benefit Analysis Panel, the London School of Economics and Political Science, or any other institution with whom the authors are affiliated.

Contents

1. Introduction **5**
1.1. Background to this Review . . . 5
1.2. The Technical Annexes . . . 6

2. Summary of Findings and Recommendations **8**
2.1. Principles . . . 8
2.2. Terms of Reference questions . . . 8

3. How We Have Undertaken this Review **12**
3.1. The survey . . . 12
3.2. Observer Groups . . . 15
3.3. Expert panels . . . 15
3.4. Green Book framework . . . 16

4. Social Discounting across Time and Countries **18**
4.1. The evolution of Green Book discounting . . . 18
4.2. The evolution of US social discounting . . . 20
4.3. Other international social discounting guidance . . . 21
4.4. Other UK social discounting . . . 21

5. The Foundations of Discounting **25**
5.1. Introduction . . . 25
5.2. Present Values . . . 25
5.3. Stochastic Discount Factors . . . 26
5.4. When to avoid discounting . . . 27
5.5. Dual discounting . . . 29
5.6. Marginal Cost of Public Funds . . . 32
5.7. Discount rate sensitivity analysis . . . 32

6. STPR or SOC? **35**
6.1. Introduction . . . 35
6.2. STPR . . . 36
6.3. SOC . . . 38
6.4. Best Unbiased Estimators . . . 40
6.5. Surveys and international guidance . . . 41

7. Discounting in a Deterministic World **42**
7.1. The Ramsey Rule and asset pricing . . . 42
7.2. The SOC risk-free rate . . . 43
7.3. The Ramsey Rule in welfare economics . . . 47
7.4. Sustainability and the Ramsey Framework . . . 52
7.5. Interpreting the Green Book . . . 57

8. Choosing the Ramsey Parameters **59**
8.1. Introduction . . . 59
8.2. Interpreting the Green Book Ramsey Rule . . . 59
8.3. The current calibration and evidence not used . . . 61
8.4. The pure rate of time preference: δ . . . 64
8.5. The elasticity of marginal utility: μ . . . 70

8.6. Expected growth, g . . . 75
8.7. Concluding remarks . . . 77

9. Risk-free Rates under Uncertainty **79**
9.1. Introduction . . . 79
9.2. Prudence . . . 79
9.3. Declining Discount Rates . . . 80

10. Average systematic risk under STPR **84**
10.1. Certainty equivalents . . . 84
10.2. The CCAPM . . . 85
10.3. Risk in the Green Book . . . 86
10.4. The L term . . . 86
10.5. The Arrow-Lind Theorem . . . 88

11. Systematic risk under SOC **90**
11.1. The CAPM . . . 90
11.2. International SOC guidance . . . 91
11.3. Estimates of the equity risk premium . . . 93
11.4. Project-varying risk premia . . . 95

12. Rare disasters **100**
12.1. Introduction . . . 100
12.2. Rare disasters and the risk-free rate . . . 100
12.3. Rare disasters and the risk premium . . . 101
12.4. Illustrative numerical examples . . . 102

13. Declining Discount Rates: Empirical Evidence in in the UK **109**
13.1. Introduction . . . 109
13.2. A general stochastic discount factor . . . 110
13.3. Model classes . . . 110
13.4. Data and empirical methods . . . 112
13.5. Results . . . 114
13.6. Conclusion . . . 115

14. Catastrophic Risk, Risk-Free Discounting and Project-Specific Risk Premia **118**
14.1. Introduction . . . 118
14.2. Theory: risk-free disaster risk and project systematic risk . . . 118
14.3. Barro and Pindyck–Wang catastrophe specifications . . . 119
14.4. Data and empirical implementation . . . 120
14.5. Empirical disaster estimates . . . 120
14.6. Main results for $\mu = 1.25$ . . . 120
14.7. Exposure required for target premia . . . 121
14.8. Conclusion . . . 121

15. Relative prices, commodity specific discount rates and sustainability **126**
15.1. Introduction . . . 126
15.2. Foundations . . . 126
15.3. A two-good Ramsey framework . . . 127
15.4. Substitutability, cross-partial effects, and environmental inequality aversion . . . 129
15.5. Empirical evidence . . . 130
15.6. Production, natural-capital stocks, and environmental drag . . . 131
15.7. Relative prices and the social cost of carbon . . . 132
15.8. Relative-price appraisal and natural-capital sustainability tests . . . 133
15.9. Government appraisal practice . . . 139
15.10.A recommendation for relative-price adjustments in CBA . . . 142

15.11. Conclusion and final recommendations . . . 144

16. Place-based objectives and welfare weights 146
16.1. Introduction . . . 146
16.2. Welfare weights . . . 147
16.3. A welfare-metric accounting theorem . . . 148
16.4. The mean-consumption Ramsey denominator . . . 149
16.5. CBA with interpersonal welfare weights . . . 150
16.6. Reference units and dynamic equity corrections . . . 151
16.7. Mean and median reference paths . . . 154
16.8. From individual weights to group weights . . . 156
16.9. Diagnosis and recommendations for the Green Book . . . 158
16.10. Conclusion . . . 162

17. High complexity 163
17.1. Introduction . . . 163
17.2. Private finance models . . . 163
17.3. Ultra-long-term projects . . . 167
17.4. Transformational projects . . . 168

Appendices 173

A. Survey responses 174
A.1. Characteristics of respondents . . . 174
A.2. Qualitative comments . . . 175

B. Appendix to Chapter 8 216

C. Appendix to Chapter 13 219
C.1. Raw recent-sample permanent-state calibrations . . . 219
C.2. Robustness: alternative time samples . . . 219
C.3. Diagnostics and model selection . . . 219
C.4. Regime-switching estimates . . . 219
C.5. Data-boundary diagnostic . . . 220
C.6. Bai–Perron and low-frequency block robustness . . . 220
C.7. AR(1)-GARCH robustness . . . 222

D. Appendix to Chapter 14 224

E. Environmental Discount Rate Review 226
E.1. Background . . . 226
E.2. Relative prices and uprating: Brief formal analysis . . . 227
E.3. How to view the Green Book Health Discount Rate . . . 227
E.4. *E* in aggregate versus its components . . . 228
E.5. Simple Examples of Relative Price Effects . . . 230

F. A Place-based Ramsey Rule 232

1. Introduction

1.1. Background to this Review

1.1.1. HM Treasury's Green Book is the UK government guidance on assessing the costs, benefits and risks of different options for achieving government objectives. The Green Book provides a structured framework for developing evidence-based, objective and impartial advice to decision-makers on the best way of achieving their objectives.

1.1.2. A key parameter within the Green Book is the "discount rate". This allows for costs and benefits that arise at different points in time to be compared. HM Treasury previously reviewed the approach to discounting in the Green Book in the 2010s. This culminated, in 2018, with a published report by Michael Spackman (LSE) and ourselves (Freeman et al. 2018). HM Treasury undertook a further external review to examine the application of the discount rate specifically in relation to environmental impacts. The conclusions from this review were published in 2021 (HM Treasury 2021). We provide a more detailed description of the evolution of Green Book discounting in Section 4.

1.1.3. In a speech in January 2025, the Chancellor of the Exchequer, the Rt Hon Rachel Reeves MP, announced a further formal review of the Green Book. She said:

> "…we will review the Green Book and how it is being used to provide objective, transparent advice on public investment across the country, including outside London and the Southeast. This means that investment in all regions is given a fair hearing by the Treasury that I lead."

1.1.4. The government's ongoing commitment to this agenda was re-affirmed in the Chancellor of the Exchequer's Spring Forecast 2026 speech, where she said:

> "...and the choice to give people in all parts of our country the opportunities that they deserve by reforming the Treasury Spending Rules in the Green Book to unlock investment in all of our urban, rural and coastal communities".

1.1.5. The results of the 2025 Green Book review were announced by the Chancellor as part of her Spending Review statement in June 2025. She said:

> "…today I am publishing the conclusion of the review of the Treasury Green Book, which is the Government's manual for assessing value for money. Our new Green Book will support place-based business cases, and make sure that no region has Treasury guidance wielded against it. I said that we would do things differently, and that we wanted growth in all parts of Britain, and I meant it."

1.1.6. The Green Book Review report outlined six major issues raised by stakeholders. One of these relates to the Green Book's perceived "…ineffectiveness at assessing transformational change." To help address this issue, the review called for the following action:

> "HM Treasury will improve the Green Book guidance on appraising transformational change. HM Treasury will commission an independent review of the Green Book discount rate to make sure that the government is taking a fair view of the long-term benefits that arise from transformational investments."

1.1.7. This statement gathered broader interest, with Paul Johnson writing in the Times:

> "Selecting the appropriate discount rate is one of the knottiest, most intellectually demanding and impossibly fascinating problems in all of economics. The Treasury is going to commission an independent review. Good. Such a review must be properly robust and independent."

1.1.8. HM Treasury published the Terms of Reference for this Green Book Discount Rate Review in December 2025, for which we were appointed as lead authors by the Chief Secretary to Treasury at that time, the Rt Hon James Murray MP.

1.1.9. The Terms of Reference required us to address eight subtle questions, each of which gets to the heart of the knotty problems alluded to by the Times.

- What evidence should the UK government consider for remaining with the current STPR approach or instead changing to an SOC approach?
- Do the current estimates for the parameters of the STPR (excluding L) remain broadly valid for short-to-medium term risk-free social discounting? What new evidence on these parameter values has emerged since the previous review in 2018?
- What rate of return should be used for short-to-medium term risk-free social discounting if the UK government chooses to switch to an SOC approach?
- What evidence should the UK government consider for and against the declining discount rate (DDR) regime currently in use? Does the existing guidance on intergenerational wealth transfers and social discounting remain appropriate?
- Does the L term within the STPR continue to correctly account for 'catastrophic' and 'systemic' risk? Should the discount rate adjustment for risk depend on whether the government chooses an STPR or SOC approach?
- What are the main conceptual issues and considerations when using discounting in appraisals involving private finance models?
- Should the discount rate be adjusted for issues such as place-based objectives and environmental scarcity?
- Should the social discount rate be adjusted for transformational projects, and if so, how?

1.1.10. The Terms of Reference also stated that "The lead authors will have full authority over the contents of the published report. HM Treasury will have discretion about whether to accept the recommendations within the report, and about how to incorporate those recommendations into the Green Book where necessary". This stresses the independent nature of our work

1.1.11. Within this independence, though, HM Treasury made three key requirements through the Terms of Reference:

- "The lead authors will be expected to liaise with the wider academic community on behalf of HM Treasury and fairly reflect the breadth of views in the final report".
- "The review should consider the UK context for these questions. The review should also look at international comparisons to better inform judgements on the UK approach."
- "HM Treasury will convene an observer group. The observers will represent organisations across the UK public sector whose decisions involve an understanding of discounting. HM Treasury will convene two meetings of the lead authors and the observer group".

1.2. The Technical Annexes

1.2.1. The Terms of Reference state that "The main output of the review will be a short report, co-authored by the lead authors" with a deadline of the end of June 2026. This is being delivered as a separate document (the *Green Book Discount Rate Review: Summary of Findings and Recommendations*).

1.2.2. The purpose of these Technical Annexes is to give greater insight into how we have reached our recommendations. This is a complex and evolving field, with significant evidence available both from the academic literature and from social discounting practice internationally. As an indication, *Google Scholar* revealed over 30,000 relevant documents on a search in late June 2026 for the single term "Social Discount Rate" (in quotation marks), with over 9,000 of these being published since the start of 2018; the time of the last detailed review of the Green Book discount rate.

1.2.3. This is also a topic that, despite its technicalities, leads to strong, and often conflicting, opinions. There are a number of different ways in which certain questions raised by the Terms of Reference could reasonably be handled. In these Technical Annexes we explain why our chosen recommendations are robust, consistent, and practical, thus directly responding to the Green Book Review's action: "HM Treasury will radically simplify and shorten the Green Book and the accompanying business case guides ... mak(ing) clear the level of detail that is proportionate for business cases of different levels of cost and complexity".

1.2.4. The remainder of these Technical Annexes proceed as follows. In Chapter 2 we summarise our findings and recommendations. Each subsequent chapter can then be viewed as falling into one of four broad categories:

1.2.5. **Background**. We begin by providing some relevant background and context. In Chapter 3 we describe the process that we have taken to writing this review, and specifically how we have met the requirements laid out in paragraph 1.1.11. In Chapter 4, we provide a brief history of social discounting within HM Treasury and an overview of international social discounting practice. We show that a number of questions asked in the Terms of Reference have been of concern to policy makers for many decades across different parts of the world. This makes clear that, despite different governmental bodies considering these matters robustly and deeply, they frequently come to different policy recommendations on the discount rate. We also briefly describe situations in the UK public and regulatory sectors where the Green Book discount rate does not apply. This chapter is particularly relevant for the "STPR or SOC" question, which we return to explicitly in Chapter 6. Before then, in Chapter 5, we review the conceptual foundations of social discounting. This sets out some very general principles that overarch the more specific questions raised by the Terms of Reference.

1.2.6. **Theory**. While this is a highly applied review, directly addressing the policy-relevant questions set in the Terms of Reference, this cannot be done without a firm theoretical understanding of the Ramsey Rule and its extensions. In Chapter 7 we derive the Green Book formula, $r = \delta + \mu g$ (excluding the L). This can emerge from a number of different conceptual frameworks and specific assumptions. In all cases, this requires that the future is known with certainty. We extend this to economies with macroeconomic uncertainty in Chapter 9, where declining discount rates become directly relevant. Chapter 10 then introduces microeconomic, project-specific, risk into the STPR discount rate model, while Chapter 11 moves this discussion into an SOC setting. In Chapter 12 we consider how the probability of rare disasters can affect the risk-free rate, the average social risk premium, and why this risk premium might be negative for projects with very strong social insurance properties.

1.2.7. **Calibration**. Having established a strong theoretical and conceptual base for determining the Green Book social discount rate, we then turn to how such models might best be calibrated in an STPR environment. Chapter 8 gives numerical value recommendations for the three parameters of the Ramsey Rule (δ, μ, g), while Chapter 13 calibrates how macroeconomic uncertainty and rare disaster risk influences the term structure of declining discount rates based on up-to-date UK consumption data. In Chapter 14 we bring all these theoretical and empirical concepts together to derive an STPR social discount rate.

1.2.8. **Complexity**. We complete this review by considering the application of discount rates in more complex situations. In Chapter 15 we focus specifically on how the environment should be handled in an economic appraisal. In Chapter 16 we consider place-based economic appraisals. Our recommendation is that such effects should be considered in welfare weights, and we provide some detailed thoughts on how current Green Book guidance might be enhanced. In Chapter 17 we discuss three of the most difficult cases; (i) public finance models, (ii) projects with ultra-long maturities, and (iii) transformational projects. In the first case, the social discount rate should reflect the risk that the public sector bears in such partnerships. The latter two cases offer such complexity that formal welfare analysis, rather than discount rates themselves, is likely to be the most proportionate way of the government have confidence in the accuracy of the economic appraisal.

1.2.9. We start with a definition which we use throughout these Technical Annexes. We define *public project* (or sometimes just *project*) as a general term to cover any proposed public investment, policy, or regulation that falls under Green Book guidance.

2. Summary of Findings and Recommendations

We begin by summarising the findings and recommendations that we have reached following from the analysis that we describe in subsequent chapters. These underpin the document *Green Book Discount Rate Review: Summary of Findings and Recommendations* which we have submitted separately to HM Treasury, with that being the document that formally fulfils our requirements under the Terms of Reference.

2.1. Principles

We lay out four principles that underpin our specific recommendations.

1. That the Green Book discount rate should account for maturity and risk only. This is discussed primarily in Chapter 5 and also Section 10.1.

2. That Green Book discounting guidance should be applied consistently, without separate discount rates for health and the environment as examples, is introduced in Section 5.5, with then a much more detailed discussion in Chapter 15. That there must exist a unique risk-free rate that should be applied to all projects of the same maturity is formally established in Section 5.3.

3. That Green Book discounting principles should be applied proportionately follows directly from the Findings and Actions of the Green Book Review (2025) and the Green Book (HM Treasury 2026*c*, #1.8). We discuss this briefly in Section 3.4.

4. That the Green Book should not allow for discount rate sensitivity analysis is discussed in Section 5.7 and the associated academic working paper (Nesje et al. 2024).

2.2. Terms of Reference questions

2.2.a. STPR or SOC?

2.2.1. That the discount rate should conceptually be viewed as an STP rate is set more generally by the framework of the Green Book itself, which makes explicit that it is concerned with welfare economics (HM Treasury 2026*c*, #2.4). The majority of respondents to our survey (see Section 3.1 for a description of this survey) supported this approach: see Table 6.1 and the qualitative comments in paragraph A.2.2.

2.2.2. That the Ramsey Rule itself has a clear foundation in welfare economics, and that it reflects HM Treasury's essential view of the welfare objective of government as the discounted sum of well-beings, is explained in Chapter 7. However, this chapter also makes clear that the precise form is a pragmatic simplification that facilitates estimation. We extend this model to account for macroeconomic risk in Chapter 9 and then rare disaster risk in Chapter 12. While these extensions improve the theoretical model, they remain a parsimonious representation of intertemporal social welfare objectives.

2.2.3. That the Ramsey Rule and its extensions raise serious conceptual and empirical challenges to calibrate is discussed in detail in Chapter 8.

2.2.4. That calibration of the Ramsey Rule and its extensions cannot, therefore, set a 'perfect' STP rate is discussed in Chapter 6. In this chapter, we present a formal case for triangulating such an estimate with information from other sources. This includes rates of return from financial markets, survey data, and precedents set by other governments.

2.2.5. On its own the Ramsey Rule does not embody concepts such as sustainability, so the choice of parameters for the Ramsey Rule does not ensure that wider welfare objectives are met (see Chapter 5).

2.2.b. Calibrating the Ramsey Rule

2.2.6. Chapter 8 addresses the specific issue of calibrating the Ramsey Rule. From our analysis of the most up-to-date evidence presented there, as well as from our survey respondents in Tables 8.1 and 8.2 and the qualitative comments in paragraph A.2.3, we recommend that:

- The rate of pure time preference, $\delta = 0.5\%$;
- The elasticity of marginal utility of consumption, $\mu = 1.25$;
- The expected growth rate of real per capita consumption $g = 1.5\%$.

2.2.7. Based on these calibrations, the Ramsey Rule itself gives a risk-free STP rate of 0.5% + 1.25 × 1.5 = 2.375%.

2.2.8. Our main STPR risk-free rate includes the potential for macroeconomic risk (Chapter 9) and rare disasters. We discuss the impact of the latter on the risk-free rate in Section 12.2, present some stylised calibrations of such models in Section 12.4, and then a much more detailed calibration in Chapter 14. The potential for such a catastrophe reduces the risk-free rate because it gives society a stronger desire to invest now for what may be a poorer future. We estimate this effect at minus 0.5%, reducing the STPR risk-free rate to 1.875%.

2.2.9. To harmonise and simplify the guidance, we also recommend in Section 8.5 that the coefficient of intratemporal inequality aversion, ϵ, which is used in welfare weights, is set equal to $\epsilon = \mu = 1.25$.

2.2.c. The SOC risk-free rate

2.2.10. We address the question of what risk-free rate to use under an SOC framework in Section 7.2. There is broad consensus that this should be based on a mid-maturity real return to a UK Treasury security; see the qualitative survey responses in paragraph A.2.4. We recommend that the Bank of England real yield curve is used for this purpose, giving a current rate of approximately 2%.

2.2.11. Other sources of information that we use to triangulate the risk-free rate come from survey data and international guidance; see Tables 11.3, 11.4 and 4.1 for more detail. When also considering the STPR risk-free rate of 1.875%, our overall recommendation is that the risk-free component of the discount rate for maturities of 30 years or shorter should be 2%.

2.2.d. Declining discount rates

2.2.12. We introduce the concept of declining discount rates in Section 9.3, report some stylised calibrations of such models in Section 12.4 when there is rare disaster risk in the economy, and then provide more detailed calibrations in Chapter 13. We also consider survey data and international guidance; see Tables 9.1, 11.3 and 4.1 for more detail.

2.2.13. Our recommendation from this analysis is that the risk-free forward rate (period-to-period) should decline by 0.5% at 31 years and then a further 0.25% at 76 years.

2.2.14. For ultra-long maturities of greater than 125 years, we recommend that HM Treasury undertakes more fundamental welfare analysis. We explain our reasoning for reaching this conclusion in Section 17.3.

2.2.e. The L term

2.2.15. In Chapter 10 we consider the current representation of L in the Green Book and how this compares with the standard STPR approach for risk premia; the Consumption Capital Asset Pricing Model (CCAPM). We recommend to HM Treasury that this term is reconceptualised as an adjustment for systematic risk and not specifically for catastrophic and systemic risk.

2.2.16. Our recommendation is that considerations of individual mortality risk should not influence this term, and that the risk of absolute societal collapse is overstated in the current value of L (Chapter 8). By contrast, considerations of obsolescence should be accounted for in the numerator of the present value equation (Section 10.4).

2.2.17. The CCAPM gives a very small value for L, thus partially supporting the Arrow-Lind Theorem (Section 10.5) that there should be no social risk premium. However, once rare disaster risk is included in the STPR model, then the risk premium becomes material. We discuss this in Section 12.3, present some stylised calibrations of such risk premiums in Section 12.4, and then give a much more detailed calibration in Chapter 14. Overall, an STPR risk premium of 0.5% is supportable based on this analysis.

2.2.18. This section makes explicit that projects that protect society from rare but severe disasters have negative risk premia. Such arguments have been critical in the economics of climate change (Section 12.3). We recommend to HM Treasury that, in exceptional circumstances, projects that have such social insurance benefits (i.e. that pay off in bad states of the world) have a risk premium of minus 0.5% applied to them. Pandemic protection programmes and defence may have these properties.

2.2.19. Since our approach is to triangulate different information sources, in Chapter 11 we consider the SOC risk premia. By considering a range of sources of evidence, our central estimate of this value for the average public project is 1.9%. Public projects have a low beta, the equity premium is currently at c.5%, and the government does not face the taxes or frictions of financial investors. The gap between the STPR and SOC risk premia is narrower than has often previously been implied by the literature.

2.2.20. The survey data given in Tables 11.3, 11.4, together with the international guidance summarised in Table 4.1, all suggest that, for an average public project, the risk premium should lie in the range 0%-2%. Having appraised all this evidence our recommendation is that HM Treasury applies a risk premium of 1% for the average public project.

2.2.21. In Section 11.4 we consider whether the risk premium should vary between projects (assuming that they don't meet the social insurance criteria). While the theoretical case to do so is strong, and many of our survey respondents thought this should be our recommendation (Table 11.2), we have decided not to do so. We feel that, for many projects, such an adjustment is disproportionate and that the evidence is at present not available to do this practically.

2.2.22. There are two caveats to this recommendation, however. First, if there are reasons to believe that making such a project-specific risk adjustment is practical, proportionate, and evidence based, then a case should be made on an exceptional basis to HM Treasury. Second, we recommend that HM Treasury keeps this matter under review, improving the evidence base on the systematic risk of public projects, and considering whether ignoring such project-specific effects is leading to sub-optimal economic analysis.

2.2.23. There is a literature on the term structure of the risk premia; see, for example, Gollier (2014). Some of our empirical models, particularly in Section 12.4, display such characteristics. However, the economic drivers of this term structure also impact on the expected benefits; see Figure 12.2. At present, we believe that making suitable adjustments to both the numerator and denominator of the present value equation is challenging because of the evidence base available and also on considerations of proportionality. We therefore recommend that risk premia are fixed at 1% (and minus 0.5% for social insurance projects) for all maturities.

2.2.24. When combined with the term structure of the risk-free rate, this implies a forward discount rate for standard projects of 3.0% up to 30 years, 2.5% from 31–75 years, and 2.25% from 76–125 years. For social insurance projects, the forward discount rate is 1.5% up to 30 years, 1.0% from 31–75 years, and 0.75% from 76–125 years. For maturities longer than 125 years, formal welfare analysis is recommended.

2.2.f. Private finance models

2.2.25. We consider private finance models in Section 17.2. Two additional complexities arise compared to public financed projects.

2.2.26. First, the SOC-STPR divide becomes of greater significance because the unitary charge is generally set based on a private sector cost of capital to give the private partner an acceptable return. However, under current guidance, this is discounted at the standard Green Book rate. If market and social rates differ significantly, this introduces a problem. However, as we will show in subsequent chapters, this divide is currently much narrower than it has been historically and therefore this is less of a point for concern than it has been in recent years.

2.2.27. Second, private finance models split risk in specific ways between the public and private sector. Generally, the private sector has high initial risk (often during a construction phase) but then receives very low risk unitary charge payments from the public sector.

2.2.28. From the public sector perspective, this arrangement with the private sector reduces initial risk but increases longer-term risk as the unitary charge acts like gearing, giving the public sector something resembling (economically) an equity claim on the project.

2.2.29. Our recommendation is that both the cost of capital that is used to set the unitary charge and the social discount rate used to calculate the present value reflect the risk to each party. The extent to which this moves the discount rate away from the standard 3% rate is a matter of proportionality depending on the nature of the risk and the economic significance of the project.

2.2.g. Place-based and environmental objectives

2.2.30. We have discussed environmental objectives in Principle 2: we support the conclusions of the Environmental Discount Rate Review (HM Treasury 2021) that changes to environmental scarcity and limited substitutability of non-market goods should be reflected in relative pricing in the numerator of the present value equation. We discuss this in some detail in Chapters 5 and 15.

2.2.31. We discuss place-based objectives in Chapter 16. Our recommendation, which also reflects the strong majority opinion of our survey respondents (see Table 16.1), is that welfare weights in the numerator of the present value equation should adjust for such matters. As with environmental effects, we give some detailed thoughts on how this might best be done.

2.2.32. A reconsideration of the way in which welfare weights are constructed is also recommended, both their form and whether they should vary over time.

2.2.h. Transformational projects

2.2.33. We hold off an explicit discussion of transformational projects until the end of these Technical Annexes (Section 17.4). These are likely to involve all the complexities that we have considered so far; regional effects, environmental effects, the use of private finance models, and potentially very long maturities. In addition, they may have other complexities that we cover in Section 5.4. They may involve Real Options — the opportunity to expand, contract or withdraw from a project. They may also have very high levels of uncertainty associated with them (Knightian uncertainty). They may even affect the growth rate of the economy itself.

2.2.34. Because of this complexity, their economic significance, and the detail of the estimated costs and benefits that public bodies produce, applying a standard centralised discount rate fails the principle of proportionality in our opinion. Our recommendation is that major transformational projects are afforded detailed welfare analysis aligned with their scale and macroeconomic impact.

2.2.35. The analysis in many of the chapters that follow will, we believe, be of value to HM Treasury as it considers how such welfare analysis might be undertaken.

3. How We Have Undertaken this Review

We start the main body of these Technical Annexes by describing the *process* we have taken to writing this review to be consistent with the requirements laid out in paragraph 1.1.11 above. We have:

- Considered the evolution of the Green Book discount rate over time, and compared and contrasted that with discount rates both in other countries and in other public sector work in the UK; for example, discount rates used by utility regulators. This has also required us to consider relevant consultancy advice given to public bodies; for example, Oxera (2026);
- Reviewed the academic literature, consolidating the background given in HM Treasury's last review of the discount rate (Freeman et al. 2018) and updating this for relevant evidence since 2018;
- Appraised new empirical evidence on, for example, expected economic growth rates, bond yields, risk premia, and how the declining discount rate schedule might best be modelled and calibrated in the current environment;
- Distributed a survey to allow for input into this review from a wide range of interested parties;
- Met with the governmental Observer Groups;
- Met with specific academic Expert Panels to discuss particular parts of the Terms of Reference;
- Held regular discussions with, and provide updates to, HM Treasury. HM Treasury chose to take the working version of these Technical Annexes on the 31st March 2026 to meet the requirement for an interim, non-public, report. We presented interim findings to HM Treasury in a face-to-face meeting on the 19th May 2026 and via an Executive Summary document that was shared with the Observer Groups for feedback on the 29th May 2026.

3.1. The survey

3.1.1. On the 19^{th} January 2026, we launched a Qualtrics survey through the University of York that allowed people to input into this review process. We will refer to the results of this survey ("our survey") in subsequent sections.

3.1.2. We publicised the survey in a number of ways, focusing on the academic community because of the Terms of Reference requirement cited in paragraph 1.1.11 above. However, our survey was open to all; both academic and non-academic. Specifically:

- We posted the survey link on our social media accounts (LinkedIn, BlueSky, Substack), asking people to share more widely.
- We wrote to the Royal Economic Society (RES), the British Accounting and Finance Association (BAFA), and the Society of Benefit-Cost Analysis (SBCA) to ask them to distribute their surveys more widely. BAFA posted about this on LinkedIn, and the RES publicised it via its "Conference of Heads of University Departments of Economics" newsletter, for example.
- We wrote to a member of both the Economics department and the Accounting & Finance group in major research-active Universities in the UK, asking for them to forward on the survey link to their colleagues. "Research intensive" here meant the Russell Group, London Business School, and other high-performing departments in the 2021 Research Excellence Framework (City, Bath, UEA, Lancaster, Surrey, Strathclyde, Cranfield, Loughborough, Sussex, St Andrews for Accounting, Finance & Economics, plus Birkbeck, Brunel, Essex, Kent, and Royal Holloway College for Economics only). In total we sent out 70 emails via this route — in a very limited number of cases, one contact was sufficient to contact all relevant members of faculty in that University.
- We sent a link to the survey by email to all members invited to participate in the academic Expert Panels.

3.1.3. While the survey text was completely identical in each case, we ran six separate versions of the survey, sending different links to each group. This allowed us to identify which route respondents arrived by: social media / email to departments / RES / BAFA / SBCA / identified experts.

3.1.4. The preamble to the survey ran as follows, with relevant hyperlinks to external sources where appropriate throughout the survey: "HM Treasury is currently reviewing its Green Book discount rate. This rate is used by the UK government and other public bodies to compare costs and benefits occurring over different periods of time on a consistent basis. This survey gives you the opportunity to input into this review. We are requesting responses by 27th February 2026.

An overview of current discounting practices is given in Section A6 of the Green Book. Further details can be found in an associated background paper.

The survey is being run by Professors Mark Freeman (University of York) and Ben Groom (University of Exeter) under appointment from HM Treasury.

We believe there are no significant risks associated with your participation in this survey. We are not collecting any personally identifying information apart from some professional information within very broad categories. Because of this, all responses will be fully anonymous. For further details, please read the Participant Information Sheet. This survey has received ethics approval from the University of York.

Analyses of this survey will be published, including potentially on the gov.uk website and in academic channels (e.g., working papers, conference presentations, and publications in scientific journals). These may refer to individual anonymous responses, and a dataset of responses may be made publicly available.

Before proceeding, please confirm that you understand that, by completing this survey [Yes/No]

- Your responses will be used in the aggregated analysis of this survey,
- Your responses may be anonymously quoted in any subsequent report that results from this survey, and
- Your responses may be anonymously archived in a publicly available dataset.

Please also confirm that you understand that you have the right to have your data removed from this study at any time before 31 March 2026. You can withdraw from the research at any point up to 31 March 2026 without giving a reason. Below is a textbox where you can enter a unique code - please ensure this retains your anonymity. If you wish to subsequently withdraw your data, please send this unique code to mark.freeman@york.ac.uk. Your response will then be identified and deleted.

The thirteen discount rate related questions in this survey reflect the Terms of Reference as set out by HM Treasury for this review. Please note that:

- All rates should be given in real (inflation-adjusted) annualised terms for the UK context.
- The discount rate must reflect best evidence while also being applicable in practice across a wide range of public projects.
- We use "public project" as a general term to cover any project, policy, or regulation that falls under Green Book guidance.
- We use "project risk" as a general term to cover all types of uncertainty over the future outcomes from any given public project. This includes, but is not limited to, systemic and catastrophic risk.

An overview of current discounting practices is given in Section A6 of the Green Book. Further details can be found in an associated background paper. The discount rate is currently given by a Social Time Preference Rate (STPR) approach based on fundamental considerations of social welfare. Excluding systemic and catastrophic risk, the Green Book follows the "Ramsey Rule". It is comprised of the rate of pure time preference (a measure of societal impatience, currently set at 0.5%) plus a wealth term. The wealth term is the product of the elasticity of marginal utility of consumption (currently set equal to 1) and the expected annualised real per capita growth rate of consumption (currently set at 2%). Excluding any adjustment for project risk, the "risk-free" STPR therefore equals 0.5% + 1*2% =2.5%. In addition to this, there is a 1% premium for "systemic" and "catastrophic" risk, which increases the discount rate to 3.5%. In what follows, we will separate questions about the "risk-free" and "risk-adjusted" discount rate".

3.1.5. After this preamble, the list of questions was as follows:

- Do you think that the UK government should retain its current STPR approach, or should it instead change to a Social Opportunity Cost (SOC) approach based on the rates of return available to capital? [STRP/SOC/Unsure]. Please explain your answer
- Do the current Ramsey Rule values in the Green Book for the ""risk-free"" STPR remain valid for discounting a 10-year public project in the UK (excluding any project risk adjustment)? These are: Rate of pure time preference = 0.5%, Elasticity of

marginal utility of consumption = 1, Expected average annual real per capita consumption growth rate of 2%. [Yes/No]. Please explain your answer.

- If you were advising HM Treasury, what single values would you recommend for the following components of the Ramsey Rule. Where appropriate, please give your answers in average real (inflation-adjusted), per capita, annualised percentages specifically for the UK (e.g., "x" = "x% real per year"): (i) Rate of pure time preference (%), (ii) Elasticity of marginal utility of consumption, (iii) Expected consumption growth over the next 10 years (%), (iv) Expected consumption growth over the next 100 years (%).

- If the UK government chooses instead to switch to an SOC approach, which financial asset, private project, or return to capital should be used to determine the "risk-free" discount rate for a 10-year public project?

- Should the discount rate (excluding any project risk adjustment) vary with the time horizon or be constant for all time horizons? [Vary with the time horizon/Stay constant for all time horizons/Unsure]. Please explain your answer.

- The Green Book currently applies a single premium of 1% to all public projects for "systemic" and "catastrophic" risk. If HM Treasury decides to keep a single risk premium for all public projects, what value do you think this should take? [Risk premium (%)]. Please explain your answer.

- Would you advise that the Green Book recommends a project risk premium that varies according to the public project being appraised? [Yes/No/Unsure]. Please explain your answer.

- Please give values of the discount rate that you think should be applied in the UK for public projects of different project risks and maturities, expressed as annualised real percentages (e.g., "x" = "x% real per year") [4 × 3 boxes for responses: Risk-free project, Low risk project (lower quartile), Median project, High risk project (upper quartile) × 10-year maturity (%), 50-year maturity (%), 100-year maturity (%)]

- What are the main conceptual issues and considerations when using social discounting in appraisals involving private finance models, such as public-private partnerships?

- HM Treasury is currently working towards building place-based business cases. These "... will bring together the different projects that are needed to achieve the objectives of a particular place. Place-based business cases will make sure that the government properly assesses the complementarities between different projects, such as housing and transport". Should the discount rate be adjusted for place-based objectives? [Yes/No/Unsure]. Please explain your answer.

- Following a prior review, HM Treasury concluded that "... the Green Book should not change the discount rate for environmental impacts" and, instead, "... favoured improved valuation for environmental impacts and updating these estimates to reflect latest evidence", such as using relative price change adjustments. HM Treasury is now again seeking guidance on this question. In your opinion, should the social discount rate be adjusted for environmental scarcity and the limited substitutability of nature? [Yes/No/Unsure]. Please explain your answer.

- This discount rate review is being actioned "...to make sure that the government is taking a fair view of the long-term benefits that arise from transformational investments". The current version of the Green Book defines "transformational change" as "... a radical permanent qualitative change in the subject being transformed, so that the subject when transformed has very different properties and behaves or operates in a different way". Should the social discount rate be adjusted for transformational projects? [Yes/No/Unsure]. Please explain your answer, and say how you think such an adjustment might be implemented if HM Treasury proceeds with this.

- Is there anything further that you would wish HM Treasury to take into consideration as it reviews its discount rate?

- Do you currently reside in the UK and/or have UK citizenship? [Yes/No]

- What is your highest academic qualification? [School level or equivalent, Undergraduate degree or equivalent, Masters degree or equivalent, Doctorate].

- Which of these best describes the field of your highest academic qualification? [Economics / Finance & Accounting / Business & Management / Law / Philosophy / Another Social Science / Another Arts & Humanities Discipline /Natural Sciences, Healthcare, or Engineering].

- For how many years have you been working since leaving full-time education? [Less than 5 years, 6-15 years, More than 16 years].

- And in relation to your experience (0 = None, 100 = A great deal): (i) The use of discount rates in your daily work, (ii) Your knowledge of the Green Book, (iii) Your knowledge of social discounting in other countries, (iv) Your knowledge of private sector discounting".

3.1.6. We report, discuss and analyse summary statistics for the quantitative and binary responses in subsequent chapters where most directly relevant. In Appendix A we report the characteristics of the respondents and give the full text of all qualitative comments received.

3.2. Observer Groups

3.2.1. HM Treasury organised three meetings between us and the government Observer Groups. These were hybrid meetings (face-to-face in HM Treasury) held with us between mid-March and mid-April 2026. The following government departments contributed to these sessions:

- Home Office;
- UK Export Finance;
- HMRC;
- West Midlands Combined Authority;
- Welsh Government;
- Transport for London;
- Homes England;
- Ministry of Defence;
- Government Actuaries Department;
- Department for Transport;
- Scottish Government;
- Ministry of Housing, Communities & Local Government;
- Department for Energy Security & Net Zero;
- Department for Work & Pensions;
- Department for Culture, Media & Sport;
- Greater London Authority;
- Liverpool City Region Combined Authority;
- Department for Environment, Food & Rural Affairs;
- Ambition North Wales;
- Department of Health & Social Care.

3.2.2. We have also received emailed input from the Department for Science, Innovation & Technology and spoken with the Department for Transport separately.

3.2.3. HM Treasury also sent links to the survey to governmental departments, with a response deadline of 30th April 2026. We report the outcome from their responses separately in Appendix A and will discuss and analyse these results within the body of these Technical Annexes in the relevant sections.

3.2.4. We also organised a joint call with Oxera and the Financial Conduct Authority (FCA) concerning the FCA's recent review of its use of the Green Book discount rate (Oxera 2026).

3.3. Expert panels

3.3.1. We identified a relatively small number of academic experts in fields directly related to specific parts of the Terms of Reference. Many of these experts have either worked with governments in the UK, EU and/or the US to help set discounting policy, and/or have extensive knowledge of the application of this policy in particular contexts. We sent each a survey link and an invitation to join one of our Expert Panel online meetings, held between early April and early May 2026, with each meeting having a focus on different elements of the Terms of Reference. The following academics participated in one of these seven online meetings:

- Professor Seth Armitage, University of Edinburgh;
- Dr. Roberto Cardinale, University College, London;
- Professor Karl Claxton, University of York;
- Professor Maureen Cropper, University of Maryland;
- Professor Simon Dietz, London School of Economics and Political Science;
- Professor Moritz Drupp, ETH Zurich;
- Professor Eli Fenichel, Yale University;

- Professor Massimo Florio, University of Milan;
- Professor Christian Gollier, Toulouse School of Economics;
- Professor Paul Kelleher, University of Wisconsin-Madison;
- Professor Bengt Kristrom, Swedish University of Agricultural Sciences;
- Professor Ulrike Kornek, Kiel University;
- Professor Deborah Lucas, Massachusetts Institute of Technology;
- Professor Peter Mackie, University of Leeds;
- Professor Frikk Nesje, University of Copenhagen;
- Professor David Newbery, University of Cambridge;
- Professor Aude Pommeret, Professor at Université Savoie Mont Blanc; Scientific Advisor at France Stratégie;
- Professor Lisa Robinson, Harvard University;
- Professor Roger Vickerman, Univesity of Kent;
- Professor Herman Vollebergh, Tilberg University.

3.3.2. Lessons learned from the Observer Groups and Expert Panels will be brought into these Technical Annexes in the relevant sections under the Chatham House Rule. We greatly thank all the participants for their thoughtful, incisive, and constructive contributions to this Review.

3.3.3. Towards the end of this review and as we were reaching our conclusions, we consulted members of the academic community for additional input. In particular, to get opinions from beyond economists, Freeman ran a workshop at the Centre for Society, Nature and Organisations at the University of Nottingham, hosted by Professor Mihaela Kelemen, on matters related to the Terms of Reference. Groom attended an event in Toulouse hosted by Professor Christian Gollier (member of the expert panel), where he also spoke to Professor Aude Pommeret (member of the expert panel), Professors James Hammitt and James Stock (both Harvard University), Alain Quinet (SNCF Group), Professor Geoff Heal, (Columbia University), Professor Antony Millner (UC Santa Barbara), Professor Rick van der Ploeg (University of Oxford), and Professor Nicolas Treich (Toulouse School of Economics) amongst others on questions directly raised by this review. Again, their input is reflected in our final recommendations and we thank them for their guidance.

3.3.4. We also greatly thank Ryland Thomas from the Bank of England's 'millenium of data' project for providing us with the most up-to-date time series of UK consumption, GDP and population date ahead of the planned release. This we used to help calibrate the Ramsey Rule and estimate the term structure of the social discount rate. We also remember Roger Guesnerie who passed away in the Spring of 2026, whose work on relative price adjustments in CBA influenced some of our recommendations.

3.3.5. This document represents our best considered analysis and conclusions at the time of publication. As with all technical guidance, we reserve the right to update these Technical Annexes to incorporate additional evidence, Green Book user experience, and constructive feedback received following publication. Any future revision is intended to enhance clarity and completeness and should not be taken to imply any change in the fundamental conclusions, and should be regarded as part of the normal evolution of technical guidance.

3.4. Green Book framework

3.4.1. This review of social discount rates is explicitly framed within a Green Book context. This gives three key principles that we lay out before getting into more detailed analysis in subsequent chapters.

3.4.a. The Ramsey Rule

3.4.2. Current Green Book discounting guidance (HM Treasury 2026*a*) is framed around the following equation for the discount rate, r. We will refer to this equation as the 'Ramsey Rule' (occasionally 'Simple Ramsey Rule') throughout these Technical Annexes

$$r = \delta + \mu g + L$$

δ will be referred to as the 'utility discount rate' or the rate of 'pure time preference'. μ is the 'elasticity of marginal utility of consumption' while g is the real growth rate in per capita consumption. L is referred to in the Green Book as "Catastrophic risk, also known as 'systemic' risk, this represents an allowance for unpredictable risks that are not normally included in appraisal'. Over the course of these Technical Annexes, this term will become more closely associated with *systematic risk*, which is how we recommend it is described in any future edition. The Green Book introduces a variable $\rho = \delta + L$ which represents 'time preference'. We do not use ρ in our own analysis because δ and L are conceptually distinct concepts that should be treated individually.

3.4.b. Zero sum game

3.4.3. The Green Book says that:

> "… most government proposals are funded from pre-determined departmental budgets. Decisions about the overall level of public spending are made separately from, and in advance of, individual spending decisions." (HM Treasury 2026*c*, #6.44)

3.4.4. Whatever values HM Treasury ultimately decides upon for its discount rates following this review, there will be no more and no less money in government, or in specific departments, as a consequence. When it comes to investment decision-making, altering the Green Book discount rate is a zero sum game. A raising or lowering of this rate will favour some projects but such a decision must also, therefore, simultaneously disfavour others when spending decisions are being taken within a fixed budget.

3.4.5. The purpose of a discount rate review, therefore, is to apply relevant economic appraisal methods in order to help the government fund the most socially beneficial projects while diverting resources away from those that are less beneficial. This is a relative, not absolute, exercise.

3.4.6. Readers who feel that government spending levels are too low or too high, or believe that HM Treasury allocates these resources sub-optimally between different public bodies, will not find the solution to these issues in this review. Discounting is concerned with selecting the best individual projects under pre-set resource constraints, not setting the broader spending envelope.

3.4.c. Proportionality

3.4.7. The 2025 Green Book Review stressed the importance of proportionality when Green Book guidance is applied by its users. This principle is now captured in the Green Book itself:

> "Green Book guidance should be used proportionately. Business cases should contain sufficient detail to support sound decision making, but they should not take up more time and resources than is necessary. The effort involved in developing a proposal should correlate to its scale, cost, complexity and risk". (HM Treasury 2026*c*, #1.8)

3.4.8. The Observer Groups made clear to us that the Green Book is applied in an extremely wide range of settings that offer widely differing degrees of scale, maturity, risk, and complexity; from local projects under regional authorities, to the extremely long timescales of nuclear decommissioning, to the very large economic commitments of the High Speed 2 rail link and other mega-projects.

3.4.9. Under current Green Book guidance, all these very different types of project apply the same schedule of discount rates irrespective of their "… scale, cost, complexity and risk". By contrast, considerations of proportionality will lie at the heart of this review. Many of the issues that we raise in these Technical Annexes involve complex considerations of what constitutes social welfare, and how such considerations should best be incorporated into a UK social discount rate (or, indeed, whether a discount rate is the right framework at all). Many of these complexities will need to be carefully addressed for certain types of projects where they are most relevant. But recommending a Green Book discount rate schedule that captures all theses issues for every single project does, in our opinion, risk violating this principle of proportionality. As one respondent to our survey noted (paragraph A.2.5), "Don't let the ideal be the enemy of the useful".

3.4.10. The current Green Book setting of one discount rate, declining over time, is, in our opinion, suitable for many project appraisals, but considerations of proportionality requires more analytical depth for the most significant projects. In some cases, this will involve moving into the broader and more general framework of welfare economics from which discounting is derived as one specific case. We will return to this matter in our final Chapter 17.

4. Social Discounting across Time and Countries

We now turn to the social discount rate itself. In this Chapter, we:

- Briefly describe how social discounting in the UK has evolved over time;
- Contrast this with the evolution of social discount rates in the United States;
- Begin the comparison of current Green Book discounting practices with those applied by governments in other countries;
- Briefly describe how discounting is undertaken in some other UK public and regulatory environments.

The purpose of this chapter is to show that many of the questions raised by the Terms of Reference have been carefully considered by HM Treasury and many other international governments & public organisations over long periods without any clear consensus having yet emerged.

4.1. The evolution of Green Book discounting

4.1.1. We begin by sketching the evolution of social discount rates in the UK prior to the 2026 Green Book Review. Our purpose is not to give a full history — we refer the reader to the excellent review by Michael Spackman (2013) — but instead to show the fluidity of HM Treasury policy and thinking on this matter, particularly prior to 2003.

4.1.2. The foundations of the current Green Book discount rate can be traced back to at least as early as the 1967 HM Treasury document *Nationalised Industries: Review of Economic and Financial Objectives*. This "…set a standard 'test discount rate' (TDR) for nationalised industry investment appraisal. This (pre-tax) rate was set initially at 8 per cent in real terms and was presented as 'broadly consistent . . . with the average rate of return in real terms looked for on low risk projects in the private sector . . . [which] does not include allowance for the risks of individual investments'" Spackman (2013, p.191).

4.1.3. In 1969 the TDR was increased upwards to 10% real "...partly in response to evidence about the appraisal rates used in the private sector and partly because of the pressure on resources following sterling devaluation in 1967." Spackman (2013, p.191–192).

4.1.4. It is interesting to note that "There was at this time, in the academic literature, much debate about the relative merits of two approaches to public sector discounting, described then and now as 'social opportunity cost' (SOC) and 'social time preference' (STP) ... A contemporary note by a Treasury official (Carr 1966), criticising an academic suggestion that STP had a role, suggests that the Treasury had not at that stage mastered all the technical issues." Spackman (2013, p.191–192)

4.1.5. By 1973, HM Treasury focus had moved away from nationalised industries alone to a more general consideration of cost-benefit analysis. The document *Use of Discounted Cash Flow and the Test Discount Rate in the Public Sector* "... was the first of the Treasury guides known from 1980 as 'the Treasury Green Book', defining, from an economics perspective, the principles of UK government appraisal methodology" Spackman (2013, p.193).

4.1.6. For things other than lease-or-buy decisions, HM Treasury (1973, p.3) says: "The discount rate used throughout the public sector at present, (and since 1969), known as the Test Discount Rate or TDR, is 10% 'in real terms'—this phrase is explained below. The TDR was first developed in 1962 for analysing nationalised industry investment. The government has now extended its use to all investment in the public sector." Therefore the wider use of the TDR did not, at this stage, influence its numerical value that had been derived from economic considerations of nationalised industries.

4.1.7. By 1980, the rate had been reduced: "This ideal discount rate for the whole of the public sector has recently been reassessed at 5% (see Cmnd 7131) 'in real terms'—paragraph 18 gives an explanation of this term. However, in practice Departments have agreed to use a 7% TDR as a normal rule for public service appraisals. This is to allow for the tendency of appraisals in the public services sector to face special difficulties; market prices are not always available to value costs and benefits, and there may be a leaning towards appraisal optimism ie a tendency to overestimate benefits and underestimate costs. However, 5% may be used (a) where it can be shown by experience that the assessment of costs and benefits can be made with a good degree of objectivity; and (b) When the problem is one of simply comparing different techniques of production for meeting a given output." HM Treasury (1980, p.3). Here, the 5% baseline "... was a projection of the pre-tax rate of return on assets *achieved* by private companies. . . ." Spackman (2013, p.192).

4.1.8. A new version of the Green Book was published in 1982. This set the discount rate for public services to 5%, removing the 2% premium for optimism bias HM Treasury (1982, pp.12–13). As we understand the organisational history, this was in the middle of a period where the *public sector economists* in HM Treasury were advocating a STP rate while those *overseeing the nationalised industries* still preferred an SOC approach. However, with strong economic growth at the time, 5% could be defended under both SOC and STPR and therefore there was no particular need to take a stand on this choice.

4.1.9. The 1984 version of the Green Book (HM Treasury 1984) retains the 5% real test discount rate. It explains that "In choosing this rate as the opportunity cost of capital consideration is given to pre-tax rates of return achieved by private companies, the likely growth in income and consumption, and other factors affecting or indicative of the appropriate time preference rate for government" (p.14).

4.1.10. By the end of the 1980s, HM Treasury was finding it harder to remain neutral on the choice of whether to take an SOC or STPR approach. The Required Rate of Return (no longer TDR) for the nationalised industries was increased to 8% but for public sector bodies it increased only to 6%.

4.1.11. The 1991 version of the Green Book re-asserted this choice of 6% as the appropriate rate for social discounting. The annexes to this version of the Green Book make clear the distinction between the SOC and STPR approaches. They also discuss potential adjustments for *systematic risk* through the use of *certainty equivalent* net benefits, as well as (implicitly) the *equity premium puzzle*, both of which we will discuss in later sections. It concludes that "the systematic risk of ... public sector activities is small" HM Treasury (1991, p.83). The 1997 version of the Green Book (HM Treasury 1997) retained these discounting choices.

4.1.12. The Green Book of 2003 (HM Treasury 2003) introduced social discounting in what is highly recognisably its current form, with 3.5% real as the recommended rate. It explicitly based the rate on STPR considerations, with the components of the Ramsey Rule parameterised with their current values. The adjustment for 'catastrophic risk' was articulated, and a declining term structure of discount rates for the long-term were introduced for the first time not only in the UK but in social discounting globally.

4.1.13. In 2008, the Treasury published a supplementary report on very long-term discounting (Lowe 2008). This introduced a 'reduced rate' declining term structure, with the rate of pure time preference set equal to $\delta = 0\%$ rather than the standard $\delta = 0.5\%$. This brought consistency with the discounting choices made by the *Stern Review on the Economics of Climate Change* (Stern 2007). For very long-term projects, the 2008 guidance stated that business cases should report two net present values based on the standard rates and reduced rates respectively.

4.1.14. The discount rate was reviewed again in the mid-2010s culminating in a published report we co-authored with Michael Spackman (Freeman et al. 2018). This led to the 2018 version of the Green Book (HM Treasury 2018) retaining the Ramsey Rule parameter values and overall standard discount rate of 3.5%. This version gave more detailed justification of the discount rate and made two changes of note. First, health effects were excluded from the 3.5% discount rate, with 1.5% being recommended instead. We will discuss this in more detail below. Second, for very long-term projects, discounting should only be applied for 125 years using both the standard and reduced rates. Beyond that, considerations should be given more broadly to implied intergenerational wealth transfers.

4.1.15. In 2021, HM Treasury explicitly considered the discount rate in relation to the environment. To do this, it called together an "an expert, external review", which was an informed meetings of experts in the field chaired by one of the authors of this report (Groom). The conclusions from this review were published in 2021, and stated that "... the Green Book should not change the discount rate for environmental impacts. This was deemed to be an imprecise way of accounting for effects such as lack of substitutability and increasing scarcity in the environment. The interaction with valuation approaches was also considered and how these effects can already be adequately accounted for in the Green Book methodology through relative price adjustments and the uprating of values over the appraisal period". (HM Treasury 2021). Note that no specific changes for environmental prices were made despite this conclusion. In Appendix E we publish for the first time the informal, but more detailed, notes that Groom presented to HM Treasury following the meetings of experts. These informed the final 2021 conclusions, and we publish the notes here as they remain highly relevant to a number of issues raised in the 2026 discount rate review Terms of Reference.

4.1.16. During the time that we have been undertaking this review, HM Treasury has issued a new Green Book (HM Treasury 2026*c*), with separate discounting supplementary guidance (HM Treasury 2026*a*). This is somewhat ambiguous about discounting beyond 125 years (excluded in Table 3.A. but included in Table 3.B. for the *Stern Review* consistent value of $\delta = 0\%$). Otherwise, this version of the Green Book did not bring in any significant new changes to discounting practice given the timing of its publication in relation to this ongoing review.

4.2. The evolution of US social discounting

4.2.1. As the Terms of Reference require us to "... look at international comparisons to better inform judgements on the UK approach", it is useful to compare the evolution of the Green Book discount rates with those in the United States. There the Green Book equivalent is *Circular A-4* for regulatory impact analysis and *Circular A-94* for the allocation of Federal grants.

4.2.2. The first version of Circular A-94 in 1969 "... bases the discount rate calculation on an average of past current yields on Government bonds and is less than 5 percent at the present time. This should be compared with the actual current yield on Government bonds of 6% percent" (Bureau of the Budget 1969, p.18278). However, "The Bureau of the Budget will request specific higher rates for particular projects or program evaluation efforts", with discount rate tables given for $4\frac{7}{8}\%$ and 10%.

4.2.3. By 1972, Circular A-94 had fixed on the higher discount rate of 10% real. This represented "an estimate of the average return on private investment before taxes but after inflation" (OMB 1972, p.4) .

4.2.4. Circular A-94 was next updated in 1992. This reduced the real discount rate to 7%, which again "...approximates the marginal pretax rate of return on an average investment in the private sector in recent years" OMB (1992, p.9). However, it also notes that "... in analyzing a regulatory proposal whose main cost is to reduce business investment, net present value should also be calculated using a higher discount rate than 7 percent".

4.2.5. The Circular A-94 guidance of 7% was applied in the first version of Circular A-4 in 1996 (OMB 1996) This explicitly noted that the discount rate should *not* be adjusted for risk and uncertainty of future costs and benefits.

4.2.6. Well in advance of HM Treasury's Environmental Discount Rate Review in 2021, the 1996 version of Circular A-4 stated: "... increasing scarcity of certain environmental resources could increase their value over time relative to conventional consumer goods. In such a situation, it is inappropriate to use current relative values for assessing regulatory impacts. However, while taking into account changes over time in relative values may have an effect similar to discounting environmental impacts at a lower rate, it is important to separate the effects of discounting from the effects of relative price changes in the economic analysis. In particular, the discount rate should not be adjusted for expected changes in the relative prices of goods over time. Instead, any changes in relative prices that are anticipated should be incorporated directly in the calculations of benefit and cost streams".

4.2.7. The 1996 version of Circular A-4 also considered intergenerational discounting effects without drawing clear conclusions, stating that agencies "...should consult with OMB prior to conducting special analyses of regulations having substantial intergenerational effects".

4.2.8. Circular A-4 was updated in 2003 (OMB 2003). This brought about an important distinction between whether government interventions displaced capital spending or consumption. In the former case, the existing 7% real rate was retained. However, for activities that displaced consumption, a 3% rate was recommended. This represented a 30-year average value of the real yields on 10-year Treasury securities and was argued to be a market representation of the STPR. The guidance given was that two net present values should be reported, one calculated at 7% and the other at 3%, with potentially a higher value in some cases. This version of Circular A-4 also considered the case for using a declining term structure of discount rates, but ultimately left this as a sensitivity analysis issue.

4.2.9. The issue of intergenerational discounting increased in importance when an Interagency Working Group calculated the Social Cost of Carbon in 2010. For this, the discount rate was reduced to a central value of 3% real, with 2.5%, and 5% as sensitivity analysis rates (Greenstone et al. 2013, p.38).

4.2.10. In 2023, new versions of both Circular A-4 and Circular A-94 were published (OMB 2023*a*,*b*). The detail of the discounting guidance is given in Circular A-4. By considering the 30-year average of the yield on 10-year Treasury securities, and adding 0.3% as a small inflation adjustment between the inflation index for personal consumer expenditure and the consumer price index, the discount rate was reduced from 3% to 2%. The 7% rate was discarded entirely. For the first time, the Appendix of this version of Circular A-4 explicitly articulated a declining term structure of discount rates.

4.2.11. In terms of dealing with project benefit risks, Circular A-4 expressed a preference for adjusting the net benefits to *certainty equivalents*. However, Appendix D of Circular A-94 allows instead for an average risk premium of 1.1%, raising the discount rate to 3.1%.

4.2.12. Alongside these were new estimates of the Social Cost of Carbon that applied new approaches to intergenerational discounting (USEPA 2023, Newell et al. 2022, Rennert et al. 2022). These used a relatively complex approach that estimated the Ramsey Rule parameters based on observed short-term bond yields in order to determine the term structure of discount rates, thus intertwining the SOC and STPR approaches. The USEPA present values for the social cost of greenhouse gas emissions based on short-term discount rates of 1.5%, 2% and 2.5%.

4.2.13. When President Trump came to power for his second term in January 2025, he issued Executive Order 14192 "Unleashing Prosperity Through Deregulation" that null-and-voided OMB (2023*a*,*b*): "The Director shall revoke OMB Circular No. A-4 of November 9, 2023 (Regulatory Analysis), and all accompanying appendices, guidelines, and documents, and shall reinstate the prior version of Circular A-4, issued on September 17, 2003". This puts the 3% and 7% rates back into force.

4.2.14. The President also issued Executive Order 14154 "Unleashing American Energy", which stated that the "Interagency Working Group on the Social Cost of Greenhouse Gases (IWG), which was established pursuant to Executive Order 13990, is hereby disbanded, and any guidance, instruction, recommendation, or document issued by the IWG is withdrawn as no longer representative of governmental policy". This removes the 1.5%, 2% and 2.5% short-term rates applied in USEPA (2023).

4.2.15. The 7% and 3% rates in the 2003 version of Circular A-4 are based on long-term averages: "Over the last thirty years, (the real rate of return on long-term government debt) has averaged around 3 percent in real terms on a pre-tax basis" OMB (2003, p.33). Therefore, data from before 1975 is being used to inform the current social discount rate in the United States, while economic information from any date after 2003 is not.

4.3. Other international social discounting guidance

4.3.1. UK and US social discounting guidance point to the two very different routes that governments can take to this issue: the UK's firmly rooted in the STPR approach and the US's firmly rooted in the SOC approach. It cannot be reasonably concluded that these differences arise because one country has considered these issues for longer, or more carefully and deeply, than the other. The long evolution of discounting practice in both countries just points to the different ways that the pros and cons of each approach have been summarised. This again makes clear that there are no single "correct" answers to a number of questions that have been posed to us in the Terms of Reference.

4.3.2. Perhaps unsurprisingly, the role of specific individuals has played a major role in the evolution of practice in each country, and more generally internationally (Groom & Hepburn 2017). The ultimate choice of discount rate may never be purely scientifically clinical.

4.3.3. In Table 4.1, we very briefly summarise some of the other international guidance that is in force at present in North America, Europe, and Australasia. This information is adapted from Rijksoverheid (2025) and Groom et al. (2022). This table concentrates on:

- Whether the rate is predominantly based on an SOC or STPR approach;
- If the STPR approach is preferred, what the choice of δ, μ for the Ramsey Rule are;
- The actual rate or rates applied, and whether a risk premium is explicitly included;
- Whether the country applies a declining term structure of social discount rates.

4.3.4. In specific sections of these Technical Annexes, we will return to the most relevant parts of these international guidelines. For the time being, we would just like to stress the variety of intellectual frameworks, parameter choices, and numerical values used in developed countries across the Western world, with the UK and US broadly acting as the two book-ends to these approaches.

4.4. Other UK social discounting

4.4.1. The Green Book discount rate is used more widely in the UK public and regulatory sectors than the specific situations for which it has been developed. Notably, the Financial Conduct Authority (FCA) uses this rate for the cost-benefit analyses of financial regulations, which it is required to undertake by the Financial Services and Markets Act (2000) (FCA 2024). The FCA has recently undertaken an independent review of its decision to use this rate (Oxera 2026), accepting the recommendation that it should retain the use of the Green Book values.

4.4.2. Yet in a number of other areas of the public sector, different rates entirely are applied, with very different intellectual justifications from that used for the STPR in the current version of the Green Book.

4.4.3. An example of this is the public sector pensions (SCAPE) discount rate. HM Treasury recently published a review of a major consultation on this rate (HM Treasury 2023):

> The consultation sought views on two possible methodologies for the SCAPE discount rate. One method was based on long-term future GDP growth expectations set by the OBR. The second was based on the Social Time Preference Rate (STPR) prescribed in HM Treasury guidance for project appraisal … The government concluded that a methodology based on long-term future GDP growth expectations, as set by the OBR, best met the balance of objectives. The consultation response notes that "this methodology will best provide intergenerational fairness by ensuring that pension promises are made in a way that is sustainable and affordable to future taxpayers, which will support the long-term stability of the public service pension system" … (because this rate) is more directly related to the expected growth in the tax base, which is the source of income from which future pension payments are ultimately funded." (Review website).

On this basis, the SCAPE rate currently sits at 1.7% above CPI, based on "…the expected long-term GDP growth figures, published by the OBR in July 2022." This is clearly well below the current STPR Green Book discount rate.

4.4.4. A different approach is taken to setting the Government Actuaries Department Personal Injury Discount Rate (PIDR), which was last reviewed in 2024. The rate set here is 0.5%, real based on an expert review. Here the core principle is based on the investment returns that claimants might reasonably expect based on a well-balanced portfolio of market assets; a clear SOC approach.

4.4.5. The Government Financial Reporting Manual (HM Treasury 2026*b*) (FReM) takes its own approach to discounting. These rates are laid out in Public Expenditure Systems (PES) papers, but these (as we understand it) are not generally in the public domain. However, current rates are reported in Chapter 4 (Annex 7) of the Department of Health and Social Care group accounting manual. This gives a range of nominal general provision discount rates ranging from 3.64% for maturities of ≤ 5 years, to 5.32% for maturities between 10 and 40 years. By contrast, the post-employment benefits discount real rate is set at 2.95%. These rates vary annually.

4.4.6. A further example of where the Green Book discounting framework is not applied is found in HM Treasury's Financial Transaction Control Framework (FTCF). This prescribes rules for when the public sector acquires or sells financial assets, including setting target rates of return for such financial transactions (FTs). This rate is effectively set to a gilt yield because this ensures fiscal neutrality, although it does note that "… generating a rate of return at or above the government's cost of borrowing may not be sufficient for an FT intervention to be welfare-generating in an economic sense". This again recognises the STPR/SOC distinction while coming down on the side of a risk-free SOC.

4.4.7. The value-for-money model of the British Business Bank's (BBB) Interim Evaluation of British Patient Capital, Table 10-3, appears to blend the SOC and STPR approaches. It both prescribes the Green Book discount rate of 3.5%, and has separate opportunity costs of capital of 14.1% for the private sector (based on venture capital fund returns) and 1.5% for the public sector (based on average gilt yields).

4.4.8. We stress that we are not experts in the SCAPE, PIDR, FReM. FTCF, or BBB discount rates and costs of capital, or other discount rates in the public sector such as the Existing Use Value for Social Housing (EUV-SH) discount rate. We have not been involved in any discussions around these as they lie outside the Terms of Reference for this review. We merely note here that, even within HM Treasury and the Government Actuaries Department, there are different intellectual approaches to discounting depending on the matter at hand and whether there is a requirement to follow standard accounting approaches. As far as we can tell, these matters are resolved by the individual teams working in each area rather than a broader, coordinated, HM Treasury approach to how and why different discount rates and costs of capital might by applied in different contexts under a single unified view of the social purpose of these calculations.

4.4.9. Perhaps most relevant for the discussions that follow later in these Technical Annexes, regulated utilities in the UK work under price controls to:

> "... protect consumers from excessive prices and to incentivise companies to invest, innovate, deliver cost efficiencies, and provide a decent quality of service ... (while) ensuring regulated companies have a reasonable expectation of earning returns which compensate capital providers for the risk of investing in these businesses." (UKRN 2023, p.6).

4.4.10. The approach taken to balancing the needs of consumers and investors is underpinned by the regulated cost of capital that utility companies can employ. The approach specified by the UK Regulator Network is based very heavily on a private-sector weighted-average cost of capital (WACC) method. The current values are given in UK Regulators Network Costs of Capital 2025.

4.4.11. Since the potential for utility re-nationalisation remains an active topic of debate (Johnson et al. 2025), it is interesting to note that any change in ownership structure would significantly alter the discount rates used in the industry from WACC values to the Green Book STP rate. We will return to this question in some detail when we consider the choice of STPR vs. SOC (Chapter 6), and also the treatment of project-specific risk in the social discount rate (Chapter 11).

Country	Year	SOC/STPR	Parameters	Rate	Term structure
Australia	2025	SOC		7%	5.4% $31-75$; 4.8% $76-125$; 4.3% $126-200$; 4.0% $201-300$; 3.7% ≥ 301
Canada	2022	SOC STPR	$\delta = 0\% - 0.5\%$ $\mu = 1.3$	SOC 7% STPR 3%	Flat
Denmark	2021	SOC		3.5% $(r_f = 2.0\%+$ $r_p = 1.5\%)$	$2.5\%(1.75\% + 0.75\%)$ $36-70$; $1.5\%(1.5\% + 0\%)$ > 70
France	2023	STPR		3.2% $(r_f = 1.2\%+$ $r_p = 2.0\%)$	Rate valid to 2070
Germany	2025	SOC		1.2% nominal	Maturity matched
Ireland	2019	STPR	$\delta = 1\%$ $\mu = 1 - 1.5$	4.0% $(r_f = 4.0\%+$ $r_p = 0\%)$	3.5% $31-60$; 3.0% $61-100$; 2.5% $101-175$; 2.0% $176-275$; 1.5% ≥ 276
Netherlands	2025	SOC		Central 2.8% Low 2.2% High 4.2%	Central 1.8% > 35 Low 1.4% > 35 High 2.7% > 35
New Zealand	2024	SOC STPR	$\delta = 0\% - 2\%$ $\mu = 0.25 - 1.5$	SOC 8% STPR 2%(2% + 0%)	SOC Flat STPR 1.5% $31-100$ STPR 1.0% ≥ 100
Norway	2021	SOC		4.0% $(r_f = 2.5\%+$ $r_p = 1.5\%)$	$3.0\%(2.0\% + 1.0\%)$ $41-75$; $2.0\%(2.0\% + 0\%)$ > 75
USA	2003	Mainly SOC		Both 7% and 3%	Flat except for intergenerational effects
USA	2023	Mainly SOC		3.1% $(r_f = 2.0\%+$ $r_p = 1.1\%)$	Long-term decline: See Appendix, Circular A-4
Sweden	2026	STPR	$\delta = 0.5\%$ $\mu = 1$	3.5% $(r_f = 2.5\%+$ $r_p = 1.0\%)$	Flat

Table 4.1.: Current international discounting guidance. All rates are real except for Germany. Based on Rijksoverheid (2025, p.81) and Groom et al. (2022), which provides information for additional countries. France and USA (2023) allow for the risk premium (r_p) to vary by project. Australia applies a rate of 4% for nonmarket goods, while Canada uses a rate of 3% for environmental and health goods. France and the Netherlands both use a relative price adjustment for environmental goods. Germany uses different rates for envionmental and transport effects. The USA 2003 guidance is currently under application by the second Trump administration.

5. The Foundations of Discounting

5.1. Introduction

5.1.1. In Chapter 4, we emphasised that many international governments have deeply considered, over long periods of time, how best to estimate the social discount rate. Despite all this work, no overarching consensus has emerged that HM Treasury can draw upon.

5.1.2. In this chapter, we therefore turn to the core foundations of discounting. This will help us better understand why such differences in international practice continue to prevail. Our focus remains on the practical questions raised in the Terms of Reference. However, we will show that these questions cannot be properly addressed without going back to such fundamental core principles.

5.1.3. Specifically, in this chapter, we concentrate on explaining the conditions under which social discount rates and present values have policy-relevant meaning. This chapter first emphasises that adjustments that allow for place-based, transformational, health and/or environmental projects can be made in either the numerator or denominator of the present value equation. There is no single correct approach because the two are (mathematically) equivalent. However, there are clear reasons to prefer adjusting the net benefits over changing the discount rate. Second, we show that social discounting cannot be applied in all cost-benefit analysis situations. Sometimes no discount rate exists at all.

5.2. Present Values

5.2.1. The reason why we need a social discount rate is to calculate the present value (PV) of the costs and benefits of public projects:

> "The Green Book recommends using discounting to compare social costs and social benefits occurring over different time periods on a consistent basis. Discounting enables profiles of benefits and costs stretching over many years to be expressed in 'present value' terms". HM Treasury (2026*c*, #6.52)

5.2.2. We will use the notation P_i to denote the per capita present value of any public project, i. The total present value is therefore MP_i, where M is the total population. This distinction between total PV and per capita PV will become of relevance below when we consider place-based discount rates, but we introduce it here for completeness.

5.2.3. To start, we make three key assumptions that are necessary to have a solid foundation for discounting in worlds of both certainty and uncertainty:[1]

A.1. If we scale up a public project to k times its original size, then the PV of this larger project is just k times the PV of the original project: $P_{k\times i} = kP_i$.

A.2. If we combine two different projects, then the PV of this is just the sum of the PVs of the two projects individually: $P_{i+j} = P_i + P_j$.

A.3. The possible values that all (stochastic) variables might take at any point in the future can be identified today, and we are able to assign a probability to each of these possible outcomes.

Assumptions A.1. and A.2. assure that present values are *linear* (more precisely, a "linear functional"), which will be important in the discussions that follow. Assumption A.3. allows us to calculate expectations, variances, and covariances (assumed finite). For discounting under certainty in the sense we define it in Chapter 7, there is only one possible outcome for all variables at any given future date, and this will occur with 100% probability. We extend this to situations with uncertainty in Chapters 9, 10 and beyond.

5.2.4. First apply assumption A.2. Consider any project i with maturity T. Denote its real per capita annual net benefits by $b_{i1}, b_{i2}, \ldots b_{iT}$, expressed in monetary units (more formally, 'the consumption numeraire'). Denote the present value of these individual net benefits by $p_{i1}, p_{i2}, \ldots p_{iT}$. Assumption A.2. allows us to view the per capita present value of the total project as just the sum of the present values of each of the annual real per capita net benefits: $P_i = p_{i1} + p_{i2} + \ldots + p_{iT}$. For discounting purposes, we can consider any public project as a portfolio of individual projects, each of which gives benefits in just one year. This process will be familiar to readers who price Treasury bonds by considering them as if

[1] A formal treatment of this issue would phrase these assumptions more precisely.

they are a portfolio of zero-coupon bonds of different maturities.

5.2.5. The standard PV equation is then given by one of the two following equations depending on whether discrete or continuous discounting is being used:

$$p_{it} = \begin{cases} \dfrac{E[b_{it}]}{(1+r_{it})^t} & \text{Discrete} \\ E[b_{it}]\exp(-r_{it}t) & \text{Continuous} \end{cases}$$

where $E[\cdot]$ denotes the mathematical expectation, which is well-defined by Assumption A.3. The discount rate, which is the focus of this review, is denoted by r_{it}. Notice that this discount rate can be both $i-$ and $t-$ dependent. There is no theoretical reason why we should expect different public projects with the same maturity to have the same discount rate. And there is also no reason to theoretically assume that, for a long-horizon project, the same discount rate should be used in the early years as in the later years. The Green Book uses discrete discounting but for the next three chapters it will be more convenient to work with continuous discounting; we will return to this distinction below.

5.2.6. The Green Book uses discounted values in a number of different contexts. To understand these, the Green Book defines the following two variables (HM Treasury 2026*c*, #6.93):

- **Present value of monetisable social benefits** (*PVB*): The sum of the monetisable social benefits of a particular option, which have been converted into real terms and present value terms.
- **Present value of monetisable social costs** (*PVC*): The sum of the monetisable social costs of a particular option, which have been converted into real terms and present value terms

5.2.7. From here, the two main measures of social value (HM Treasury 2026*c*, Table 10) are:

- **Net Present (Social) Values: (NPVs or NPSVs):** $PVB - PVC$. Equivalently, we can see this as the sum of the expected discounted net benefits. This is a measure of the expected change in welfare from undertaking a new project.
- **Benefit Cost Ratio (BCR):** PVB/PVC. Note that this is equal to $NPV/PVC+1$, and therefore there is a clear link between the BCR and NPV measures. The BCR measures the expected social benefit generated for each pound of social cost.

5.2.8. Discounting is also used in cost-effectiveness analysis, which the Green Book describes as when "...the real discounted social costs of the proposal are considered against the unmonetisable benefits" (HM Treasury 2026*c*, #2.26).

5.2.9. As the Green Book itself notes, more care has to be taken when working with BCRs than NPVs: "For example, some options may increase carbon emissions, while others may reduce them. One approach would be to define 'reduced emissions' as a benefit and include them in the numerator of the BCR. An alternative would be to define 'increased emissions' as a cost and include them in the denominator. This choice affects the BCR, but not the NPSV." (HM Treasury 2026*c*, #6.96). We would also note that BCRs are not additive in the way that NPSVs are. In addition, in a world of uncertainty, $E[x-y] = E[x] - E[y]$, but $E[x/y] \neq E[x]/E[y]$. The ordering of the expectation therefore matters for BCRs (which are defined in $E[x]/E[y]$ terms), but not NPSVs.

5.2.10. For our purposes in these Technical Annexes, we work generically with present value calculations, which lie at the heart of all these social value measures, without drawing conclusions on which metric is most appropriate for any economic evaluation. This is an issue that lies outside the Terms of Reference.

5.3. Stochastic Discount Factors

5.3.1. The theory of discounting was put on a rigorous footing by a series of papers that have become seminal in asset pricing (Harrison & Kreps 1979, Chamberlain 1983, Chamberlain & Rothschild 1983), and now widely covered in advanced-level textbooks including Cochrane (2009) and LeRoy & Werner (2014). However, it has not yet deeply permeated into the social discounting discourse. This, we believe, helps explain why there continues to be deviations in international government discounting guidance that HM Treasury can resolve through this review.

5.3.2. Central to this rigorous foundation was the observation that Assumptions A.1.-A.3. are more mathematical than economic — all we are formally assuming is that present values are linear functionals within a well-defined single probability space. When viewed this way, it is clear that mathematics alone can be expected to tell us important things about discounting.

5.3.3. Specifically, Assumptions A.1.-A.3. guarantee (through the Rietz Representation Theorem) that there must exist a variable π_t such that, for any public project, i:

$$p_{it} = E[b_{it}\pi_t]$$

5.3.4. This is a highly powerful equation — it is known as the “Fundamental Theorem of Asset Pricing” — which again has no economic content beyond that contained in Assumptions A.1.-A.3.. The variable π_t is known as either the *stochastic discount factor (SDF)* or the *pricing kernel*. What is key to note about this variable is that it *does* depend on the maturity, t, but it *does not* depend on the individual public project being considered, i. For any given maturity, the same SDF can be used for all public projects that public bodies are economically appraising.

5.3.5. By comparing the two expressions we have presented for p_{it} (using continuous discounting) in offset equations, simple re-arrangement shows that:

$$r_{it} = \underbrace{-\frac{1}{t}\ln\left(E[\pi_t]\right)}_{\text{Risk-free rate}} \underbrace{-\frac{1}{t}\ln\left(1 + \text{Cov}\left[\frac{b_{it}}{E[b_{it}]}, \frac{\pi_t}{E[\pi_t]}\right]\right)}_{\text{Project-specific risk premium}}$$

5.3.6. We make an important distinction at this point. “Discounting under certainty" means there is no stochasticity in the economy at all. “Risk-free discounting" means that we are discounting a public project whose benefits are non-stochastic but where there may still be macroeconomic uncertainty. In Chapter 7, we concentrate on the former, while returning to the risk-free rate in the presence of macroeconomic uncertainty in Chapter 9.

5.3.7. As we will show below, all the credible social discount rates that HM Treasury may consider implementing — as well as the discount rates used in the private sector — have interpretation in this form. Essentially, every question in the Terms of Reference is asking us to better refine the choice of SDF, π_t, that HM Treasury will ultimately use.

5.3.8. This formulation also makes clear that the discount rate is concerned with time (the risk-free rate) and project uncertainty (the risk premium). As the Dutch government’s recent report notes:

> “The discount rate central to this report is the so-called risk-weighted discount rate (in real terms). This consists of two parts: the risk-free discount rate and the risk premium. The risk-free discount rate is the required real return on a risk-free project, for which costs and revenues are fully certain. The risk premium is the socially required compensation for project risk, the fact that the costs and benefits of the project are surrounded by uncertainty..." (Rijksoverheid 2025, p.15, Google Translate).

5.3.9. We recommend that HM Treasury follows this approach:

R.1. The discounting accounts for the maturity and systematic risk of projects

We recommend that HM Treasury uses the discount rate to account for project maturity and project systematic risk only in its economic appraisals.

5.4. When to avoid discounting

5.4.1. It is critical that HM Treasury recognises that, under some circumstances, either no discount rate exists at all, or that discounting is not the best economic technique available to appraise proposed public projects. In all cases, there are broader types of welfare economics that can allow for an appraisal of the public project to be undertaken, but discounting and present values are not the appropriate tools for doing so.

5.4.2. Consistent with this, Lucas (2023, Abstract) makes clear her view that discount rates cannot be applied for very long-term environmental problems: “For policies with long-term impacts, intergenerational concerns become paramount, projections of cash-flows and discount rates become highly uncertain, and present value calculations are an intrinsically unreliable measure of value. No approach to discount rate selection can overcome those problems; alternative decision criteria need to be established." A similar comment was made by one of the respondents to our survey “But for very long-term costs and benefits, a discounted present value approach is too noisy to be a reliable and decision-relevant tool" (paragraph A.2.2). We return to this point in Chapter 17.

5.4.a. Knightinan uncertainty

5.4.3. Key to discounting is Assumption A.3.; that there is a well defined probability space around which we can calculate expected net benefits, variances and covariances.

5.4.4. However, as famously noted by then US Defense Secretary Donald Rumsfeld in 2002:

> “... as we know, there are known knowns; there are things we know we know. We also know there are known unknowns; that is to say we know there are some things we do not know. But there are also unknown unknowns—the ones we don’t know we don’t know. And if one looks throughout the history of our country and other free countries, it is the latter category that tends to be the difficult ones"

Knightian Uncertainty economically formalises the concept of 'unknown-unknowns' by saying that the probability space is unknown, ill-defined, or unknowable. There are either states of the world we are blind to, or to which we cannot meaningfully assign probabilities. Examples might include extreme climate change outcomes, future pandemics with very high morbidity rates, catastrophic AI or nanotechnology disruption, nuclear war, asteroid strikes or major gamma-ray bursts, or a mega-volcano eruption.

5.4.5. If HM Treasury wishes to incorporate the consideration of such Knightian uncertainty into its cost-benefit appraisals, discounting is toothless. Instead, a different type of welfare analysis should be applied, such as "minimax regret discounting" (Iverson 2013). While theoretically robust, such techniques are highly complex and largely untested in practice.

5.4.6. However, we also believe that Knightian uncertainty has frequently been mis-used in practice to avoid situations where probability estimation is hard but not impossible. We would advise HM Treasury to only allow for Knightian uncertainty arguments to be used in very limited cases.

5.4.7. In terms of these Technical Annexes, we use 'risk' and 'uncertainty' interchangeably to give some linguistic variation. In both cases, we are in an economy described by a well-defined probability space because this is the environment in which the Green Book is currently set. We use 'Knightian uncertainty' explicitly when this is not the case.

5.4.b. (Dis-)economies of scale and additivity

5.4.8. Assumption A.1. requires that, if we increase the scale of a project, then the PV increases proportionately. While this assumption is innocuous for marginal projects, it does not necessarily hold once projects get sufficiently large.

5.4.9. Specifically, transformational projects may offer economies, or dis-economies, of scale. If these are significantly large to be of practical significance, then Assumption A.1. is violated and there is a requirement to go back to non-marginal welfare economics when undertaking cost-benefit analysis. Again, no discount rate, in the normal meaning of the term, exists. In Chapter17 we examine this issue in more detail.

5.4.10. Another issue to consider in terms of transformational projects is additivity, as given by Assumption A.2.: $P_{i+j} = P_i + P_j$.

5.4.11. There is a danger that if the economic appraisals of very large projects are silo-ed into individual government departments, then the synergies that arise from transformational effects get lost in the gaps between departmental analyses. The whole will then be greater than the sum of the estimated departmental parts.

5.4.12. This clearly violates Assumption A.2. Because of this, no adjustment to the discount rate can properly correct for this weakness. As before, we are in a world where the most basic assumptions about discounting do not hold.

5.4.13. Instead, this is a question about ensuring that all benefits are correctly identified and included in the analysis of one, and only one, department. We urge HM Treasury to ensure this is the case because, without it, any economic appraisal will be fundamentally flawed. We again return to this issue in Chapter 17.

5.4.c. Real Options

5.4.14. We have assumed that an investment today in a public project will bring in a future stream of net benefits, $b_{i1}, \ldots, b_{iT}$. While these net benefits are, in general, uncertain at the time that the investment decision is taken, we have not introduced the possibility that a future management decision could further influence them.

5.4.15. For example, it may be that if a project performs better than expected, then the government can either scale it up or use it to develop other similar projects. By contrast, if it performs poorly, then the government may be able to either scale it back or stop it altogether. Allowing for such future management flexibility in economic appraisals is known as *Real Options Analysis*.

5.4.16. While it is technically possible to frame Real Options within a stochastic discount factor framework (Edge 2011), this is not the most natural framework for doing so. Instead, it would be much easier for HM Treasury to apply standard techniques in Real Options analysis. There are a wide range of books including Guthrie (2009) and Trigeorgis (1996) that provide an introduction to such approaches.

5.4.17. In his 2022 Presidential Address to the American Finance Association, John Graham evidenced the use of investment appraisal techniques amongst for-profit firms in the US (Graham 2022). This showed a significant increase in Real Options usage by both small and large corporations since 2001, with well over a third now "always or almost always" using this technique.

5.4.18. The potential use of Real Options is mentioned in the 2026 Green Book (HM Treasury 2026*c*, A.6.76), with a somewhat more detailed discussion in the 2018 version (HM Treasury 2018, A5.17–A5.20 and Box 23). However, we note that the current version of the Green Book says: "... real options analysis typically requires estimating probabilities of different scenarios. This can sometimes introduce spurious accuracy into the analysis". We stress that the estimation of probabilities is, under Assumption A.3., also a requirement for accurate present value models. Without a well-defined probability space, we are not in a position to calculate expected net future benefits, which prohibits the use of the PV formula in a world with uncertainty.

5.4.19. The extent to which the application of Real Options methods in the UK public sector is practical is a question that lies outside the Terms of Reference for this review. However, the theoretical case, particularly for large public projects, is strong. We understand that such methods have been successfully applied in CBA by the government in the Netherlands; for example around the recent development of the Rotterdam harbour.

5.4.20. We should stress that the value of Real Options increases with future uncertainty. Unlike standard present values, where the discount rate increases with the systematic risk of the project, the more uncertain we are about future states of the world, the more value there is in having the flexibility to adjust management decisions as things evolve. Greater systematic risk is no longer necessarily associated with lower economic values.

5.5. Dual discounting

5.5.1. We now turn to the question of whether adjustments to allow for specific elements of a project — be they environmental, health-related, transformational, or place-based — are better handled in the numerator or denominator of the PV equation. This section excludes consideration of project risk, which we return to in Chapter 10.

5.5.2. A number of points in the Terms of Reference question whether using "usual" discount rates lead to a mis-estimate of the "true" social present value, p^*_{it}, of certain types of public projects. This includes transformational projects, place-based projects, and environmental projects. The implication is that, by choosing a "better" discount rate, r^*_{it}, then government departments would make more accurate economic evaluations. This process of changing the discount rate for certain types of projects is commonly referred to as "dual discounting":

$$E[b_{it}] \exp(-r^*_{it}t) = p^*_{it} \neq p_{it} = E[b_{it}] \exp(-r_{it}t)$$

5.5.3. Rather than changing the discount rate, the alternative is to change the estimates of the net benefits to b^*_{it}:

$$E[b^*_{it}] \exp(-r_{it}t) = p^*_{it}$$

These two approaches are entirely equivalent if:

$$r^*_{it} = r_{it} + \frac{1}{t} \ln \left[\frac{E[b_{it}]}{E[b^*_{it}]} \right]$$

and so any adjustment that increases the estimated expected net benefits, $E[b^*_{it}] > E[b_{it}]$, is equivalent to the choice of reducing the discount rate, $r^*_{it} < r_{it}$ for certain types of projects.

5.5.4. Given this equivalence, there is no objectively correct way of deciding whether to adjust the discount rate or to adjust the net benefits for issues such as transformational, place-based, or environmental projects. You can get to the same, "correct", value either way if you are careful. Generally, the preference in government departments has been to change the benefits rather than the discount rate, for example through welfare weights in the Green Book, and through changes to relative prices in respect to environmental matters.

5.5.5. In relation to the risk-free component of the discount rate, where the covariance between b_{it} and π_t is zero, the theory above says adjusting the benefits should be the preferred approach. This risk-free rate is a function of π_t only. As explained above, π_t can vary by maturity but not by asset, and therefore the risk-free rate should be the same for all assets.

5.5.6. We see this in practice in the private sector, where the use of the market-based Capital Asset Pricing Model (CAPM) prevails for setting the discount rate: $r_i = r_f + \beta_i E[r_m - r_f]$. Here, r_f is the risk-free rate, β_i is the 'systematic risk' of the project, and $E[r_m - r_f]$ is the 'equity premium' — the difference between the expected return to a broadly diversified portfolio of risky assets and the risk-free rate. While there are many different types of projects across the economy, all companies will use a Treasury security for the r_f term; the risk-free rate does not depend on the asset being appraised.

5.5.7. Intuitively this is what we might expect. The risk-free discount rate exists to compare net benefits expressed in the same monetary units that arise at different points in time, not to compare different types of net benefit at any given time.

5.5.8. Given that the Ramsey Rule, excluding the L term, in the Green Book is a risk-free rate that does not allow for the systematic risk of any given public project, the Fundamental Theorem of Asset Pricing tells us that the correct approach, theoretically, is to never adjust this rate for different types of project. As one of the respondents to our survey requested "Please embed as little as possible in the discount rate." (paragraph A.2.13).

R.2. Do not overwork the discount rate

We recommend that HM Treasury does not *overwork* the discount rate by putting project-specific issues into this variable that, more rigorously, should be handled elsewhere in the analysis. The Green Book should, instead, focus on providing a schedule of consumption discount rates that can be applied consistently across all different types of public project. This schedule may allow for variations in project maturity and risk.

5.5.9. In 2024, a large number of international experts in environmental cost-benefit analysis combined to write a paper in *Science* to consider this issue (Drupp et al. 2024). This also concluded that adjusting the benefits should be preferred to dual discounting, consistent with HM Treasury's findings in HM Treasury (2021).

R.3. The findings of the Environmental Discount Rate Review remain robust

We recommend that HM Treasury retains the conclusions drawn from its Environmental Discount Rate Review (HM Treasury 2021).

5.5.10. In Table 5.1 we report the responses from our survey on whether the discount rate should vary for environmental projects; see Appendix A for a description of the different responder groups represented. Apart from among the respondents who came to our survey from social media, there was a clear view that the discount rate should not be adjusted for environmental issues. Not a single one of the academic experts who responded to this survey advocated for this approach. For example, one respondent to our survey said "Environmental scarcity and limited substitutability are better reflected through improved valuation methods, such as relative price adjustments and updated shadow prices, rather than through changes to the social discount rate. Adjusting the discount rate risks conflating time preference with valuation issues and reduces transparency and comparability across projects", with many similar comments (paragraph A.2.11).

5.5.11. Next turn to the 1.5% health discount rate applied by the Green Book. The reader may feel that this is an example of dual discounting already being successfully used by HM Treasury. But this is to misunderstand current discounting guidance. We have stated that the net benefits, b_{it}, are expressed in the consumption numeraire. This, though, is not current practice with health discounting. Instead, the recommendation is to discount Quality Adjusted Life Years (QALYs) and then convert these discounted QALYs into a monetary value using the current cost/QALY figure. QALYs are not, themselves, a monetised value but instead reflect utility.

> "The recommended discount rate for risk to health and life values is 1.5%. This is because the 'wealth effect', or real per capita consumption growth element of the discount rate, is excluded. As set out in Annex 2, health and life effects are expressed using welfare or utility values, such as Quality Adjusted Life Years (QALYs), as opposed to monetary values. The diminishing marginal utility associated with higher incomes does not apply as the welfare or utility associated with additional years of life will not decline as real incomes rise." (HM Treasury 2018, A6.16).

Therefore, HM Treasury currently uses a utility discount rate, rather than a consumption discount rate, for QALYs. This is not dual discounting in the normal sense of that term.

5.5.12. An issue that was raised about health discounting in more than one of our Expert Panels was that participants were not clear exactly how this lower rate should be applied given the existing wording of the Green Book. Was it only for QALYs or also for monetary values associated with health? They recommended that, should HM Treasury remain with utility discounting for health, then the application should be clarified. They also wondered about the extent to which health effects delivered outside the Department for Health & Social Care, including environmental impacts for example, were covered by this rate. On this latter point, the Observer Groups were more satisfied with the current consistent rate of lower health discount rates across different departments.

5.5.13. An alternative for HM Treasury to consider is to monetise the QALYs in each future year and then discount these at a consumption discount rate. In this case, the arguments above would state that the appropriate risk-free rate to apply would be the same as for every other type of net benefit. There would be nothing special about health from a discounting

	Count	Yes	No	Unsure
UK Academic	20	30%	65%	5%
Social Media	19	58%	32%	11%
SBCA	11	27%	64%	9%
Experts	11	0%	73%	27%
Total (non-government)	61	33%	56%	11%
Government	12	33%	50%	17%
Total (whole sample)	73	33%	55%	12%

Table 5.1.: Should the discount rate be adjusted for environmental projects?

perspective. The Green Book would then set the consumption discount rate for use across all departments, and those departments would then have a responsibility for expressing the future stream of net benefits in terms of the consumption numeraire in a way that would be overseen by HM Treasury.

5.5.14. To see how this might work, imagine the present value of one QALY, Q_t, that a public project manages to deliver at time t, with a current cost-per QALY of C_{Q0}. Its present value under existing Green Book guidance is then given by: $Q_t C_{Q0}/(1+\delta+L)^t$. Notice that this continues to involve the risk adjustment term, L, which we return to in Chapter 10. The alternative would be to increase the price of the QALY at time t to C_{Qt} and discount this at the standard discount rate $\delta+\mu g+L$. Now the present value is $Q_t C_{Qt}/(1+\delta+\mu g+L)^t$

5.5.15. Let r_Q represent the average rate of growth in the price of the QALY over time. Then $C_{Qt} = C_{0t}(1+r_Q)^t$. Comparing these equations, the present value is the same under both the existing Green Book approach and this new approach if:

$$\begin{aligned} \frac{Q_t C_{0t}}{(1+\delta+L)^t} &= \frac{Q_t C_{0t}(1+r_Q)^t}{(1+\delta+\mu g+L)^t} \\ \Rightarrow \quad 1+\delta+\mu g+L &= (1+\delta+L)(1+r_Q) \end{aligned}$$

5.5.16. Since δ, L, r_Q are small, the right band side can be well approximated by (this is an exact equality under exponential discounting, with no approximation error):

$$\begin{aligned} 1+\delta+\mu g+L &\approx 1+\delta+L+r_Q \\ \Rightarrow \quad r_Q &\approx \mu g \end{aligned}$$

5.5.17. Therefore the proposed approach would give an identical outcome to current practice if the net benefit ascribed to one QALY increased at $r_Q \approx \mu g = 2\%$ per year (under the current Green Book values). If QALYs are expected to grow in monetary terms at a faster or slower rate than that, then the different approaches will alter the estimate of the present value of health outcomes.

5.5.18. There was general consensus among the Expert Panels that consistent consumption discounting, without dual discount rates, should be HM Treasury's preferred approach. One of the respondents to our survey also made this point: "If we have a time based preference parameter and an income/consumption elasticity parameter, there is no reason for the discount rate to vary by context. The fact that health ignores the income effect is already a significant error in current practice; we should avoid reconsidering the question which was posed to the environmental discount rate a few years ago" (paragraph A.2.10).

5.5.19. This continues to reflect our own opinion which we expressed in relation to health and the environment, but also specifically to the the Value of Transport Time Saved, in a published report commissioned by the Department for Transport (Freeman & Groom 2021*a*). Relative price adjustments rather than alterations to the discount rate is, in our opinion, the better alternative.

R.4. Remove the health discounting schedule

We recommend that health costs and benefits are discounted at the standard rate, with relative price adjustments used to convert health utility units into the consumption numeraire.

5.5.20. Yet the Expert Panels also stressed in the strongest possible terms the *requirement* to also have a relative price schedule for environment, health and other benefits from a public project where real values are likely to change over time. Considerable attention should be given to estimating the true value of r_Q and its equivalent for the environment and other areas of public policy where relative prices are changing. There was a material concern among the Expert Panels that such schedules do not exist and that, therefore, a single schedule of Green Book consumption discount rates would be applied to an inappropriate schedule of benefits. This view was also expressed in the survey responses; for example, "While we agree with the problem outlined in the paper [Drupp et al. (2024)] and the theoretical approach described to correct for the issue of environmental degradation in policy appraisal, often the data required for such adjustments is unavailable to policy makers and appraisal practitioners." and "Even if the review concludes against the adjusted discount rate for environmental impacts specifically, we think explicit discussion around this issue is required, with further practical guidance to relative price adjustments for environmental impacts included in the resulting Green Book guidance. This is particularly important now that the Green Book has been shortened and detailed discussions around complexities of environmental impacts have been removed from the main text." (paragraph A.2.11). Such practical guidance falls outside our Terms of Respondents, but we fully agree that this task is required.

5.5.21. We appreciate that this requirement may introduce transition issues while such price schedules are constructed. In the meantime, we recommend that HM Treasury takes a more heuristic approach in the numerator of the PV equation — for example, by increasing the value of QALYs by the revised estimate of $r_Q \approx \mu g$ (2% in the 2026 Green Book and, as we will see below, 1.875% under the parameters that we are recommending to HM Treasury through this review, although 2% is likely to remain a good heuristic). Since this is a pricing issue, any heuristic adjustment should not be made in the discount rate. Because of the mathematical equivalence of numerator (b^*_{it}) and denominator (r^*_{it}) adjustments, accurately adjusting the net benefit stream is no more complex in practice than accurately adjusting the discount rate.

R.5. Relative price adjustments

We strongly recommend that HM Treasury undertakes a body of work to evaluate future relative price changes for a range of social, health, and environmental benefits. This is a necessary condition for the successful application of a single schedule of consumption discount rates.

5.6. Marginal Cost of Public Funds

5.6.1. The Marginal Cost of Public Funds (MCPF) is an issue that has historically played a key role in the determination of social discount rates; see, for example, the discussion in Spackman (2024). This is an also a matter that has been raised with us during the Expert Panels and Observer Groups as part of this review process.

5.6.2. The role of the MCPF is explicitly discussed in the Green Book, extending the quotation that we introduced in paragraph 3.4.3:

> "There are costs associated with raising public funds: taxes create distortions in the economy, and the government typically pays interest on its borrowing. Practitioners should not generally include these costs in appraisal. This is because most government proposals are funded from pre-determined departmental budgets. Decisions about the overall level of public spending are made separately from, and in advance of, individual spending decisions. The costs of raising public funds are therefore not relevant for the appraisal of different options." (HM Treasury 2026*c*, #6.44)

5.6.3. For this reason, considerations of the MCPF are explicitly outside the Terms of Reference for this review.

5.6.4. We would note, though, that even if HM Treasury did want to incorporate MCPF into its economic appraisals, then adjustments to the discount rate are not the appropriate way to do this. The cost of raising funds is exactly that — a cost. As such, any adjustment is better handled in the estimate of net costs and benefits; the numerator of the present value equation. As Spackman notes: ".. the cost of public funding can be usefully framed as an absolute cost (≥ 1) relative to consumption, but not as a rate of return" (Spackman 2024, p.229).

5.7. Discount rate sensitivity analysis

5.7.1. We next consider the question of whether a range of discount rates should be recommended by the Green Book for discounting any given benefit, b_{it}, in a practice that is commonly referred to as "present value sensitivity analysis".

5.7.2. This is current Green Book practice for cases with maturity beyond 50 years, when sensitivity analysis should be undertaken with $\delta = 0\%$ and $\delta = 0.5\%$ (HM Treasury 2026*a*, Table 3.B). The choice of $\delta = 0\%$ is consistent with the value chosen by Lord Stern in his review of the economics of climate change (Stern 2007); see Lowe (2008).

5.7.3. It is also recommended in the 2003 version of Circular A-4, which is currently under use by the

Trump administration in the United States. This says "For regulatory analysis, you should provide estimates of net benefits using both 3 percent and 7 percent" (OMB 2003, p.34).

5.7.4. A further recent policy example of the application of multiple discount rates to discount the same stream of net benefits is given in the 2023 estimates of the Social Cost of Greenhouse Gas emissions by the US Environmental protection agency (USEPA 2023), following academic work in Newell et al. (2022) and Rennert et al. (2022).

5.7.5. More generally, it is just very tempting to use multiple discount rates in the appraisal of public projects because the task of identifying the single, 'correct', discount rate is just so challenging. In practice, present value sensitivity analysis around both project net benefits and the discount rate is extremely widely applied in both the public and private sectors.

5.7.6. In recent work, with co-authors Drupp and Nesje, we have argued strongly against this practice (Nesje et al. 2024). Our argument is straightforward. The present value of any future benefit is given by the Fundamental Theorem of Asset Pricing as explained in paragraph 5.3.3: $p_{it} = E[b_{it}\pi_t]$. The present value is then, in a mathematical sense, an *expectation*. But mathematical expectations are single numerical values (scalars). Say, '20'. And it is not possible to report a sensitivity analysis around the number 20.

5.7.7. Assumption A.3. lies behind this finding. We exist in a well-defined probability space where all potential future outcomes are known and to which a single probability can be assigned. Therefore, quantifiable uncertainty over the numerical values of π_t and b_{it} can be integrated out under the expectations operator in a way that should be incorporated in a single Green Book risk-free discount rate and project-specific risk premium for any given maturity.

5.7.8. A more complex situation arises when there is uncertainty over π_t because people cannot agree on the intellectual basis on which this should be determined. For example, some members of HM Treasury may advocate for an STPR approach and others for an SOC approach. To resolve this dilemma, one possibility would be for HM Treasury to give two different discount rates — one for each approach — and require analysts to undertake sensitivity analysis with both. Such type of 'model uncertainty' cannot be integrated out under the expectations operator. Alternatively, people may not be able to agree on the parameters of the Ramsey Rule and therefore introduce sensitivity analysis over δ, μ, as is currently the case for very long-term projects in the Green Book. Again, such disagreement cannot be integrated out under the expectations operation.

5.7.9. However, because π_t should be the same for all public projects, it is unreasonable to expect individual project analysts to know what to do with such sensitivity analysis. Similarly, HM Treasury may struggle to ensure that decisions from such sensitivity analysis are consistently taken across different appraisals. For this reason, however difficult it may be to reach a final decision, we recommend that HM Treasury chooses the best single discount rate it can for any given project at any given maturity because it is better placed to do so than anyone else.

5.7.10. Additional to these arguments are the fact that present value sensitivity analysis is not 'well behaved' in the way that the central present value is. In particular:

- Present value sensitivities are not additive. You can sum the means of two stochastic distributions but this property does not hold for most other features of the probability density function. Present value sensitivity values therefore violate Assumption A.2. Because of this, how parts of any project are aggregated within the economic analysis affects the results.
- There is a danger that project risk is double-counted; first, through the central social discount rate, and then through the sensitivity analysis. This potentially makes decision-makers over-cautious. If the correct discount rate is applied, then risk is already fully accounted for.
- There is no clear decision-making rule on exactly how an analyst should use the present value sensitivity analysis results.
- The correct discount rate is specifically derived to calculate the central present value. There is no basis to believe that the same discount rate can be applied to non-central values, and neither is it clear how the discount rate should be adjusted for sensitivity analysis.

5.7.11. The reader of this section may believe that this argument is insufficiently pragmatic. The counter-argument may be that, while these intellectual arguments may be correct, in practice, present value sensitivity analysis is *useful*. We have searched carefully for evidence to support such claims without success.

5.7.12. At the heart of this issue is again the danger of overworking the discount rate. Users of present value calculations need to be clear about what it does, and does not, tell the decision-maker. It quantifies the *expected* change in welfare created by benefits (PVB) and costs (PVC). It does not reveal further information about what the ex-post distribution of welfare might be: that is not the purpose of the discount rate. Present value analysis therefore gives one, but only one, important input into the economic evaluation of public projects.

5.7.13. The Institute and Faculty of Actuaries kindly sent us a report to consider as part of this Review (Joslin et al. 2026). We will adopt a recommendation from this report in the next chapter. However, for reasons that we have explained here and in more detail in the associated academic paper, the IFoA makes two recommendations with which we do *not* agree: "For all projects consider applying a sensitivity removing the wealth effect from appraisal calculations. This would ensure that decision makers for all projects understand how BCRs would be impacted if future growth in wealth is much lower than the central assumption assumes" and "For transformational projects consider removing the social time preference rate from baseline calculations, or, as a less preferred option, introduce its removal as a sensitivity calculation" (p.3). The SDF is the same for all projects, including transformational ones, and discount rate sensitivity analysis lacks rigour.

5.7.14. More general quantitative sensitivity analysis may still add useful additional information. For example, the identification of states with very poor outcomes for the project may be politically unacceptable and therefore the project may be rejected on this basis. Sensitivity around the discount rate itself, though, is non-informative for such broader analysis.

5.7.15. Specifically, quantitative sensitivity analysis may be useful as part of the process when undertaking economic appraisals. By showing the variables which have greatest effect on the outcome, departments can both concentrate their resources on estimating these impacts more accurately, and on considering how the project design might be changed so society might best profit from these sensitivities. Therefore, quantitative sensitivity analysis can help in the policy and estimation processes. This is distinct from quoting sensitivity analysis around the final present value, which we believe there is little evidence to support.

5.7.16. Should HM Treasury disagree with this recommendation, at the very minimum we believe that it should be clear to departments about how precisely this information should be economically interpreted, and also how it should be consistently applied for decision making purposes across different departments and appraisals.

5.7.17. For example, in relation to the $\delta = 0$ rate of pure time preference sensitivity analysis for very long-term projects, the discounting supplementary guidance says "Practitioners should explain the different appraisal results clearly in the associated business case or impact assessment. The difference between these results provides an estimate of the wealth transfer that is attributable to pure social time preference" (HM Treasury 2026*a*, #3.5). However, it gives no explicit guidance on how this sensitivity analysis should be used for decision-making purposes.

5.7.18. A similar point was made by one of the governmental respondents to our survey: "While the current Green Book instructs practitioners to conduct sensitivity analysis with pure time preference element removed for these cases, we question to what extent this is applied consistently and effectively in practice across projects to sufficiently address the issue. We would therefore welcome this issue to be revisited and addressed in full, with clear and explicit guidance" (paragraph A.2.11).

5.7.19. Our recommendation can also be seen in light of longstanding cautionary remarks about the potential ills and abuse of sensitivity analysis. (Little & Mirrlees 1974, p. 135) warned that sensitivity analysis can be abused where it becomes "an excuse not to try to quantify things" or where it just presents a range of switching values without being clear about how these can be interpreted or decisions made. Our recommendations align with this view.

R.6. Do not undertake discount rate sensitivity analysis

We recommend that, for any specific net benefit, b_{it}, only one consumption discount rate should be given by the Green Book. This does not prohibit the risk-free discount rate varying by maturity, t, or the risk premium varying by project, i. HM Treasury (2026*a*, Table 3.B) should be reviewed in light of this recommendation.

6. STPR or SOC?

6.1. Introduction

6.1.1. In this chapter, we address the question in the Terms of Reference: "What evidence should the UK government consider for remaining with the current STPR approach or instead changing to an SOC approach?"

6.1.2. We begin by formalising what is meant by SOC ("Social Opportunity Cost") and STPR ("Social Time Preference Rate"). In the Terms of Reference, the SOC approach is "... based on the rate of return that would be expected on funds left in the private sector, rather than raised by government borrowing or taxation". By contrast, the STPR is "the rate at which society values the present compared to the future" (HM Treasury 2003, #5.49). We gave fuller definitions in our review of international social discounting:

> The SOC approach is primarily a production-side approach that measures the social cost of public investment against the social cost of public funding. The SDR stemming from this approach reflects the cost of capital derived from the different sources of government funding and reflects the return on the private or corporate investment that it displaces. The argument is that public in vestments should have rates of return that compete with returns available elsewhere . We define the SOC approach as an SDR that reflects this cost of capital, and for simplicity assume that it reflects the returns available in the private sector. It is appropriate for discounting costs and benefits when the numeraire is units of investment." (Groom et al. 2022, p.470).

and

> "The STPR approach is primarily a consumption-side approach in which the STPR reflects the rate at which society is willing to trade off consumption today for consumption tomorrow. The STPR reflects society's intertemporal preferences and the trade-offs that increase or decrease welfare. Public appraisal using the STPR as the discount rate requires that public interventions compensate deferred consumption at a rate that ensures welfare is increased. The STPR can be reflected by the market savings rates, reflecting the consumption rate of interest, as in the United States, or otherwise the STPR arising from a calibrated social welfare function, e.g., the Ramsey Rule as in the United Kingdom. It is appropriate for discounting costs and benefits when the numeraire is units of consumption" (Groom et al. 2022, pp.470–471).

6.1.3. Here, we will define them in slightly different ways that align with our Recommendation R.1:

- **STPR**. The discount rate is derived from a model that explicitly considers the maximisation of a social welfare function.
- **SOC**. The discount rate is taken from financial market returns, the rates of return observed in the private economy for real investment, or financial economic models of investor welfare, for projects of equivalent maturity and systematic risk.

6.1.4. Under these revised definitions, the US approach of using market savings rates is viewed as SOC rather than STPR because there is no explicit consideration of a social welfare function. For the SOC rate, because discount rates vary with systematic risk and maturity, it is essential that the private sector comparator is matched on this basis. It would be a simple error to compare a long-term, safe, public project against a short-term, high risk private project under the SOC approach.

6.1.5. The argument we will present throughout these Technical Annexes, influenced by considerations of the theoretical & empirical evidence, is that the question in the Terms of Reference is over-dichotemised by framing it as a binary choice. Instead, the argument that we will present, in this and future chapters, is:

- Neither the STPR nor SOC approaches can be 'perfected' by theoretical extensions or better calibration to give an error-free estimate of the social discount rate. There are strengths and weaknesses to both approaches and both, therefore give relevant, yet imperfect, information.

- Because of this, our recommendations are based on a **triangulation** approach. Rather than relying on a single source, we will look to extract the most relevant information from four main sources of evidence: (i) social welfare functions that are both theoretically extended beyond the Ramsey Rule used in the current Green Book and calibrated using updated UK data, (ii) financial market returns and evidence from financial economics, (iii) survey data, where experts have balanced the strengths and weaknesses of different arguments to deduce their own optimal values, and (iv) international governmental guidance, where experts in other countries have gone through similar processes to our own.

- An STPR model that is both theoretically and empirically robust should result in discount rates that are, at least broadly, consistent with SOC rates. While it is well-established that SOC and STP rates will always differ for a wide range of reasons — see, for example Groom et al. (2022) — the SOC/STPR divide is currently overstated.

6.1.6. In future chapters, we will present evidence for all four key information sources we have used in this review: STPR models with updated calibrations, SOC rates, survey data, and other international guidance. We will show that, while these lead to a range of different values, these ranges are narrower than some readers of these Technical Annexes might have anticipated. This gives us confidence that the discount rates that we recommend are robust.

6.1.7. That both STPR and SOC information is relevant for setting social discount practice is consistent with findings in Drupp et al. (2018). While that survey focused on very long-term (≥ 100 year) discounting, it asked economic experts: "What relative weight (summing up to 100 percent) should the governmental body place on the following rationales for determining the social discount rate: (a) Normative issues, involving justice towards future generations [*X* percent], and (b) Descriptive issues, involving forecasted average future returns to financial assets [*X* percent]?", which broadly correspond to STPR and SOC approaches respective in an intergenerational context. The responses had a mean / median / mode value of weight on STPR of 61.5% / 70% / 50% respectively, with the full range of responses from 0% to 100%. This suggests, at least in the context raised by Drupp et al. (2018), that economics experts see value in using a range of source to derive the optimal social discount rate. In a subsequent smaller survey of philosophers with expertise in social discounting (Nesje et al. 2023), there was a much stronger leaning towards an STPR-only approach: the mean and median responses for the weight on STPR in this case were 78.5% and 80% respectively. Nevertheless, even philosophers (as a group) were not entirely willing to fully discard SOC information.

6.1.8. In our survey, we set this as a binary question that did not allow respondents to choose a mix. In Table 6.1 we report the percentage of our respondents who recommended an SOC approach, an STPR approach, or were unsure about which approach to prefer. Via all routes, our respondents clearly signalled a preference for STPR over SOC. However, amongst the expert academic group, who are arguably the most qualified non-government group to comment, just under half of the respondents were unsure which approach to prefer.

6.1.9. The qualitative survey comments around this issue, as given in paragraph A.2.2, are particularly enlightening. Many respondents expressed a clear preference for one over the other, but some advocated for the balanced approach that we take here: e.g.:

- "A composite rate may be more appropriate (i.e., STPR & SOC)"
- "You frame the question as either/or. I would agree with the idea of making distinctly more use of market based measures of discount rates...",
- "Either approach has unique upsides but also flaws. I don't think that a pure form of one or the other approach is the optimal solution.".

6.2. STPR

6.2.1. The broader purpose of the Green Book, much wider than the discount rate itself, is captured in the following statement:

> "Appraisal involves estimating the social value of different options and selecting the option that represents best value for money. This concept of social value is based on the principles of welfare economics." (HM Treasury 2026*c*, #2.4)

6.2.2. From this statement of purpose (and similar, but separate, legislation such as the Public Services (Social Value) Act 2012), it appears completely uncontentious to state that, in a UK context, the government's preferred conceptual basis for Green Book economic appraisal is STPR and not SOC, and therefore any discount recommendations must fall within this wider framework. There is no triangulation of what it is we are trying to estimate — it is an STP rate.

SOC or STPR	Count	SOC	STPR	Unsure
UK Academic	32	9%	66%	25%
Social Media	26	15%	73%	12%
SBCA	17	6%	82%	12%
Experts	14	14%	43%	43%
Total (non-government)	89	11%	67%	21%
Government	20	5%	70%	25%
Total (whole sample)	109	10%	68%	22%

Table 6.1.: The percentage of our respondents who recommended the SOC or STPR approach to discounting for the Green Book.

6.2.3. Therefore, when, in later chapters, we introduce evidence on SOC rates, this is not because we believe that there should be a conceptual 'blend' of SOC and STPR. This, in our opinion, is unambiguously a purely STPR framework. Instead, we will be using SOC rates because we believe they give relevant information content even within a purely social welfare context.

6.2.4. In theory, this should not be necessary. If three tasks could all be completed perfectly, then an STPR approach alone would be sufficient to set the social discount rate without any further triangulation from other sources:

- The social welfare function that the government is aiming to maximise is explicitly specified, and
- The conceptual basis on which this welfare function should be calibrated is explicitly specified, and
- Highly detailed empirical evidence is available to calibrate the social welfare function on that conceptual basis.

6.2.5. In later chapters, we will argue in some detail that none of these individual conditions are met, let alone all of them. In Chapter 7, we derive the currently-used Ramsey Rule to show the very large number of assumptions that are needed to get to this functional form. It seems self-evident that governmental social welfare objectives are not fully captured by this model. In Chapters 9 and 10 we respectively extend this model to allow for macroeconomic and project risk respectively, with a model with 'rare disaster' risk given in Chapter 12. These theoretical discussions, and associated illustrative calibrations, show how sensitive the estimated STPR discount rate is to the precise modelling choice of both social welfare function and the underlying dynamics of the economy. In reality, we can identify neither of these things precisely enough to be confident that our model of social welfare is sufficiently well understood to give the true social discount rate.

6.2.6. In Chapter 8, we consider the difficulties associated with calibrating any social welfare model. First, there are different conceptual bases underlying social welfare. At its most reduced level, the question can be seen as asking whether social welfare should reflect the aggregated preferences of individual citizens themselves, or should it take a more "Government House" approach to what is best for society and the people within it. These do not necessarily lead to the same choice of parameter values.

6.2.7. Once the conceptual approach is decided upon, evidence is needed to accurately calibrate the parameter values. For certain components of the Ramsey Rule (particularly δ), the evidence base is surprisingly sparse. While, in principle, it *may* be possible to intergrate out some of this uncertainty in the social welfare function optimisation, this is unlikely to be realistic in practice and we know of no related policy context that takes such an approach.

6.2.8. Beyond this, the "true" STP rate is never observable, even with the benefit of hindsight. This means that, in contrast to models in financial economics, there is little chance to refine welfare models over times based on the evolution of direct evidence.

6.2.9. For these reasons, the STPR discount rate calculated from the Ramsey Rule and its extension cannot be viewed as the "true" STP rate. To paraphrase a member of one of the expert academic panels: "You can believe in the Ramsey Framework without believing in the Ramsey Rule". This view was re-enforced by some of the qualitative survey responses (paragraph A.2.3): "Where to begin? All three of the assumptions are contentious, and measured with significant error", and " ... none of these [Ramsey Rule] quantities reflect valuation in the real world. Rather, they are derived from estimates of parameters in highly stylzed models of the macroeconomy that are informative about basic mechanisms but not about how society evaluates the effects of time and risk on value." Drupp et al. (2018) demonstrated that experts do not appear to use the Ramsey Rule when setting their own recommendations for the social discount rate.

6.2.10. Similar opinions were made by a number of members of the academic expert panels. Academic models were seen as being extremely useful ways of conceptualising the problem of intertemporal welfare optimisation. One member also stressed how helpful such models are for explaining to practitioners why discounting is needed at all, and why the rate does not necessarily always equal zero. The intellectualisation and communication of the Ramsey Rule and its extensions were seen as being fundamental to the task of setting the social discount rate.

6.2.11. But that was seem as being entirely distinct from using the Ramsey Rule as if it were an accounting identity. The highly parsimonious nature of the model and the chasm between economic theory and the real world were emphasised by a range of academic experts.

6.2.12. But there is a danger of overstating this critique. Calibrating the Ramsey Rule and its extensions as accurately as possible gives a direct estimate of what the "true" STP rate might be. We see STPR models as being both conceptually correct and useful, but not perfect.

6.2.13. This means that there remains an estimation error of the "true" STP rate even if we carefully calibrate the best available STPR model. The relevant question for HM Treasury is then whether SOC-based discount rates can help reduce this error if used carefully and in partnership with STP rates. We will argue that they can.

6.3. SOC

6.3.1. In Chapter 7, we will make the case that the best estimate of the risk-free component of the social discount rate, based on an SOC approach, is a Treasury bond yield. A key advantage of this is that it is observable. There is no need for the development of any economic models or empirical methods to calibrate them. Almost instantaneously, HM Treasury can access the Bank of England's real yield curve and take a rate off the shelf. This gives it a major advantage over the Ramsey Rule. This, of course, would be irrelevant if the SOC rate is entirely uninformative for the STP rate. But there are several reasons to believe that this is not the case.

6.3.2. In standard asset pricing theory, the baseline model for the risk-free SOC rate shares an identical functional form with the Ramsey Rule; see, for example, Cochrane (2009, Equation 1.7). This should not be surprising. STPR concerns the maximisation of social welfare, financial economics concerns the maximisation of investor welfare. Investors are also citizens and, in the representative agent economy that underpins the Ramsey Rule (see Chapter 7), they are the same agent. As a consequence the SOC risk-free rate is exactly the same as the STP rate within this highly stylised theoretical framework. This equivalence forms the basis for more subtle models that incorporate market imperfections; see, for example, Sandmo & Drèze (1971). Given this, depending on the approach we take to calibrating the Ramsey Rule under an STPR approach, there is at least some degree of theoretical equivalence between the SOC and STPR risk-free rates.

6.3.3. In addition, if the STP rate is substantially below the SOC rate, then financial markets underprice Treasury bonds according to the government perspective. Repurchasing gilts is therefore welfare enhancing according to Green Book methods. For example, assume inflation is zero and that there exists a Treasury bond paying £2 in coupons semi-annually for 10 years before redemption. If this has a market semi-annual yield of 5%/2, then its market price is £92.21. But, if the government repurchases this bond, then it saves £2 every six months up to 9.5 years, and then a final benefit of £102. Discounting this at an annual Green Book discount rate gives a value of £104.45. After subtracting the purchase price of £92.21, this gives economic value of £12.24 per bond that is repurchased. Conversely, if STPR > SOC, then the government should — at least applying Green Book appraisal methods — be issuing more debt to fund its projects. While the Green Book makes clear that its purpose is not to set overall governmental spending levels (see paragraph 3.4.3), this simple arbitrage argument suggests there must be some barrier to the extent to which SOC and STPR can reasonably deviate.

6.3.4. In a general equilibrium framework Barrage (2018) makes some of these arguments concrete. If the social discount rate is calibrated to be below the rate of return on capital, the entire market equilibrium is distorted from the planners perspective. This introduces the need for further policy interventions to correct these distortions. Primarily these are capital subsidies to essentially align market returns with the planners perspective on intertemporal allocations. In the world of the second best, peculiar policy prescriptions can emerge too. In the case of Barrage (2018) carbon taxes would have to be lower in the absence of capital subsidies in order to achieve the welfare maximising outcome. The conclusion is not that the social discount rate has to be the same as the SOC, rather that there are further policy implications of such decisions.

6.3.5. This view was expressed by some of the respondents to our survey; e.g., “The taxpayer does not care about rates that come from macroeconomic theory models. HM Treasury must show that these project create value after accounting for the opportunity cost of capital, and the clearest opportunity cost is simply to cut public debt" (paragraph A.2.9).

6.3.6. In contrast to the risk-free rate, SOC systematic risk premia are not observable. We can calculate observed average returns to risky assets, but the discount rate is determined by future expected returns. As expectations are not directly observable, and because the future can differ from the past, discount rate risk premia must be estimated. There are a range of models in asset pricing theory that help investors make these estimates, and we will cover some of these briefly in Chapter 11.

6.3.7. Again, the Consumption Capital Asset Pricing Model (CCAPM) that underpins the risk premium in an STPR environment, and which we discuss in more detail in Chapter 10, features centrally in the asset pricing literature; see, for example, Cochrane (2009, Equation 1.16). The reason is as before; both social discounting and asset pricing are centred around optimising welfare. The distinction between the theory of social discount rates and the theory of asset returns is much narrow than is commonly believed.

6.3.8. However, the CCAPM performs very poorly in explaining certain properties of observed financial market behaviour. Perhaps most famously, the observed average return to equity markets has historically been much higher than this model can explain. This is known as the “equity premium puzzle", and we will return to this topic in detail in Chapters 11 and 12.

6.3.9. The conclusion that we reach in these chapters is that this puzzle is less relevant for setting social discount rates than has previously been suggested. On the one hand, STPR refinements to the CCAPM increase the predicted equity premium. On the other hand, (i) the equity premium has declined over recent decades, (ii) social projects have lower risk than private projects and therefore the SOC rate — when matching maturity and systematic risk — is considerably lower than the equity premium itself, and (iii) the government does not pay taxes or suffer financial market frictions, meaning it can accept a lower gross return than private sector investors for the same risk and maturity. The gap between SOC and STPR estimates of the social discount rate, even in the presence of systematic risk, may strike readers as being surprisingly narrow.

6.3.10. This, though, is what we should expect. A well specified model of welfare should both form the basis for both an STPR social discount rate (social welfare) and have some ability to partially explain asset returns (investor welfare). We reiterate — investors are also citizens.

6.3.11. This does *not* imply that we think STPR and SOC rates should be identical. There are many factors that will drive a wedge between them which has been well documented in the literature; e.g. Groom et al. (2022). But it is our opinion that there must be some type of tie between the two, meaning that the SOC rate contains information content even when setting a purely STPR social discount rate.

6.3.12. If this were not the case, then again significant conceptual issues arise if the STP rate is too far below the SOC rate. If the government has a significantly lower cost of capital than the private sector then — taken to its logical and somewhat extreme conclusion — this, on its own, would mean that the government should nationalise all private industries. As with the gilt repurchase example earlier, any nationalisation would be likely, based on the discount rate alone, to pass an economic appraisal. Of course the Green Book does not exist for such policy contexts, but it shows the logical consequences of having a large STPR–SOC divide.

6.3.13. This returns us to the history of discounting in the Green Book. The “test discount rate", described in Section 4.1 was introduced precisely because nationalised industries were producing rates of return below their private sector counterparts and that was seen, by the government at the time, as being of material concern. The need, even in social discounting, to be somewhat constrained by the discipline of the markets was a point made to us in the expert panels.

6.3.14. This view is advocated by (Lucas 2023, p.118), which we cite here at some length:

> “A commonly held view among economists is that market prices are the best available aggregators of information about preferences, scarcity, and expectations about the future. In fact, most resource allocation, in both the private and public sectors, is mediated by the market price system. There is no need for analysts to agree about preferences, whether people are rational, or to what extent markets are complete. Estimates are disciplined

because market prices often are observable and auditable, and when they are not available, valuation experts can be called upon to verify the quality of analysts' estimates. By contrast, utility function–based inferences require assumptions about preference parameters, the weights to assign to different individuals or groups, the completeness of markets, and so forth. Those quantities are unobservable and must be assumed or inferred from models that are unfamiliar to most policy makers and the public and that are viewed by many economists as insufficiently robust for use as the basis for government valuations. Model-based discount rates effectively delegate an important part of the decision-making process to analysts that have limited accountability."

6.3.15. This is why we believe that the SOC rates of return prescribed by utility regulators, and which we return to in Chapter 11, are so informative. Scottish Water, being outside England, is nationalised and therefore falls under the Green Book discount rate. Thames Water, being in England, is privatised and therefore is constrained in what it can charge consumers based on costs of capital dictated by the Ofwat. If discount rates reflect maturity and risk only, then ownership structure (private or nationalised) should not in theory make any difference.

6.3.16. This point was also raised by comments received to our survey: "I would emphasise in particular the need for consistency with the way expected returns and hence discount rates feed into investment decisions in regulated utilities - all potentially part of the public sector. How would investment decisions change if they were renationalised? Does this make sense?" (paragraph A.2.9).

6.3.17. Foreshadowing results from Table 11.1, the regulated real vanilla cost of capital for water utilities is currently around 4%. While this is not identical to the current Green Book 3.5% rate, neither is the gap as large as might be inferred from a direct reading of the equity premium puzzle literature. Again, this is evidence of the narrowness of the STPR-SOC divide.

6.4. Best Unbiased Estimators

6.4.1. We now formalise our argument concerning triangulation. Let r_{STPR} be the "true" social discount rate that we would calculate if we could perfectly calibrate the precise welfare function that the government is aiming to optimise. Let $\hat{r}_{\text{STPR}}$ be the estimate of the "true" social discount rate, taken from what in practice is the best possible calibration of the Ramsey Rule or one of its extensions. Because neither the model nor the calibration is perfect, this has an estimation error, $\hat{r}_{\text{STPR}} = r_{\text{STPR}} + \epsilon_{\text{STPR}}$, where it is assumed that the error term is a normally distributed and unbiased: $E[\epsilon_{\text{STPR}}] \sim N(0, \sigma^2_{\text{STPR}})$.

6.4.2. Let $\hat{r}_{\text{SOC}}$ be an estimate of the social discount rate taken from consumer interest rates, financial markets, or private real investments. This differs from the "true" STPR social discount rate for two reasons. First, it will, like the STPR estimate, just be imperfect. There will be a simple forecasting error from it. In addition, it may be that such an approach gives a biased estimate, B_{SOC}, where it is assumed that $B_{\text{SOC}} \sim N(\mu_B, \sigma^2_B)$. A 'mean bias adjusted SOC' can be constructed by $\hat{r}^*_{\text{SOC}} = \hat{r}_{\text{SOC}} - \mu_B$. This unbiased estimator from an SOC approach has a forecasting error $\hat{r}^*_{\text{SOC}} = r_{\text{STPR}} + \epsilon_{\text{SOC}}$, where it is again assumed that the error term is normally distributed: $E[\epsilon_{\text{SOC}}] \sim N(0, \sigma^2_{\text{SOC}} + \sigma^2_B)$. There are now two terms in the variance which are assumed to be independent of each other: a simple forecasting error and an error from there potentially being an unknown bias from using an SOC basis to estimate what is conceptually an STPR social discount rate.

6.4.3. The errors from the STPR and SOC approaches may not be independent, not least because the same theoretical models lie behind both. Let $\sigma_{\text{STPR,SOC}}$ denote the covariance between ϵ_{STPR} and ϵ_{SOC}.

6.4.4. HM Treasury wishes to use all available information to construct the best estimate it can of r_{STPR}, the "true" social discount rate. It does this by combining $\hat{r}_{\text{STPR}}$ and $\hat{r}^*_{\text{SOC}}$ with weights w and $1 - w$ respectively to minimise the least squares forecasting error:

$$\min_w w^2\sigma^2_{\text{STPR}} + (1-w)^2(\sigma^2_{\text{SOC}} + \sigma^2_B) + 2w(1-w)\sigma_{\text{STPR,SOC}}$$

6.4.5. This results in the Best Linear Unbiased Estimator (BLUE) of the true STPR discount rate, $\hat{r}_{\text{BLUE}}$:

$$\hat{r}_{\text{BLUE}} = w\hat{r}_{\text{STPR}} + (1-w)(\hat{r}_{\text{SOC}} - \mu_B)$$
$$w = \frac{\sigma^2_{\text{SOC}} + \sigma^2_B - \sigma_{\text{STPR,SOC}}}{\sigma^2_{\text{STPR}} + \sigma^2_{\text{SOC}} + \sigma^2_B - 2\sigma_{\text{STPR,SOC}}}$$

6.4.6. The central message from this equation is that, under BLUE, the SOC estimate becomes irrelevant to HM Treasury ($w = 1$) only if one or both of two conditions holds:

- The STPR estimate is perfect: $\sigma^2_{\text{STPR}} = 0$, and/or
- The SOC estimate gives precisely no information content about an STPR "true" social discount rate, $\sigma^2_{\text{SOC}} + \sigma^2_B \to \infty$

Since there are strong reasons to believe that neither of these conditions hold, even if HM Treasury considers that an SOC approach gives biased and highly imperfect estimates of the "true" STPR discount rate, it should still give some weight to evidence from the private sector.

> R.7. The discount rate is STPR but SOC evidence is relevant.
>
> We recommend that HM Treasury views the social discount rate as a purely STPR concept. However, when establishing a numerical value for it, STPR and SOC evidence are both relevant.

6.4.7. To empirically calibrate the value of w in equation 6.4.5 in order to formally use a BLUE estimator, the standard approach in other fields would be to undertake an econometric evaluation on previous forecasting errors of $\hat{r}_{\text{STPR}}$ and $\hat{r}_{\text{SOC}}$ when compared to r_{STPR}. Yet the "true" social discount rate is never observed, even with the benefit of hindsight. Therefore it is not possible to empirically establish how different approaches have performed in practice.

6.4.8. Because of this, there is no objective scientific way of determining what w should be. Instead this will reflect how any given individual subjectively assesses the relative merits of each case. Ultimately this comes down to matters of professional judgement. This is one explanation for why it is just so challenging to reach broad consensus on matters relating to the social discount rate, particularly under an STPR framework.

6.5. Surveys and international guidance

6.5.1. Different experts will assign different values to w. This is due to the lack of empirical evidence on which to calibrate this parameter in order to reach consensus. This gives us two possible channels that we could follow throughout this review.

6.5.2. The first choice would be for us to rely entirely on our own professional judgement. This, broadly speaking, was the approach taken by Lord Stern in the Stern Review of the Economics of Climate Change (Stern 2007). He, as a world-leading expert in the field, assessed the competing cases of different arguments and reached conclusions based on this. As Groom & Hepburn (2017) emphasise, governmental social discounting decisions are, in practice, directly influenced by the specific individuals who undertake the writing.

6.5.3. The alternative would be for us to attempt to establish some average of what the community of experts feel is the right weight to place on w and then reflect this in the final recommendations irrespective of our own personal views. This is, at least in part, encouraged by the Terms of Reference, which require us to '... liaise with the wider academic community on behalf of HM Treasury and fairly reflect the breadth of views in the final report".

6.5.4. The approach we take is a blend of these. We recognise that it is impossible for any scientist to completely disassociate their personal views from such a process; see, for example, Van Eck et al. (2024). However, expert survey responses — both from the survey we have run as part of this exercise and previous published surveys — and the way that other international governments have grappled with issues closely related to those in the Terms of Reference, all add significant value to the process of setting an STPR social discount rate in the UK.

6.5.5. This explains the triangulation approach that we will apply throughout the rest of these Technical Annexes. Conceptually we are clear; it is an STP rate that we are endeavoring to estimate. But simplified theoretical STPR models are imperfect, and therefore we look for other information sources that we believe can be incorporated into the analysis in order to make the recommended values more accurate. The exact way in which this is done ultimately reflects our own professional judgements as the appointed lead authors for this review. But we stress that the range within which many of these estimates lie is relatively narrow, resulting in fairly low sensitivity of the recommended discount rate to the weights that we place on each source.

7. Discounting in a Deterministic World

In this chapter, we concentrate on the foundations of discounting under certainty, which we define to be discounting in a world where the values of all variables at all future points in time are known today — more formally, this is an economy with no stochasticity. This setting captures both the Ramsey Rule as currently recommended in the Green Book, and considerations of the appropriate choice of the risk-free rate under an SOC approach.

This chapter makes clear that there are conceptual and empirical choices to be made irrespective of whether the STPR or the SOC approach is taken to social discounting. In the case of the STPR, there are several theoretical framings from within financial and welfare economics that lead to the mathematical form that is generally called the Ramsey Rule. In fact Ramsey (1928) had something specifically in mind when thinking about optimal saving. Not all approaches share this view. Similarly, with regard to the SOC, choices need to be made about the conceptually appropriate empirical quantity (market data) and temporal reference to guide social decision-making.

7.1. The Ramsey Rule and asset pricing

7.1.1. We start from the stochastic model that was derived in Chapter 5 and move from here towards the Ramsey Rule. This requires the existence of a representative agent economy:

A.4. The economy can be modelled as if there is just one 'representative agent' who lives indefinitely and consumes aggregate real per capita consumption, c_t, at all times, t.

To justify this assumption in a world of uncertainty, there is a requirement that all personal shocks to consumption are insured against so that everyone in the economy experiences the same consumption growth rate as the aggregate (it does not require that everyone has the same consumption level). Put more formally: "Full consumption insurance implies that heterogeneous consumers are able to equalize their marginal rates of substitution state by state and, at least for consumers with von Neumann-Morgenstern preferences, that the equilibrium of a heterogeneous-consumer, full-information economy is isomorphic in its pricing implications to the equilibrium of a epresentative-consumer, full-information economy" (Constantinides & Duffie 1996, p.220).

7.1.2. Society is assumed to receive intertemporal welfare, W, by summing over time the 'utility', $U(c_t, t)$, that the representative agent receives each period, t, from consuming c_t. The objective of the government is to maximise expected intertemporal welfare:

A.5. The objective of the government is to maximise $E[W]$:

$$\text{Objective} \quad \max E[W] = \sum_{t=0}^{\infty} E[U(c_t, t)]$$

7.1.3. If the government enters into a new project that gives a per capita benefit of b_{it} at time t (only) at a per capita investment cost of p_{it} today, then this changes the total expected welfare, $E[\Delta W]$, over time as follows:

$$U(c_0 - p_{it}, 0) - U(c_0, 0) + E[U(c_t + b_{it}, t) - U(c_t, t)]$$

Utility today is reduced because we are investing for the future, but expected utility increases in the future when the project comes to fruition.

7.1.4. We now introduce a further assumption which emphasises that social discounting only applies in the case of marginal projects:

A.6. p_{it}, b_{it} are "small" relative to c_0, c_t.

7.1.5. This assumption allows us to apply a first order Taylor's series approximation to the change in total welfare over time, which becomes:

$$E[\Delta W] = -p_i U'(c_0, 0) + E[b_{it} U'(c_t, t)]$$

where $U'(\cdot)$ denotes the first derivative of the utility function with respect to consumption.

7.1.6. The highest price, p_{it}, that the government should be prepared to pay for the project is the one that leaves total expected welfare over time unchanged: $E[\Delta W] = 0$. Rearranging the previous offset equation then gives:

$$\begin{aligned} p_{it} &= E[b_{it}\pi_t] \\ \pi_t &= \frac{U'(c_t, t)}{U'(c_0, 0)} \end{aligned}$$

where, as above, π_t, is a stochastic discount factor. While not generally explained this way in the social discounting literature —although it is in the asset pricing literature, and one respondent to our survey said “Why not moving to a stochastic discount factor approach?" (paragraph A.2.13) — this firmly places the Ramsey Rule within the much broader framework of the Fundamental Theorem of Asset Pricing. It is clear that this expression for p_{it} is consistent with Assumptions A.1. and A.2. provided that scaling and combining projects does not violate the marginality Assumption A.6. In addition, because π_t depends on aggregate consumption only, the SDF is the same for all public projects, as the Fundamental Theorem tells us it must be.

7.1.7. The Ramsey Rule only considers the risk-free component of the discount rate, with no adjustment for project systematic risk; we refer to this as r_{ft}. Therefore, from above, we know that:

$$r_{ft} = -\frac{1}{t}\ln(E[\pi_t]) = -\frac{1}{t}\ln\left(E\left[\frac{U'(c_t,t)}{U'(c_0,0)}\right]\right)$$

7.1.8. It is next necessary to make an assumption about the form of the utility function.

A.7. The utility function takes the form:

$$U(c_t,t) = \begin{cases} e^{-\delta t}\ln(c_t) & \mu = 1 \\ e^{-\delta t}\dfrac{c^{1-\mu}-1}{1-\mu} & \mu \neq 1 \end{cases}$$

where δ is the rate of pure time preference (or the ‘utility discount rate’) which is used for the health discounting of QALYs under current Green Book guidance, and $\mu \geq 0$ is the elasticity of marginal utility of consumption. This form of utility function is known as ‘isoelastic utility’ and is commonly applied in many fields of economics. L’Hôpital’s rule can be used to prove that the bottom equation tends to the top equation as $\mu \to 1$.

7.1.9. Under this utility function, $U'(c_t,t) = e^{-\delta t}c_t^{-\mu}$ for all $\mu \geq 0$ and therefore:

$$r_{ft} = \delta - \frac{1}{t}\ln\left(E\left[\left(\frac{c_t}{c_0}\right)^{-\mu}\right]\right)$$

7.1.10. Let g_t be the annual logarithmic growth rate in real per capita consumption:

$$g_t = \frac{1}{t}\ln\left(\frac{c_t}{c_0}\right)$$

then:

$$r_{ft} = \delta - \frac{1}{t}\ln E\left[\exp(-\mu t g_t)\right]$$

7.1.11. To progress from here, we need a further assumption:

A.8. g_t is known with certainty at the current time (it is ‘non-stochastic’).

The implication of this is that the analysts appraising any public project know, with certainty, the level of future real per capita consumption at all points into the future.

7.1.12. This assumption makes the expectations operator, $E[\cdot]$ redundant, and the expression for the risk-free discount rate collapses to something close to its form in the Green Book:

$$r_{ft} = \delta + \mu g_t$$

7.1.13. With two more assumptions, we finally arrive at the standard ‘Ramsey Rule’ (Ramsey 1928), as applied in current Green Book guidance:

A.9. $g_t = g$ for all t. The annualised consumption growth rate is known, with certainty, by analysts at the current time to take the same value for all horizons into the future. This means that, with $\delta = 0.5\%$, $\mu = 1$, and $g = 2\%$, the Ramsey Rule derived risk-free discount rate of 2.5% should be applied at all horizons (before we get to the Green Book’s arguments for declining discount rates, which we will return to in the Chapter 9).

A.10. The discount rate derived this way can be used by government departments as a discrete rate, even though the formula has been derived as a continuous rate. This point is of very little practical significance as continuous and discrete rate are so similar. The risk-free Ramsey Rule rate currently in the Green Book of 2.5% corresponds to a discrete rate of 2.532%. Nevertheless, for formality, we include it here to make clear the set of implicit assumptions that underlie current Green Book discount rate guidance.

7.2. The SOC risk-free rate

7.2.1. Having laid out one single derivation of the Ramsey Rule for the risk-free component of the current Green Book social discount rate (excluding declining discount rates and the L adjustment) under an STPR approach, we now turn to how the risk-free component of the social discount rate might be estimated under an SOC environment.

7.2.2. We had a number of qualitative survey responses on this matter; see paragraph A.2.4 in the Appendix. Consistent with a wide range of international guidance (Groom et al. 2022), the very strong preference is that the real yield on some UK Treasury security should be used. As one respondent noted: "At last a question with a clear-cut answer! Implied zero coupon yield to maturity on a 10 year indexed bond."

7.2.3. Such opinions have strong theoretical support as gilts are the lowest risk investment opportunity that is available in the UK economy, although, as one respondent noted "there are no risk-free assets, projects or returns to capital in the real world". Some respondents preferred a US Treasury Bond yield, presumably on the grounds that these are lower risk, but these returns are not denominated in sterling and therefore, from a UK perspective, risk is present because of changes in the USDGBP exchange rate.

7.2.4. In the private sector, which commonly uses the CAPM as the basis for its weighted-average costs of capital, it is again usual practice to use a Treasury security for the risk-free rate. However, as private sector cost-benefit analysis is generally undertaken in nominal terms, with the benefits inflated over time, a real Treasury rate is not used.

7.2.5. Despite this broad preference of respondents for UK gilt yields, which we agree with on theoretical grounds and noting the parallel with the private sector (while accepting that such rates are not perfect), a number of more detailed questions still need to be addressed:

7.2.a. Vanilla or index-linked gilts?

7.2.6. Most gilts in the UK offer fixed coupons and redemption payments in nominal terms, therefore not offering a risk-free real rate. If using a vanilla gilt yield, there is therefore a requirement to deduct both expected future inflation and the risk-premium that is implicit in gilt yields for future inflation uncertainty.

7.2.7. Given this, the case seems stronger to use index-linked gilts. While the inflation hedge is not perfect, there is much greater protection against future changes in price levels. However, the market for index linked gilts has greater distortions caused by, for example, illiquidity, demand spikes caused by policy decisions including pension reforms under the Brown administration, and fixed-income supply choices by the Debt Management Office.

7.2.8. A further complication is that inflation is handled in the Green Book through the GDP deflator. By contrast index-linked gilts are protected, with an indexation lag, to the Retail Price Index. On average over the period 1949–2025, the GDP deflator minus RPI inflation averaged -0.2% with a range of -5.9% to +3.5%. This difference in the treatment of inflation is therefore not trivial.

7.2.9. The ex-post differences in realised inflation between these two measures is likely to be larger at certain times than the ex-ante forecasts on which public sector NPV analysis is undertaken. Inflation forecasts, which are more relevant for Green Book purposes, are given by the Office of Budget and Responsibility (OMB). In Table 7.1, we present the OMB's forecasts in March 2026 for the period up to 2030. It is clear from this that the distinction between these two inflation measures is likely to require some correction if HM Treasury decides to use an SOC approach.

Year	GDP Deflator	RPI
2026	2.2%	3.9%
2027	1.9%	2.9%
2028	1.9%	2.9%
2029	1.9%	2.9%
2030	1.9%	2.3%

Table 7.1.: Office for Management and Budget forecasts for inflation for 2026–2030. Forecast date: March 2026.

7.2.10. HM Treasury could decide to engage with the academic literature on fixed income securities to find a path through these issues. While this is not recommended by any of our survey respondents, we believe instead that the Green Book should call on existing available sources where this is possible.

7.2.11. The Bank of England publishes a real yield curve every day on its website. The rate that is of relevance to the Green Book is the spot rate (the rate between now and time t) and not the forward rate (the instantaneous rate that exists between time t and $t+\Delta t$ for very small Δt). We believe that this provides the best option for the Green Book, should HM Treasury decide to take an SOC approach.

7.2.12. In describing its estimation of real discount rates, the Bank of England notes:

> "A real debt market provides information on the ex ante real interest rates faced by borrowers and lenders who want to avoid the effects of inflation. In practice, there are factors that mean index-linked gilts do not offer complete inflation protection, and the UK index-linked gilt market is not as liquid as that for conventional UK gilts. Nevertheless, this market allows us to calculate

real spot and forward rates analogous to the nominal spot and forward rates described above".

7.2.b. Maturity of the Gilt

7.2.13. While there may be some difference between the yield inferred from vanilla gilts adjusted for inflation and the yield on index-linked gilts, this is likely to be small compared to the difference in yield on any given date between Gglts of different maturities.

7.2.14. As an example, in Figure 7.1, we present the real spot yield curve, as given by the Bank of England, for the 31st March 2026. On this date, yields varied between -0.48% at the short end of 2.5 years, and 2.17% at 25 years. Such a variation would make a significant difference to any cost-benefit analysis calculation.

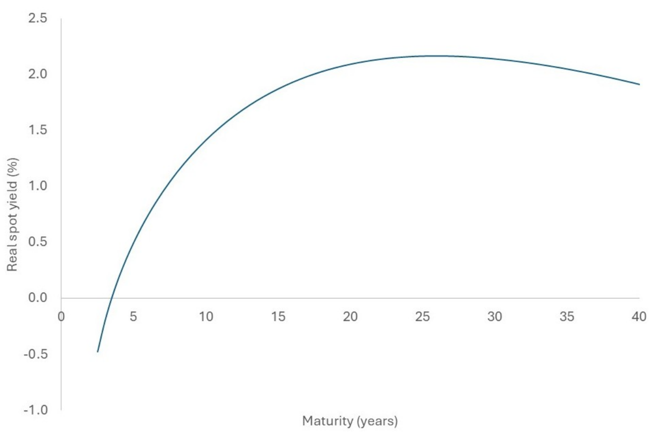

Figure 7.1.: This figure presents the Bank of England's real yield curve (spot rate) for the date of 31st March 2026 for maturities over 2.5 years.

7.2.15. Many of our survey respondents stated that a 10-year maturity should be used, but, simultaneously, in Table 11.3 below, we also see that our respondents felt that the discount rate should vary with maturity. In addition, the Green Book already varies the social discount rate with maturity.

7.2.16. Theoretically, the net project benefit each year should be discounted by the yield at that given maturity, but this is unlikely to be pragmatic. To overcome this issue, we adopt a recommendation made in the recent report by the Institute and Faculty of Actuaries that was referred to us as part of this review (Joslin et al. 2026). It recommends that the Green Book should "Enhance the evidence base for discounting impacts by requiring users of the Green Book to calculate the duration of annual costs and benefits respectively for their projects and use this to assess and monitor any bias toward short termism on an ongoing basis" [p.3]. This is consistent with guidance in HM Treasury (2026*c*, #A.17) for private finance models, where discount rate sensitivity analysis is required based on the gilt rate, where "The precise gilt rate to use should be the one that most closely matches the weighted average duration of the private finance model contract".

7.2.17. 'Duration' is a technical term in fixed income analysis that gives information about the sensitivity of a bond price to a small change in interest rates. It can be loosely interpreted as the weighted-average maturity of discounted net benefits and therefore gives a single date that might be considered relevant for determining the appropriate point on the term structure.

7.2.18. Formally, to estimate a Modified or Macaulay Duration, a discount rate is needed for the calculation. Here, though, the duration is being suggested to estimate the discount rate itself. For pragmatism, therefore, Joslin et al. (2026) recommends, effectively, that for the purposes of the duration calculation, the discount rate is set equal to zero. Under this assumption, the duration of the net benefits of any project, D_i is given by:

$$D_i = \frac{\sum_{t=1}^{T} t b_{it}}{\sum_{t=1}^{T} b_{it}}$$

7.2.19. For example, for a project that gives the same net benefits each year, $b_{it} = b_i$, then this simplified duration is just $(T+1)/2$.

7.2.20. Consider a 10-year project that will pay £1 each year starting in Year 3 and finishing in Year 12. The simplified duration of this project is 7.5 years, and using the Bank of England real spot interest rate for the 31st March 2026 of 1.04%, this would give a PV of £9.25. If, instead, we had discounted each net future benefit by the real discount rate of equal maturity (which is theoretically more accurate), the PV would be £9.20: a valuation difference of approximately 0.5%. For the 30-year project with fixed annual real benefits, the valuation error is under 2% (£21.43 against £21.82).

7.2.21. We stress that this approach is a heuristic without formal theoretical support. There are situations where it will work poorly, particularly when the Bank of England term structure is U-shaped. However, it has a further advantage. The Bank of England term structure only goes out to 40 years, which may be insufficient for a long-term project. However, 40 years is the simple duration of an 79 year project with fixed net benefits in each year. This approach, therefore, allows for the evaluation of many long-term projects that the 40-year term structure itself would not.

7.2.22. We note that the expert review for the PIDR advised against dual or multiple discount rates by term, arguing that the technical benefits of doing so are outweighed by the practical frictions of doing so.

7.2.c. Stability of the real yield curve

7.2.23. In the previous two subsections, we have taken a relatively pragmatic approach to the real yield curve. This is because the most substantive issue that HM Treasury must face does not concern the identity of the correct yield on any given date, but the fact that this rate moves rapidly within the time-frame of a standard policy cycle.

7.2.24. In Figure 7.2, we plot the evolution of real yields since 1985 at maturities of 5, 10, 20, 30 and 40 years. From this it is clear, that the maturity difference on any given date is much smaller than the change in yields at any given maturity over time.

Figure 7.2.: This figure presents the Bank of England's real yield (spot rate) estimates for maturities of 5, 10, 20, 30 and 40 years from January 1985 to March 2026. Uninterrupted longer maturity data starts at dates later than 1985.

7.2.25. Given the infrequency with which Green Book discount rates are reset, this presents a significant challenge. One approach to help solve this issue is given in the 2003 version of Circular A-4 which has been re-activated by the second Trump administration:

> "If we take the rate that the average saver uses to discount future consumption as our measure of the social rate of time preference, then the real rate of return on long-term government debt may provide a fair approximation. Over the last thirty years, this rate has averaged around 3 percent in real terms on a pre-tax basis. For example, the yield on 10-year Treasury notes has averaged 8.1 percent since 1973 while the average annual rate of change in the CPI over this period has been 5.0 percent, implying a real 10-year rate of 3.1 percent." (OMB 2003, pp.33–34).

7.2.26. While this is pragmatic, we find it difficult to support this approach. Taking a 10-year bond yield from 30 years ago, 1996, contains implicit economic forecasts that bond traders thought at the time were relevant for the 1996–2006 period. Why that should be informative for decisions made for the 2026–2036 horizon is unclear to us. Inherent in this approach appears to be an implicit view that real bond yields revert to a given mean over a long period of time, but Circular A-4 does not test that assumption. Certainly Figure 7.2 does not appear to be stationary, while some authors have argued authoritatively that real interest rates appear to have been on both a multi-decade (Obstfeld 2023) and multi-century (Schmelzing 2020) downward trend.

7.2.27. One of the issues that has been repeatedly raised with us by the Observer Groups and Expert Panels is that the growth rate used in the Ramsey Rule (2%) is not compatible with the growth rates that are currently used to inform the future flow of expected net benefits; see also very many of the qualitative survey responses reported in paragraph A.2.3, perhaps most pithily expressed as "Sadly no-one expects real per capita consumption growth of 2% p.a. anymore". There is a *temporal inconsistency* between assumptions in the numerator and denominator of the PV equation. A similar problem has the potential to arise in an SOC environment; information is taken from the bond yield curve that is formed under different market expectations than that used to forecast future project benefits.

7.2.28. In principle, the solution to this problem is straightforward. For any given project, an "evaluation date" should be set that is consistently applied to take interest data from the Bank of England real spot yield curve, real growth forecasts from the OBR/OECD (see Chapter 8) and to form the macroeconomic background for project forecasts. In principle, this date could vary for different projects, with the most up-to-date information used in each case. However, we recognise that this may open the opportunity for tactical decision-making by departments over when to pick the evaluation date, particularly if bond yields are distorted from true economic fundamentals by market frictions or policy interventions.

7.2.29. An alternative, under the SOC approach, would be for HM Treasury to periodically reissue the rates that should be used by Green Book users based on the bond yields that prevail at the time. For example, the Government Actuaries Department re-estimates the Personal Injury Discount Rate every five years under Part 2 of the Civil Liability Act 2018 while regulators generally re-review the Weighted Average Cost of Capital to be used by utility companies once a year during the price review process. The Dutch government states:

> "The working group advises establishing the above numerical values for a period of five years ... In addition, the working group recommends reviewing the established numerical values if 1) the risk-free real interest rate – measured over a continuous period of one year – changes by more than 1.5 percentage points within four years; or 2) if the adjustment of the required returns by the next Parameters Committee, scheduled for 2027, warrants this. The CPB has the responsibility to signal in a timely manner whether an interim revision is warranted" (Rijksoverheid 2025, Google Translate, p.9)

7.2.30. As one of the Governmental respondents to our survey noted (paragraph A.2.3), "A regular (say every 5 to 10 years), pre-announced review of the expected average annual real per capita consumption growth rate should be undertaken. It is currently hard to justify this 2% based on recent and forecasted trends", while another noted somewhat directly (paragraph A.2.13) "The discount rate changes over time. Just accept it".

> R.8. The social discount rate should be regularly reviewed.
>
> We recommend that HM Treasury reviews both the expected growth rate component of the Ramsey Rule, g, and the real rates of interest from the Bank of England yield curve every five years, or sooner if there is a significant change in real yields.

7.2.d. Negative real interest rates

7.2.31. As shown by Figure 7.2, there was an extended period in the 2010s and early 2020s when the Bank of England real yield curve presented negative real interest rates. A relevant question under an SOC framework is whether the real risk-free social discount rate should ever be lower than 0%, which makes net benefits in the future *more* valuable than the same real net benefits today.

7.2.32. To address this question, it is necessary to determine the extent to which these negative rates reflect true economic fundamentals against whether they instead are a consequence of distortions in financial markets.

7.2.33. There is no reason in principle while "true" real interest rates should always be positive. This is most easily seen through the Ramsey Rule with negative g, but similar arguments can also be made within an SOC framework. Theoretically, therefore, there is no reason to always fully preclude the potential use of negative real STP rates.

7.2.34. This is a question that was addressed by the Dutch government in its previous 2020 review of discount rates, which are based on an SOC approach. In this review, they did not shy away from negative real rates.

> "The recommended value of -1 percent for the risk-free portion is based on observed real risk free interest rates with maturities of 10 and 30 years. The working group sees no reason to deviate substantially from this. This entails a fairly substantial adjustment of the current value of approximately 0 percent. The 2015 working group set the risk-free component higher than market rates suggested at the time, in light of interest rate volatility." (Rijksoverheid 2020, Google Translate, p.19).

7.2.35. However, we argue that the negative real gilt yields of the late 2010s and early 2020s were heavily driven by structural factors operating within the gilt market, notably the Bank of England's large-scale asset purchases and regulatory-induced demand from defined-benefit pension schemes; e.g., Bailey et al. (2020), House of Lords (2021). For this reason, they may be unsuitable for use in a social discounting context.

7.2.36. But, whether driven by economic fundamentals or policy / central bank responses, low real yields may indicate to HM Treasury that the growth rate requires reviewing. Major changes in the risk-free SOC rate is, in our opinion, a strong basis to decide whether or not an early review of the STP rate is required.

7.2.37. Our recommendation is that, if real rates go negative for a sustained period of time, or if otherwise the STPR–SOC divide becomes too wide, then the Green Book discount rate should be reviewed.

7.3. The Ramsey Rule in welfare economics

7.3.1. The Ramsey Rule also emerges from the Discounted Utilitarian (DU) social welfare function (SWF) that is a cornerstone of welfare economics and intertemporal cost-benefit analysis. In this section we briefly summarise the welfare theoretical underpinnings of the DU SWF and the Ramsey Rule. In doing so it is important to draw a distinction between the form of the Ramsey Rule and the substance. There are different normative justifications for the form of the Ramsey Rule but not all of them stem from the impartial Utilitarian framework envisaged by Ramsey.

7.3.2. This section starts with a summary of Ramsey (1928) before moving on to discuss some welfare theoretical interpretations of the discounted utilitarian social welfare function and the mathematical form typically taken in applied welfare analysis and CBA.

The Ramsey (1928) framework

7.3.3. In the original paper by Frank Ramsey (Ramsey 1928) the equation $r = \delta + \mu g$ was stated as a normative goal for society when thinking through society's essential long-run saving-consumption decision. Dasgupta (2005), using a term coined by Sen & Williams (1982, p16), describes the Ramsey's approach as reflecting a useful objective for 'Government House' economics, meaning a suitable normative Utilitarian representation of intertemporal welfare for the government to make intertemporal decisions for the good of society. It is Utilitarian in motivation — that is, the outcome of utility is described as 'enjoyment' Dasgupta (2005, p151) — indicating individual experiences. Ramsey (1928) provided no representation theorem (axiomatic motivation) for this structure, and famously argued that the discount factor on the function $U(\cdot)$ ought to equal 1, or rather than $\delta = 0$ (more on this later).

7.3.4. This essential normative framing is sometimes mistaken for the kind of representative agent model discussed earlier in this chapter. This misconception probably stems from reference to Ramsey (1928) in the wider context of positive neoclassical growth theory, in which models are referred to as Ramsey-Solow-Koopmans-Cass models. Typically though, optimal growth models have a representative agent in mind with the intention of simultaneously understanding optimal allocations and decentralised behaviour in the economy. Ramsey (1928) made no mention of a representative agent.

7.3.5. Dasgupta (2005) explains that there are various different interpretations of the mathematical form. Typically the DU welfare function is written as follows,

$$W = \sum_{t=0}^{\infty} e^{-\delta t} U(c_t).$$

where the parameter δ is the utility discount rate. Yet the precise approach taken by Ramsey (1928) is:

$$W(\mathbf{c}) = \sum_{t=0}^{\infty} U(c_t),$$

which explicitly embeds Ramsey's view that $\delta = 0$. The approach implies a number of assumptions including time separability (utility in the future is not a function of utility in the past) and exponential utility discounting (constant δ). The summation is taken over an infinite horizon. Again, utility is generally assumed to be isoelastic, implying that marginal utility is given by $U'(c) = c^{-\mu}$. With δ in the utility discount rate and g the growth of per capita consumption, the Ramsey Rule is that the Social Discount Rate (SDR):

$$SDR = \delta + \mu g$$

7.3.6. While the expression for the STRP in the Green Book is frequently described as the Ramsey Rule, strictly speaking the Ramsey Rule is an optimal savings condition that accounts for the rate of return from capital, r:

$$r = \delta + \mu(c_t)g$$

where $\mu(c_t) = -c_t U''/U'$. This equation reflects the solution to the essential problem of interest to Ramsey: given an inherited capital stock with rate of return r, what is the optimal path of consumption and saving for society? Dasgupta (2005, p163) explains that this equation as being 'fundamental to intergenerational ethics' where (with emphasis in the original):

> "...the right hand side reflects the percentage rate at which, at the margin, it is ethically *permissible* to exchange consumption among the successive generations, t and $t + 1$".

whereas the left hand side is described as:

> "... the percentage rate at which consumption can *feasibly* be exchanged among successive generations..."

The right hand side is, therefore, the rate at which future consumption should be discounted.

7.3.7. The social welfare function approach assumes cardinal utility and interpersonal comparability in changes (albeit not levels) (Kelleher 2025*b*), allowing comparisons between different proposed consumption paths. The mathematical structure of the DU approach arises from several different conceptual underpinnings to Ramsey's original conception, not all of which assume this property.

7.3.8. Since Ramsey (1928) different justifications for the discounted utilitarian social welfare function and its mathematical form have been discussed. Dasgupta (2005) notes the basic structure embodies a Rawlsian social welfare function as $\mu \to \infty$ and can be thought of as capturing Koopmans conception of intertemporal preferences if $\delta > 0$. In the following paragraphs we briefly explain some of these works and their relevance to social discounting. However, it is our view that the social welfare framing of the Green Book and the associated Ramsey Rule social discount rate is best understood as a normative social-planner framework, rather than a representative agent framework. This framing allows us to think clearly about the eligible sources of information for estimating the parameters of the SWF and Ramsey Rule.

Samuelson's welfare-representative agent

7.3.9. Samuelson (1956) shows that the Ramsey framework and the DU SWF form can be understood as a normative anaolgy to the representative agent framework, a welfare-representative agent. This differs from the positive representative agent outlined above, the proof of which relied on frictionless, maximising behaviour, decentralised markets and perfect insurance. The analogous device is welfare optimisation. The same essential approach is taken by Heal & Millner (2014*a*, 2013), who investigate the existence of a representative agent in the intertemporal context.

7.3.10. Samuelson (1956) starts with a welfare function of the Bergson–Samuelson form,

$$W = \mathcal{W}(U_1, \ldots, U_n),$$

where each individual i's utility U_i at a point in time depends on that individual's allocation of goods. The existence of a representative agent in this context depends on how welfare is distributed in society and whether this distribution is optimal. The result is as follows. For any aggregate of, say, consumption c, define the associated indirect social welfare value as

$$U(c) = \max_{\{c_i\}_{i=1}^n} \mathcal{W}(U_1(c_1), \ldots, U_n(c_n)) \quad \text{s.t.} \quad \sum_i c_i = c.$$

7.3.11. In words, the single function $U(c)$ gives a social valuation of the aggregate bundle once that bundle has been distributed in the welfare-maximising way. Welfare maximisation (not perfect markets) means that the first-order condition is

$$\mathcal{W}(U_1, \ldots, U_n) U_i'(c_i) = \lambda \qquad \text{for all } i,$$

where λ is the marginal (shadow) value of aggregate consumption. By the envelope theorem it follows that,

$$U'(c) = \lambda.$$

Thus Samuelson provides a representative social utility function over aggregate consumption. Rather than relying on assumptions about perfect markets this *indirect* welfare function relies on welfare maximisation, which implies an optimal distribution of utility across households. This is not the utility function of a literal representative household. It is the value of aggregate consumption after the allocation of that consumption has been chosen according to the Bergson–Samuelson social welfare function.

7.3.12. Samuelson (1937) looks at welfare at a particular point in time, but this proof is relevant to the discounted utilitarian framework because it shows one way in which a period utility term $U(c_t)$ can be given a welfare-economic foundation. If, at each date, aggregate consumption is distributed optimally according to the social welfare function, then intertemporal welfare can be written as

$$W = \int_0^\infty e^{-\rho t} U(c_t)\, \mathrm{d}t.$$

We discuss the precise assumptions underpinning the representative agent in an intertemporal context, but note here that the assumptions are strong. In particular, the assumption that the economy has an optimal redistribution of consumption is unlikely to hold, in the same way that perfect markets are unlikely to exist. Nevertheless, with additional assumptions (time separability and exponential discounting) Samuelson (1956) provides the foundation for the interpretation of the DU SWF as a representative agent in which the distribution of consumption in society can be ignored.

7.3.13. It is helpful to relate this to the form of the Ramsey Rule Green Book. If aggregate consumption is $c_t = \sum_{i=1}^n c_{it}$ and the weighted utilitarian Bergson–Samuelson welfare function has common isoelastic individual utility:

$$\mathcal{W}(c_1, \ldots, c_n) = \sum_{i=1}^n a_i \frac{c_i^{1-\mu}}{1-\mu},$$

$$a_i > 0, \qquad \mu > 0, \qquad \mu \neq 1.$$

Samuelson's representative social utility over aggregate consumption is the indirect welfare function obtained by distributing c_t at time t optimally:

$$U(c_t) = \sum_{i=1}^n a_i \frac{\left(c_t a_i^{1/\mu}/A\right)^{1-\mu}}{1-\mu}.$$

This simplifies to

$$U(c_t) = A^\mu \frac{c_t^{1-\mu}}{1-\mu}.$$

where $A = \sum_{j=1}^{n} a_j^{1/\mu}$.

7.3.14. Up to an irrelevant affine transformation, the representative social utility has the same form as the underlying individual utility. If this SWF is then placed in an intertemporal discounted utilitarian form:

$$W = \int_0^\infty e^{-\rho t} U(c_t)\ \mathrm{d}t,$$

the Ramsey Rule logic follows directly by calculating the welfare preserving trade-off of consumption over time, as shown in the previous section.

7.3.15. The assumptions underpinning this proof are strong. We interrogate this further below, but to give a flavour of the strength of the assumptions required, note that in addition to optimal distribution, if individual utilities have different curvatures, the representative utility need not be isoelastic and may not be a stable function of aggregate consumption alone.

7.3.16. Thus Samuelson provides a coherent route from a Bergson–Samuelson welfare function to a representative period utility function. In general the resulting representative welfare function is a social-planner object and not strictly a representative agent in the *behavioural* sense. In the Utilitarian case with equal weights one can infer the structure of utility through revealed preference. In this case the indirect welfare function can essentially be thought of as a representative agent.

7.3.17. In any event, Samuelson (1956) provides a welfare theoretic way in which to interpret the DU SWF underpinning the Ramsey Rule in the Green Book. The Ramsey period utility $U(c_t)$ is an indirect welfare function of aggregate consumption if aggregate consumption is being evaluated after optimal distribution, or if distributional concerns can otherwise be summarised by aggregate or mean consumption.

7.3.18. As we now explain, this representative agent interpretation is not the interpretation that we think is most relevant to this Green Book.

Harsanyi's aggregation theorem

7.3.19. A useful first step in understanding the principles that underpin the DU SWF is to consider the aggregation result of Harsanyi (1955). Harsanyi showed that, under von Neumann–Morgenstern expected-utility assumptions for individuals and for the social evaluator, together with a Pareto-type condition linking social and individual preferences, social welfare can be represented as a weighted sum of individual vNM utilities (see also Boadway & Bruce 1984, Kelleher 2025*b*):

$$W = \sum_i a_i U_i.$$

Harsanyi (1955) is an explicit aggregation theorem, aggregating over individuals in society. This does not by itself deliver the discounted utilitarian Ramsey framework, which requires additional assumptions to determine the aggregation over time. Nevertheless, this theorem provides a basis for welfare to be measured by a sum of utilities.

7.3.20. This provides a basis for an additive structure. It is possible to interpret each utility as time dated (e.g. $U(c_t)$) and imagine that the weights are utility discount factors: $a_i = e^{-\delta t}$. However, to move from this intra-temporal interpretation to the inter-temporal one requires numerous further assumptions on the intertemporal preference ordering.

The Koopmans and Diamond approach

7.3.21. To get a flavour of the additional assumptions required to arrive at the DU SWF used by Ramsey, it is useful to consider the approach taken by Koopmans and Diamond in their seminal contributions in intertemporal decision-making (Koopmans 1960, 1972, Diamond 1965).

7.3.22. Koopmans (1960, 1972) and Diamond (1965) consider an ordering (sometimes referred to as a preference relation $\succeq$) over infinite future consumption streams, e.g. to dertemine which of $\mathbf{c} = (c_0, c_1, c_2, \ldots)$ versus $\mathbf{c}' = (c_0', c_1', c_2', \ldots)$ is preferred/ranked highest. They impose successive ethical properties or axioms on the ordering and these determine how the ordering can be represented by a function.

7.3.23. Regularity conditions of continuity, transitivity, and reflexivity allow the numerical representation of the preference relation via a general functional:

$$W(\mathbf{c}) = \mathcal{W}(c_0, c_1, c_2, \ldots),$$

Montonicity then establishes that more consumption is strictly better than less.

7.3.24. *Time separability* is then imposed, capturing that welfare depends on outcomes at each point in time only, through period-specific utilities rather than through cross-date interactions. This yields an additive representation of the form:

$$W(\mathbf{c}) = \sum_{t=0}^{\infty} U(c_t, t),$$

where $U(c_t, t)$ captures the dependence on time via the argument t.

7.3.25. Multiplicative time-separability then yields:

$$W(\mathbf{c}) = \sum_{t=0}^{\infty} D_t\, U(c_t),$$

where $U(\cdot)$ is period utility and D_t is a sequence of weights that depend on calendar time t.

7.3.26. A constant utility discount rate follows from imposing *stationarity* (or time invariance). This can be understood as follows:

$$(\alpha, c_1, c_2, \ldots) \succeq (\alpha, c_1', c_2', \ldots) \iff$$
$$(c_1, c_2, \ldots) \succeq (c_1', c_2', \ldots).$$

that is, the ordering of consumption streams is not altered by a like shift in timing. In the additive setting this implies $D_{t+s} = D_t D_s$, so that, up to normalization,

$$D_t = \beta^t = e^{-\delta t}, \qquad 0 < \beta \leq 1,\ \delta \geq 0.$$

where δ is the utility discount rate.

7.3.27. The welfare function therefore becomes

$$W(\mathbf{c}) = \sum_{t=0}^{\infty} \beta^t U(c_t),$$

This has the same mathematical form as the original Ramsey paper, only now it emerges from the application of axioms of intertemporal ethics. The fact that the approach starts with a preference relation $\succeq$ suggests that this is a preference based approach, but the outcome is often interpreted as a social welfare function for a unitary social decision maker.

7.3.28. A remaining question here is whether the utility discount rate, δ ought to be zero or greater than zero. A remarkable result in this literature coming from Koopmans (1960) and Diamond (1965) shows that the axioms discussed so far imply that for an ordering to exist in which one infinite stream of consumption can be ranked against another, it has to be the case that $\delta > 0$. We return to this in the discussion of parameterisation of the Ramsey Rule.

7.3.29. While Koopmans' dealt with ordinal not cardinal preferences, these results indicate possible axiomatic requirements of a preference ordering that lead to the time separable exponentially discounted intertemporal interpretation of the Harsanyi aggregation, and the mathematical form found in Ramsey (1928) and the Green Book. Koopman's (Koopmans 1960, 1972) approach was not specifically concerned with the aggregation of utilities across individuals *within* time periods. It does not provide a proof of the existence of a representative agent. It rather commences with a single agent, which is often interpreted as a social evaluator.

Harsanyi-Broome

7.3.30. (Kelleher 2025*b*, part II) explains that Broome (2004*a*) identified a potential shortcoming that arises from the infinite horizon setting of Koopmans (1960) and Diamond (1965). It turns out that continuity of the welfare function leads to a situation in which a consumption stream of approximately zero forever is preferred to a positive level of consumption today followed by precisely zero thereafter (see Kelleher 2025*b*, Ch3). This seems ethically dubious.

7.3.31. Kelleher (2025*b*) argues that the approaches taken in Harsanyi (1955) and Broome (2004*a*) offer a way out of this conundrum. The Harsanyi (1955) use of vNM utility functions but over life-time utilities as in Broome (2004*a*), coupled with the consideration of finite, rather than infinite, time horizons leads to a *Prioritarian* SWF of the following form:

$$W(\mathbf{c}) = \sum_{t=0}^{T} \beta^t G\left(U(c_t)\right),$$

where population is constant and aggregation is undertaken over identical individuals in society. $G(.)$ is a concave transformation of utility that characterises a Prioritarian social welfare function. This is different to the DU SWF because it prioritises additional utility to people, time periods or generations with lower utility (e.g. Adler 2012). Nevertheless, Broome (2004*b*) argues that the most appropriate form for this particular Social Welfare Function stems from a linear transformation, indicating risk neutrality over lifetime well-beings (Kelleher 2025*a*, p111-120). With additional assumptions on time separability and utility discounting this approach leads to a Discounted Utilitarian SWF of the following form:

$$W(\mathbf{c}) = \sum_{t=0}^{T} \beta^t U(c_t),$$

An intertemporal welfare-equivalent representative agent

7.3.32. The DU form is often interpreted as a *representative agent.* a single agent that represents the preferences or well-beings of the constituent individuals in society.

7.3.33. Ramsey (1928) posited the intertemporal social welfare function in 7.3.27 as a normative Utilitarian construct representing a intertemporal societal welfare goal (Dasgupta 2005). In general this is not seen as being a representative agent approach. Neither, in the context of CBA, is it seen as being a 'welfare representative agent' of the kind imagined by Samuelson (1956).

One issue here is that CBA does not typically assume that the distribution in society has been optimised.

7.3.34. Heal & Millner (2013) clarify what is involved in treating a heterogeneous society as if it had a representative intertemporal preference. Their starting point is the social-planner problem in which agents have heterogeneous pure rates of time preference. The planner allocates consumption efficiently across agents and over time, solving

$$\max_{\{c_{i,t}\}} \sum_i \alpha_i \int_0^\infty e^{-\rho_i t} U_i(c_{i,t})\ \mathrm{d}t$$

subject to an aggregate resource constraint. The first-order conditions equate weighted, discounted marginal utilities across individuals through the shadow value of aggregate resources.

7.3.35. In this framework there are two representative objects. A *welfare-equivalent* representative agent reproduces the group's welfare for a given aggregate consumption stream. This is the dynamic analogue of Samuelson (1956) inwhich for a given C_t, welfare in time t is obtained by optimally distributing C_t across heterogeneous individuals. Heal & Millner (2013, 2014*a*) define the welfare representative agent as U^W

$$U^W(C_t, t) = \max_{\{c_{i,t}\}} \sum_i \alpha_i e^{-\delta_i t} U_i(c_{i,t}) \quad \text{s.t.} \quad \sum_i c_{i,t} = C_t.$$

7.3.36. In addition define the *policy-equivalent* representative agent. This agent need not reproduce society's welfare level for every possible consumption stream. It is only required to generate the same optimal aggregate consumption path as the social planner. Importantly, the social discount rate in each of these cases is the same, so marginal projects evaluated around the chosen aggregate reference path are treated in the same way in this respect.

7.3.37. A key question in this context concerns the form of the intertemporal welfare function representative agent, and whether this in any circumstance coincides with the DU SWF used in the Green Book. It turns out this is true only under very restrictive, probably unrealistic circumstances. To see this, note that by definition of U^W, even if the utility functions are identical the welfare representative discount factor becomes:

$$D(t) = \sum_i \alpha_i e^{-\delta_i t}.$$

In general, this object is *not* exponential, i.e. there does not exist a constant $\bar{\delta}$ such that $D(t) = e^{-\bar{\delta} t}$ for all t, unless $\delta_i = \bar{\delta}$ for all i. Instead, the implied discount rate,

$$\delta^W(t) = -\frac{D'(t)}{D(t)} = \frac{\sum_i \alpha_i \delta_i e^{-\delta_i t}}{\sum_i \alpha_i e^{-\delta_i t}},$$

is time-varying and typically declining to the lowest value held in society. The Green Book social discount rate will not typically reflect that of a representative agent in a society with hetergeneous preferences.

Summary

7.3.38. In the forgoing paragraphs we have tried to place the DU SWF used in the Green Book within a broader welfare-economic tradition, and illustrate that the essential mathematical form has different possible interpretations. The interpretation is important because it will guide how we parameterise the social welfare function that HM Treasury uses for CBA.

7.3.39. Samuelson provides a welfare-theoretic analogue of the representative-agent construction at a single point in time, providing the beginnings of a welfare-representative agent interpretation for the Ramsey Rule. Yet, Ramsey (1928) is not explicitly a representative agent approach. The DU SWF is rather asserted as a normative societal objective. Harsanyi provides an axiomatic argument for additive aggregation of individual (vNM) utilities, while Koopmans provides axioms under which a single ordinal preference over consumption streams has a discounted utility representation. Broome (2004*a*) provides an alternative interpretation of DU SWF form by extending Harsanyi to addresses difficulties in Koopmans' approach. In short, several conceptually different approaches lead to the same mathematical form as the DU SWF.

7.3.40. Finally, Heal & Millner (2014*a*, 2013) show that when intertemporal preferences are heterogeneous among the agents in an economy, the welfare-equivalent representative agent will typically not have the simple constant-rate discounted-utilitarian form.

7.3.41. This coverage of social welfare functions that could underpin CBA is very far from exhaustive. Prioritarian (Adler 2012), Rank ordered Utilitarianism (Zuber & Asheim 2012, Asheim & Zuber 2014), and altruistic Utilitarianism (Asheim & Nesje 2016, Nesje 2024), to name but a few, are all eligible frameworks in principle.

7.4. Sustainability and the Ramsey Framework

7.4.1. The Ramsey rule is derived from a discounted-utilitarian social-welfare framework. Often discussions of the appropriate discount rate for public appraisal or welfare analysis revolve around the idea that higher discount rates bias against longer term returns and hence sustainability in the sense of non-declining well-being or non-declining ecological and environmental stocks.

7.4.2. Unfortunately, if sustainability is the objective then the Ramsey framework is the wrong framework in its simplest form. Undertaking CBA within the Ramsey Framework presented thus far, and using the discount rate arising from the Ramsey Rule places no such constraints on the evaluation of marginal projects or the optimal programme at the macro scale. Discussions about the parameters of the Ramsey Rule are not directly relevant in this sustainability sense.

7.4.3. Similarly, changing relative prices for flows of goods and services, including ecosystem services, that determine utility, which we ultimately recommend (rather than sectoral discount rates), also do not solve issues of sustainability in and of themselves. We now provide a brief and semi-formal exposition of these points. The point we want to illustrate is that the Ramsey Rule is not, by itself, a rule of sustainability, hence changing the parameters of the discount rate will not build in sustainability to CBA.

7.4.a. What the Ramsey Framework does and doesn't do

7.4.4. The Ramsey Rule is derived from a discounted-utilitarian social-welfare framework; it is not, by itself, a rule of sustainability. In this framework, utility depends on both consumption, C_t and the environment, E_t; $U(C_t, E_t)$. Let

$$V^R(\mathbf{K}_t, t) = \sup_{\mathcal{A}_t^R} \int_t^\infty U(C_s, E_s) e^{-\delta(s-t)}\, ds,$$

where $\mathbf{K}_t$ collects produced, natural, human, and other relevant capital stocks. The feasible set $\mathcal{A}_t^R$ contains the allocations permitted by the technologies, ecological laws of motion, resource constraints, and institutions represented in the model. The superscript R denotes the Ramsey welfare problem. It does not mean that the economy is free of physical or institutional constraints; it means that no additional sustainability criterion has been imposed unless that criterion is already part of $\mathcal{A}_t^R$.

7.4.5. In the familiar one-good, deterministic benchmark, the welfare objective implies the Ramsey consumption discount rate

$$SDR_{C,t} = \delta + \mu_{CC,t} g_{C,t}, \qquad \mu_{CC,t} = -\frac{U_{CC,t} C_t}{U_{C,t}}.$$

The pure-rate-of-time-preference term δ and the wealth-effect term $\mu_{CC,t} g_{C,t}$ determine how marginal consumption-equivalent benefits at different dates are compared (Ramsey 1928, Gollier 2013). The rule is therefore an intertemporal valuation condition, conditional on the social objective, the expected consumption path, and the risks represented in the model. It does not identify which environmental assets must be preserved, which ecological states are inadmissible, or which distribution of well-being across people and generations is fair.

7.4.6. In cost–benefit analysis, the Ramsey framework is normally used to express a project's heterogeneous impacts in consumption-equivalent units and to compare those units across time. For a marginal change in the relevant capital stocks, the Ramsey accounting price of stock $K_{i,t}$ and the associated local welfare change can be written as

$$q_{i,t}^R = \frac{1}{U_{C,t}} \frac{\partial V^R(\mathbf{K}_t, t)}{\partial K_{i,t}}, \qquad \frac{\Delta V_t^R}{U_{C,t}} \simeq \sum_i q_{i,t}^R \Delta K_{i,t} \equiv NPV_t^R.$$

A positive NPV_t^R means that the project raises discounted welfare, at the margin, relative to the specified baseline and within the modelled feasible set. It does not establish that the baseline or the project maintains the productive asset base, respects an ecological threshold, or secures a non-declining level of well-being for future generations.

7.4.7. The distinction arises because discounted-utilitarian welfare permits trade-offs across dates, people, and goods. Unless the model says otherwise, a loss of natural capital or ecosystem services may be accepted when the resulting gain in consumption or other valued outcomes more than compensates for that loss in discounted utility. A marginal appraisal may also be conducted around a baseline that is itself unsustainable. A positive NPV therefore cannot, on its own, certify compliance with ecological thresholds, safe-minimum standards, intergenerational commitments, or other sustainability requirements.

7.4.b. Exhaustible resources and sustainable well-being

7.4.8. The classic analysis of Dasgupta & Heal (1974) and Dasgupta & Heal (1979), together with Solow (1974*b*) and Stiglitz (1974), makes the distinction especially clear. In its simplest form, the economy contains produced capital K_t and a non-renewable resource stock S_t. Extraction R_t is an input to production, while output can be consumed or invested:

$$\dot{K}_t = F(K_t, R_t) - C_t, \qquad \dot{S}_t = -R_t, \qquad S_t \geq 0, \quad R_t \geq 0.$$

Consumption growth is no longer an exogenous number inserted into the Ramsey formula. It is the endogenous result of decisions about extraction, consumption, investment, technology, and substitution between produced capital and the exhaustible resource.

7.4.9. Under a discounted-utilitarian objective, the chosen depletion path maximises the present value of utility subject to these stock dynamics. Depending on technology, initial stocks, substitutability, and the welfare parameters, the optimal consumption path may be front-loaded and can rise initially before peaking and subsequently declining as resource scarcity tightens. This is not a mechanical consequence of one particular discount rate; it is the outcome of the entire dynamic optimisation problem. Equally, lowering δ need not by itself produce a path that preserves a resource stock, maintains consumption, or meets a specified intergenerational standard.

7.4.10. The alternative question posed by Solow (1974*b*) is whether the economy can sustain a constant positive level of consumption, rather than which path maximises a discounted sum. In the canonical resource model, the answer depends on the production technology and on whether resource scarcity rents are converted into other productive assets. Hartwick (1977) showed that, under restrictive conditions, investing all rents from exhaustible-resource depletion in reproducible capital can support constant consumption. Hartwick's rule is thus a wealth-maintenance or asset-substitution condition, not a modification of the Ramsey discount rate. It also illustrates why sustainability conclusions depend on the management of stocks and investment, rather than only on the ethical parameters used to discount future utility.

7.4.11. Figure 7.3 illustrates paths of consumption in some stylized economies. Neither discounted utilitarianism nor a sustainability constraint generates one universal consumption path. The comparison instead highlights that the criteria answer different questions. Discounted utilitarianism selects the feasible path with the largest discounted welfare sum. A maximin or non-declining-well-being criterion restricts the admissible intergenerational profile before choices are ranked.

7.4.12. An important point for this review is that the path of non-declining well-being cannot generally be recovered by a small adjustment to δ or μ_{CC} in the Ramsey Rule. Sustainability is an additional welfare criterion.

7.4.13. A further foundation is provided by Heal (1998), who allows the stock of the environment itself to enter utility. A schematic representation is

$$V^R(\mathbf{K}_t, t) = \sup_{\mathcal{A}_t^R} \int_t^\infty U(C_s, N_s) e^{-\delta(s-t)}\, ds,$$

where N_s denotes the environmental or natural-capital stock. This gives the stock a direct welfare role, in addition to any contribution it makes to production or future ecosystem-service flows. It broadens what can be valued within a Ramsey welfare function, but it still does not guarantee that the stock will be maintained. If substitution in utility is permitted, discounted welfare may continue to accept a decline in N_s in exchange for sufficiently large gains elsewhere.

7.4.c. Scarcity prices, wealth accounting and sustainability constraints

7.4.14. The relative-price approach developed above can be situated within a general Ramsey value function. Assuming some feasible optimum exists, let

$$V^R(\mathbf{K}_t, t) = \max_{a \in \mathcal{A}_t^R(\mathbf{K}_t)} \int_t^\infty U(C_s, E_s) e^{-\delta(s-t)}\, ds,$$

where $\mathbf{K}_t$ collects produced, natural, human, and other relevant capital stocks, and $\mathcal{A}_t^R(\mathbf{K}_t)$ is the set of allocation programmes permitted by the technologies, ecological laws of motion, resource constraints, and institutions represented in the model. The superscript R denotes the ordinary Ramsey welfare problem. It does not imply an absence of physical or institutional constraints, but it does indicate that no additional sustainability requirement has been imposed unless it is already contained in the feasible set.

7.4.15. Within this welfare problem, the contemporaneous shadow price of an environmental service flow in units of consumption is

$$p_{E,t}^R = \frac{U_{E,t}}{U_{C,t}}.$$

This is the relative price examined in the preceding sections. Its level and evolution reflect the social demand for environmental services, their physical scarcity, their substitutability with consumption, and the production and ecological conditions represented in the model. The relative-price adjustments recommended in this report are projections of how $p_{E,t}^R$ changes as these underlying conditions evolve.

7.4.16. For cost–benefit analysis, the first step is to identify the physical ecosystem-service changes caused by an intervention. Let $\Delta E_s(P)$ denote the change in an ecosystem-service flow at date s caused by project P. Ignoring other project effects for the

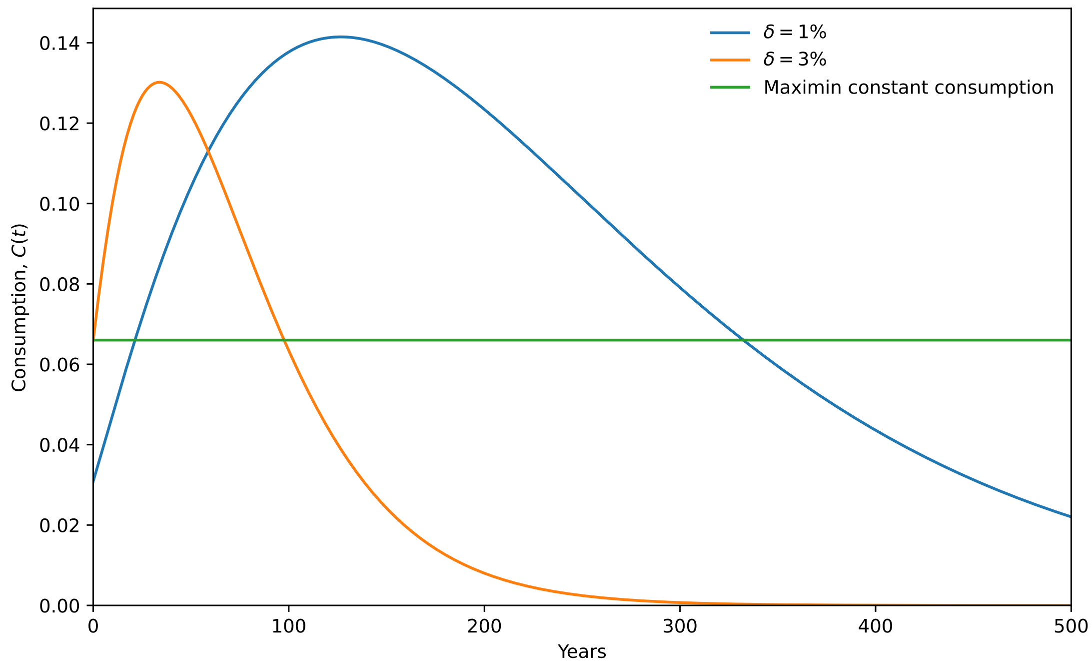

Figure 7.3.: Discounted-utilitarian consumption paths in the canonical Dasgupta–Heal–Solow–Stiglitz exhaustible-resource economy under two positive utility discount rates, labelled $\delta = 1\%$ and $\delta = 3\%$. The illustration assumes logarithmic utility, Cobb–Douglas production with capital share $\alpha = 0.6$, initial produced capital $K_0 = 1$, and initial resource stock $S_0 = 0.1825$. The horizontal path is the corresponding maximin constant-consumption level. Lowering δ shifts consumption towards later generations, but both discounted-utilitarian paths ultimately fall below the sustainable constant and converge towards zero.

moment, its marginal contribution to project value, expressed in date-t consumption units, is

$$\frac{\Delta V^R_{E,t}(P)}{U_{C,t}} \simeq \int_t^\infty D_C(t,s)\, p^R_{E,s}\, \Delta E_s(P)\, ds,$$

where

$$D_C(t,s) = e^{-\delta(s-t)} \frac{U_{C,s}}{U_{C,t}}$$

is the consumption discount factor. The appraisal procedure is therefore to forecast the physical service flow, value each future unit using its period-specific environmental shadow price, convert the effect into consumption-equivalent units, and discount those units using the consumption discount factor. Direct effects on consumption, health, time, and other outcomes can be included in the same way using their corresponding shadow prices.

7.4.17. The price of an environmental service flow should be distinguished from the accounting price of the natural-capital stock that generates it. Let N_t denote a natural asset and suppose that its current state affects ecosystem services at future dates. Its Ramsey accounting price is

$$q^R_{N,t} = \frac{1}{U_{C,t}} \frac{\partial V^R(\mathbf{K}_t, t)}{\partial N_t}.$$

This price measures the marginal contribution of an additional unit of the natural asset to current and future welfare. Schematically, it may be written as

$$q^R_{N,t} = \int_t^\infty D_C(t,s)\, p^R_{E,s} \frac{\partial E_s}{\partial N_t}\, ds + q^{R,\text{prod}}_{N,t} + q^{R,\text{res}}_{N,t} + q^{R,\text{risk}}_{N,t},$$

where the additional terms represent effects operating through production, resilience, risk, and other welfare channels not already included in the ecosystem-service flow. Non-use values may also be included where they are not represented directly by E_s.

7.4.18. The projected environmental relative-price path is therefore an input into the accounting price of natural capital, but it is not itself the price of the natural-capital stock. Constructing $q^R_{N,t}$ also requires the ecological law of motion, the persistence of the asset, the complete stream of services it generates, and its effects on production, resilience, and risk. Natural assets whose services are expected to become scarcer or more difficult to replace will see rising environmental relative prices and capital accounting values.

7.4.19. A marginal project can consequently be evaluated through either of two representations, provided that each is comprehensive and internally consistent. One can value the complete stream of ecosystem-service changes caused by the project, or one can value the project's change in the underlying natural-capital stock using its accounting price. The same welfare effect should not be included in both forms. If the full stream of service losses has already entered the NPV, the capitalised value of those same losses should not be added again as a separate natural-capital cost.

7.4.20. The accounting-price construction applies to the entire productive base. Let

$$\mathbf{K}_t = \left(K_t^P, K_t^H, N_t, \ldots\right)$$

contain produced capital, human capital, natural capital, and other assets that contribute to present and future well-being. The Ramsey accounting price of asset i is

$$q_{i,t}^R = \frac{1}{U_{C,t}} \frac{\partial V^R(\mathbf{K}_t, t)}{\partial K_{i,t}}.$$

For a marginal project whose effects can be represented as changes in these stocks,

$$\frac{\Delta V_t^R(P)}{U_{C,t}} \simeq \sum_i q_{i,t}^R \Delta K_{i,t}(P) \equiv NPV_t^R(P).$$

A positive $NPV_t^R(P)$ means that the intervention increases discounted welfare relative to the specified baseline and within the modelled feasible set. Where a project also has direct flow effects that are not captured by changes in the measured asset stocks, those effects must be added separately using the corresponding flow prices.

7.4.21. The same accounting prices can be used to evaluate changes in comprehensive wealth. Differentiating the Ramsey value function through time gives

$$\frac{\dot{V}^R(\mathbf{K}_t, t)}{U_{C,t}} = \sum_i q_{i,t}^R \dot{K}_{i,t} + \frac{1}{U_{C,t}} \frac{\partial V^R(\mathbf{K}_t, t)}{\partial t}.$$

The first term is genuine investment:

$$GI_t^R = \sum_i q_{i,t}^R \dot{K}_{i,t}.$$

The explicit-time term captures changes in technology, population, institutions, or other determinants of welfare that are not represented fully by the measured capital stocks. In a stationary setting,

$$GI_t^R = \frac{\dot{V}^R(\mathbf{K}_t, t)}{U_{C,t}},$$

so that non-negative genuine investment is equivalent, under comprehensive and correctly constructed accounts, to non-declining continuation welfare (Dasgupta & Mäler 2000, Arrow et al. 2003, 2004).

7.4.22. Ramsey accounting prices thus perform two related but distinct functions. In project appraisal they are applied to project-induced changes in the capital stocks,

$$\sum_i q_{i,t}^R \Delta K_{i,t}(P),$$

in order to compare alternative programmes. In comprehensive-wealth accounting they are applied to changes in the productive asset base through time,

$$\sum_i q_{i,t}^R \dot{K}_{i,t},$$

in order to assess whether continuation welfare is rising or falling. Project appraisal differentiates welfare across programmes, whereas genuine investment differentiates welfare through time along a programme. The prices can be the same even though the calculations answer different questions.

7.4.23. The ordinary Ramsey framework therefore has an important role in sustainability analysis. It does not merely provide a consumption discount rate. Once environmental, human, produced, and other assets have been assigned comprehensive accounting prices, the framework can be used both to evaluate marginal welfare changes and to test whether comprehensive wealth and continuation welfare are non-declining. In particular, the environmental relative-price adjustments developed in this report should feed into the accounting prices assigned to natural-capital stocks.

7.4.24. Nevertheless, a positive Ramsey NPV does not by itself establish that the resulting development path is sustainable. A project can raise welfare relative to its baseline while the baseline and resulting programme continue to exhibit negative genuine investment. It can also increase comprehensive wealth while reducing a particular natural asset if the welfare model treats accumulation of other assets as sufficient compensation. A separate question must therefore be asked about what is to be sustained: aggregate continuation welfare, comprehensive wealth, a particular natural-capital stock, an ecological condition, or compliance with a physical or legal threshold.

7.4.25. High Ramsey scarcity prices should likewise not be confused with the shadow value of a separately imposed sustainability constraint. The formulation of Drupp (2018), for example, causes the relative price of environmental services to rise sharply as those services approach a subsistence requirement. This captures essentiality and declining substitutability within the welfare function. It does not, by itself, require natural capital, comprehensive wealth, or well-being to be non-declining.

7.4.26. A sustainability requirement can instead be incorporated directly into the feasible set. Let

$$\mathcal{A}_t^S(\mathbf{K}_t) \subseteq \mathcal{A}_t^R(\mathbf{K}_t)$$

contain only programmes satisfying the relevant sustainability, ecological, or legal constraints, and define

$$V^S(\mathbf{K}_t, t) = \max_{a \in \mathcal{A}_t^S(\mathbf{K}_t)} \int_t^\infty U(C_s, E_s) e^{-\delta(s-t)}\, ds.$$

The corresponding accounting prices are

$$q_{i,t}^S = \frac{1}{U_{C,t}} \frac{\partial V^S(\mathbf{K}_t, t)}{\partial K_{i,t}}.$$

When the additional constraint is slack, these prices coincide with the ordinary Ramsey prices. When it binds, they also reflect the marginal opportunity cost of the sustainability constraint. Chapter 15 examines these sustainability-constrained prices, and the alternative of combining ordinary Ramsey CBA with a separate wealth-based or physical sustainability test, in greater detail.

7.4.27. The central point is therefore not that Ramsey prices are irrelevant to sustainability. Environmental relative prices value ecosystem-service flows; those flow prices, together with ecological dynamics, help determine the accounting prices of natural-capital stocks; and the resulting system of Ramsey accounting prices can be used to value changes in natural, human, produced, and other forms of wealth. What the ordinary Ramsey framework does not do automatically is require the programme selected by CBA to satisfy a particular sustainability criterion. That requirement must either be tested separately or incorporated explicitly into the feasible set.

7.5. Interpreting the Green Book

7.5.a. The Green Book concept of welfare

7.5.1. Several possible welfare-theoretic underpinnings of the discounted utilitarian social welfare function that underlie the Ramsey rule have been reviewed above. Taken together, this non-exhaustive review suggests that the Green Book Ramsey rule should not be interpreted as a representative agent approach. Rather, it is best understood as a pragmatic attempt to reflect a societal objective. In either case, the assumptions underpinning the precise functional form of utility are restrictive even when we move away from the representative agent framing, so the Ramsey Rule and its underlying SWF can be interpreted as an approximation for social welfare that captures some essential aspects of concern: well-being is determined by consumption (broadly defined), there is diminishing marginal utility and society might be impatient. The Ramsey Rule reflects a normative social-planner framework for valuing marginal changes in social welfare. As applied to CBA, government appraisal evaluates policy- or project-induced marginal changes in welfare around a chosen reference path of mean real per capita consumption.

7.5.2. Social Cost-Benefit Analysis seeks to express social costs and benefits in monetary terms and then discount them to obtain net present social value. The inputs to CBA are usually monetary evaluations of individual welfare changes: willingness to pay, willingness to accept, compensating variation, equivalent variation, resource costs, or consumption-equivalent impacts. In this respect, CBA has an inherent welfarist or consumer-sovereignty element: individual choices and valuations are treated as important evidence about how individual welfare changes. The Green Book maintains that tradition.

7.5.3. However, the intertemporal social welfare function in the Green Book plays a different role. It specifies how monetary welfare changes at different dates are socially valued. This does not make the approach non-welfarist. Rather, it means that the Green Book combines welfarist measurement of individual impacts with a normative social-planner rule for intertemporal aggregation. In some quarters, this kind of public welfare calculation has been associated with "Government House" utilitarianism (Sen & Williams 1982). The term was originally used critically, reflecting the concern that a utilitarian elite might make decisions for a wider public using utilitarian principles that the public itself may not share. Others have sought to reclaim the idea, arguing that utilitarian reasoning may be an appropriate public philosophy for government decision-making (see e.g. Goodin 1995). Nevertheless, a central concern remains: how should the parameters of the social welfare function be chosen, which sources of information are relevant, and what are the consequences of alternative choices (Dasgupta 2005, 2008, Nordhaus 2007)?

7.5.4. Dasgupta (2008) suggests that one can avoid the worst aspects of the "Government House" if the selection of parameters is not rooted in the opinions of government officials or so-called philosopher-kings, but is rather informed by careful reflection on their ethical foundations, empirical sources and practical consequences. This bears some resemblence to the idea of the *reflective equilibrium* proposed by Rawls (1971): the choice of parameters should be disciplined by both principles and implications. That is the approach taken in this review.

7.5.b. Sustainability and appraisal and the Green Book

7.5.5. Current Green Book practice is broadly consistent with treating natural-capital condition as a parallel appraisal consideration rather than attempting to encode it entirely in the social discount rate. The Green Book requires natural-capital screening, consideration of physical changes to natural assets, valuation of the resulting welfare consequences, and attention to uncertainty. It also recognises legal and environmental requirements as potential constraints on admissible options and states that a single summary metric such as NPSV or a benefit–cost ratio is not sufficient on its own (HM Treasury 2026*c*, paras. 5.5–5.9, 6.96–6.101, and 8.76–8.81). In practical terms, this creates a reasonable dual test: assess monetised social value, while also checking natural-capital stocks, condition, targets, and material non-monetised effects.

7.5.6. This approach is more defensible than trying to express sustainability by reducing the pure rate of time preference or making another small adjustment to the parameters of the Ramsey Rule. A change in δ changes the weight on all future consumption-equivalent impacts, regardless of which asset, person, or threshold is at issue. It neither identifies the natural-capital stock that must be monitored nor specifies the minimum condition that must be met. A separate stock or feasibility check is more transparent about the substantive policy commitment and avoids loading an economy-wide discounting parameter with a task for which it was not designed.

7.5.7. The same conclusion applies to measures of macroeconomic performance. GDP or consumption growth can remain positive while the economy runs down natural, human, or other productive assets. A serious sustainability framework therefore requires balance-sheet information on the quantities and condition of capital stocks, their accounting prices, the distribution of gains and losses, and any critical thresholds or irreversibilities. Comprehensive-wealth and inclusive-investment approaches are designed to address these questions (Arrow et al. 2003, Dasgupta 2021). They cannot be replaced by selecting a different value of δ in a Ramsey consumption rule.

7.5.8. The conclusion for this review is therefore deliberately two-part. Within the existing cost–benefit framework, we retain a common Ramsey consumption discount schedule and use explicit relative-price adjustments to reflect the changing scarcity value of environmental and other non-market outcomes. That is a useful and transparent correction to monetary appraisal, but it does not guarantee sustainability or intergenerational fairness. Those concerns depend on the social objective, the distribution of well-being, production and substitution possibilities, stock dynamics, investment, thresholds, and the definition of macroeconomic performance. They must be addressed directly through a broader wealth- and sustainability-oriented framework, including explicit natural-capital checks or constraints alongside NPV, rather than through marginal recalibration of the Ramsey Rule alone.

7.5.c. Concluding remarks

7.5.9. We interpret the Green Book social welfare function as a Utilitarian social-planner appraisal framework rather than as the preference of a literal representative agent. This interpretation preserves the welfarist foundations of CBA while recognising that the parameters of the Ramsey Rule should reflect public, not private, social valuation parameters.

7.5.10. Concerns about sustainability are not explicitly embedded in the social welfare function and the shadow prices deployed in CBA. Discussion of the calibration of the Ramsey Rule has only an indirect bearing on this issue. Sustainability is captured to some extent by the recommendation in Section 5 to evaluate impacts on natural capital via the ENCA guidelines. Other forms of national wealth are not specifically addressed. Relative price adjustments proposed reflect scarcity not necessarily sustainability in the Green Book (Section 6).

7.5.11. The task of calibration is therefore to choose those parameters transparently, drawing on both evidence and normative reasoning, and with explicit attention to the consequences of alternative choices, noting that sustainability issues are better handled via explicit accounting for national wealth impacts (as intended by Green Book guidance and ENCA) or shadow prices that reflect sustainability constraints.

8. Choosing the Ramsey Parameters

8.1. Introduction

8.1.1. This chapter considers the choice of the two preference parameters in the Green Book Ramsey Rule: the pure rate of social time preference, δ, the elasticity of marginal utility of consumption, μ, as well as the expected rate of growth of real per capita consumption, g. Before deciding which estimates are relevant, it is necessary to be clear about what the intertemporal social welfare function represents and, therefore, what kinds of information can legitimately be used to calibrate it.

8.1.2. The preceding chapter interpreted the Green Book as adopting a pragmatic discounted-utilitarian social-planner framework, rather than the preferences of a literal representative agent or a complete general-equilibrium model of the economy. Cost–benefit analysis evaluates the welfare effect of policy- or project-induced deviations from a chosen reference path, expressed in consumption-equivalent units. The reference path is the expected path of mean real per capita consumption. This interpretation matters directly for calibration. The parameters of the Ramsey Rule are public social-valuation parameters. They need not reproduce private impatience, individual saving behaviour or any single market return, although all of these may contain information that is relevant to this social choice, of for no other reason than the welfarist principle that individuals are the best judges of their own welfare.

8.1.3. We begin by reporting, in Table 8.1, the percentage of respondents to our survey who considered the current Green Book values of δ, μ and g to remain broadly appropriate. By a ratio of approximately two to one, respondents considered that the parameter values should change. This view was particularly strong among government respondents, while the expert academic group was evenly divided.

8.1.4. Table 8.2 reports the median recommendations for δ, μ as well as for growth, g, at a 10-year and 100-year horizon. The number of respondents varies slightly by question; the count shown is for the question on δ. Despite the appetite for change shown in Table 8.1, two of the median recommendations for the whole sample are the same as the current Green Book values: $\delta = 0.5\%$ and $\mu = 1$. Forecasts for growth, however, at 1.35% median for 10 years (1.5% from non-governmental respondents) and 1% median for 100 years are well below the current value of $g = 2\%$. This apparent tension is informative rather than contradictory. It reflects disagreement about the separate parameters, the term structure and the rationale for the existing calibration, as well as a concentration of responses around the present values. The qualitative responses also contain a substantial body of support for increasing μ, an issue to which we return below.

8.2. Interpreting the Green Book Ramsey Rule

8.2.a. Discounted utilitarianism and the reference path: Recap

8.2.1. As discussed in Chapter 7, the discounted-utilitarian social welfare function underlying the simple Ramsey Rule can be written as

$$W(\mathbf{c}) = \sum_{t=0}^{\infty} \beta^t U(c_t),$$

where c_t denotes mean real per capita consumption, broadly defined, and $U(c_t)$ is the flow of social well-being associated with it. For the present discussion, population and the intratemporal distribution of consumption are held fixed. Distribution within a period is addressed separately in Chapter 16.

8.2.2. The equation in 8.2.1 should not be read as the lifetime utility of one immortal consumer. Nor does its use in appraisal require the economy to be on a competitive equilibrium path in which the Ramsey rate is equal to a market interest rate. It is a tractable social-planner representation of how marginal welfare changes at different dates are aggregated, albeit over infinite time. The reference path is denoted by $\bar{\mathbf{c}} = \{\bar{c}_0, \bar{c}_1, \ldots\}$. A policy or project generates a stream of consumption-equivalent net benefits, x_t, so that the relevant path becomes $c_t = \bar{c}_t + x_t$. Its exact welfare effect is $\Delta W = \sum_{t=0}^{\infty} \beta^t \left[U(\bar{c}_t + x_t) - U(\bar{c}_t)\right]$, which has the following first order approximation

$$\Delta W \simeq u'(\bar{c}_0) \sum_{t=0}^{\infty} \left[\beta^t \frac{U'(\bar{c}_t)}{U'(\bar{c}_0)}\right] x_t.$$

The term in square brackets is the social discount factor for a consumption-equivalent impact at date t. Normalising by current marginal utility, net present social

	Count	Yes	No
UK Academic	26	31%	69%
Social Media	21	33%	67%
SBCA	15	33%	67%
Experts	12	50%	50%
Total (non-government)	74	35%	65%
Government	15	20%	80%
Total (whole sample)	89	33%	67%

Table 8.1.: The percentage of respondents to our survey who believe that the Ramsey Rule parameter values (δ, μ and g_c) reported in the 2026 Green Book remain broadly appropriate.

	Count	δ	μ	g (10 years)	g (100 years)
UK Academic	22	0.50%	1.00	1.50%	1.00%
Social Media	17	0.25%	1.00	1.00%	1.00%
SBCA	12	0.35%	1.10	1.50%	1.00%
Experts	9	0.10%	1.40	1.50%	1.00%
Total (non-government)	60	0.45%	1.00	1.50%	1.00%
Government	8	0.50%	1.00	1.00%	1.00%
Total (whole sample)	68	0.50%	1.00	1.35%	1.00%

Table 8.2.: Median recommendations for the Ramsey Rule parameters δ, μ, and g in our survey.

value is therefore

$$NPSV = \sum_{t=0}^{\infty} D_t x_t, \qquad D_t = \beta^t \frac{U'(\bar{c}_t)}{U'(\bar{c}_0)}.$$

Thus CBA evaluates deviations from the reference path, rather than selecting or optimising the whole path anew for every appraisal. Of course $D_t = e^{(-STPR*t)}$ where in the isoelastic case this leads to

$$r \simeq \delta + \mu g.$$

The Green Book adds the separate term L and writes

$$STPR = \delta + L + \mu g.$$

This chapter is predominantly concerned with δ and μ, returning to g in Section 8.6. The parameter δ changes the social weight attached to utility solely because of its date. The parameter μ determines how the marginal social value of a unit of consumption changes as the reference level of consumption changes.

8.2.b. Principles for choosing the parameters

8.2.3. The social-planner interpretation does not imply that parameter choice should proceed by ethical introspection alone. Nor does the welfarist basis of CBA imply that every observation of private behaviour is admissible. We use four principles to organise the evidence.

Normative and positive evidence are both eligible. The parameters have normative content because they determine how welfare is aggregated over time and consumption levels for multiple generations. Estimates require empirical analysis and some introspection. At the same time, CBA is grounded in the welfarist idea that people are ordinarily important judges of what contributes to their own welfare (e.g. Boadway & Bruce 1984). Observed behaviour, stated judgements and actual social choices may therefore discipline parameter selection. The relevant question is not whether evidence is "positive" or "normative", but whether we can learn something about the measurement of social welfare for the purposes of the public project appraisal.

The evidence should reflect social rather than private decision making. Private choices and public judgements are not the same object. An individual deciding how much to save for retirement, or choosing between two small experimental payments, may appropriately give special weight to themselves and their family. The same individual, considering a public policy as a citizen, may adopt a more impartial position. This is the sense in which individuals can wear different "hats" (Harsanyi 1982, Dietz et al. 2009). Following Harsanyi (1982), (Harsanyi 1982, Dietz et al. 2009) distinguish manifest preferences from the preferences that would be formed with adequate information, careful reasoning and a context appropriate to the social question. They also emphasise that revealed private behaviour may identify personal ethics rather than a view clear view on the preferred political or social arrangements.

8.2.4. This principle does not exclude all empirical evidence. Choices made explicitly on behalf of society, public decisions, socially framed experiments and informed stated judgements are eligible. Aggregate market evidence may also be informative, particularly because markets aggregate many decisions and constrain what governments can implement. Nevertheless, aggregation in a market does not by itself transform private preferences into a social ethic. Market outcomes reflect the distribution of wealth, institutions, taxes and constraints, as well as preferences. In sum, markets prices probably do not price in the wellbeing of unborn future generations. Yet it remains difficult to design credible experiments to elicit social pure time preference for very long horizons (Grijalva et al. 2014)

Parameter choice as a reflective equilibrium. Following Dasgupta (2005, 2008), parameter values should be tested against both their underlying principles and their implications. This is close to the reflective-equilibrium approach associated with Rawls (1971). A principle that initially appears compelling may generate implications that are difficult to accept; conversely, an observed outcome may rest on assumptions that are ethically unattractive. The purpose of examining thought experiments, model implications and empirical calibrations is therefore not to replace ethics with consequences, but to have a balanced consideration of the relevance of both.

8.2.5. This process is also one way of avoiding the worst implications of "Government House" utilitarianism. Public officials should not select ethical parameters as philosopher-kings and then conceal their judgements within a technical formula. As Dasgupta (2008) puts it,

> "social ethics contains an irremediably democratic element."

The reasons for the choice, the evidence considered, the implications of alternatives and the remaining disagreement should therefore be made explicit. Expert opinion is relevant because the problem is technically and ethically difficult, but experts provide evidence and argument rather than political authority (e.g. Pindyck 2013, Sunstein 2014).

Market rates provide triangulation, not complete identification. A normative calibration need not coincide with observed interest rates. The Ramsey Rule in CBA is not an equilibrium equation and many social benefits are not tradable assets. Even so, a very large and persistent divergence between social discount factors and market prices can create practical difficulties. Where a project's cash flows can be replicated in as-if markets, inconsistent discounting can imply opportunities for fiscal arbitrage. Where public expenditure displaces private investment or must be financed through borrowing and taxation, market returns affect opportunity cost even if they do not determine social ethics.

8.2.6. Equally, introducing a STPR that is lower than the rate of return in the market place can introduce several 'second best' issues, which in turn means that interlinked countries need to consider further, complimentary policies, such as capital subsidies, to achieve a welfare maximising environment (Barrage 2018).

8.2.7. While there are many reasons why the STPR may not coincide with SOC rates, ranging from argument relating to the isolation savings paradox Sen (1967), population/evolutionary dynamics (Robson & Szentes 2014), or even the unexpected effects of altruism and bequest (Nesje 2024). The real gilt yields, safe rates and broader asset returns are therefore useful checks on a calibration and may require adjustments elsewhere in appraisal. They should not be treated as direct estimates of δ or μ.

8.3. The current calibration and evidence not used

8.3.a. The Green Book calibration; recap

8.3.1. The 2026 Green Book retains the decomposition

$$STPR = \delta + L + \mu g,$$

with $\delta = 0.5\%$, $L = 1\%$, $\mu = 1$ and $g = 2\%$ for the first 30 years. The terms δ and L were added together and expressed as ρ giving, we think, the wrong inpression that these terms were part of a general utility discount rate, In any event, summing together this gave an initial STPR of 3.5% (HM Treasury 2026*a*). The adoption of a social-time-preference approach in 2003 marked a significant change in UK practice. Earlier guidance had placed greater weight on the social opportunity cost of capital and used substantially higher rates.

8.3.2. Pearce & Ulph (1995) is a useful reference here to understand the theoretical and empirical background behind the parameterisation. Moreover, it was here that the arguments for the use of the Ramsey Rule STPR are laid out.

8.3.3. The remaining subsections explain why three sources sometimes used to motivate the utility discount rate ρ in the current guidance are not used in the calibration here: experimental evidence on private impatience including on hyperbolic discounting (relating to (δ), individual mortality risk, and a generic existential hazard (relating to ρv via L in the previous guidance). The exclusion follows primarily from the interpretation of the welfare function as a social concept and not a representative agent. These sources of evidence are interesting in their own right, but we see them as not fitting the principles outlined above.

8.3.b. Private impatience and hyperbolic discounting

8.3.4. There is a large literature estimating individual time preferences from observed behaviour and from hypothetical or incentivised experiments (Frederick et al. 2002, Andreoni & Sprenger 2015, Andersen et al. 2014, Cohen et al. 2020). Estimates vary greatly with the elicitation method, horizon, stakes, payment mechanism and treatment of risk. Present bias and apparently hyperbolic discounting are common, and annualised short-horizon rates can be inthe order of 30%, so extremely high for a social discount rate. This introduces a tension between the principle of consumer sovereignty in CBA, which states that individuals are the best judges of their own well-being, and the much lower values of the social discount rate typically used in practice, bet they SOC or STPR.

8.3.5. This point is forecefully put by Henderson & Bateman (1995), who remind us of the importance of individual behaviour and positive inputs into CBA:

> "To deny [observed hyperbolic behaviour] and to insist that the social discount rate be expressed only as an exponential function is to risk ignoring positivist evidence by substituting a normative judgment as to how people ought to discount."

(emphasis our own)

8.3.6. Nevertheless, our view is that certain types of behavioural evidence is unsuitable in the present setting. The behavioural evidence on hyperbolic discounting generally concerns small, private allocations over hours, days or months. The Green Book parameter concerns a government acting for society over years and generations. Even a well-identified private discount function would therefore identify the wrong object. The empirical/experimental difficulties reinforce this conclusion: apparently hyperbolic behaviour can reflect limited credibility, complexity, liquidity constraints or the confounding of time and risk preferences (Andreoni & Sprenger 2015, Epper & Fehr-Duda 2023, Cavatorta & Groom 2025, Enke et al. 2023).

8.3.7. Socially framed experiments are not ruled out by this argument. Studies asking respondents to make decisions for society, with a horizon and informational setting relevant to public policy, are potentially admissible. The available evidence suggests that socially framed rates lead to much lower social discount rates being expressed than private rates (Howard 2013, Venmans & Groom 2021). Such exercises remain sensitive to framing and incentive compatibility, particularly at long horizons (Grijalva et al. 2014, 2018), but they at least frame the questions as social rather than private behaviour. There are difficulties in design.

8.3.8. There is also a practical reason not to embed private hyperbolic discounting in the social welfare function. A hyperbolic public discount rule can generate preference reversals and time-inconsistent appraisal: a project accepted for the future may be rejected when the future becomes the present. It can also expose government to self-defeating sequences of decisions or Dutch-book-type losses (Karp 2005, Hepburn et al. 2010).

8.3.9. These issues were discussed at length prior to the introduction of the 2003 Green Book update where declining social discount rates were introduced for the first time. Importantly, time varying social discount rates, based on uncertainty of future growth, can be applied through a time-consistent valuation rule and do not introduce parallel behavioural biases in social decision making as an happen at he individual level due to hyperbolic discounting. Time consistency is in principle maintained (Gollier 2013, p65). See also Groom et al. (2005) for a discussion.

8.3.c. Individual mortality risk

8.3.10. For similar reasons, our view is that individual mortality risk is also unsuitable as a component of social pure time preference or as part of an extended utility discount rate ρ. Mortality can in principle explain why an individual places less weight on consumption that they may not live to enjoy. It does not follow that a social planner should place less weight on welfare at that date, since other people will be alive and may receive the benefits. This distinction was already recognised by Pearce & Ulph (1995), who questioned whether an individual's increasing risk of death should determine the government's weighting of consumption at different dates in the background paper to the 2003 revision to the Green Book (see also Oxera 2002).

8.3.11. The only case in which a mortality-like probability enters the discount factor as a social concept is as the survival probability of the welfare-bearing entity: society and human-kind as a whole. As we discuss in more detail below, if the probability that society itself continues to exist to date t is $S(t)$, then *expected* (probabiity weighted) social welfare is weighted by $S(t)$; with a constant existential hazard, this is algebraically equivalent to an addition to the utility discount rate. But the annual proportion of current individuals who die, which is the empirial motivation in Pearce & Ulph (1995), is not that hazard. It is a rate of demographic turnover. In the per capita Green Book social welfare function, background deaths and births do not by themselves alter the reference path of per capita welfare. In a total-population social welfare function, the relevant demographic object would be the path of population N_t. This welfare measure is influenced by births, deaths and migration together, not the crude death rate alone. This leads to rather different welfare assessments and an entirely new literature on population ethics which has not entered into the Green Book to date (Dasgupta 2005, Millner 2013, Broome 2004*b*, Parfit 1986).

8.3.12. This matters for interpreting the older estimates. The calculation in Newbery (1992) referenced in the current guidelines is supposed to reflect a rough hazard of the end of mankind or society. Pearce & Ulph (1995) report it in that spirit, but also note that a similar number can be obtained by dividing UK deaths by population. However, the numerical similarity is misleading. A 1% annual extinction hazard would imply roughly a 63% probability that society has ceased to exist within a century. A 1% crude death rate merely says that one per cent of the current population dies each year. If births or migration replace those deaths, future welfare-bearing persons remain. The statistic therefore cannot be used as a generic reason to discount all future social benefits. Apart from anything, it is a rather large number for an existential risk.

8.3.13. A simple example makes the normative problem clear. Consider two stationary societies with constant populations and identical paths of per capita consumption and well-being. In Society A, each person dies with certainty at age 65 and is immediately replaced by a birth. In Society B, death and replacement occur at age 70. Assume, for the purpose of the example, that well-being at every date is otherwise identical. An anonymous discounted-utilitarian social welfare function assigns the same value to the two streams: the identity of the person receiving a unit of well-being does not matter. A discount rate based on individual survival would nevertheless discount the future more heavily in Society A. It would treat a unit of future well-being differently solely because it is more likely to accrue to a replacement person rather than to someone alive today. That violates the anonymity or person-neutrality implicit in the social-planner interpretation, which we see as problematic.

8.3.14. This does not imply that mortality is irrelevant to CBA. Mortality matters when a project changes survival probabilities, life expectancy, health or the number and quality of life-years enjoyed. Those effects belong in the benefit stream itself. Subject to whatever distributional, age-related or health-state conventions are adopted, an anonymous social CBA does not treat a life-year differently merely because it is enjoyed by a different person.

8.3.15. So there are problems with using individual mortality risk as a generic discounting term. Using the probability that members of the current cohort will die reduces the weight attached to all future net-benefits, including benefits enjoyed by replacement persons who will be alive at the relevant date. It therefore turns a social welfare calculation into a current-cohort survival calculation. The 65/70 example exposes this clearly: the per capita welfare streams may be identical, but a mortality-adjusted discount rate values them differently because the identities of future beneficiaries change more quickly in one society than in the other. Mortality should therefore enter as an outcome affected by the project, not as a generic reduction in the value of future social welfare.

8.3.16. There is a separate behavioural reason for caution. Heimer et al. (2019) show that subjective survival beliefs are systematically distorted: younger people tend to underestimate survival while older people tend to overestimate it, and these beliefs help explain under-saving among the young and slow decumulation among retirees. Private financial behaviour therefore reflects perceived mortality, bequest motives, annuity-market imperfections and other constraints, not simply an actuarial survival probability.

8.3.17. For all of these reasons we find individual mortality risk ineligible as a basis for informing the social discount rate, despite some fascinating work in this area (Kula 1985, 1987, Addicott et al. 2020).

8.3.d. Existential hazard and disaster risk

8.3.18. A constant, exogenous probability that society ceases to exist can be represented algebraically as an additional exponential survival factor. If the survival probability of the welfare-bearing entity is $S(t) = e^{-ht}$, expected welfare contains the term e^{-ht}, which is equivalent to adding h to the utility discount rate (Yaari 1965). This observation explains the long-standing use of catastrophe risk in decompositions of the Green Book rate. It does not make the hazard an element of pure time preference. It is a probability of existence, not a judgement that future well-being matters less merely because it occurs later. Note that this approach is different from work by Dasgupta & Maskin (2005) and Sozou (1998) who model the hazard of payoffs disappearing as an explanation for hyperbolic discount rates among individuals.

8.3.19. The numerical implications are also worth making explicit. If a survival probability over horizon T is translated into a constant hazard, the implied annual rate is

$$h = -\frac{\ln S(T)}{T}.$$

Thus the suggestion in Rees (2003) that humanity faced roughly a 50–50 chance of surviving the next century implies

$$h = -\frac{\ln(0.5)}{100} \simeq 0.0069,$$

or about 0.7% per year. This is below Newbery's 1% estimate, reported as consistent with the perceived risk of the end of mankind in 100 years, but much significantly higher than the Stern Review's 0.1% assumption (Newbery 1992, Pearce & Ulph 1995, Stern 2007). The differences are large. A 0.1% annual hazard implies about a 9.5% probability of extinction within a century; a 0.7% hazard implies 50%; and a 1% hazard implies about 63%. These are substantive beliefs about existential risk, not innocuous calibration details.

8.3.20. This comparison also reinforces the point made above about mortality. 1% existential risk is not defensible as the annual proportion of current individuals who die. Ordinary demographic turnover leaves future welfare-bearing persons in place; extinction removes the future welfare stream itself. The relevant probability for an additive hazard term is therefore the survival probability of the social welfare-bearing entity, not the death rate of current individuals.

8.3.21. There is a further modelling issue. The Yaari transformation is algebraically convenient only because the hazard is assumed to be constant, exogenous and memoryless. As Dasgupta (2008) notes in this context, folding such a probability into an exponential discount factor effectively turns an uncertain future existence problem into a deterministic survival-weighted welfare stream. That may be a useful shortcut for a small, policy-invariant background hazard, but it is not a satisfactory representation of risks whose probability changes over time, is affected by policy, or is correlated with consumption and social conditions. In those cases the hazard should be modelled directly.

8.3.22. Most catastrophic risks relevant to appraisal are not literal extinction risks in any case. They are risks of large losses to consumption, output, capital, health, security or social functioning. In such cases the effect on discounting can have the opposite sign from a simple positive addition to δ. With diminishing marginal utility and prudence, low-consumption states increase the value of marginal future benefits, and can lower the appropriate risk-free social discount rate. This is the approach taken in macroeconomic catastrophe models such as Pindyck & Wang (2013), and it is also the spirit of Martin & Pindyck (2015), where multiple potential catastrophes interact and cannot always be evaluated by treating each hazard as an isolated exponential adjustment.

8.3.23. Our view is to recommend that we do not use existential risk to calibrate δ (or ρ via L in the Green Book currently). A strictly existential and policy-invariant background hazard could, in principle, be recorded as a separate survival adjustment. But the plausible numerical value is deeply uncertain, the exponential representation is restrictive, and the relevant risks for appraisal are usually better modelled as explicit stochastic shocks to consumption, growth, capital or survival probabilities as in Barro & Ursúa (2011), Barro (2006). Such shocks actually reduce the discount rate for precautionary reasons as we discuss in Chapter 9. Catastrophic and systematic risks are considered in the risk analysis elsewhere in this Review, including through models in which rare disasters reduce consumption rather than eliminate all future welfare; see Chapter 12.

8.3.24. Deeper considerations of population ethics are always possible, but for the purpose of choosing δ, individual mortality and a generic existential hazard are set aside.

8.4. The pure rate of time preference: δ

8.4.1. Once individual mortality and generic catastrophe risk are removed, δ has a precise interpretation: it is the rate at which the social welfare function reduces the weight attached to a unit of utility solely because it occurs later. This is an ethical quantity, but it can be informed by positive evidence. We consider aggregate and market calibrations first, then the normative arguments for zero and positive pure time preference, and finally the evidence from expert and social judgements.

Table 8.3.: Qualitative quotations on δ existential risk, risk eligibility in δ and the parameterisation in general

Use	Quote	Note
Against putting risk adjustment in the discount rate	"I don't think risk adjustment should be through the discount rate. It should be through the quantity of costs and benefits."	Strongest general quote against embedding risk in δ or the discount rate.
Against using PRTP for environmental/nature risk	"I am in general in favour of lower discount rates for environmental projects because the time preference should really not apply, given how dependent society is on nature, and growth rate is also not appropriate."	Useful for long-run environmental risks; not explicitly existential risk.
Direct extinction-risk mention, sceptical	"One could account for extinction risk as part of the PRTP, e.g. the Stern Report used 0.1% based on that motivation, although it did seem like that number was sort of pulled out of thin air."	Best direct extinction-risk quote; sceptical of the basis for including it in PRTP.
Against current 0.5% despite extinction-risk possibility	"I would prefer a PRTP = 0%, or perhaps 0.1%, or somewhere in between, over the 0.5%."	Comes from the same response as the extinction-risk quote; supports near-zero rather than current PRTP.
Risk should be modelled through bad-outcome states	"For longer time horizons pure time preference should be lower, and small chances of bad outcomes, discount rate should be adjusted by probability and level of consumption in bad cases."	Supports explicit treatment of bad states rather than relying only on a central discount rate.
Against changing PRTP for global uncertainty	"Global uncertainty might edge pure time preference up a bit – but it is at best a guestimate and should probably be left alone."	Useful against incorporating broad global uncertainty into PRTP without a clear basis.
Against Ramsey parameters as real-world valuation of risk	"None of these quantities reflect valuation in the real world. Rather, they are derived from estimates of parameters in highly stylized models of the macroeconomy that are informative about basic mechanisms but not about how society evaluates the effects of time and risk on value."	Useful caution against treating the parameters as direct evidence on time and risk valuation.
Risk preferences not fully captured by Ramsey/CRRA	"The Ramsey rule is an incomplete description of the sources of the STRP: what if risk preferences are more Epstein-Zin than CRRA?"	Useful if arguing that risk treatment requires more than a simple Ramsey parameterisation.

8.4.a. Positive and aggregate approaches

Calibration to market returns

8.4.2. The most developed positive approach chooses δ as a residual of the equilibrium/optimal version of the Ramsey Rule's. In its simplest form $\delta = r - \mu g_c$, where r is the risk-free market rate of return to capital and delta is the difference between that and the wealth effect. The more general idea is that δ is chosen to calibrate the model of aggregate consumption and asset returns to selected market moments, such as the observed risk free rate of return. This is an eligible approach according to our principles and under the social-planner interpretation, provided the identifying assumptions are made explicit. The attraction of the positive approach, compared to a fully normative alternative is expressed forcefully by Nordhaus (2007):

> "[The positive] approach does not make a case for the social desirability of the distribution of incomes over space or time of existing conditions, any more than a marine biologist makes a moral judgment on the equity of the eating habits of marine organisms in attempting to understand the effect of acidification on marine life."

8.4.3. Barrage & Nordhaus (2024) provide a prominent recent example in DICE-2023. Their objective is explicitly a social welfare function, with δ and μ described as normative parameters, but the calibration is disciplined by observed economic outcomes, in this case the risky return in the market.

8.4.4. DICE-2023 sets the pure rate of social time preference to 1% and the elasticity of marginal utility to 1.5. They introduce a precautionary adjustment for uncertainty in consumption growth and a climate-risk adjustment based on a climate beta. With consumption-growth volatility of approximately 1% per year, a climate beta of 0.5, a market risk premium of 5% and a near-term real risk-free rate of 2%, the model produces a near-term required return on climate investment that is consistent with this market calibration. The role of the exercise is not to infer ethics directly from markets, but to select a set of welfare and risk parameters whose combined implications are not grossly at odds with observed returns. This precise case focuses in on the climate application, with the climate beta making the calibration of δ here slightly problematic.

8.4.5. A similar strategy is used by Emmerling et al. (2025) in the context of an equity-adjusted social cost of carbon. Their calibration is not specifically focused on the climate context though, since their calibration organises mainly around the risk free rate, not the risky market rate. They do, however, use a more flexible preference structure: Epstein-Zinn preferences, which allows risk and intertemporal substitution preferences to be disentangled.

8.4.6. Conditional on an elasticity of intertemporal substitution of 2/3, so that the coefficient of aversion to intertemporal fluctuations is 1.5, the authors use macroeconomic volatility and rare-disaster risk to calibrate relative risk aversion to 7.313, matching a macroeconomic risk premium of about 2.75%. They then back out a pure rate of time preference of 1.546% so that the model matches a real risk-free rate of 1%. This is close in motivation to the DICE calibration: δ is a residual that reconciles a specified welfare-and-risk model with aggregate asset-market moments. Its numerical value is therefore conditional on the model, the separation of risk aversion from intertemporal substitution and the market moments selected for matching.

8.4.7. These exercises demonstrate why market evidence can be used for triangulation rather than identification. A residual calibration cannot distinguish a high pure rate of time preference from misspecification of growth, risk, preferences, taxes or market frictions. More broadly, observed asset prices may not give unborn generations or the future in general the level of voice that a social planner may want to give them. Unpriced externalities, e.g. from environmental pollution or positive spillovers, are not fully captured by market rates either. In the end the object around which the calibration takes place may not be correct either.

8.4.8. Nevertheless, a proposed public calibration should be able to explain why its implications differ from safe rates and returns on capital, especially for benefits and costs that can be replicated or financed in markets. Market calibration is therefore a useful stress test on a normative choice, not an independent moral argument for it.

Saving behaviour and cross-dynastic altruism

8.4.9. As an example of when the social discount rate may be different to the decentralised one, Nesje (2024) looks at the implications of altruistic behaviour, which is often argued to be embedded in market rates as a result of bequests. For this reason, aggregate saving is sometimes treated as a revealed social judgement about the relative value of present and future consumption. Even at the aggregate level, this inference is incomplete.

8.4.10. Nesje (2024) distinguishes concern for one's own descendants from concern for members of the next generation in other dynasties. When people care about the next generation as such, saving for one's own descendants creates preference externalities because it also benefits people who care about those descendants across dynastic lines. Decentralised saving then understates the efficient social weight on future generations. In Nesje's numerical illustration, cross-dynastic concern only 10–20% above the altruism inferred from saving reduces the efficient discount rate by approximately 20–40% relative to a Nordhaus-style calibration. The precise numbers are model-dependent, but the conceptual point is robust: observed saving need not reveal the full social concern for future people. This is akin to the isolation paradox of Sen (1967).

Socially framed stated preferences

8.4.11. Stated-preference evidence can also be relevant where the elicitation is explicitly social. Respondents may be asked to choose between public benefits accruing to different generations, rather than between private payments to themselves. Such studies tend to find substantially lower rates than the private time-preference literature, often close to 1% or below (Howard 2013, Venmans & Groom 2021).

8.4.12. The chief advantage of experimental work of this kind is its contextual relevance. The weaknesses are sensitivity to framing, limited incentives and the difficulty respondents face in reasoning over very long horizons (Grijalva et al. 2014, 2018). Note that in a similar context to Venmans & Groom (2021)z, Disque et al. (2026) find values of pure time preference in the order of 4%, which is extremely high when compared to expert opinion for instance. We therefore treat these studies as supporting evidence rather that social pure time preference is typically lower than private, than as a direct estimator of δ.

8.4.b. Normative arguments

The case for zero pure time preference

8.4.13. The central normative argument for $\delta = 0$ is equal treatment of utilities through time. If two units of well-being are otherwise identical, their date does not by itself appear to provide a reason for treating them differently, so the argument goes. This is the impartial-utilitarian position and has a formidable intellectual pedigree. In the aftermath of the Stern Review (Stern 2007) most economists became familiar with several quotes concerning the general Utilitarian tradition on pure time preference. For instance, Ramsey (1928) wrote that discounting later enjoyments relative to earlier ones is

> "ethically indefensible and arises merely from the weakness of the imagination" (Ramsey 1928).

8.4.14. Pigou's formulation is more concise. Positive pure time preference implies that

> "our telescopic faculty is defective" (Pigou 1930).

8.4.15. Harrod was more forceful still, describing positive pure time preference as

> "a polite expression for rapacity and the conquest of reason by passion" (Harrod 1948).

8.4.16. Solow later expressed the public rather than private version of the claim:

> "in solemn conclave assembled, so to speak, we ought to act as if the social rate of time preference were zero" (Solow 1974*a*).

8.4.17. The qualification in Solow's wording is important. Individuals may be impatient, but a government considering the interests of successive generations can be asked to adopt a more impartial stance. This is precisely the distinction between the private and public "hats" described above. It also explains why the fact that individuals exhibit positive time preference does not answer the social question.

8.4.18. Böhm-Bawerk is also frequently quoted in this literature, but for a different reason. He described present preference as being

> "most frankly expressed in children and savages".

This is revealing as intellectual history, but it does not provide a satisfactory rationale for a social pure time preference from an ethical perspective. It moves from a contentious description of particular groups to a conclusion about social valuation, which seems at least procedurally unhelpful as a means of determining social pure time preference.

8.4.19. A further argument for a low long-run value comes from Millner (2020). Millner considers social planners who are not completely certain that their current ethical theory is correct and who assign positive probability to revising it in the future. Although such planners begin with different welfare criteria, nondogmatism causes them to agree on the long-run social discount rate. Under the additional condition that zero pure time preference remains among the admissible views, the most patient view dominates asymptotically and the pure-time-preference component tends to zero. This is not a proof that δ must be zero at all horizons. It is a result about how not assuming that future generations share the current generation's pure time preferences affects the way the planner today see the distant future. It provides a principled reason for placing particular weight on a zero-δ sensitivity at long horizons, provided that that is the lowest pure time preference considered for the future, which seems reasonable.

8.4.20. The case for zero is therefore strongest where δ is interpreted narrowly, as it is here. It follows from equal treatment across dates, not from a claim that the consumption discount rate itself should be zero. Even with $\delta = 0$, the wealth effect μg generally gives a positive rate when future people are expected to be richer.

The case for positive pure time preference

8.4.21. **Arrow's argument** (Arrow 1999): Impartial utilitarianism is not the only defensible ethical stance when it comes to pure time preference. Agent-relative views allow a generation to give some priority to its own members while retaining duties to future people. Arrow (1999) captures the tension by quoting the Rabbi Hillel:

> "If I am not for myself, then who is for me? If I am not for others, then who am I? If not now, when?" (Arrow 1999, p. 15)

8.4.22. The point is not simply that present people happen to be impatient. It is that a the ethics of the social pure time preference may recognise a legitimate claim of the present generation on its own resources.

8.4.23. There is clearly a tension here between the two positions on offer: $\delta = 0$ and $\delta > 0$ from an ethical perspective. A respondent to the expert survey in Nesje et al. (2023) captures the difficulty succinctly:

> "Essentially, you allow people to choose between being utilitarian with a zero utility discount rate or utilitarian with a positive utility discount rate. I disagree with both these criteria, as the first one essentially gives zero weight to the present generation, even if it is the worst-off, while the second one treats generations at different points in time unequally. There are attractive alternatives to both these criteria."

Alternatives include the explicit accounting for sustainability discussed in previous sections, which can solve some intergenerational fairness issues, irrespective of the rate of pure time preference. Many other alternatives were proposed in Nesje et al. (2023).

8.4.24. Dasgupta (2005, 2008) show why the discussion of pure time preference can be material in growth models. With positive returns and an impartial objective extending indefinitely into the future, a very low δ can imply extremely high saving by the present generation, even when future generations are expected to be richer. The resulting "hair-shirt" implications are a challenge to the view that temporal impartiality is sufficient for intergenerational fairness.

8.4.25. Olson & Bailey (1981) make a related argument for positive time preference. In economies with productive investment and a sufficiently long horizon, zero pure time preference can make the claims of an indefinitely expanding future dominate present consumption. Their conclusion is conditional on the production opportunities and horizon assumed; it is not a general logical proof that all rational agents or governments must be impatient. As with the example from Dasgupta (2008) above, its continuing relevance is as a consequence test. A social objective that appears impartial at the level of utility weights may impose a highly asymmetric burden on the first generation when combined with compound returns and, crucially, unbounded future production.

8.4.26. **Koopmans and Diamond** (Koopmans 1960, 1972, Diamond 1965): The axiomatic literature discussed in the previous chapter provides another route to positive discounting. Under conditions including stationarity, continuity and separability, Koopmans (1960) obtains a recursive ordering of infinite consumption streams but only with a positive pure time preference. the contributions do not use what one could describe as ethical or moral arguments. Rather these contributions show that selecting $\delta = 0$, no matter the apparently compelling nature of impartial Utilitarianism, is not costless from the standpoint of coherent ordering over an infinite horizon.

8.4.27. Simple thought experiments at the policy coal face offer a simple illustration of the tension. In a stationary economy with no growth, risk or other differences, $\delta = 0$ makes the social planner indifferent between an otherwise identical hospital benefit delivered today and one delivered in a 140 years, 1000 years, or indeed at any point in the distant future. That implication can seem anomalous in public investment appraisal, which would clearly choose a hospital now compared to the alternatives above.

8.4.28. It is also worth noting that a positive value for δ is a blunt solution to this problem. With $\delta = 0.5\%$, one hospital today has approximately the same present value as two identical hospitals in 140 years. The decision here also seems obvious, other things equal. The examples do not decide the issue. In fact a parallel process is likely required to handle such issues of fairness and strategy (see below). The value of these examples is simply a reflective one, allowing use to consider the consequences of each ethical position.

8.4.29. One response is to retain $\delta = 0$ within an abridged utilitarian appraisal and handle anomalous burdens on the present through a parallel decision process, as proposed by Kelleher (2025*b*). This preserves impartiality in the welfare calculation while allowing a separate fairness or feasibility check. The practical difficulty is institutional. If the parallel process is not reliable, a modest positive value of δ may act as a rough safeguard against some extreme implications. This approach is probably an imperfect safeguard for issues of fairness and sustainability though.

8.4.c. Expert and survey evidence

8.4.30. The disagreement in principle is reflected in expert evidence. Table 8.4 summarises several recent surveys. In Drupp et al. (2018), the mean recommendation for δ is approximately 1.1%, the median is 0.5% and the mode is zero. The later survey of philosophers by Nesje et al. (2023) finds a median of zero among philosophers. Yet, when philosophers choose a paired DICE calibration, the median pair contains a very low δ of 0.075% and $\mu = 1.25$. These results does not settle the ethical issues. They show, however, that a modest positive value lies near the centre of considered expert judgement while $\delta = 0$ has considerable support also. This is also the finding of our survey for the Review as shown in Table 8.4.

8.4.31. Our own survey points to the same central

Table 8.4.: Survey evidence on the pure rate of time preference

Study	Sample	Object	Main finding
D18	Econ.	Direct δ	Mean 1.1%, median 0.5%, mode 0%; 38% of responses in $[0, 0.1\%]$.
N23-E	Econ.	Direct δ	Mean 1.1%, median 0.5%.
N23-P	Phil.	Direct δ	Mean 0.9%, median 0%.
N23-M	Phil. pair	DICE pair	"Median philosopher view": $\delta = 0.075\%$, $\mu = 1.25$.
GVPZ23	Econ.	Discount rates	No direct δ. Mean 10-year risk-free rate 2.30%; mean climate rate 2.28%. Non-risk-adjusters: mean 10-year rate 2.53% and mean 100-year rate 2.26%.

Notes: D18 and N23 directly elicit the pure rate of time preference δ. GVPZ23 elicits discount rates and views on risk adjustment rather than δ itself. Abbreviations: D18 = Drupp et al. (2018); N23-E and N23-P = economists and philosophers in Nesje et al. (2023); N23-M = the median philosopher calibration in that study; GVPZ23 = Gollier et al. (2023).

value. The median for the whole sample is 0.5%, as is the median among government respondents and UK academics. The expert subgroup is more patient, with a median of 0.1%, while the overall distribution contains substantial support for both zero and positive values. The qualitative responses should be read alongside the numerical median, particularly where respondents distinguish intragenerational public appraisal from very long-run intergenerational problems.

8.4.32. Table 8.5 provides some selected quotes from our survey to illustrate that reasonable people may differ on this matter, and there are argument for what we see as two separate positions: $\delta = 0$ and $\delta > 0$.

8.4.33. Earlier UK evidence provides historical corroboration rather than a separate contemporary estimate. As noted above, Pearce & Ulph (1995) used a central pure-time-preference value of 0.3% and an upper value of 0.5%. The relevance here is that a small positive value of pure time preference has a long history in UK discounting guidance. This is in line with the survey evidence, but solid stated or revealed preference evidence on the social pure time preference is hard to come by.

8.4.d. Conclusion on δ

8.4.34. There is no uniquely estimable value of pure social time preference. The most direct normative argument favours $\delta = 0$: time itself does not identify a welfare-relevant difference between otherwise identical people or benefits. Nondogmatism strengthens the case for zero in the distant future if the impartial view is held by someone in society. Against this, agent-relative ethics, axiomatic arguments and the potentially severe burden placed on early generations by an indefinitely extended impartial objective provide reasons for a modest positive value. Market calibrations commonly imply values around 1–1.5%, but these are residuals from particular models and do not identify the public ethical parameter. Cross-dynastic altruism (Nesje 2024), isolation paradox arguments (Sen 1967) interactive population dynamics (Robson & Szentes 2014) provide further reasons to expect private market rates to differ from social rates of discount.

8.4.35. Our judgement is to retain $\delta = 0.5\%$. This is the median in the principal survey of economists, the median in our own survey and the current Green Book value. It is also low enough to limit temporal discrimination to a degree, while allowing an ordering of projects in principle. The recommendation rests on triangulation across normative argument, expert judgement, socially relevant evidence and the implications of alternative values. It also recognises that some of the chief concerns, e.g. sustainability, are better addressed by parallel decision criteria. Importantly, this estimate does not include individual mortality or societal extinction risk.

8.4.36. Reasonable disagreement remains, particularly for intergenerational horizons. However, we do not recommend retaining the current sensitivity analysis of $\delta = 0$. As we have argued, if the concern about $\delta > 0$ relates to sustainability then changing the pure time preference discount rate for CBA will not solve that issue. Separate measures are required. If the concern is that long-horizons should be treated differently, the risk free rate now contains a component of (minus) -0.5% to reflect any rare disaster risks that they may face; see Chapter 12. Furthermore, a declining term structure change reflects possible persistent growth states in the future. In short there are other, more directed ways to address some of the concerns often discussed in relation to pure time preference.

R.9. The pure rate of social time preference should remain $\delta = 0.5\%$

$\delta = 0.5\%$ should represent social pure time preference alone. Individual mortality risk and private experimental impatience should not be included. Catastrophic and systematic risks should be handled explicitly in the risk analysis.

Table 8.5.: Qualitative responses on δ / pure rate of time preference

Use	Lines / strength	Quote	Note
Pro $\delta = 0$	212–218; direct	"Based on the classical impartial Utilitarian philosophy, Pigou, Ramsey, Harrod, Stern, and others argue for a pure time preference of $\delta = 0\%$."	Most direct philosophical case for zero.
Pro $\delta = 0$	394; direct	"In the latter, I'd argue that our welfare is no more important than that of future generations, suggesting a pure time preference rate of 0 for intergenerational questions."	Strong intergenerational-equity justification.
Against $\delta = 0$	268–270; direct	"The value of 0.5% for the rate of pure time preference is at the median of the survey by Drupp et al. (Discounting Disentangled) so it is quite justifiable by this assessment."	Direct defence of a positive current value.
Numerical δ	626; direct	"I adopt an SRTP framework with a pure rate of time preference of 0.5% and an elasticity of marginal utility of 1."	Simple, explicit adoption of 0.5%.

8.5. The elasticity of marginal utility: μ

8.5.a. The interpretation of μ in the Green Book

8.5.1. The interpretation of μ follows from the same social-planner framework. In the discounted-utilitarian welfare function

$$W(\mathbf{c}) = \sum_{t=0}^{\infty} \beta^t U(c_t),$$

μ is the elasticity of the marginal social value of reference consumption:

$$\mu = -\frac{c_t U''(c_t)}{U'(c_t)}.$$

With isoelastic U, a one per cent increase in consumption reduces its marginal social value by μ per cent. In the Ramsey Rule, this curvature determines the wealth effect. If future per capita consumption is higher, a marginal unit of consumption is worth less in welfare terms and is discounted more heavily.

8.5.2. It is useful to distinguish two local elasticities. First, the curvature of the period social-utility index is

$$\mu_t \equiv -\frac{\partial \log U'(c_t)}{\partial \log c_t} = -\frac{c_t U''(c_t)}{U'(c_t)}.$$

For isoelastic period utility, $\mu_t = \mu$. This is the elasticity of marginal period utility with respect to consumption.

8.5.3. Second, in a more general intertemporal Bergson–Samuelson aggregator,

$$\mathcal{W} = B(U_0, U_1, \ldots), \qquad U_t = U(c_t),$$

there may also be curvature in the aggregator B over dated utility levels. Let

$$B_t \equiv \frac{\partial B}{\partial U_t}$$

denote the marginal social welfare weight on utility at date t. The welfare-level curvature can be written as the elasticity of this marginal weight with respect to consumption at date s:

$$\chi_{ts}^B \equiv -\frac{\partial \log B_t}{\partial \log c_s}.$$

8.5.4. Thus χ_{tt}^B measures the own-date curvature of the welfare aggregator, while χ_{ts}^B, for $s \neq t$, measures how consumption at one date changes the marginal social value of utility at another date. Equivalently, when one wants the elasticity directly in utility space,

$$\tilde{\chi}_{ts}^B \equiv -\frac{\partial \log B_t}{\partial \log U_s}.$$

8.5.5. The marginal social value of consumption at date t is

$$\lambda_t \equiv \frac{\partial \mathcal{W}}{\partial c_t} = B_t U'(c_t).$$

Therefore its own elasticity with respect to consumption is

$$-\frac{\partial \log \lambda_t}{\partial \log c_t} = \underbrace{\mu_t}_{\text{curvature of } U(c_t)} + \underbrace{\chi_{tt}^B}_{\text{curvature of } B \text{ over dated utilities}}.$$

For $s \neq t$,

$$-\frac{\partial \log \lambda_t}{\partial \log c_s} = \chi_{ts}^B.$$

8.5.6. Thus a general intertemporal Bergson–Samuelson welfare function can have two sources of curvature: the curvature of the period social-utility index, μ_t, and the curvature of the intertemporal welfare aggregator, χ_{ts}^B.

8.5.7. In the standard Ramsey discounted-utilitarian case,

$$\mathcal{W}_R = \sum_{t=0}^{\infty} \beta^t U(c_t),$$

so

$$B(U_0, U_1, \ldots) = \sum_{t=0}^{\infty} \beta^t U_t.$$

The marginal welfare weight is therefore

$$B_t = \beta^t,$$

which does not vary with consumption at any date. Hence

$$\chi^B_{ts} = 0 \qquad \text{for all } s, t.$$

It follows that

$$-\frac{\partial \log \lambda_t}{\partial \log c_t} = \mu_t.$$

8.5.8. In the Ramsey case, the relevant curvature is therefore the curvature of the period social-utility index $U(c_t)$, governed by μ, not an additional curvature of the welfare aggregator over dated utilities.

8.5.9. This distinction matters because μ should not be described uncritically as the preference parameter of a representative consumer. In this setting, μ is the curvature chosen for the social valuation of consumption along an intertemporal path. A more general Bergson–Samuelson welfare function could add a separate welfare-level curvature, χ^B_{ts}, reflecting how the social marginal value of utility at one date depends on utility or consumption at the same or other dates. This would be a Prioritarian approach (see e.g. Adler 2012). But discounted utilitarianism rules out that extra welfare-level curvature by making B linear in dated utilities.

8.5.10. One further clarification is necessary for interpersonal equity-fairness comparisons. For distributional appraisal within a period, let ϵ denote the elasticity used to construct welfare weights. With an isoelastic intratemporal social valuation, a group's marginal weight relative to the national median can be written as

$$a_i = \left(\frac{m_i}{m_N}\right)^{-\epsilon}, \tag{8.1}$$

where m_i is median consumption or income for group i and m_N is the national median. This is precisely how the Green Book proposes welfare weights be calculated for group i (HM Treasury 2026c). In principle, ϵ and μ are different parameters: the former concerns distribution across people at a date, while the latter concerns the marginal value of reference consumption across dates. We recognise this conceptual difference but recommend using the same value for each type of inequality to simplify application in the Green Book. We have further to say about this in Chapter 16.

8.5.b. Eligible evidence for μ

8.5.11. The principles in Section 8.2.b apply to μ as they did to δ, but the evidential balance is different. There is more empirical information on the curvature of utility and social valuation than there is on pure social time preference. The evidence nevertheless varies in how directly it identifies the parameter required by the Green Book.

8.5.12. The most direct positive evidence comes from social or political choices over the distribution of consumption, such as tax schedules, and from exercises in which respondents make explicitly social trade-offs, such as Okun's leaky bucket experiments. Subjective-wellbeing evidence is also relevant to the welfarist interpretation because it attempts to identify how experienced well-being changes with income. Evidence from household saving, demand for insurance and private risk taking is less direct. Such approaches identify individual substitution or risk preferences under parametric model of preferences and then typically use a one-parameter structure (e.g. isoelastic utility) to determine the curvature of utility for social evaluation. Such evidence can be used to test the coherence of the isoelastic specification of a social welfare function (Venmans & Groom 2021, Disque et al. 2026) and for triangulation. They are but one set of approaches for determining the appropriate value of μ for social discounting and in some cases inter-personal equity weighting.

8.5.13. Expert judgements are also eligible. The task involves functional-form restrictions, the interpretation of evidence across several domains and normative implications that are not readily elicited from the general public. Again we take a reflective view of the source of information and the implications, normative and positive, of different choices of the μ parameter.

8.5.14. Part of this reflection is to consider the views of our survey respondents. As there is a wide variety of viewpoints that we need to navigate. Table 8.6 shows a selection of those who prefer the status quo, those who want changes to μ, and those who have particular views on what μ should represent.

8.5.c. Positive evidence for μ

Inequality aversion and equal sacrifice

8.5.15. The equal-absolute-sacrifice approach infers social curvature from the progressivity of the tax schedule. Let pre-tax income be y and tax liability be $T(y)$. If the tax system is organised so that each taxpayer bears the same utility sacrifice, then

$$u(y) - u(y - T(y)) = k.$$

Table 8.6.: Qualitative responses on μ / elasticity of marginal utility

Position	Quote
Keep $\mu = 1$	"It is safe to keep the current assumption of logarithmic utility ($\mu = 1$)."
Keep $\mu = 1$ but frame it as inequality aversion	"Elasticity of 1 seems reasonable but I would cast this as inequality aversion for social choices and appeal to that evidence rather than individual preference or some supposed individual utility function."
Keep $\mu = 1$ / broadly acceptable	"Inequality aversion or consumption smoothing motive at 1 also seems to be in line with what people prefer from my reading of the evidence."
Keep $\mu = 1$ despite uncertainty	"The elasticity of marginal utility of income has always been more or less more supposition than rooted in sound evidence, we know more about the relationship now and that cross-sectional evidence differs from longitudinal impacts of general increases in income. But it remains a rather arbitrary number, and changing it significantly would make the Ramsey formula either more cumbersome to apply or dictate unrealistic parameters for the time preference and growth elements. So leave it at one."
Increase μ above 1	"I would set the elasticity of marginal utility of consumption higher than 1 and the expected average annual real per-capita consumption growth rate lower than 2%."
Increase μ above 1	"The elasticity of marginal utility of consumption should be revised higher."
Increase μ above 1	"The elasticity of consumption is too low and inconsistent with the latest empirical evidence and its use in other parts of the Green Book."
Increase μ to 1.5	"g is more like 1%. elasticity of MU income more like 1.5."
Increase μ to 1.5, citing evidence	"From Groom and Maddison (2018) it seems difficult to argue that the elasticity is anything other than 1.5, rather than 1. .... per capita consumption growth rate has been on a clear downward trend and should be revised downwards"
Increase μ modestly above 1	"I think there is evidence that the MU of income elasticity is somewhat above unity. Overall, I think those two things roughly balance out, so the STPR should remain at around its current level."
μ may be too low	"I think the pure rate of time preference and μ are a bit low. Whereas real income growth at the median, which I prefer to the mean for social questions, is below 1% for the UK."
Lower μ below 1	"The elasticity of marginal welfare wrt consumption might be a tad high at 1, but it's in the right ballpark. Eg, something closer to 0.8 is probably more accurate."

Differentiating and imposing isoelastic utility gives

$$\mu = \frac{\ln(1 - T'(y))}{\ln(1 - T(y)/y)},$$

where $T'(y)$ is the marginal tax rate and $T(y)/y$ is the average tax rate. The approach has a long history in the UK, including Stern (1977), Evans & Sezer (2002) and Groom & Maddison (2019).

8.5.16. This evidence has an attractive social interpretation. Tax schedules are collectively produced, persist through political processes and reveal more than an isolated private choice. Groom & Maddison (2019) obtain a preferred contemporary equal-sacrifice estimate of 1.515 and a long-run historical estimate around 1.57 for the period 1949 - 2010. These estimates are close despite the very different data used.

8.5.17. The interpretation is not literal. Actual tax schedules reflect revenue requirements, labour-supply incentives, avoidance, administrative constraints, transfers and political bargaining, as well as views about equal sacrifice. Moreover, the equal-sacrifice hypothesis cannot be tested independently of the utility function imposed. Tax evidence should therefore be treated as a socially salient benchmark rather than as a clean revelation of a universal moral parameter (Del Campo et al. 2024, Groom & Maddison 2019).

The elasticity of intertemporal substitution

8.5.18. Under time-additive isoelastic utility, the Euler equation links consumption growth to the real interest rate. In a simple deterministic representation the Ramsey Rule can be rearranged as follows,

$$\Delta \ln c_{t+1} \simeq \frac{1}{\mu}(r_t - \delta), \qquad (8.2)$$

so $1/\mu$ is the elasticity of intertemporal substitution. A low willingness to shift consumption between dates implies a high μ. Using UK aggregate consumption and interest-rate data, Groom & Maddison (2019) estimate μ at approximately 1.6.

8.5.19. This estimate is informative about the curvature needed to describe aggregate consumption dynamics under the assumed model. It is less direct evidence about social valuation. Aggregate consumption reflects heterogeneous households, borrowing constraints, demographics, taxes and expectations; the Euler equation also imposes time-additive expected utility, which if it is untrue affects the estimates of μ. The estimate is therefore a useful consistency check on the isoelastic form rather than a decisive social parameter.

Risk aversion and insurance

8.5.20. Under expected utility, the same μ is the coefficient of relative risk aversion. Groom & Maddison (2019) infer this parameter from UK wealth and non-health, non-life insurance data. For a given loading factor, greater utility curvature increases the fraction of wealth insured against loss. Their estimate is approximately 2.19.

8.5.21. The risk evidence points towards curvature above one, but it is among the less direct sources for present purposes. Insurance demand is affected by background risk, liquidity, contract design, default, information and behavioural biases. More fundamentally, the social aversion to unequal consumption across dates need not equal private risk aversion. The estimate is best viewed as a test of whether the one-parameter model is grossly inconsistent across domains.

Subjective well-being

8.5.22. Subjective-wellbeing evidence is closer to the literal idea of diminishing marginal utility. If reported life satisfaction is treated as an observable transformation of latent utility, the slope and curvature of the income–well-being relationship can be used to estimate μ. Layard et al. (2008) obtain an estimate of approximately 1.32 for the British Household Panel Survey, with a confidence interval that includes one. The approach is attractive because it speaks directly to how well-being changes with income. It requires strong assumptions about interpersonal comparability, reporting scales, adaptation and omitted determinants of well-being.

Synthesis of the UK evidence

8.5.23. The important contribution of Groom & Maddison (2019) is not that the distinct concepts are theoretically identical. It is that the available UK estimates are not statistically irreconcilable. Their fixed-effects pooled estimate is 1.528, with a 95% confidence interval of approximately 1.44–1.61; the random-effects estimate is 1.594, with a wider interval of approximately 1.36–1.83. Parameter homogeneity across the methods cannot be rejected. Omitting the insurance estimate leaves a pooled estimate of approximately 1.51. Table 8.7 summarises the evidence and its relationship to the social-planner parameter.

8.5.24. Taken at face value, the UK revealed-preference evidence favours a value close to 1.5 and rejects one in the pooled analysis. The force of that conclusion depends on accepting the one-parameter interpretation and giving private risk and intertemporal behaviour the same standing as social distributional choices. Under the principles adopted in this Review, the result is strong evidence that the current value of one is at the low end of the plausible range, but not a requirement to adopt 1.5 exactly.

Table 8.7.: Selected evidence on the elasticity of marginal utility

Source or method	Estimate	Interpretation for Green Book calibration
Equal sacrifice, contemporary UK tax schedule	1.515	Relatively direct evidence from a socially produced distributional choice, conditional on equal sacrifice.
Equal sacrifice, historical UK tax schedules	$\simeq$ 1.57	Long-run political benchmark; affected by changes in tax institutions and constraints.
Euler equation	1.584	Private/aggregate intertemporal substitution under additive isoelastic utility; useful for triangulation.
Insurance demand	2.19	Private risk aversion under expected utility; indirect for social intertemporal valuation.
Subjective well-being	1.32	Evidence on diminishing marginal well-being from income, subject to comparability assumptions.
Groom–Maddison pooled estimate	1.53	Empirical synthesis; supports $\mu > 1$ but retains the one-parameter restriction.
Drupp et al. expert survey	Mean 1.35; median 1	Considered expert judgement rather than behavioural identification.
Our survey	Median 1; expert median 1.4	Direct recommendation for public appraisal; qualitative responses often favour an increase.

8.5.d. Normative and expert evidence

8.5.25. The normative implications of μ are easier to inspect than those of δ. Suppose one person or generation has twice the consumption of another. With isoelastic utility, the ratio of the marginal social value of an additional unit to the richer and poorer recipient is

$$\frac{U'(2c)}{U'(c)} = 2^{-\mu}. \tag{8.3}$$

For $\mu = 1$, an additional pound to the richer recipient is worth one half as much at the margin. For $\mu = 1.25$, the ratio is approximately 0.42; for $\mu = 1.5$, it is approximately 0.35; and for $\mu = 2$, it is 0.25. Equivalently, in an Okun-style leaky-bucket transfer from the richer to the poorer recipient, the maximum loss consistent with a marginal welfare improvement is approximately 50%, 58%, 65% and 75%, respectively. These implications provide a transparent way to test candidate values against considered views about redistribution.

8.5.26. Expert evidence lies between the current Green Book value and the UK pooled estimate. Drupp et al. (2018) reports a mean recommendation of approximately 1.35 and a median and mode of one. The paired choices of philosophers in Nesje et al. (2023) have a median μ of 1.25. In our survey, the whole-sample median is one, but the expert subgroup median is 1.4, and the qualitative responses contain repeated arguments for raising the parameter. This pattern is consistent with a central value between one and 1.5, rather than a sharp consensus at either endpoint.

8.5.27. International and sector-specific estimates span a wider range, in part because they identify different concepts. Estimates based on domestic tax schedules commonly lie between one and two, while estimates based on international aid or donor–recipient decisions can be below one. Experimental risk and inequality estimates can be considerably higher. This variation shows how context matters for the elicitation of inequality aversion parameters, so no method is the definitive estimate for social decision-making (Del Campo et al. 2024). For UK government appraisal, UK social choices and informed public recommendations should receive the greatest weight, with private and international evidence used as checks.

8.5.e. Conclusion on μ

8.5.28. The evidence for μ is more substantial than the evidence for δ, but it does not identify one uncontested number. The strongest UK empirical synthesis favours approximately 1.5. Expert recommendations and our own survey centre closer to one, while the philosopher calibration and several forms of normative reasoning support an intermediate value. The interpretation adopted in this chapter gives greatest weight to evidence from social choices, subjective well-being and informed social judgement, while treating private EIS and risk-aversion estimates as triangulation.

8.5.29. On this basis, we recommend $\mu = 1.25$. This is a deliberate social calibration rather than the point estimate from any single method. It raises the elasticity above the current value of one, reflecting the UK evidence that marginal social value declines somewhat more rapidly with consumption. At the same time, it stops short of imposing the full pooled revealed-preference estimate of 1.5, since several of the inputs identify private risk or substitution behaviour and rely on the restrictive equality of preferences across domains. The value also coincides with the median paired calibration of philosophers and lies close to the centre of the expert evidence.

8.5.30. Figure 8.1 shows where this recommendation falls among those found in the literature. Appendix B provides further evidence by concept: risk, inequality and fluctuation aversion.

> R.10. The elasticity of marginal utility of consumption should be $\mu = 1.25$
>
> $\mu = 1.25$ should be used in the Ramsey Rule. It represents a moderate increase in the curvature of social valuation relative to current guidance and is supported by triangulation across UK distributional choices, subjective well-being, expert judgement and the broader empirical evidence. Private risk and intertemporal-substitution estimates inform, but do not determine, the value.

8.5.31. Reflecting on this value, we note that Buchholz & Schumacher (2010) provide some normative reasons for moving away from the current value of $\mu = 1$. Starting with the equilibrium/optimal interpretation of the Ramsey Rule, where r is the rate of return to capital, $r = \delta + \mu g$. They point out, like Dasgupta (2008), that there is a substitutability between δ, and the elasticity of marginal utility, μ. They argue that in the Ramsey Framwork, sustainable non-decreasing consumption paths can be selected by an appropriate choice of μ. With regard to the level of consumption, they also note that productivity shocks are only shared among all generations (what they call 'circumstance solidarity') if and only if $\mu > 1$. Values below one induce larger proportional saving for the earlier generation, leaving them worse off relatively speaking. This is true even when $\delta > 0$. An alternative nor-

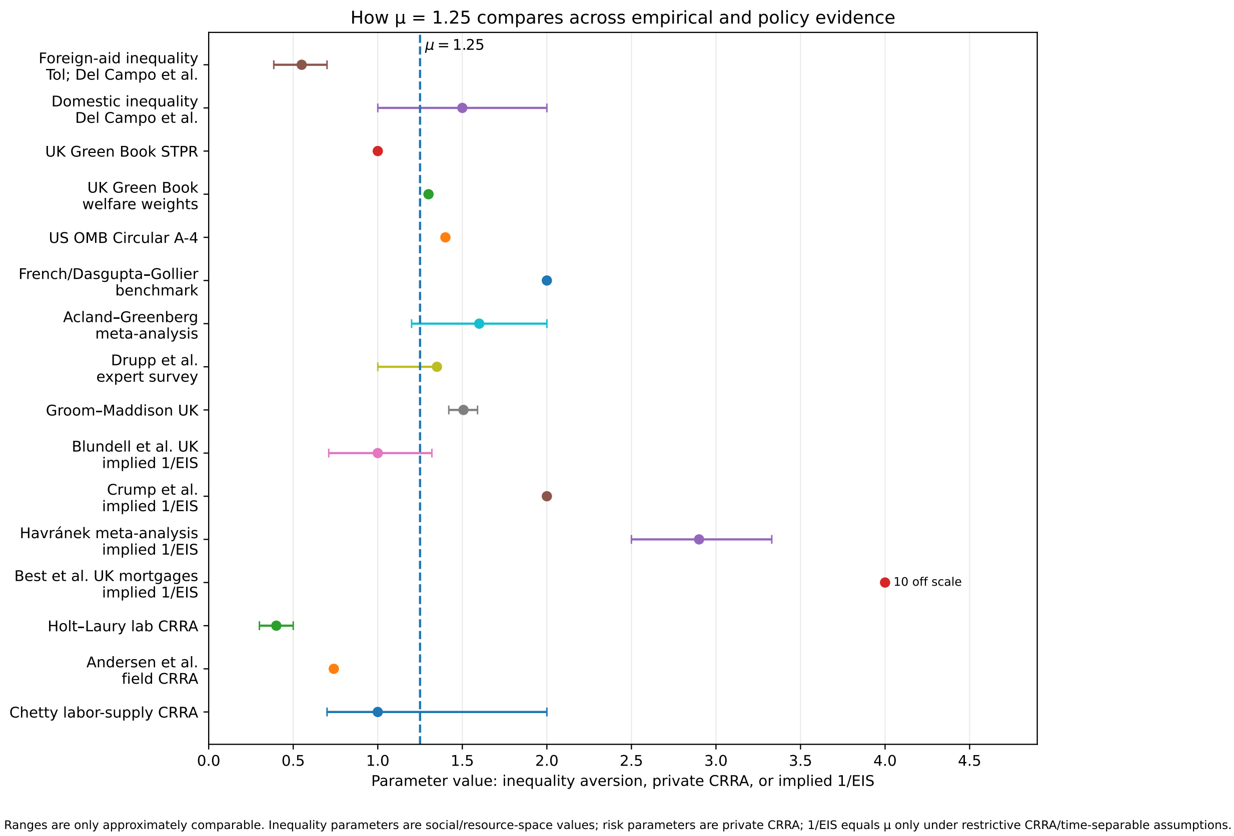

Figure 8.1.: Evidence on the magnitude of μ from studies of inequality aversion, risk aversion, intertemporal substitution and subjective well-being estimates. The proposed value of $\mu = 1.25$ is shown by the vertical dotted line.

mative criterion is the ‘no envy’ criterion. Absolute productivity-adjusted no-envy supports $\mu = 1$, while relative no-envy, where generations compare their relative status after adjusting for productivity, supports $\mu = 2$. Our choice satisfies some of these normative principles even for $\delta = 0.5\%$.

8.5.32. On the inter-personal and inter-temporal distinction, there is also a purely practical case for harmonising the curvature used in discounting with that used for distributional welfare weights with that used for discounting. The two parameters are conceptually distinct, but maintaining different values creates avoidable complexity and can produce an unexplained discontinuity between valuation across dates and valuation across people. A common value of 1.25 is close to the current welfare-weighting evidence and would make the ethical content of Green Book appraisal more transparent. The harmonisation should be described as a simplifying public convention, not as an empirical finding that intertemporal and intratemporal inequality aversion are identical.

R.11. Welfare weights should be harmonised at $\epsilon = \mu = 1.25$

Although ϵ and μ represent different concepts, Green Book welfare weights should use $\epsilon = 1.25$ alongside $\mu = 1.25$ in the Ramsey Rule. Appraisals should report the conventionally unweighted NPSV and, where distributional analysis is material, a welfare-weighted NPSV using $\epsilon = 1.25$.

8.6. Expected growth, g

8.6.1. We now turn to the final parameter in the Ramsey Rule, g. Conceptually this is a much simpler parameter value to estimate than δ and μ since it is entirely positive in nature; with the benefit of hindsight we can observe the true value. As a consequence, many of the ethical issues raised in previous sections are not relevant in this context.

8.6.2. When we considered the SOC risk-free rate in Section 7.2, we recommended that HM Treasury does not explicitly estimate this value itself but instead relies on an external source; in that case, the real yield curve from the Bank of England. For growth, we make the same recommendation. There are a number of external sources that give long-term macroeconomic growth forecasts, and these provide a strong foundation for estimating g. We can also look at historic growth rates achieved in recent years by the UK economy, which is closer to the current Green Book approach.

8.6.3. We start by presenting a variety of recent long-term forecasts for UK real GDP growth. There are two caveats to make around these numbers. First, these figures are not (generally) given on a per capita basis. Second, the core macro-variable of relevance to the Ramsey Rule is real per capita consumption, not GDP, growth. Nevertheless, these forecasts are informative for thinking about the far-horizon term structure of the social discount rate as well as the value of g that should be applied for the short-to-medium term.

8.6.4. In Figure 8.2, we present long-range real GDP growth forecasts for the UK from the OECD global long-run economic scenarios: 2025 update

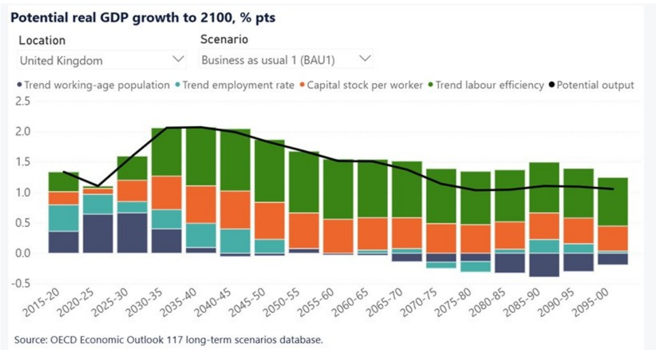

Figure 8.2.: This figure presents the OECD's 2025 forecast of long-run real GDP growth for the UK based on the Business as Usual 1 model.

8.6.5. Long-range real GDP growth forecasts for the UK are also given in Goldman Sachs' Global Economy Outlook 2075. This gives projections of 2.0% for 2024–2029, declining to 1.9% over the next decade, 1.6% in the 2040s and then reducing to 1.2% for the 2070s.

8.6.6. The Centre for Business and Economic Research World Economic League Tables 2026 predicts that, compared to 2025, real UK GDP will be 8% higher in 2030 and 27% higher in 2040, implying fairly steady annual growth of 1.5-1.6% over the next fifteen years.

8.6.7. In Figure 8.3 we present the long-run assumptions that the Office for Budget Responsibility's Long-term economic determinants – March 2025 Economic and fiscal outlook. This predicts annualised real growth of 1.8% to 2030, 2.1% in the 2030s, 1.8% in the 2040s, 1.6% in the 2050s, and 1.5% in the 2060.

8.6.8. But these figures are all on a total, not per capita, basis. The Office for National Statistics presents population forecasts for the UK to adjust for this. This predicts that population will grow from 69.3 million in mid-2024 to 71.0 million in mid-2034, an annualised growth rate of 0.24%. Between then and 2054, the population is expected to grow by 0.1% per annum before declining thereafter.

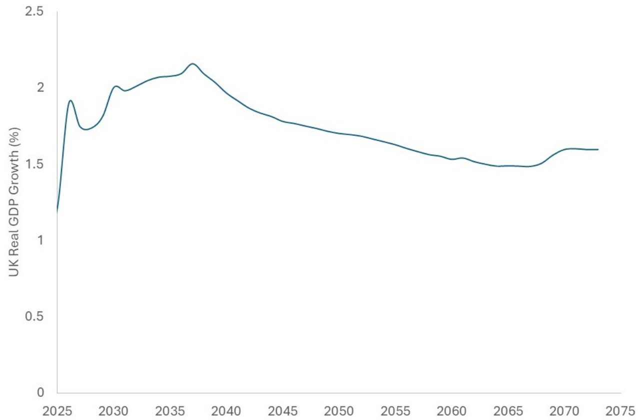

Figure 8.3.: This figure presents the long-run real GDP growth rate assumptions in the Office for Budget Responsibility's Long-term economic determinants – March 2025 Economic and fiscal outlook

8.6.9. Therefore, in terms of the term structure of consumption growth per capita, there are two offsetting effects. Absolute growth is forecast to be stronger in the early years, but population growth is also faster.

8.6.10. The Office for Budget Responsibility's Long-term economic determinants – March 2025 Economic and fiscal outlook gives some further information on this. On a per capita basis, it forecasts real GDP growth of 1.4% over the next five years, 1.7% over the 2030s, and 1.4% thereafter. This is a population adjustment of 0.4% in the early years, declining to 0.15% by the 2060s.

8.6.11. Overall, our opinion is that a forecast of 1.75% for real GDP growth over the 0–30 year period, with a downward adjustment of 0.25% to account for population growth, gives a solid basis for an estimate of real per capita consumption growth in the short-to-medium term of $g = 1.5\%$. This is also the median value given by non-governmental respondents to our survey as reported in Table 8.2. While those from government and social media argued for a lower growth rate, we believe this is not supported by the balance of evidence presented above.

8.6.12. That respondents to our survey overwhelmingly view the current value of $g = 2\%$ as too high is very strongly evidenced in the qualitative comments reported in paragraph A.2.3 in the appendix.

8.6.13. There is also evidence, both in the external forecasts of GDP growth and population, together with the survey responses in Table 8.2, that growth is expected to decline in the longer term. We will return to this in Chapter 9 when we discuss declining discount rates.

8.6.14. We stress, though, that HM Treasury should keep these growth rate forecasts under regular review, together with the real yield curve. Any significant changes should be cause for an early reassessment of the Ramsey Rule parameters, and particularly g.

8.6.15. An alternative approach, more in keeping with current Green Book practice, is to calibrate from historic UK real per capita consumption growth. In Figure 8.8, we present trends in g between 1831 and 2025 in 5-year rolling windows. According to these data, the period from 1976 to the present day (the past 50 years) experienced growth of 1.84%. This is an average of 2.9% pre-2000 and 0.8% between 2001 and the present day. The post-Brexit/Covid years have been the worst in recent times for consumption growth, with this declining to 0.32%. This can also be seen in Figure 8.4.

8.6.16. These per capita consumption growth rates can be compared to the GDP growth rates. To be consistent with the 1.5% annual growth of per capita consumption recommended above, it is necessary to select the period 1988-2025, the past 38 years, which encompasses periods of rapid growth in the late 1980s and 2000s bookended by deep recessions in the early 1990s, 2008, and in the post-Brexit and Covid-19 pandemic period.

R.12. per capita real consumption growth should be 1.5%.

We recommend that, for the parameterisation of the Ramsey Rule, $g = 1.5\%$. This reflects a growth rate in real consumption of 1.75% and a downward adjustment of 0.25% to account for population growth over the next 30 years.

8.7. Concluding remarks

8.7.1. The choice of the Ramsey parameters follows from the interpretation of the Green Book social welfare function. CBA evaluates marginal deviations from a reference path of mean real per capita consumption using a discounted-utilitarian social-planner rule. The parameters are therefore public social-valuation parameters, not the private preferences of a literal representative agent.

8.7.2. Both normative reasoning and positive evidence are eligible, but they must identify the relevant social object. Private hyperbolic discounting and individual mortality risk do not. Market returns, aggregate saving, tax schedules, subjective well-being, socially framed experiments and expert judgements can all provide information, but none is mechanically decisive. Parameter choice should instead proceed through reflective equilibrium and triangulation: ethical principles should be tested against their implications, empirical evidence should be examined for contextual relevance, and the resulting calibration should be checked for consistency with observable social and market outcomes.

8.7.3. Applying these principles, we recommend retaining a pure rate of social time preference of $\delta = 0.5\%$, and increasing the elasticity of marginal utility to $\mu = 1.25$. These values do not eliminate reasonable disagreement. They provide a transparent and practicable calibration that lies near the centre of considered expert judgement, is consistent with the strongest parts of the UK evidence, and avoids treating private impatience, mortality or any single market model as the social ethic of government appraisal.

8.7.4. By contrast, there are a range of sources that are consistent with a recommendation of $g = 1.5\%$. While there may be some reasonable disagreement over the the precise numerical value that we have recommended, the conceptual basis for setting this parameter value is much less contentious than for δ and μ.

Table 8.8.: Growth statistics by historical period, GB+NI, 1831–2025.

Period	Years	g_C mean	g_N mean	g_{pc} mean	g_C CAGR	g_N CAGR	g_{pc} CAGR
2016–2025 (last 10 years)	10	0.973	0.656	0.317	0.977	0.658	0.317
2010–2025 (post-financial crisis)	16	1.348	0.688	0.660	1.357	0.690	0.662
2001–2025 (last 25 years)	25	1.451	0.663	0.788	1.462	0.665	0.792
1988–2025 (last 38 years)	38	2.052	0.531	1.521	2.073	0.532	1.533
1976–2025 (last 50 years)	50	2.244	0.424	1.820	2.269	0.425	1.837
1976–2000	25	3.037	0.185	2.852	3.083	0.185	2.893
1926–1975	50	2.021	0.442	1.579	2.041	0.443	1.592
1876–1925	50	1.277	0.890	0.387	1.285	0.894	0.388
1831–1875	45	2.009	1.090	0.919	2.029	1.095	0.924

Notes: Growth rates are annual log growth rates in percent. The underlying consumption and population data are from the Bank of England's *A millennium of macroeconomic data for the UK* research dataset, extended here to the GB+NI series through 2025. g_C denotes real aggregate consumption growth, g_N denotes population growth, and $g_{pc} = g_C - g_N$ denotes real per-capita consumption growth. CAGR denotes the compound annual growth rate computed from the level of the relevant series over the period. The first row reports the last 50 years as a whole; the next two rows split that same interval into two 25-year sub-periods.

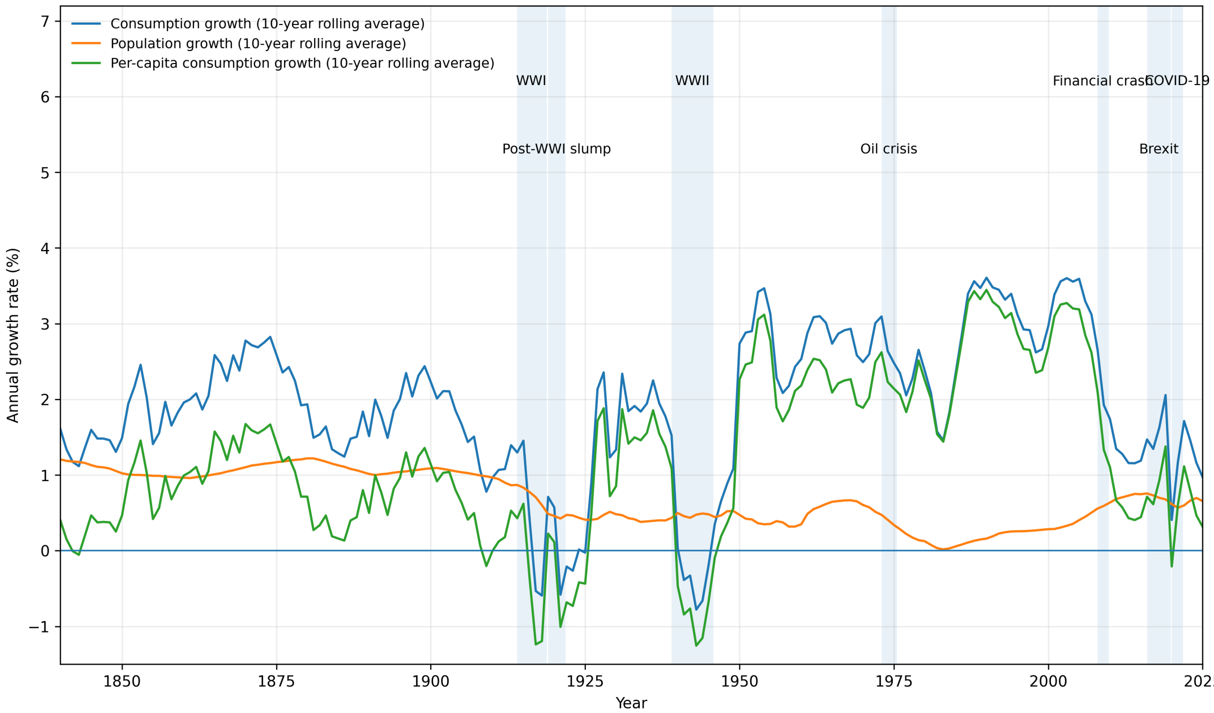

Figure 8.4.: Ten-year rolling average growth rates of aggregate consumption, population and per capita consumption, GB+NI, 1831–2025.

Notes: Growth rates are annual log growth rates in percent. g_C denotes real aggregate consumption growth, g_N denotes population growth, and $g_{pc} = g_C - g_N$ denotes real per capita consumption growth. The plotted series are trailing ten-year rolling averages, so the value shown in year t is the average annual growth rate over years $t - 9$ to t. The underlying consumption and population data are from the Bank of England's *A millennium of macroeconomic data for the UK* research dataset, extended here to the GB+NI series through 2025. Shaded bands mark selected historical episodes and are included for orientation only.

9. Risk-free Rates under Uncertainty

9.1. Introduction

9.1.1. In Chapter 7 we considered the social discount rate in a world of perfect foresight about the future of the economy and all public projects within that economy.

9.1.2. In this chapter, we extend the analysis to consider stochasticity at the macroeconomic level. In terms of the analysis of Chapter 5, macroeconomic risk relates to π_t, which is the same for all projects, while microeconomic risk relates to b_{it}, the benefit from a given project. It is crucial not to confuse these two types of risk. We will return to project-specific, microeconomic, risk in Chapters 10, 11 and 12.

9.1.3. The effects of uncertainty around π_t on the risk-free discount rate will be referred to 'prudence', caused by a 'precautionary savings motive'. We will define these terms more formally below. The effects of risk around b_{it} on the risk-adjusted discount rate will be referred to as 'risk-aversion'. This is a separate economic concept.

9.1.4. A longer and very clear discussion of many of the issues covered in this and subsequent chapters is given in Christian Gollier's book "Pricing the Planet's Future" (Gollier 2013).

9.2. Prudence

9.2.1. To understand prudence, return to the equation for the risk-free rate from paragraph 7.1.7, and assuming that current consumption, c_0, is known even in a stochastic environment:

9.2.2. Macroeconomic risk in this environment refers to the fact that c_t is unknown. We introduce the notation $c_t = E[c_t] + \epsilon_{ct}$, where ϵ_{ct} is the forecast error that has zero expectation $E[\epsilon_{ct}] = 0$. Then taking a second-order Taylor's series approximation of the numerator in the large brackets of previous equation:

$$r_{ft} = -\frac{1}{t} \ln \left(\frac{U'(E[c_t], t) + 0.5\text{Var}[\epsilon_{ct}]U'''(E[c_t], t)}{U'(c_0, 0)} \right)$$

9.2.3. We now have an additional term that depends on the third derivative of the welfare function: this is known as the 'prudence' term following Kimball (1990*a*,*b*).

9.2.4. For quadratic welfare — which is sometimes used to justify the Capital Asset Pricing Model which we return to in Chapter 11 — the third derivative is zero and therefore, however uncertain we are about future consumption growth, this has no impact on the risk-free rate. More usually, we assume that the third derivative of the welfare function is positive and that therefore risk about future consumption levels reduces the risk-free rate compared to an economy with no uncertainty.

9.2.5. Intuitively, this is relatively clear. If our future has either a very wealthy or very poor outlook, then we would increase our savings level today to protect against the worst outcomes. This is a consequence of decreasing marginal utility — we lose more from large consumption drops than we would benefit from consumption rises of the same magnitude. That greater risk raises savings levels is known as the 'precautionary savings motive'. Higher savings levels then result in lower interest rates.

9.2.a. Normally distributed risk

9.2.6. To progress from here, it is common to assume that the probability density function of logarithmic consumption growth is normally distributed: $g_t \sim N(m_{gt}, s_{gt}^2/t)$.

9.2.7. With isoelastic utility, $U'(c_t, t) = \exp(-\delta t)c_t^{-\mu}$, then:

$$r_{ft} = \delta - \frac{1}{t} \ln \left(E[\exp(-\mu g_t t)] \right)$$

9.2.8. In this case, $-\mu g_t t \sim N(-\mu m_{gt} t, \mu^2 s_{gt}^2 t)$. The expression for the risk-free rate in the previous offset equation now includes the moment generating function of a normal distribution which has a well-known analytical form. Using this expression gives:

$$r_{ft} = \delta + \mu m_{gt} - 0.5\mu^2 s_{gt}^2$$

9.2.9. The final term in this expression is the adjustment for prudence and the precautionary savings motive. Because the third derivative of isoelastic utility is positive, this effect reduces the risk-free rate as described above. This expression, and slight variations of it derived under mildly different assumptions, are known as the *Extended Ramsey Rule*.

9.2.10. In practice, consumption volatility is low. Taking a long-term proxy measure for consumption, we estimate $m_{gt} = 2.4\%$ and $s_{gt} = 2.8\%$.[1]. This means that, for $\mu = 2$, double the current Green Book value, the prudence adjustment is only 0.16%, which is generally considered to be too small to be practically relevant. For this reason, while HM Treasury has long been aware of the Extended Ramsey Rule, it has decided previously not to adjust the SDR for any precautionary savings motive except in relation to declining discount rates.

9.3. Declining Discount Rates

9.3.1. We are now in a position to consider the declining schedule of discount rates given in the current version of the Green Book (HM Treasury 2026*a*, Table 3.A) and its extensions to other potential environments. In a *Policy Forum* article in *Science* in 2013, Nobel Laureates Ken Arrow and William Nordhaus, together with many other eminent co-authors, argued that 'there are compelling arguments for using a declining discount rate schedule' (Arrow et al. 2013, p.350); see also Cropper et al. (2014) for an overview of this literature.

9.3.2. HM Treasury was pathbreaking in international social discounting guidance when it introduced declining discount rates in the 2003 version of the Green Book. Following this lead, such terms structures have been widely incorporated in social discount rates across a range of developed countries, including the US under the Biden administration (but not the second Trump administration) — see Groom et al. (2022).

9.3.3. In Table 9.1 we report the percentage of respondents to our survey who recommended that HM Treasury should vary the discount rate with the maturity of the project. There is very large agreement between all respondents that the term structure of discount rates should not be flat, consistent with current Green Book practice; see also the qualitative survey responses in paragraph A.2.5. However, there was some limited dissent, including (see paragraph A.2.3) "Longer term discount rates should be higher, driven by increasing time preference rate and the elasticity of marginal utility of consumption that also should increase".

[1]Here we have used the Office for National Statistics series ABJR: "Household final consumption expenditure: National concept CVM SA - £m" which includes durables and is not per capita, but does allow us to estimate means and variances on annualised data from 1948–2024.

9.3.a. Prudence and declining discount rates

9.3.4. Fundamental to this framework is the observation that $r_{ft} = -1/t \ln(E[\exp(-x_t t)])$ for some random variable x_t. Declining discount rates come from uncertainty over the future value of x_t as, under Jensen's inequality, $E[\exp(-x_t t)] \geq \exp(E[-x_t t])$ because exponential functions are convex. Greater uncertainty over x_t and greater convexity of the exponential function (higher t) both increase the magnitude of this inequality.

9.3.5. There are a range of interpretations that can be placed on x_t in order to derive a declining term structure of risk-free discount rates:

- Within the context of the Extended Ramsey Rule, $x_t = \delta + \mu g_t$, There is uncertainty here over future economic growth and, as shown by, for example, Gollier (2002), this uncertainty can lead to a declining term structure of discount rates.
- Within an SOC context, x_t reflects future risk-free interest rates on Treasury securities, which are also not known at the current time. This approach was pioneered by Weitzman (1998).
- In the setting of Gamma Discounting (Weitzman 2001), the distribution of x_t reflects the extent of expert disagreement on what the $t-$period risk-free discount rate should be. The expectations operator is then not a sum over values multiplied by probabilities, but instead a sum of weighted expert opinions. In Gamma Discounting, all opinions were given equal weight, but this weighting choice has subsequently been questioned (Jouini & Napp 2014, Heal & Millner 2014*b*, Freeman & Groom 2015).

9.3.6. To provide a simple illustration of declining discount rates, first assume that $x_t = 3\%$ or 1% with equal probability for all t. Then, by simply applying the formula $r_{ft} = -1/t \ln(E[\exp(-x_t t)])$, it is straighforward to verify that r_{ft} declines from 2% for very short horizons, to 1.95% at $t = 10$, 1.76% at $t = 50$, and 1.57% at t=100. At a horizon of 1000 years, $r_{ft} = 1.07\%$ and, as shown by Weitzman (1998), as $t \to \infty$, so $r_{ft} \to 1\%$, the lowest possible value of x_t. This asymptotic result is true for all possible well-defined distributions for x_t not just the simple two-state example given here.

9.3.7. Now suppose that $x_t t$ is the sum of t random variables $x_t t = \xi_1 + \xi_2 + \ldots + \xi_t$ and also assume that ξ_t has a constant mean, m_ξ and variance s_ξ^2 for all t. Then $x_t t$ has mean $t m_\xi$ and variance $1' \Sigma_\xi 1$, where 1 is the $t-$vector of 1s and Σ_ξ is the $t \times t$ variance-covariance matrix of the ξ_is.

	Count	Yes	No	Unsure
UK Academic	25	72%	12%	16%
Social Media	19	89%	5%	5%
SBCA	12	83%	17%	0%
Experts	12	83%	8%	8%
Total (non-government)	68	81%	10%	9%
Government	12	75%	8%	17%
Total (whole sample)	80	80%	10%	10%

Table 9.1.: Should the discount rate vary by project maturity?

9.3.8. If the ξ_is are all independent, then Σ_ξ is a matrix with s_ξ^2 on the diagonals and zeros elsewhere. In this case, $1'\Sigma_\xi 1 = s_\xi^2 t$.

9.3.9. If the ξ_is are, instead, perfectly correlated, then all elements of Σ_ξ equal s_ξ^2 and $1'\Sigma_\xi 1 = s_\xi^2 t^2$. In contrast to the previous paragraph, the variance of $x_t t$ is now proportional to t^2 and not t.

9.3.10. This distinction is critical for understanding declining discount rates. In the former case, where the ξ_is are independent, there is generally not a declining term structure of discount rates driven by the precautionary savings motive. In the latter case, where the ξ_is are perfectly correlated, then we get a strongly declining term structure of discount rates that tends to the lowest possible value as $t \to \infty$. Partial correlation lies in between these two cases.

9.3.11. In the case of the Extended Ramsey Rule, $g_t t = \ln(c_t/c_0) = \ln(c_1/c_0) + \ln(c_2/c_1) + \ldots + \ln(c_t/c_{t-1})$. The $t-$period consumption growth rate is just the sum of each individual period's growth rate. Therefore, it is the autocorrelation of real consumption growth that, in an STPR framework, influences prudence-driven declining discount rates.

9.3.12. Notice that, in paragraph 9.2.6 we defined the variance of g_t by s_{gt}^2/t. If $s_{gt} = s_g$ for all t, this leads to a linear relation between $\mathrm{Var}(g_t t)$ and t: the case of zero autocorrelation. In this case, the Extended Ramsey Rule becomes $r_{ft} = \delta + \mu m_{gt} - 0.5\mu^2 s_g^2$, and there is no prudence-driven declining discount rates.

9.3.13. If, instead, consumption growth were perfectly autocorrelated instead, then $r_{ft} = \delta + \mu m_{gt} - 0.5\mu^2 s_g^2 t$ and the prudence term increases linearly with time, and $r_{ft} \to -\infty$ as $t \to \infty$.

9.3.14. The core question then, is an empirical one. To what extent is the real annual consumption growth rate in the UK autocorrelated? Using the same annual consumption data for the UK as above from 1948–2024, the estimate autocorrelation is slightly positive at 0.12, although the $t-$statistic is not statistically significant ($p-$value greater than 0.3). On first inspection, this would appear to imply that the term structure should be much flatter than that recommended in the Green Book. We will calibrate this model much more carefully in Chapter 13.

9.3.15. What is of note is that, despite its STPR approach, this is not the way that HM Treasury estimated its current term structure of discount rates. Instead, it frames x_t in an SOC, interest rate way. As HM Treasury (2003) makes clear, the schedule of discount rates is based on Oxera (2002), which states:

> "This report concludes that the results from Newell and Pizer's [Newell & Pizer (2003)] modelling of interest rates could be used as the basis for policy guidance on a social time preference rate for the UK. Although the results were based on US interest rates rather than UK rates, and although the report indicates that some future refinements to the method are justified, Newell and Pizer's data offers an empirically based path for the discount rate. The results are shown in the figure below—a discount rate that declines to 1.0% in the long term. The discount rate has been set to start at 3.5%, in line with the draft Treasury Green Book guidance." (Executive Summary).

9.3.16. Newell & Pizer (2003) base their estimate of declining discount rates on US interest rate data from 1798 — 1999. That the term structure of discount rates in the Green Book is based on such historic US interest rate data seems somewhat incongruous given the general approach to discounting that it takes.

9.3.17. Since Newell & Pizer (2003) there have been a large number of different estimates of the term structure of discount rates based on interest rate data, consumption data, and the divergence of expert opinion. We will return to the question of how HM Treasury might best empirically update its declining discount rate schedule based on this literature in Chapter 13.

9.3.b. Other motivations for declining discount rates

9.3.18. Much of the literature on declining discount rates, and the application of this literature in the Green Book, is based on arguments about the variance of x_t, as sketched in the previous subsection. However, there are two other justifications for such a term structure within the STPR framework of the Extended Ramsey Rule.

9.3.19. First, the mean of x_t may decline over time. In the context of the Ramsey Rule, m_{gt} may be lower at long horizons than for shorter ones. We see evidence that many people believe this may be the case in our survey responses. Table 8.2 showed a clear view across respondents that economic growth will slow over the next century, and there was also evidence of this in Section 8.6. This gives an alternate justification for lower long-term risk-free discount rates.

9.3.20. This was raised by some of the respondents to our survey (paragraph A.2.5). For example, one noted that: "According to the standard Weitzman/Gollier reasoning, the term structure might rise or fall with the time horizon. But even setting that aside, the mean growth rate is likely to vary over time. The growth rate in the UK has been declining, which according to the Ramsey rule should lead to a declining term structure".

9.3.21. An alternative is that the rate of pure time preference, δ, might itself have a declining terms structure, δ_t. A common justification for this is that, while we can discount our own utility ($\delta_t > 0$), we cannot ethically discount the utility of future generations ($\delta_t = 0$); see Chapter 8. This is captured by two of the qualitative responses to our survey:

- "One issue I have is that the theory stems from infinite time-horizon utility maximisation, whereby you're essentially expanding the intertemporal behaviour of one representative agent to that of how society thinks about multiple generations' intertemporal importance. We as individuals discount our own future wellbeing but I consider that to be a very different phenomenon to discounting the welfare of others who do not yet exist. However, in both case, this points towards declining discount rates. In the former, evidence shows we tend to apply declining discounting wrt our own welfare. In the latter, I'd argue that our welfare is no more important than that of future generations, suggesting a pure time preference rate of 0 for intergenerational questions. Perhaps in practice this means a discontinuity from 0.5 to 0 after 15-20 years?"
- "If δ is initially greater than 0%, it should approach 0% as the time horizon increases. Short-term positive δ can reflect impatience within a generation, while $\delta = 0$ embodies the ethical argument of intergenerational impartiality" (both paragraph A.2.5):.

9.3.22. A further case for using a value of δ that declines over time is presented by Stefanski & Trew (2024). They present evidence that society has just become more patient over time. If this trend continues, then δ will be lower in a century than it is today.

9.3.23. Despite the appeal of these arguments, they introduces a difficulty. Take a simple example. Suppose HM Treasury applies $\delta_t = 0.5\%$ for the current generation — say, up until $t = 30$ years. Beyond that point $\delta_t = 0$ to reflect our ethical responsibilities to future generations. Using current Green Book parameters of $\mu = 1$, $m_{gt} = 2\%$ for all t and ignoring the L term, the forward risk-free discount rate is 2.5% up to 30 years and then 2% beyond that date.

9.3.24. Now consider two public projects that the government might invest in for the same initial outlay of £35. The first returns £100 with certainty after 40 year and the second returns £125 with certainty after 50 years. The effective spot discount rate for the former is 2.375% and for the latter is 2.300%, giving a higher present value for the latter than the former: £39.58 against £38.67. The government would invest in the longer-term project.

9.3.25. Now move forward 30 years to when the new generation are running HM Treasury. When reassessing these projects, they now have maturities of 10 and 20 years and the appropriate discount rate is 2.5% for both of them. In this case, the present value is £77.48 for the former and £75.82 for the latter. The new generation which actually benefits from this project would perceive greater value from the shorter-term project.

9.3.26. Therefore a declining term structure of δ_t, while ethically appealing, leads to *time-inconsistent decision making*. We make decisions for future generations that they would not make themselves even under exponential (not hyperbolic) discounting.

9.3.27. In Chapter 13 we will return in much more detail to the precise calibration of the schedule of DDRs for use in the UK. For the moment, we stress following three points:

- The theory and empirical evidence on DDRs has evolved significantly since 2003 and remains strong, with the majority of survey respondents supporting its continued use.

- Persistency in UK consumption growth data appears, at first sight, to be low, suggesting that the current schedule of DDR rates may be too steep.
- If long-term growth forecasts are below short-term growth forecasts, this also justifies the use of DDRs, but a declining term structure for δ raises issues around time-inconsistency in decision making.

10. Average systematic risk under STPR

In Chapter 9, we considered how macroeconomic uncertainty affects the risk-free rate, which we will return to again in Chapter 13 for more accurate calibration. We now extend this by introducing microeconomic risk that relates to risk over the benefits from any given public project, b_{it}, within the standard STPR framework. We extend this to the SOC framework in Chapter 11 and to STPR model extensions in Chapter 12, where the potential for rare disasters is shown to have a significant influence on both the risk-free rate and risk premia. In this chapter, we concentrate on the average public project rather than considering whether any specific risk adjustment to the discount rate should be made for individual projects. We also defer that discussion to Chapters 11 and 12.

10.1. Certainty equivalents

10.1.1. When setting out the purpose of the social discount rate in paragraph 5.3.9, we stated that it should account for project maturity and project systematic risk only.

10.1.2. However, looking back at paragraph 5.5.3, we know that any adjustment to the discount rate has the same effect on the present value as making a mathematically equivalent adjustment to the net benefits.

10.1.3. Adjusting the numerator, rather than the denominator, of the PV equation for systematic risk is a common policy choice. In this case, discounting takes two separate steps. First, the government decides the highest certain net benefit it would be willing to trade for the risky net benefit that the project is actually delivering. This is known as the *Certainty Equivalent* value. Once this value has been calculated, then because it is risk-free, it can be discounted at the risk-free discount rate and any potential adjustment to the discount rate becomes fully redundant.

10.1.4. Notable academic support for the use of certainly equivalents is given by the internationally preeminent scholars in the field Zeckhauser & Viscusi (2008, p.96), who state:

> "Economists generally agree that whoever is the decision maker, the discount rate should not be adjusted for risk. The preferred approach, roughly speaking, is to address risk by converting monetary payoffs to certainty equivalents, and then do the discounting".

10.1.5. This is the primary approach taken by the Office of Management and Budget in the United States in the 2023 version of Circular A-4. This states (OMB 2023*a*, pp.81–82) that "While use of certainty equivalents is the preferred method of accounting for risk, you may choose to account for systematic risk using a risk-adjusted discount rate".

10.1.6. A number of respondents to our survey also argued for the use of certainty equivalents rather than a risk-adjusted discount rate; e.g. "I think project risk should be handled by risk adjusting the project benefit and cost flows, NOT the discount rate itself" (paragraph A.2.8).

10.1.7. In the current version of the Green Book, as well as Europe more generally, it has been more common to adjust for risk in the social discount rate (Groom et al. 2022). This is also the overwhelmingly dominant approach in the private sector in all major economies.

10.1.8. Theoretically — and unlike the case we have presented against dual discounting in Section 5.5 — there is no clear reason to prefer one approach over the other. This is because the risk adjustment depends on both macroeconomic risk (π_t) and microeconomic risk (b_{it}). Our view is that there is no strong case for the Green Book to move away from its current practice of making this adjustment in the discount rate, and there are advantages from remaining with this approach in the case of projects that span the public and private sectors.

10.1.9. Our view that risk should be adjusted for in the discount rate, while relative price changes should be accounted for in the numerator of the PV equation and *not* through dual discounting is consistent with recent guidance from the Dutch government: 'The discount rate serves only to value differences in time and risk. Relative price developments serve to value (expected) changes in relative scarcity. If relevant to the CBA, relative price developments should be incorporated into the amounts to be discounted, and not into the discount rate" (Rijksoverheid 2025, Google Translate p.20)

10.2. The CCAPM

10.2.1. In an STPR setting, from paragraph 5.3.5, we know that the risk premium (R_P) component of the discount rate is given by:

$$R_P = -\frac{1}{t}\ln\left(1 + \text{Cov}\left[\frac{b_{it}}{E[b_{it}]}, \frac{\pi_t}{E[\pi_t]}\right]\right)$$

10.2.2. We first assume that the covariance term is sufficiently small to allow us to apply $\ln(1+x) \approx x$. Then:

$$R_P = -\frac{1}{t}\text{Cov}\left[\frac{b_{it}}{E[b_{it}]}, \frac{\pi_t}{E[\pi_t]}\right]$$

10.2.3. Under isoelastic utility, and continuing to assume that consumption growth, g_t, is lognormally distributed:

$$\frac{\pi_t}{E[\pi_t]} = \frac{\exp(-\mu g_t t)}{E[\exp(-\mu g_t t)]}$$

10.2.4. Now assume as before that $g_{it} \sim N(m_{gt}, s_{gt}^2/t)$, but also now that $b_{it} \sim N(m_{bit}, s_{bit}^2)$. Because b_{it} and g_{it} are both normally distributed, we can apply Stein's Lemma to significantly simply the expression for R_P. The Lemma says that, under bivariate normality, for any well-behaved function $f(\cdot)$, $\text{Cov}[b_{it}, f(g_{it})] = E[f'(g_{it})]\text{Cov}[b_{it}, g_{it}]$. Given this, the risk premium becomes:

$$R_p = \mu\text{Cov}\left[\frac{b_{it}}{m_{bit}}, g_t\right] = \frac{\mu\rho_{bg}s_{bit}s_{gt}}{m_{bit}\sqrt{t}}$$

where ρ_{bg} is the correlation between b_{it} and g_t. The square root of t in the denominator term reflects the fact that — assuming s_{gt} has close to zero autocorrelation — the *per-period* consumption growth standard deviation declines with maturity. In essence, we have time diversification for long-maturity projects with respect to macroeconomic risk.

10.2.5. When $t = 1$, this is the standard expression for the risk premium in an STPR type environment. It is commonly known as the *Consumption Capital Asset Pricing Model (CCAPM).*

10.2.6. In contrast to the prudence effect, which is driven by the third derivative of the utility function, the risk premium is driven by $\mu = -c_t U''(c_t)/U'(c_t)$; the Pratt-Arrow measure of relative risk aversion. It is the second derivative of the utility function that is of relevance here. Therefore, a quadratic utility function would have an associated risk premium but not an associated prudence term. Indeed, quadratic utility is often used to justify the standard market CAPM, which we return to in Chapter 11.

10.2.7. We can now identify the extent to which the parameter μ is *overworked* in a standard isoelastic utility function. It captures societal preferences around at least four economically distinct concepts:

1. Through the wealth term in the Ramsey Rule, μg, it captures the extent to which, under certainty, society would wish to smooth consumption across time. This, then, is about society's wish to eradicate *intertemporal* inequality.

2. Through the relative prudence term, which is formally given by $-c_t U'''(c_t)/U''(c_t) = \mu + 1$ (notice again the dependence of this term on the third derivative of the utility function), it captures the magnitude of the social precautionary savings motive; the extent to which society would hold reserves back today to protect itself against bad future macroeconomic outcomes as captured by the final term in the Extended Ramsey Rule.

3. Through the relative risk aversion term, $-c_t U''(c_t)/U'(c_t) = \mu$, it captures how averse society is to systematic risk in public projects, as captured in the CCAPM.

4. In a regional discounting context, it captures society's aversion to *intratemporal* inequality; the extent to which society wants people across the country to have similar outcomes at any given point in time.

10.2.8. These four concepts are clearly related — they all speak of society's desire to smooth consumption in some way because of the effects of the curvature of the utility function. But, fundamentally, they are intellectually distinct and there is no theoretical reason to believe that one parameter should adequately capture all of them. Our use of isoelastic utility is therefore one of convenience that balances practical purpose with reasonable properties, rather than us truly believing that all the subtleties of social welfare can be captured by a log/power utility function driven by only two preference parameters (δ, μ).

10.2.9. There are a range of theoretical extensions to isoelastic utility; for example, Kreps-Porteus utility functions (Kreps & Porteus 1978). See, for example, Starmer (2000), Backus et al. (2004), Barberis (2013) for reviews of some of the many possible extensions. These, though, are often technically complex and extremely hard to calibrate meaningfully. Within an STPR framework, we recommend that HM Treasury continues to works with isoelastic utility, but that it also remains mindful of the limitations of this welfare specification.

10.3. Risk in the Green Book

10.3.1. The Green Book (HM Treasury 2026*c*) contains an explicit section, separate to that on the discount rate, on "Risk and Uncertainty". Here the terms are defined as "Uncertainty: The assumptions that are made in an appraisal that are not fully known. Risk: The effect of uncertainty on objectives". This follows the usage in HM Treasury's Orange Book (HM Treasury 2025, #6.61).

10.3.2. These definitions of risk and uncertainty differ from those we have introduced in paragraph 5.4.7 for these Technical Annexes. We use these terms interchangeably (to introduce linguistic variation) when there are different identifiable future states of the world with associated probabilities. This contrasts with "Knightian uncertainty" which occurs when there is no well-defined probability space.

10.3.3. The Green Book section on risk and uncertainty include a discussion around: "Commercial risks: The inputs required for a proposal might be more expensive than originally assumed. Legal risks: The proposal might face additional costs due to litigation. Fraud risks: The proposal might lose public money to fraudsters" as exemplars of the broader types of risk and uncertainty covered in the Orange Book and that are relevant to a public sector economic appraisal.

10.3.4. This is very different to the type of risk that influences the discount rate because these issues are not *systematic* — that, is, they are not correlated with the SDR, π_t. For this reason, they are better accounted for in the numberator of the NPV equation and thus are not relevant to the Terms of Reference for this review.

10.3.5. Some in the government Observer Groups considered optimism bias, which is also discussed in the Green Book, as a type of risk that should influence the discount rate. We disagree with this for two reasons. First, this risk is also not systematic. Second, it reflects a mis-estimation of the probability space that underpins an NPV. Either the estimated outcomes themselves are too optimistic, or the probabilities assigned to the best outcomes are too high. In either case, the current Green Book guidance is, in our opinion, clearly correct and does not fall under the Terms of Reference for this review:

> "Practitioners must account for optimism bias by making explicit adjustments at the outset of an appraisal. This means increasing estimated costs and timeframes, and decreasing estimated benefits, to provide a more realistic view of how much the proposal will cost, how long it will take, and what benefits it will deliver" (HM Treasury 2026*c*, #6.79).

10.3.6. What is relevant for our consideration here is the "catastrophic risk" term L in the current Green Book discount rate. We turn to this next.

10.4. The L term

10.4.1. The $L = 1\%$ term in the current Green Book discount rate was introduced in the 2003 version, which was the first to take an explicitly STPR approach in a way that closely resembles current practice. It is informative to see how the definition of L has changed over time, although the numerical estimate of 1% has stayed fixed.

10.4.2. Throughout all these iterations L has been added to δ to make the term ρ, which the Green Book (2018) describes as "'time preference' — the rate at which consumption and public spending are discounted over time, assuming no change in per capita consumption. This captures the preference for value now rather than later".

10.4.3. In the 2003 version, L was defined as an adjustment for 'catastrophe risk' that "...is the likelihood that there will be some event so devastating that all returns from policies, programmes or projects are eliminated, or at least radically and unpredictably altered. Examples are technological advancements that lead to premature obsolescence, or natural disasters, major wars etc. The scale of this risk is, by its nature, hard to quantify".

10.4.4. A value of 1% was used to capture this catastrophe effect. Based on Footnote 1 on page 97 of the 2003 version of the Green Book, this estimate was based on four papers; Newbery (1992), Pearce & Ulph (1995), Kula (1987), Oxera (2002). Three of these estimates are based on personal mortality risk, while the estimate in Newbery (1992) is based on the 'perceived risk of the end of mankind in 100 years' — see Pearce & Ulph (1995, Table 5.1) and Oxera (2002, Section 3.3). For example, Kula (1987, p.170) says: "The basic argument underlying a mortality based pure time discount rate is that each member of the community discounts his/her future utility by the probability of not being alive at that date. Therefore what we need here is the average annual death rate for the representative individual", which is estimated at 1.18% for 1975. Newbery's approach, by contrast, reflects that of former Astronomer Royal Martin Rees, who observed "I think the odds are no better than fifty-fifty that our present civilisation on Earth will survive to the end of the present century without a serious setback." Rees (2003, Prologue). This raises the question about how *individual* mortality rates and total societal eradication should impact on the social discount rate derived based on the assumption of an infinitely-lived representative agent, which we discussed in Chapter 8.

10.4.5. By the 2018 version of the Green Book, the L term was defined as "an allowance for unpredictable risks not normally included in appraisal, known as 'catastrophic' and 'systemic' risk". This definition is then expanded to:

> "The risks contained in L could, for example, be disruptions due to unforeseeable and rapid technological advances that lead to obsolescence, or natural disasters that are not directly connected to the appraisal. L also includes a small premium for 'systemic risk' because costs and benefits are usually positively correlated to real income per capita" (A6.10).

10.4.6. In the 2026 version of the Green Book, L is defined as "Catastrophic risk ... also known as 'systemic' risk, this represents an allowance for unpredictable risks that are not normally included in appraisal". This is then expanded to:

> "The risks contained in L could, for example, be disruptions due to natural disasters that are not directly connected to the appraisal. L also includes a small premium for systemic risk because social costs and social benefits are usually positively correlated to real income per capita. The Green Book assumes a value of 1.0% for the probability of catastrophic risk" (HM Treasury 2026*a*, #2.11).

10.4.7. There are several elements in these definitions that need disentangling:

- "The risks contained in L could, for example, be disruptions due to unforeseeable and rapid technological advances that lead to obsolescence..." relates, as we read this, to project risk; that the *project itself* becomes obsolete. In that case, this should be reflected in the expected benefits, $E[b_{it}]$. The expectations operator needs to incorporate a state of the world, with an associated probability, that reflects this potential technological advance. The adjustment should not be made in the discount rate.
- "... or natural disasters that are not directly connected to the appraisal". Given the discussion above, we interpret this as a macroeconomic risk unrelated to uncertainties over the project outcome. Instead, it reflects either the way that individual mortality risk aggregates into a social discount rate, or as with Newbery (1992), the threat of an outcome so catastrophic as to stop human civilisation altogether. This latter interpretation has an important effect. If we believe that society will continue to survive but at very low consumption levels, we should be investing today on their behalf, pushing down the discount rate as described in Chapter 9. But if we think that society may just cease to exist, then there it is futile to save for the future. This lowers the demand for savings and therefore raises the risk-free rate. Therefore, the directional change in the risk-free rate is the opposite under this interpretation of L than the risk introduced by a catastrophe that does not wipe out society — we return to this point in subsection 12.2. In either case, such an adjustment is correctly accounted for in ρ as it is a macroeconomic risk.
- That L includes "... a small premium for systemic risk because social costs and social benefits are usually positively correlated to real income per capita". To understand this, take the example of a public project with expected net benefits at time t of £100m and where the 95% confidence bounds for this net benefit are £50m–£150m. Under the normality assumption, a 95% confidence interval corresponds to approximately ±2 standard deviations, so the standard deviation of $b_{it}/m_{bit} = 25/100 = 0.25$. Using $\mu = 1$, consistent with the current version of the Green Book, and the standard deviation of consumption growth given above of 2.8%, we can derive an upper limit on the risk premiumn from the equation in paragraph 10.2.4 (as the correlation between b_{it} and g_t cannot exceed 1):

  $$R_P \leq 1 \times 0.25 \times 0.028/\sqrt{10} = 0.22\%$$

  This is small, as the Green Book notes; see also Moore, Boardman & Vining (2017). However, it is not correct to incorporate this into ρ, which is the risk-free component when "there is no change in per capita consumption". Instead, it is the correlation between consumption risk and project net benefit risk that drives this effect. This should therefore be separated out into a separate risk premium component if considered sufficiently large to be of policy relevance.

10.4.8. We would also note that the L term, which is a consumption discount rate risk premium, is also included in the health discount rate $\delta + L = 1.5\%$ even though that is a utility discount rate. This introduces a conceptual inconsistency into the health discount rate, further re-enforcing our Recommendation R.4 that measures of health utility should be converted into monetary units and then discounted as usual.

10.4.9. We therefore conclude that L consists of (i) one term that should best be treated in the numerator of the PV equation, (ii) one term that should be in ρ but potentially has the wrong sign unless we believe that absolute human extinction is imminent or instead that individual mortality risk is relevant to the social discount rate, and (iii) a small risk premium adjustment which should be treated separately.

> R.13. Disentantling the components of L.
>
> We recommend that HM Treasury revisits the parameter L to more explicitly disentangle the separate economic effects that underlie it.

10.5. The Arrow-Lind Theorem

10.5.1. In a seminal 1970 paper, the Nobel Laureate Kenneth Arrow, together with Robert Lind (Arrow & Lind 1970), argued that the systematic risk of all public projects is zero and that therefore the CCAPM risk adjustment term is not just small but instead precisely zero. This theorem has historically been highly influential within the practice of international social discount rates.

10.5.2. As noted by Foldes & Rees (1977, p.188) the assumptions required to prove the Arrow-Lind Theorem are: (i) the government initially appropriates all benefits and pays all costs, distributing the net returns subsequently through changes in the level of taxes; (ii) the net returns are statistically independent of each person's disposable income in the absence of the project; and (iii) each person's share of the net returns tends to zero as the number of persons tends to infinity.

10.5.3. A major symposium on the continued relevance of the Arrow-Lind theorem was reported in the Journal of Natural Resources Policy Research, Volume 6, Issue 1 (2014); we would particularly refer the reader to the papers by Baumstark & Gollier (2014) and Lucas (2014), who are internationally pre-eminent experts in this field. In general, they are skeptical about its real-world applicability because the assumptions that lie behind the theorem are unrealistically restrictive; see also earlier critiques including in Foldes & Rees (1977).

10.5.4. For example, Lucas (2014, p.86) notes that "The conclusions of Arrow and Lind rest on the presumption that government investments are free of aggregate risk. That assumption is clearly violated for many if not most of the investments made by governments around the world".

10.5.5. This view, though, is not held by all the respondents to our survey. As one noted (paragraph A.2.12): "HM Treasury clearly does not believe in Arrow-Lind, I do."

10.5.6. From our perspective, this is an empirical question on which we are theoretically agnostic. Only by considering the extent to which aggregate consumption growth and public project net benefits are observed to be correlated can we correctly deduce the appropriate risk-adjustment term for the social discount rate in a CCAPM type environment.

10.5.7. We will return to consideration of the beta of public projects in Chapter 11, taking more of a markets-based approach for calculating their systematic risk. For now, in Table 10.1, we extend table 8.2 to report the median response to our survey on what the new Ramsey Rule parameter values should be, and now adding the median response for the risk premium, L.

10.5.8. In terms of the risk premium, L, non-governmental respondents gave a median response of 0.85%, while governmental responses were much lower. There was not clearly lower for those who favoured an STPR approach compared to other respondents. For example one respondent from the academic expert panel who advocated for an STPR approach simultaneously advocated for a value of L of 2%.

10.5.9. There is also a risk that some of these responses are pulled down by the percentage element. We asked for 2% to be reported as "2", but we cannot be sure this was always the case. For example, one respondent from the Society of Benefit Cost Analysis route, responded "0.011", but this may well reflect the 1.1% rate used by the Office of Management and Budget in 2023. We had a number of responses of "0.01" which may, again, in some cases be meant to be read as 1%. This is one reason why we quote medians rather than means, but may suggest that the values L in Table 10.1 are biased downwards.

10.5.10. There was significant differences in opinion on this matter, though. Taking two consecutive responses to our survey from paragraph A.2.6 gives "Systematic risk on the market is priced with a 5% risk premium. If we include catastrophic risk - and perhaps other forms of risk too - then the risk premium should be at least 6%. The current value of 1% is ridiculously small compared to the size of risk it should

	Count	δ	μ	g (10 years)	g (100 years)	L
UK Academic	22	0.50%	1.00	1.50%	1.00%	0.70%
Social Media	17	0.25%	1.00	1.00%	1.00%	0.75%
SBCA	12	0.35%	1.10	1.50%	1.00%	0.28%
Experts	9	0.10%	1.40	1.50%	1.00%	1.00%
Total (non-government)	60	0.45%	1.00	1.50%	1.00%	0.85%
Government	8	0.50%	1.00	1.00%	1.00%	0.01%
Total (whole sample)	68	0.50%	1.00	1.35%	1.00%	0.70%

Table 10.1.: This table extends Table 8.2 by giving the median response for what the Ramsey Rule parameter values should become, now including L, from our survey.

compensate for", followed by "If a single risk premium is retained, it should be set below the current 1% to avoid over-penalising long-term public investments".

10.5.11. When combined with the evidence in Chapter 11, we conclude that the government is not sufficiently well diversified for the Arrow-Lind theorem to apply and that the appropriate value for $L > 0$.

11. Systematic risk under SOC

We now turn to how project-specific risk is set in (i) countries that set their discount rate based on the SOC approach, (ii) the real private economy, and (iii) financial markets. We present a brief and stylised representation of this framework to raise issues that are of particular relevance to HM Treasury as it reviews its social discount rates. This chapter should not be read as a more comprehensive review of private sector discount rates.

11.1. The CAPM

11.1.1. To estimate risk premia in an SOC environment, there are a number of different approaches, all of which have interpretation within a stochastic discount factor framework. The most widely applied model, e.g. Graham (2022), arises in a single-period economy when investors try to maximise the expected return of their portfolio for a given level of risk (standard deviation of returns) — they are known as 'mean-variance decision makers'. This framework was developed in a portfolio context by Markowitz (1952), and then developed into the standard market-based CAPM by Sharpe (1964), Lintner (1965) and others.

11.1.2. In this case, it can be shown that the stochastic discount factor is given by $\pi = A + Br_m$ (notice there are no t subscripts here because this is a single-period economy), where A, B are constants and r_m is the rate of return on a well-diversified portfolio of risky assets, commonly referred to as the "market portfolio". While, in theory, this market portfolio should consist of all risky assets, in practice it is often taken to be a broad equity market index with dividends reinvested; perhaps the FTSE All Share Total Return Index. We will follow that convention throughout this section.

11.1.3. This mean-variance framework results in the expected return to any asset, i, being given by perhaps the most famous set of equations in financial economics:

$$\begin{aligned} r_i &= r_f + \beta_i E[r_m - r_f] \\ \beta_i &= \frac{\text{Cov}[r_E, r_m]}{\text{Var}[r_m]} \end{aligned}$$

11.1.4. There are extensions of the CAPM to a multi-period framework, perhaps most notably by Merton (1973), but these are not commonly applied in practice.

11.1.5. This shares a number of features with the STPR CCAPM approach. In each case, the discount rate comprises of three components:

- The first element is a risk-free component, which is given by the Ramsey Rule in the standard Green Book STPR approach but here is an SOC risk-free rate, such as the yield on an index-lined Treasury security as discussed in Section 7.2.
- The second term is a beta — a measure of systematic risk. In a standard CCAPM STPR setting, this is measured as a covariance between project benefits and aggregate consumption growth, but here it is measured as co-movement against the market portfolio.
- The third item is a "price of risk". How much additional return is required for each additional unit of beta. In the CCAPM this is given by the risk-aversion interpretation of μ; see paragraph 10.2.7. In the CAPM it is given by $E[r_m - r_f]$, the expected excess return on the market portfolio, since (by construction) the market portfolio has a beta exactly equal to one.

11.1.6. A key difference between the CCAPM and CAPM approaches is that, in the former, all rates of return can be calculated theoretically. For the latter, the risk-free rate and the expected return to the market portfolio are exogenous and all other rates of return are calculated using these two values as a benchmark. We will return to this issue in Section 11.3.

11.1.7. When considering the CAPM for social discounting purposes, it is essential to distinguish between betas on the equity of quoted companies, and asset betas. Most private firms are funded through a proportion $0 \leq G \leq 1$ (for 'gearing') of debt, D, and $1 - G$ of equity, E: $G = D/(D + E)$. Here, E is the market capitalisation of the equity, not the value on the balance sheet.

11.1.8. If we denote the beta of the equity of the firm by β_{Ei} and the assets of the firm by β_{Ai}, and if we assume that the corporate debt is close to risk-free, then in a tax-free environment, $\beta_{Ai} \approx G\beta_{Ei}$. This shows that corporate asset risk is lower than the equity risk. This is intuitive clear since equity is a levered claim to the assets of the firm once the corporate debt has been fully paid off. Since equity holders are the last to be paid, they experience the greatest risk on their investment.

11.1.9. For social discount rates under an SOC approach, departments are valuing the projects themselves, not levered claims on those projects, and therefore it is the asset beta that is relevant. We begin by looking at the approaches taken by other international governments in either their implicit or explicit estimates of such asset betas for social economic evaluation.

11.2. International SOC guidance

11.2.1. Looking back at SOC guidance in Table 4.1, it is clear that there are two very different levels of discount rate. In Australia, Canada, and the upper rate under the 2003 United States guidance, the rate is 7%, which implies a significant risk premium for the social discount rate. The New Zealand SOC rate of 8% is also high but its use is limited to commercial public projects.

11.2.2. By contrast, the SOC risk premium is zero under the lower USA 2003 discount rate and German Federal rate, and small (1%-2%) for the 2023 US rate, the Norwegian rate and the discount rate used in Denmark. The central SOC risk premium in the Netherlands is somewhat higher at 2.8%, yet it is still well below the risk premium implicit in some of the higher SOC discount rate values.

11.2.3. These very distinct differences in rate can be largely explained by estimates of social asset betas. We consider each in turn.

11.2.a. High SOC rates

11.2.4. The higher rate of 7% recommended in the United States in OMB (2003) is given as

> "The 7 percent rate is an estimate of the average before-tax rate of return to private capital in the U.S. economy. It is a broad measure that reflects the returns to real estate and small business capital as well as corporate capital. It approximates the opportunity cost of capital, and it is the appropriate discount rate whenever the main effect of a regulation is to displace or alter the use of capital in the private sector" (p.33).

11.2.5. Canadian and Australian guidance take a similar approach to this. The latest Canadian guidance is given in Treasury Board of Canada (2022). Its SOC discount rate of 7% is based on a 2008 presentation update of Jenkins & Kuo (2007). The rate here is derived as a weighted average of the rate of return on postponed / displaced investment, the social cost of new domestic savings, and the marginal cost of additional foreign capital inflows. Taking economic returns to capital on domestic investment as an example, this "... can be measured by the sum of the private net-of-tax returns on capital and all direct and indirect taxes generated by this capital." (p.8).

11.2.6. Australian guidance is given in Australian Centre for Evalauation (2025), with further detail about the central 7% rate, the sensitivity analysis at 3% and 10%, and the declining term structure of discount rates given in Office of Best Practice Regulation (2014). The base rate of 7% is based on Harrison (2010, Table 4.1), although the central rate is slightly lower: see Office of Impact Analysis (2023, footnote 5). Here the central rate is set by starting "... from the marginal return to capital (which contains a risk premium reflecting average market risk)". The 10% sensitivity rate is for high-risk projects, while the 3% rate is a weighted average risk-free rate.

11.2.7. What is clear from all these examples is that social projects are expected to give the same rates of return as an average private investment. This approach implicitly assumes that (i) government projects have the same asset beta, on average, as the private sector, and (ii) that the government should price each unit of risk the same as the private sector.

11.2.b. Low SOC rates

11.2.8. The lower 3% rate recommended in the United States in OMB (2003) is explained as follows:

> "When regulation primarily and directly affects private consumption (e.g., through higher consumer prices for goods and services), a lower discount rate is appropriate. The alternative most often used is sometimes called the 'a social rate of time preference'. This simply means the rate at which society discounts future consumption flows to their present value. If we take the rate that the average saver uses to discount future consumption as our measure of the social rate of time preference, then the real rate of return on long-term government debt may provide a fair approximation. Over the last thirty years, this rate has averaged around 3 percent in real terms on a pre-tax basis" (p.33).

11.2.9. Note that, under our framing of the SOC/STPR divide, long-term debt is not an STP rate since it is not directly derived from an explicit model of social welfare. Instead, we have viewed this as a risk-free SOC rate; see Chapter 6. In this case, the OMB does not add a risk premium onto this risk-free rate. This is implicitly saying that the average social asset beta is zero.

11.2.10. A similar approach is taken by German guidance, which is currently using 1.2% *nominal* for the central Federal social discount rate. This is again based on risk-free interest rates and therefore an implicit social asset beta of zero; see Section B.VII here.

11.2.11. In our opinion, the OMB (2003) does not fully reconcile why its 3% rate is so much lower than the 7% rate. Its explanation focuses on economic distortions; primarily taxes on capital. Yet, as should be clear from the discussions above, a core reason why these rates differ is becaus private capital has very different systematic risk to Treasury Bonds; see also Li & Pizer (2021, footnote 3).

11.2.12. This omission is adjusted for in the later Appendix D of Circular A-94 (OMB 2023*b*). This allowed for an average social risk premium of 1.1% for social discounting in the United States, with agencies permitted to seek OMB approval if they wish to use different rates (although there is a more general preference for certainty equivalent net benefits discounted at a fixed rate than an adjustment for risk in the discount rate; see Section 10.1). The 1.1% value is explained as follows:

> "There are a variety of approaches for calculating this rate. OMB's Office of Economic Policy uses the standard consumption capital asset pricing model. In this model, this rate is calculated by multiplying together (1) the market risk premium by (2) the correlation between market risk and net benefits from government investment. The market risk premium, which OMB arrives at using the average implied risk premium in equity markets over the 1960-2022 period, adjusted for leverage, is estimated to be 2.5%. The correlation between market risk and net benefits is estimated to be 0.45, which captures the correlation in equity markets for economic sectors closest to government investment, adjusted for the difference in nonpayment risk between equity market investments and government investments."

11.2.13. The beta estimate of 0.45 points to a fundamental difference in approach to the 'High SOC' countries. Here there is formal recognition that the asset beta of a social project is likely to be lower than the average asset beta in the private sector because much of what the government does inherently has low pro-cyclicality. If this is the case, then there is no reason to believe that, even under an SOC approach and no frictions or externalities, the public sector should offer the same rate of return as the private sector.

11.2.14. We see this clearly in the Norwegian guidance (Norwegian Ministry of Finance 2021), with discounting practice still calling on Norwegian Ministry of Finance (2012, Section 5.5 and Table 5.2). This assumes a risk-free rate of 2.5%-3% and a stock market risk premium of 3.0%-3.5%. But it then assumes that "... an ordinary public measure, being a transportation measure, has a risk profile that is somewhat closer to a government bond than to an average project funded via the stock exchange, the calculations above suggest that a risk-adjusted real required rate of return (before tax) of about 4 percent is reasonable" (p.74) implying a beta of close to 0.5, which is similar to the US value.

11.2.15. Estimates of 0.4-0.5 for the public sector asset beta, when measured against a well-diversified equity market index, can be supported from other sources: we give three here:

- A primary source of information related to equity and asset betas, required rates of return, and gearing of different industries in an international context is provided on the website of Professor Aswath Damodaran at New York University. He gives estimates of average sectoral asset betas, adjusted for cash holdings, and averaged across the period 2021–2025, across Western Europe (Column P in the "Industry Averages" worksheet). These estimates are dated January 2026. We present these estimates for a variety of sectors that might be of most relevance to HM Treasury in Figure 11.1. Because these are asset betas, the average across sectors is below one. For certain industries that are highly relevant for the public sector (hospitals, utilities including telecoms, education, renewable energy) the betas are below 0.6. This would imply that the "government beta" is well below one based on an SOC approach.
- In Table 11.1, we reproduce beta estimates, and the rates of return prescribed, by the UK Regulator's Network Cost of Capital Annual Update 2025, Table 2. These rates of return are used to determine consumer utility prices that will ensure that that private firms receive a 'fair' rate of return on their capital base. The asset betas are largely consistent with the values given in Figure 11.1.

- A more indirect approach is to examine the betas of International Public Partnerships Ltd and HICL Infrastructure PLC. These are major, long-standing listed investment trusts whose consist predominantly of core UK social infrastructure assets, including schools, hospitals, and public-sector offices. At the beginning of June 2026, both equity betas stood in the 0.50-0.55 range (source: LSEG database).

11.2.16. The Dutch guidance (Rijksoverheid 2025) also takes an SOC approach to discounting. The average risk premium that they apply is 2.8% (the risk-free rate is zero). The methodology is given in Hoogendoorn & Romijn (2020) and, rather than being based on the risk of social assets, is instead a weighted average return of Dutch households; comprising debt, housing, bonds and other assets. This, then, is a distinct approach to that of other countries that we have considered here and will be less directly relevant to the issues raised below.

11.3. Estimates of the equity risk premium

11.3.1. Having estimated an 'average' social beta, as calculated against a broad domestic equity index, in the region of 0.4-0.5, it is next necessary to estimate the price of risk. In an SOC environment, this is given by the equity premium: $E[r_m - r_f]$

11.3.2. There is a full academic literature on different empirical approaches to estimating this parameter value, a review of which lies well beyond the scope of these Technical Annexes. We refer the reader to, for example, Damodaran (2026). As with the SOC risk-free rate and forecasts of consumption growth, our recommendation is that HM Treasury does not estimate this directly but instead refers to well-established sources. We call here on four:

- Damodaran (2026, p.151) directly gives an estimate of equity risk premium for the United Kingdom of 5.0% for 2026.

- On 1/6/2026, the Starmine equity risk premium app on the LSEG database gave an equity risk premium of 5.53% for the UK.

- Based on survey data, Fernandez et al. (2025) reports an estimate for the UK for 2025 of 5.4% and 5.1% for the mean and median of 68 responses respectively.

- Recent estimate of the equity premium applied by utility regulators lie in the 4.7%–5.5% range; see Table 11.1.

11.3.3. Collectively, these imply that a market risk premium of 5%-5.5% for the UK is well supported by current empirical evidence. Even allowing for gearing, this is well above the 2.5% unlevered equity premium used by OMB (2023*b*); see paragraph 11.2.12. The estimate of the levered equity premium in the US is about 4.5% from Damodaran (2026); below the UK value. Damodaran estimates the unlevered beta of non-financial firms in the US at 0.9, implying an unlevered equity premium for the US of approximately 4% using the methods we are advocating here, well below the OMB value.

11.3.4. Beyond these empirical methods, another approach that can be used is to theoretically estimate the equity premium and the risk-free rate. Here, standard financial asset pricing reveals equations that are, in terms of their functional form, identical to the Ramsey Rule for the risk-free rate and the CCAPM for the risk premium. See, for example, equations (1.7) and (1.16) in Cochrane (2009).

11.3.5. As we have argued in Chapter 6, there should be little distinction in a perfect frictionless market between the SOC and STPR approaches to the discount rate as the same models underpin each. In a representative agent economy, such as underlies the Ramsey Rule, the investor is also the citizen and therefore the utility function may be calibrated with the same values of δ, μ — at least under very highly stylised theoretical assumption.

11.3.6. As is very widely known, though — and widely discussed in the international guidance on social discount rates —the Ramsey Rule and CCAPM have very poor ability to explain the real return to risk-free assets and the observed average returns to the stock market respectively. The CCAPM equity premium is too low and the Ramsey Rule risk-free rate is too high when compared against observed market yields and returns (for the latter, see, for example, National Audit Office (2018, Figure 5)). These are the "equity premium puzzle" of Mehra & Prescott (1985) and the "risk-free rate puzzle" of Weil (1989).

11.3.7. HM Treasury discussed this issue, and its relevance to the Green Book discount rate (which, at the time, was taking an SOC approach) as far back as HM Treasury (1991, p.83). More recently, this has been considered in a UK context by Spackman (2020, 2021, 2024), who is largely skeptical of the relevance of this issue for HM Treasury.

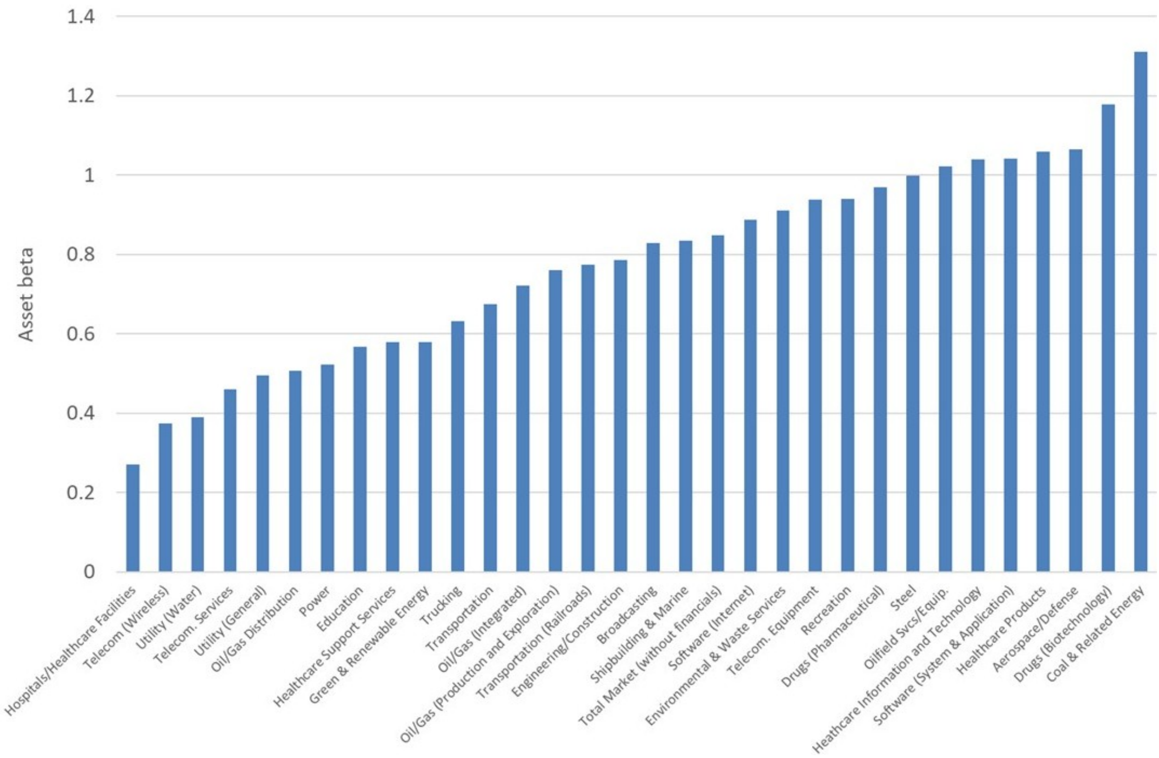

Figure 11.1.: Asset betas, calculated under an SOC approach, for Western Europe averaged over the period 2021–2026 for sectors that may be most relevant for HM Treasury as given by Ashwan Damodaran's website.

11.3.8. We, by contrast, are more positive in our opinions concerning the relevance of the equity premium for the Green Book discount rate. First, we would note that the SOC risk premium for a social asset would be less than one half of the equity premium because the asset beta of social projects is likely to be below 0.5. An SOC approach based on estimates of the equity premium does *not* imply a social risk premium of c.5%.

11.3.9. Second, we would note that the equity premium has come down somewhat from the original values used by Mehra & Prescott (1985), who cited 6.18%, particularly for the US.

11.3.10. Third, as we have stressed throughout these Technical Annexes, the Ramsey Rule and the CCAPM give STPR values only under very parsimonious theoretical assumptions. A range of different, plausible, extensions to these models give very different estimates of what a welfare-based equity premium should be. See, for example, Mehra (2007) for a review of much of the early groundwork on this topic.

11.3.11. More recently, Cochrane (2017, p.948) has provided ten possible sources of explanation for why core summary statistics for financial markets have been so difficult for standard models to explain (bibliography details of references in this quotation given in the original paper):

1. "Habits (Campbell and Cochrane, 1999a, 1999b).
2. Recursive utility (Epstein and Zin, 1989).
3. Long-run risks (Bansal and Yaron, 2004; Bansal, Kiku, and Yaron, 2012).
4. Idiosyncratic risk (Constantinides and Duffie, 1996).
5. Heterogeneous preferences (Gârleanu and Panageas, 2015).
6. Rare Disasters (Reitz, 1988; Barro, 2006).
7. Utility nonseparable across goods (Piazzesi, Schneider, and Tuzel, 2007).
8. Leverage; balance sheet; institutional finance (Brunnermeier, 2009; Krishnamurthy and He, 2013).
9. Ambiguity aversion, min–max preferences (Hansen and Sargent, 2001).
10. Behavioral finance; probability mistakes (Shiller, 1981, 2014)".

11.3.12. Particularly relevant for our arguments are papers that extend the rare disaster models of Rietz (1988), Barro (2006) and Barro (2009) that we that are the focus of our attention in Chapter 12. Gabaix (2012), Wachter (2013), and more recently Chen et al. (2024) show that allowing for a time-varying 'catastrophe' risk in extensions of the CCAPM can help resolve a wide range of financial market anomalies that include — but also go way beyond — the risk-free rate and equity premium puzzles. Such an approach has been highly influential in setting social discount rates in France although support for such models is not universal (Cochrane 2017).

11.3.13. The core point is that augmented STPR-type models that go beyond the standard Ramsey Rule / CCAPM, be they rare disaster models or alternatives, have the ability to broadly explain key properties of observed financial market returns. And this is what we would broadly expect. Governments want to optimise social welfare while individuals, who make up society, wish to more directly optimise their own welfare. While there are many reasons why these will not lead to identical values — we are not claiming that the SOC and STP rates should be identical — it is our opinion that the STPR-SOC divide cannot be as large as implied by Mehra & Prescott (1985) if the welfare model used allows for greater sophistication than the CCAPM. This is an argument we also made in Chapter 6 and will expand upon in Chapter 12.

11.3.14. Even given this, the 5%-5.5% equity premium for the UK is almost certainly too high as an estimated social cost of risk. There are a wide range of reasons for this:

- Investors need to consider their net return after tax, transaction frictions, potential premia for asset illiquidity, and compensation for informational asymmetry. Foreign investors may also require a premium for foreign exchange risk. These all mean that the realised returns that investors receive are lower than the gross equity premium.
- Investors receive returns without having to account for the externalities that their investments have on, for example, the environment. Therefore the true social returns from these investments are likely to be lower than the gross equity premium.
- There is no reason, in theory, why the government should require the same return for a given level of beta as an average investor. The government may just have a higher willingness to bear risk.

11.3.15. For these, and other, reasons, we believe that the equity premium should be treated as an upper bound for the social price of risk. While it is difficult to know exactly how to calibrate any adjustment, we believe that a social price of risk of 4.0%-4.5% is broadly consistent with market returns under an SOC approach.

11.3.16. Therefore, with a beta of 0.4-0.5 and a price of risk of 4.0%-4.5%, this would give an SOC risk premium for the average risk public project in the range 1.6%-2.25%. We take a central value, based on SOC considerations alone, of $0.45 \times 4.25\% = 1.9\%$.

11.3.17. This, again, is broadly consistent with the regulated values for utilities given in Table 11.1 and sits very much within the "Low SOC" countries.

11.4. Project-varying risk premia

11.4.1. While the average SOC risk premium for a standard public project sits at approximately 1.9%, the question arises as to whether, in an SOC environment, this should vary by project.

11.4.2. The theoretical answer is straightforward — it is clearly 'yes'. The specific beta of any project is multiplied by the price of risk to derive the SOC discount rate. And, as shown in Figure 11.1, there are clear differences in betas between industry sectors that roughly align with different government departments.

11.4.3. Recently, Gollier (2026, abstract) has argued that "The welfare loss of using a single discount rate is equivalent to a permanent reduction in consumption that lies somewhere between 15% and 45% depending on which single discount rate is used", showing how significant this issue is for policy-making purposes.

11.4.4. We also note, from Table 11.1, that UK regulators vary the cost of capital by sector depending on estimates of the asset beta.

11.4.5. That the risk premium should vary by project was firmly supported by our expert academic panels; see Table 11.2. There is a very clear consensus, both among the non-government sample as a whole, but particularly within the Expert group, that the risk premium should vary by project. However, government respondents were much less convinced about the merits of project-specific risk premia.

11.4.6. Further information on this matter resulting from our survey, and additional evidence in support for the use of declining discount rates, is given in Table 11.3. Across all respondents, we report the median recommended social discount rate for different horizons and different levels of project risk. This reveals a general preference for project-varying risk premia and a declining term structure of discount rates. For example, one respondent to our survey said (paragraph A.2.6) "There is no single risk premium for all public projects. To use a single rate despite varying risk is madness."

11.4.7. A further key survey that is relevant for this question was reported by Gollier et al. (2023). We very briefly summarise their findings in Table 11.4. This paper gives considerable detail on this question; much more so than we can consider here. The results, though, are largely aligned with the results of our own survey. Most respondents (c.75%) advocated for the use of risk-adjusted discount rates, and the majority felt that transport (pro-cyclical) should be discounted at a higher rate than healthcare (a-cyclical). Even then, the median risk premium on railroads (1%) is half of the equity premium (2%), again implying a social beta of about 0.5 and consistent with much of the evidence we have presented above. Gollier et al. (2023), however, find this small variation in discount rates anomalous because their beta estimate for railroads is above 2.

11.4.8. They call this effect the "discounting premium puzzle". Many experts agree that the social discount rate should be adjusted according to the specific systematic risk for any public project, yet the range of risk premiums that these experts propose appear much narrower than theory might suggest.

11.4.9. One possible explanation for this is that, despite their theoretical appeal, betas have low power to explain variations in observed rates of return between different finanical assets. This finding can be traced back at least as far as Fama & French (1992). They showed that, once corporate size is accounted for, beta has no ability to explain the cross-section of equity returns. This has led to a plethora of multi-factor models that add further explanatory variables into empirical asset pricing models, including the 5-factor model of Fama & French (2015) and the q−factor model of Hou et al. (2015). In a UK context, Tharyan et al. (2024) have recently shown that both these models subsume the CAPM in terms of their ability to explain the cross-section of asset returns; see also Fletcher et al. (2024).

11.4.10. Such multi-factor extensions to beta cannot be applied under an SOC social discounting framework. Many of the characteristics that are adjusted for, such as a firm's book-to-market ratio, have no clear analogy within the public sphere. Therefore, while the UK Regulator Network continues to rely on betas for setting costs of capital, it is unclear to us, based on empirical evidence, that beta alone provides a strong basis for the UK to vary the risk premium by project.

11.4.11. We also note that where countries have attempted to vary individual project discount rates according to their systematic risk, this recommendation has frequently struggled to gain traction in practice. This is true even within departments that would be allowed to use lower than average discount rates as a consequence. For example, in the case of Norway:

> "In the period 1999 – 2014 the guidelines allowed for differentiation of the discount rate based on the project's systematic risk. After 2014, this was no longer allowed. This was in part because the risk adjustment proved to have low transparency and consistency across sectors". (Nesje & Lund 2018, p.34, authors' own translation)

11.4.12. Some respondents to our survey were also concerned that the introduction of project-specific risks might be open to politicisation; see paragraph A.2.7. For example: "While I can see value in this approach, it also introduces scope for bias (conscious or otherwise) and possibly political interference", "Because it would create a very expensive assessment system, with potential political incursions", and "I think this sounds good in principle but in practice it'd add a layer of bureaucracy to the business case process and would probably just result in 'gaming' the assumptions around what type of project it is".

11.4.13. Some respondents to our survey also questioned whether there were the analytical resources within public bodies to do this effectively, e.g. "This might be too difficult and provide spurious accuracy. .... Most people struggle with basic CBA, let alone doing this sort of thing" (paragraph A.2.7).

11.4.14. One possibility, as used in the Netherlands, is to have a limited number of risk categories into which projects could be assigned (high/medium/low, for example) with a different risk premium for each. Again, such a concept received support from some respondents to our survey; e.g. "Rather than applying a single uniform risk premium to all public projects, HM Treasury could specify a small number of standard consumption betas by project type (e.g. resilience, state-independent public goods, cyclical infrastructure, demand-driven regeneration). Risk adjustments would then be derived from these betas and an STPR-consistent consumption risk premium, better reflecting systematic risk differences across projects while remaining operationally simple." (paragraph A.2.7.

11.4.15. We carefully considered this as a recommendation. However, even this simpler approach remains subject to the critiques of the poor empirical performance of single-factor beta models, the practical challenges of operationalising such an approach, and the additional difficulty of directly estimating betas for different categories of public projects rather than inferring them from the private economy.

11.4.16. Given these observations, our recommendation is that HM Treasury only adjusts the average risk premium where it is practical, proportionate, and evidence-based to do so. However, we also advise that HM Treasury keeps this recommendation under review and gathers additional evidence that will help it decide in future whether adjusting for beta can improve the practice of economic appraisal.

11.4.17. Projects that have social insurance properties present a strong case for being exceptions to this rule, and we turn to these specifically in Chapter 12.

R.14. With limited exceptions, the average social risk premium should be applied to all projects.

Except where it is practical, proportionate, and evidence-based to do otherwise, the government should apply the same risk premium to all projects irrespective of their betas. A notable exception, covered in Chapter 12, is projects with social insurance properties. More detailed evidence on the systematic risk of public projects should be gathered to review the impact of this recommendation at a later date.

Sector	Mar-21 Telecoms	Mar-21 Telecoms	May-22 Gas TSO	Nov-22 Electricity	Feb-23 Gas Distribution	Oct-23 Rail	Oct-24 Electricity TSO/DSO	Dec-24 Water
Allowed Return on debt (pre-tax)	1.5%	1.6%	1.15%	3.01%	2.52% - 3.94%	2.33%	3.82%	3.15%
Risk free rate	-1.0%	-1.0%	-1.1%	1.23%	1.77%	1.68%	1.9%	1.52%
Total market return	6.7%	6.7%	6.8%	6.5%	6.5%	6.5%	6.75%	6.68% - 6.98%
Equity risk premium (implied)	7.7%	7.7%	7.9%	5.27%	4.73%	n/a	4.85%	5.16% - 5.46%
Notional equity beta	0.88	1.05	0.76	0.759	0.69-0.82	0.7	0.69	0.59-0.65
Debt beta	0.10	0.10	0.075	0.075	0.075	n/a	0.075	0.05-0.15
Asset beta	0.53	0.62	0.35	0.349	0.35-0.41	n/a	0.35	0.32-0.35
Cost of equity (pre-tax)	7.6%	9.1%	6.08%	6.67% - 7.50%		n/a	6.97%	6.80%
Allowed return on equity (post-tax)	5.8%	7.1%	4.92%	5.23%	5.02%- 5.65%	5.05%	5.23%	5.10%
Notional gearing	45%	45%	60%	60%	55%	62.5%	55%	55%
Assumed tax rate	19%	19%	19%	19%	19%	25.0%	25%	25%
Return on capital (pre-tax)	n/a	n/a	3.12%	n/a	4.76-5.17%	3.98%	5.32%	4.79%
Return on capital (vanilla)	3.8%	4.6%	2.66%	3.90%	3.93% - 4.42%	3.35%	4.53%	4.03%
Retail net margin deduction	n/a	n/a	n/a	n/a	n/a	n/a	n/a	0.06%
WACC (vanilla)	3.8%	4.6%	2.66%	3.90%	3.93% -4.42%	3.35%	4.53%	3.97%

Table 11.1.: Regulated real (CPI/CPIH adjusted) costs of capital in the UK utility sector. Source UK Regulator's Network Cost of Capital Annual Update 2025, Table 2. "DSO" = Distribution System Operators. "TSO" = Transmission System Operators.

	Count	Yes	No	Unsure
UK Academic	25	56%	32%	12%
Social Media	18	56%	17%	28%
SBCA	10	40%	30%	30%
Experts	12	92%	8%	0%
Total (non-government)	65	60%	23%	17%
Government	12	25%	50%	25%
Total (whole sample)	77	55%	27%	18%

Table 11.2.: Should the risk premium be adjusted by project?

	No risk	Low risk	Medium risk	High Risk
10 years	2.00%	2.00%	3.00%	3.50%
50 years	1.75%	1.78%	2.50%	3.35%
100 years	1.50%	1.50%	2.40%	3.00%

Table 11.3.: What should the Green Book discount rate be across different project risks and maturities? Median value across all respondents, where the number of responses varies between 33 and 41 for the different questions.

	Mean	Median
Risk-free	2.30	2.00
Stock market	4.72	5.00
Equity premium	2.43	2.00
Railroads	3.38	3.00
Healthcare	2.79	2.50
Public debt	2.52	2.00
Climate mitigation	2.28	2.00

Table 11.4.: This gives the 10-year real discount rate (except for climate mitigation), expressed in percentages, for different public and private asset types as reported by Gollier et al. (2023, Table C.9) from a survey of expert economists who advocated for risk-adjusted social discount rates.

12. Rare disasters

12.1. Introduction

12.1.1. Within an STPR framework, we have shown in Chapter 9 that extending the Ramsey Rule to allow for volatility in consumption growth makes only a small difference to the estimated risk-free component of the social discount rate. Similarly, in Chapter 10, we demonstrated that the standard CCAPM generates only a small social systematic risk premium.

12.1.2. In Section 11.3.11, we described a number of extensions to the standard Ramsey Rule / CCAPM framework that have been influential in financial economics for explaining the equity premium, risk-free rate, and other asset pricing puzzles

12.1.3. One of these, the risk of rare macroeconomic disasters following Rietz (1988), Barro (2006), has also been influential in the social discounting literature. It is, for example, referenced in governmental documents from the Netherlands (Rijksoverheid 2025, p.68 et. sec.) and Norway (Norwegian Ministry of Finance 2012, Chapter 8), as well as being raised in our 2018 Green Book discount rate review (Freeman et al. 2018). It has been particularly influential in setting social discount rates in France.

12.1.4. In this chapter, we provide illustrative examples of such models to demonstrate the effect that, in a purely STPR environment, the risk of rare disasters can have on the estimated social discount rate.

12.1.5. Our aim here is not to precisely calibrate these models for the UK economy; that is a task that we defer until Chapter 14. Instead, by using illustrative numerical examples we make a number of conceptual points that will be central for our overall recommendations:

- The Ramsey Rule and its CCAPM extension are only one encapsulation of the STPR framework. Different choices over the form of the social welfare function or, as here, the dynamics of the underlying economy make very significant differences to the estimated STP rate. This emphasises the point made in Chapter 6 that these equations should not be treated as accounting identities in an STPR framework.

- The estimated STPR discount rate in this framework depends on three key things: (i) the probability of a rare disaster, (ii) the magnitude of such a disaster, and (iii) the persistence of any macroeconomic crash.

- The Ramsey Rule is, within an STPR environment, likely to provide an upper estimate of the risk-free component of the social discount rate. The CCAPM, by contrast, is likely to provide a lower estimate of the average social risk premium.

- Projects that offer'social insurance' properties by delivering the greatest benefits at times of societal collapse (and hence high marginal utility) may demand a much lower social discount rate than the public sector average.

12.2. Rare disasters and the risk-free rate

12.2.1. We begin by returning to the framework of Chapter 9. Suppose that, over the next t years, there is a high probability, q that $g_t \sim N(m_{gt}, s_{gt}^2/t)$ as usual. However, there is also a small probability, $1-q$, that g_t will take a highly significant negative value g_{dt}. Taking the formula for the risk-free rate from paragraph 9.2.7:

$$\begin{aligned} r_{ft} &= \delta - \frac{1}{t}\ln[q\exp(-\mu m_{gt}t + 0.5\mu^2 s_{gt}^2 t) \\ &\quad +(1-q)\exp(-\mu g_{dt}t)] \end{aligned}$$

12.2.2. Take an example. Suppose that, consistent with current Green Book guidance, $\delta = 0.5\%, \mu = 1, m_{gt} = 2\%$ and $s_{gt} = 2.8\%$. Take a horizon of $t = 10$ years where, if the 'down' state occurs, consumption will be 50% lower than current levels: $g_{dt}t = -50\%$, or $g_{dt} = -5\%$. Assume that the probability of this happening is $1-q = 2\%$. In this case, the risk-free rate is much lower than before at 2.262%.

12.2.3. There are three elements at play here:

- First, the mean is reduced from 2% to 1.86%, lowering the simple Ramsey Rule social discount rate to 2.36% if this jump risk is recognised.

- Second, this downside risk increases the standard deviation of annualised consumption growth. Within the Extended Ramsey Rule, s_{gt} should be replaced by s'_{gt} where

$$s'_{gt} = \sqrt{qs_{gt}^2 + q(1-q)(m_{gt} - g_{dt})^2 t}$$

where the final term remains multiplied by t because the jump risk is persistent at -5% each year. With these parameter values, $s_{gt} = 4.158\%$, so the Extended Ramsey Rule risk-free rate is 2.274%

- The distribution is heavily left-hand skewed, which increases the precautionary motive slightly further, reducing the risk-free rate to 2.262%.

12.2.4. All three of these different impacts of rare disaster risk on the risk-free rate become more pronounced as (i) μ increases, (ii) the probability of being in the down state increases, and (iii) the severity of the down state gets more disastrous.

12.2.5. This highlights the dangers of estimating the economic growth parameter of the Ramsey Rule over historic periods. If we have lived through relatively stable times, then such jumps will not be observed in the data. But that does not mean that such risks will not impact in the future. This is known in financial economics as a 'peso problem'; see, for example Krasker (1980).

12.2.6. Given this risk of mis-calibration, we can conclude that downside jump risk potentially has a practically significant impact on the risk-free component of the social discount rate. Such a drop in consumption may be caused by, for example, a major pandemic, natural disaster, technological disaster, or war in a way that does not hit the hurdle of Knightian uncertainty.

12.2.7. What is also crucial here is the persistence of any severe decline in consumption. If, as with Covid, the severe macroeconomic effects are relatively short-lived, then this will affect the discount rate for a few individual annual benefits, b_{it}, but not for the entire stream of benefits $b_{i1}, \ldots, b_{iT}$. However, if we expect there to be a very slow recovery from any consumption shock, then this will affect the risk-free discount rate that we should apply to many of the future net benefits, leading to a significant effect on the overall public project PV. There is therefore a need to estimate the scale, probability, *and* persistency of any shock to consumption should HM Treasury wish to apply these jump adjustments to the Ramsey Rule in future.

12.3. Rare disasters and the risk premium

12.3.1. Now introduce a risky project into this economy. With probability q, g_t and b_{it} are bivariately normally distributed, $g_{it} \sim N(m_{gt}, s_{gt}^2/t)$, $b_{it} \sim N(m_{bit}, s_{bit}^2)$, with correlation ρ_{bg}. Otherwise, with probability $1-q$, then $g_t = g_{dt}$ and $b_{it} = b_{dit}$, two scalar values, with certainty.

12.3.2. In this case, we can calculate the risk premium using the general formula given in paragraph 5.3.5. The full algebraic expression for this is now more complex, but the individual terms in this equation are given by (proof available from the authors on request):

$$\begin{aligned} X &= \exp(-\mu m_{gt} t + 0.5\mu^2 s_{gt}^2 t) \\ Cov(b_{it}, \pi_t) &= -qX\mu\rho_{bg} s_{bit} s_{gt}\sqrt{t} + [q(1-q)\times \\ &\quad (X - \exp(-\mu g_{dt} t))(m_{bit} - b_{dit})] \\ E[b_{it}] &= q m_{bit} + (1-q) b_{dit} \\ E[\pi_t] &= qX + (1-q)\exp(-\mu g_{dt} t) \end{aligned}$$

12.3.3. Notice that X is just the expected value of the stochastic discount factor, π_t, if there were no jumps. Given this, if $q = 1$, we return to the CCAPM as before.

12.3.4. To understand the relevance of this, take an example. Consider the 10 year public project ($t = 10$) that we considered above in paragraph 10.4.7. In normal times, it is expected to generate a net benefit of m_{bit} =£100m. The 95% confidence interval for this benefit is between £50m and £150m, meaning that the standard deviation of the benefit in normal times is s_{bit}=£25m as a 95% confidence interval is plus or minus two standard deviations. We assume that the correlation between the public project and economic growth $\rho_{bg} = 0.45$.

12.3.5. Now introduce jump risk as before. There is a $1 - q = 2\%$ probability of a $g_{dt} t = 50\%$ decline in consumption over the 10 years. We consider two projects. One of these is highly pro-cyclical so that, should the down state occur, $b_{dit} = -200$ because of high fixed costs associated with the project and low benefits. The other is highly counter-cyclical with $b_{dit} = +500$ because this project has provided some type of implicit social insurance against the disaster state.

12.3.6. Other parameters are as usual; $\delta = 0.5\%$, $\mu = 1$, $m_{gt} = 2\%$ and $s_{gt} = 2.8\%$. In Table 12.1, we present the risk-free rate, the project risk premiums, and the overall discount rate for the pro- and counter-cyclical projects for $\mu \in \{0.5, 1, 1.5, 2\}$. In the no-disaster case, the risk-free rate is calculated using the Extended Ramsey Rule.

12.3.7. The difference between the risk-free rate in the disaster and no-disaster cases is the impact of severe downside risk on the precautionary savings demand, as described in detail in Section 12.2. This difference increases for higher μ as explained above. The risk premium from the CCAPM remains very low because of the time-diversification of consumption risk over a 10-year horizon. However, disaster risk makes a material difference to the discount rate for pro- and counter-cyclical projects even when $\mu = 1$.

μ	0.5	1	1.5	2
		No disaster risk		
r_f	1.49%	2.46%	3.41%	4.34%
R_P	0.05%	0.10%	0.15%	0.20%
r_{it}	1.54%	2.56%	3.56%	4.54%
		Disaster risk (procyclical: $b_{dit} = -200$)		
r_f	1.41%	2.26%	3.05%	3.76%
R_P	0.32%	0.75%	1.34%	2.18%
r_{it}	1.72%	3.01%	4.39%	5.95%
		Disaster risk (counter-cyclical: $b_{dit} = 500$)		
r_f	1.41%	2.26%	3.05%	3.76%
R_P	-0.25%	-0.61%	-1.09%	-1.72%
r_{it}	1.15%	1.65%	1.96%	2.04%

Table 12.1.: The discount rate of pro- and counter-cyclical projects, plus the risk-free rate, in the presence of rare disaster risk.

12.3.8. As illustrated by Table 12.1, both in relation to the risk-free rate and the risk-premium, the small probability of a major drop in consumption can have a much more significant impact on the discount rate than more usual day-to-day fluctuations as might be modelled by a normal distribution.

12.3.9. Where such arguments have played a prominent role in public policy discussions has been around the economic case for investing in climate change. This was originally noted in a series of papers written 15–20 years ago. For example:

- "Spending money now to slow global warming should not be conceptualized primarily as being about optimal consumption smoothing so much as an issue about how much insurance to buy to offset the small chance of a ruinous catastrophe" Weitzman (2007, p.704–705);
- "To what extent does economic analysis of climate change depend on low-probability, high-impact events? The short answer is a great deal" Dietz (2011, p.537);
- "I have argued that the economic case for stringent GHG [greehnouse gas] abatement cannot be made based on 'most likely outcomes' ... (instead) any case for stringent abatement must be based on the possibility of a catastrophic climate outcome" Pindyck (2013, pp.234–235).

12.3.10. More recently, and specifically in relation to climate tipping points, Dietz et al. (2021, p.1) observe that "Collectively, climate tipping points increase the social cost of carbon (SCC) by ~ 25% in our main specification. The distribution is positively skewed, however. We estimate an ~ 10% chance of climate tipping points more than doubling the SCC". USEPA (2023) considered the impact of such tipping points on the SCC in some detail when updating Federal regulations during the Biden administration.

12.3.11. For the avoidance of ambiguity, we make these points in relation to the choice of the social discount rate in the UK and *not* the social cost of greenhouse gas emissions. The Green Book does not set these values using a discounted cash flow approach. Instead, it recognises the governmental commitment to Net Zero by 2050, and then estimates what the cost will be of reaching this target: "Since 2009, a 'target consistent' approach has been used to estimate the values, where these are calculated as the marginal abatement cost of meeting targets". Detailed guidance is given in DESNZ (2023).

12.4. Illustrative numerical examples

12.4.1. As explained in previous sections, the risk of rare disasters: (i) potentially has a significant impact on the risk-free rate and project risk premia within a STPR framework, (ii) has been shown to have a significant impacts on estimates of the social cost of carbon in a way that has been explicitly considered for policy purposes by USEPA (2023), and (iii) has been influential in asset pricing as we discussed in Chapter 11.

12.4.2. Because of the theoretical and practical relevance of these models, we extend our analysis of this issue through six illustrative multi-period examples. The results are presented in the six panels of Figure 12.1.

12.4.3. In each case, as usual, we set $\delta = 0.5\%$ and $\mu = 1$. These examples are constructed using a Monte Carlo simulation across 250,000 iterations. In each case, we model annual per capita real consumption growth, together with the growth rate in the benefits of a pro-cyclical project and a counter-cyclical project, with the benefits of both starting at £1. We use a Markov switching model with between two and four states. We consider project benefits with maturity up to $T = 100$ years.

12.4.4. Example 1 has two states of the world: an up-state ($s = 1$) and a down-state ($s = 2$). Consumption growth = 4.8% in the up-state and -0.8% in the down-state. The pro-cyclical project benefits grow at 2.30% in the up-state and -0.38% in the down-state. The counter-cyclical project benefits grow at -0.38% in the up-state and 2.8% in the down-state. In addition, the growth rate of the two projects has an idiosyncratic component drawn from a uniform distribution in the range [-3%,+3%]. The transition probability matrix is given by $T(s_1, s_2)$, where s_1 is the current state and s_2 is next-period's state. In our first example, there is no persistence in consumption growth, so $T(s_1, s_2) = 50\%$ for $s_1, s_2 \in \{1, 2\}$. This period's consumption growth is non-informative for next period's consumption growth.

12.4.5. Results are given in Panel A of Figure 12.1. This confirms what we have seen in earlier chapters when there is no persistency of consumption risk. The term structure of discount rate is flat, the risk-free rate is close to the Green Book value of 2.5% (slightly lower because of the precautionary savings motive) and the positive/negative risk premia for the pro-cyclical and counter-cyclical project are sufficiently small to be of little significance for policy purposes.

12.4.6. Example 2, presented in Panel b of Figure 12.1 is the same as Example 1 except in one regard. Now there is full persistency of states, $T(1, 1) = T(2, 2) = 100\%$ and $T(1, 2) = T(2, 1) = 0$. The initial state is set at random with equal probability and then the economy remains in that state for the next century. This is similar to the fully persistent model of Weitzman (1998). This confirms what we have already established in earlier chapters: high persistency in consumption growth risk both reduces the risk-free rate — explaining the current declining schedule of discounting rates in the current version of the Green Book — but also increases the risk premium on the pro-cyclical project; see for example Gollier (2014). The overall discount rate for the risky project is relatively flat.

12.4.7. Example 3 considers a case with highly persistent, but not fully persistent consumption growth. Now $T(1, 1) = T(2, 2) = 80\%$ and $T(1, 2) = T(2, 1) = 20\%$. Otherwise, this is the same as Examples 1 and 2. The question is whether this will look more like Panel A or Panel B. The result is presented in Panel C. While there is some decline in the risk-free rate over time, this is relatively small given the assumed high degree of persistence in the consumption process. Compared to Panel A, though, there are greater differences in the three discount rates at longer maturities.

12.4.8. Example 4 adds a third, 'rare disaster' state. In this state, consumption growth is −40%, the growth rate of the benefits of the pro-cyclical project is −30% and of the counter-cyclical project +30%. There is a 2% probability of entering this state from either the up-state or the down-state. Once in this state, the economy will return to the down-state with certainty next period. Otherwise $T(1, 1) = T(2, 2) = 60\%$ and $T(1, 2) = T(2, 1) = 38\%$. In full, the transition probability matrix is given by:

$$T(s_1, s_2) = \begin{bmatrix} 60\% & 38\% & 2\% \\ 38\% & 60\% & 2\% \\ 0\% & 100\% & 0\% \end{bmatrix}$$

Because this is a one-period consumption decline, it is a persistent shock in *levels*. The economy will never recover to its previous condition. The results are presented in Panel D. The term structures remain largely flat, but the level of the discount rate is much lower than in Panels A and C. This is the increased precautionary savings demand caused by the small probability of major consumption declines.

12.4.9. Example 5, presented in Panel E, adds a fourth, 'recovery' state. This can only be entered from the 'rare disaster' state and allows for levels of consumption to largely recover after one period. This reflects what happened following Covid. In this recovery state, consumption growth is +36%, the change in benefits from the pro-cyclical project is +25% and from the counter-cyclical project is -25%. There is a 50:50 chance that the economy will recover or not (if it does not recover after one year, it does not recover at all). In this case the transition probability matrix is given by:

$$T(s_1, s_2) = \begin{bmatrix} 60\% & 38\% & 2\% & 0\% \\ 38\% & 60\% & 2\% & 0\% \\ 0\% & 50\% & 0\% & 50\% \\ 0\% & 100\% & 0\% & 0\% \end{bmatrix}$$

This raises the discount rate compared to Panel D because of the potential for consumption recovery. Apart from some shape in the term structure at very short maturities (caused by the fact that the economy can't recover before it has fallen), the term structure is largely flat.

12.4.10. Panel F presents our final example. Now, once in the 'jump' state, there is the potential of remaining in that state as well as transitioning to either the down state or the recovery state. Very severe drops in consumption level can happen in consecutive years with 20% probability. The transition probability matrix is now given by (the only changes are in the third row):

$$T(s_1, s_2) = \begin{bmatrix} 60\% & 38\% & 2\% & 0\% \\ 38\% & 60\% & 2\% & 0\% \\ 0\% & 40\% & 20\% & 40\% \\ 0\% & 100\% & 0\% & 0\% \end{bmatrix}$$

In this case, we return to discount rates that look very similar to those in Panel D. The probability of recovery in the level of consumption growth is offset by the risk of a further significant consumption decline.

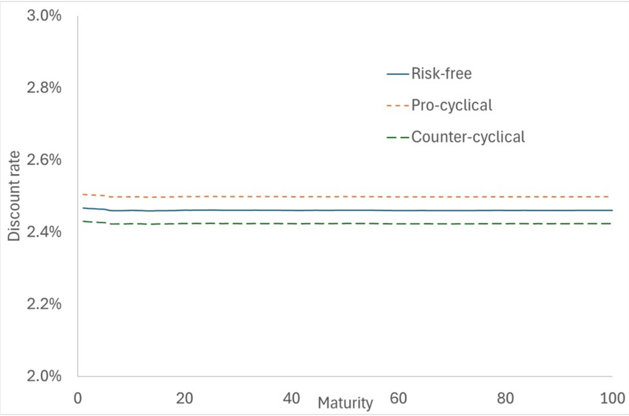

(a) No jump risk, no persistence in consumption growth

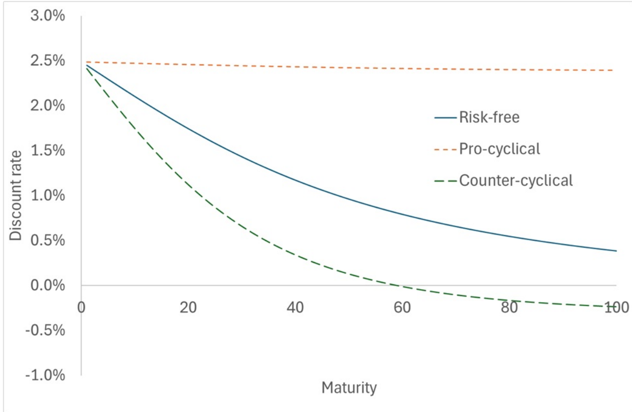

(b) No jump risk, full persistence in consumption growth

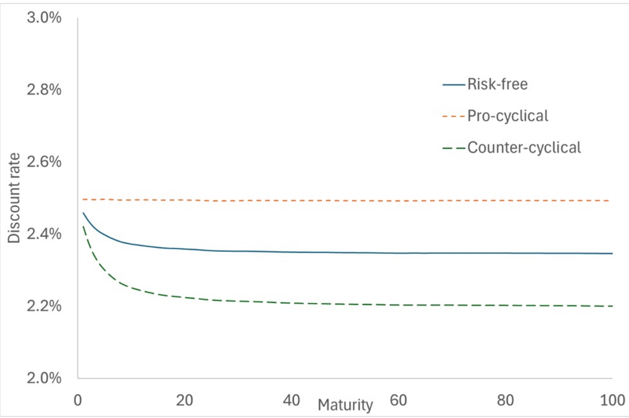

(c) No jump risk, partial persistence in consumption growth

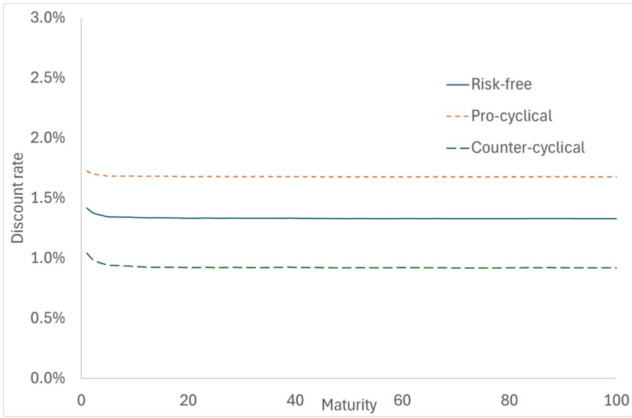

(d) One-period jump risk

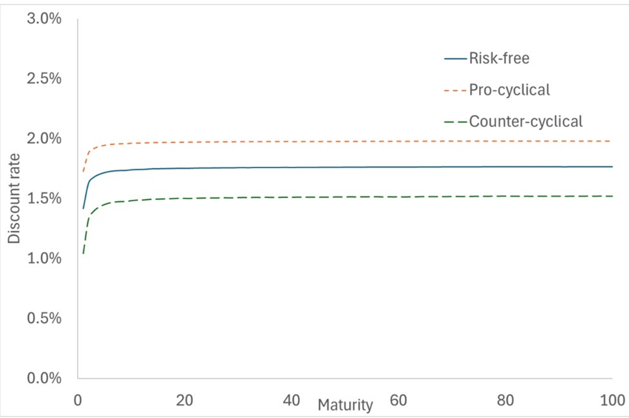

(e) One-period jump risk, with potential for recovery

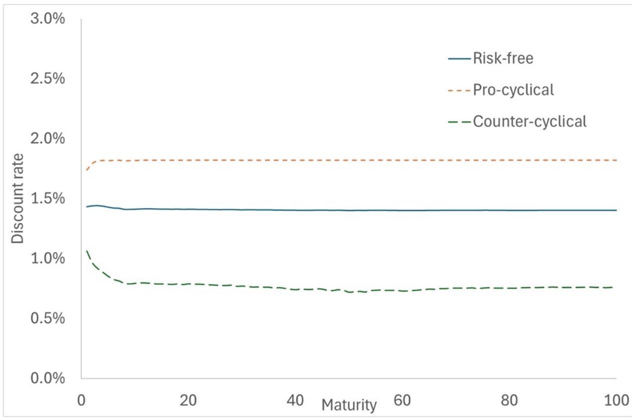

(f) Multi-period jump risk, with potential for recovery

Figure 12.1.: These figures capture three core elements described in this chapter: (i) the impact of persistency in consumption growth risk on the precautionary savings motive, (ii) CCAPM risk premia, and (iii) the impact of downside jump risk on the risk-free and risk-adjusted discount rates in an STPR environment.

12.4.11. Overall, these examples show the sensitivity of the STPR discount rate to subtle modelling choices over the way that consumption evolves over time, and how that is correlated with individual public project benefits. The Ramsey Rule, even augmented with a declining term structure of discount rates and the L term for risk, is a highly parsimonious way of modelling economic dynamics.

12.4.a. Expected benefits and PVs

12.4.12. Within these jump models, it is insufficient to consider the discount rate alone. The purpose of calculating a discount rate is not as an end in its own right, but as a part of the process to determine the project's present value. As the present value is also determined by the expected future net benefits, we also need to consider how this variable evolves in a model with stochasticity.

12.4.13. In the most basic model we have described here, the expected growth rate of the project benefits is (2.3%-0.38%)/2=0.96%. Therefore, after 100 years, with no stochasticity, the benefit would be $\exp(0.0096 \times 100) = 2.612$ for both the pro- and counter-cyclical projects.

12.4.14. However, because of Jensen's inequality, stochasticity (without changing the mean) increases this expected benefit. This is most obviously shown by the second example with full persistency in states. In this case, the benefit is either $\exp(0.023 \times 100) = 10.014$ or $\exp(-0.0038 \times 100) = 0.681$ with equal probability. Therefore, the expected net benefit is 5.347; more than twice the non-stochastic benefit.

12.4.15. Therefore, for the counter-cyclical project, there are two significant effects that both work to increase the present value. First, the discount rate declines because of the (i) precautionary savings effect on the risk-free rate, and (ii) the risk-premium becoming more negative as the insurance benefits from the project are higher. Second, the expected benefit also increases. These two effects combine to have a very large impact on the present value, particularly at long maturities.

12.4.16. Once we start to introduce the rare disaster state, then this increases the expected benefits to the counter-cyclical project, and reduces the expected-benefit to the pro-cyclical project. This further impacts on the present values of the two types of project.

12.4.17. In Figure 12.2 we plot the expected benefits for both projects for all six examples that we have described in this section, while in Figure 12.3 we show how the expected benefits and discount rate combine to give a present value. From this, it is clear that the potential for rare but severe drops in consumption significantly changes the estimated social welfare from public projects compared to the Ramsey Rule and its usual extension. This is most significantly the case for projects that offer social insurance at such times.

12.4.18. This is a feature of the Stress Discounting framework recently introduced by Cherbonnier et al. (2025). This was developed particularly for policy application in the economic appraisal of nuclear waste disposal projects in France. There are a range of other examples where such insurance characteristics might play a key part in establishing their social value. We have already highlighted projects that potentially help protect against climate and environmental disasters. Some health programmes also have this characteristic; see, for example, Rao et al. (2026). Pandemic prevention initiatives are likely to fall into this category. It was also suggested in the expert academic panels that certain types of defence spending, such as building new warships, take place for social insurance purposes.

12.4.19. For these projects, it is inappropriate to use a positive risk premium; it should instead be negative. The "L" term in the current version of the Green Book therefore risks penalising these projects 'unfairly'. Given the impact these arguments have had in setting international climate policy, we believe that similar considerations should be given to other projects that share economic properties with greenhouse gas reductions programmes. Our opinion, then, is that there is a danger that current Green Book methods understate these insurance benefits and that this should be reviewed by HM Treasury. This opinion was also expressed by a number of respondents to our survey reported in paragraph A.2.6; for example, two from government noted "In some investments such as addressing climate change, you could make the argument that this risk premium should be negative, as acting now will reduce risk in the future." and "If done properly, I would argue that risk should be captured in the appraisal through risk analysis, scenario testing etc and due to the fact this effectively discounts long-term benefits (which occur more regularly that long term costs) on the basis they might not materialise (e.g., due to war or systemic shock), even though those are precisely the states of the world where welfare would be lowest and where resilience-building investments may matter most." Also, in paragraph A.2.7 " ... project with returns that vary inversely with the state of the economy - such as a mental health facility - should have a negative risk premium added to the 'standard' discount rate".

12.4.20. Departments should not be allowed to "game"

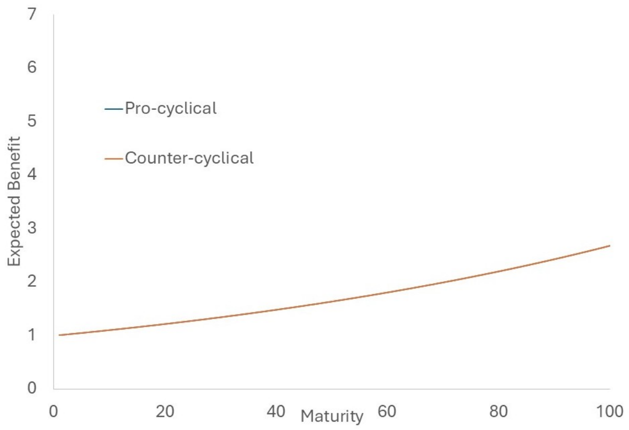

(a) No jump risk, no persistence in consumption growth

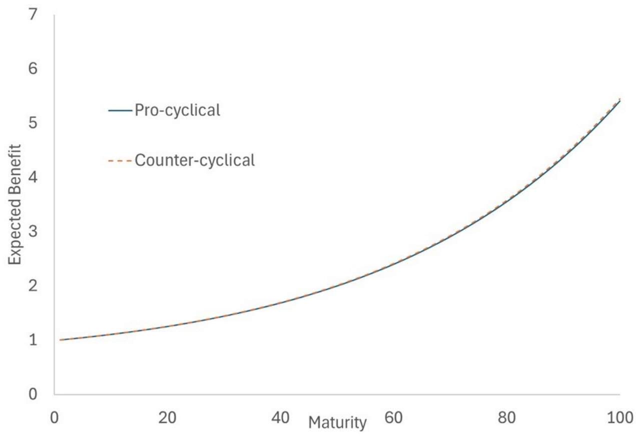

(b) No jump risk, full persistence in consumption growth

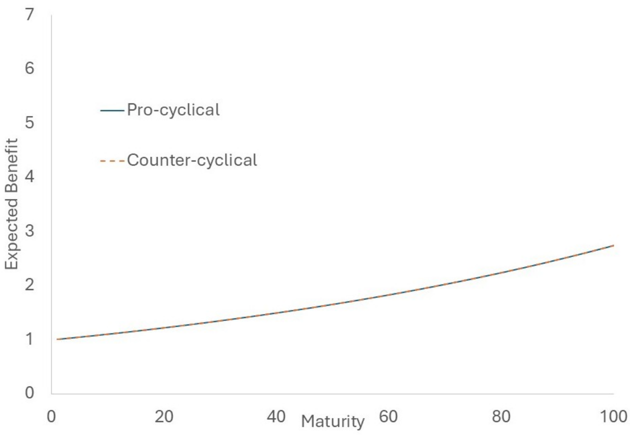

(c) No jump risk, partial persistence in consumption growth

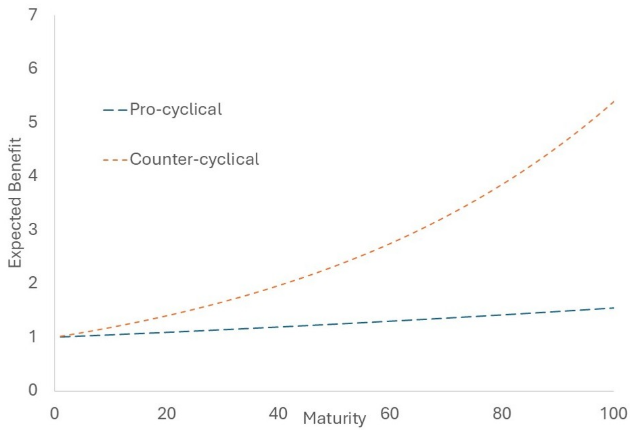

(d) One-period jump risk

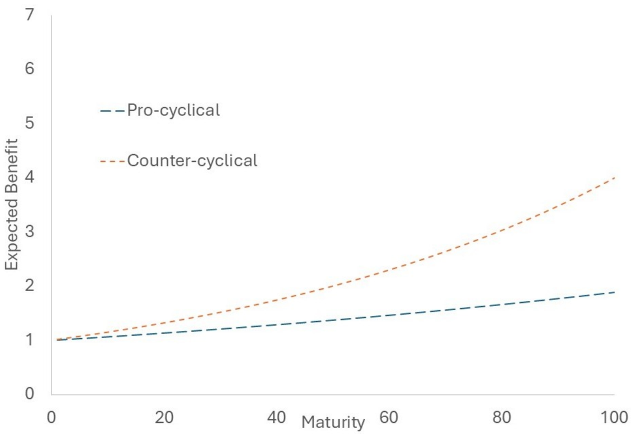

(e) One-period jump risk, with potential for recovery

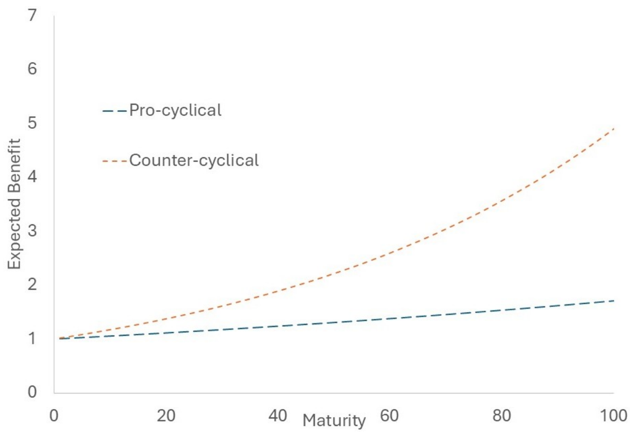

(f) Multi-period jump risk, with potential for recovery

Figure 12.2.: This shows the expected net benefits from the pro- and counter-cyclical projects for the six economies described in Figure 12.1.

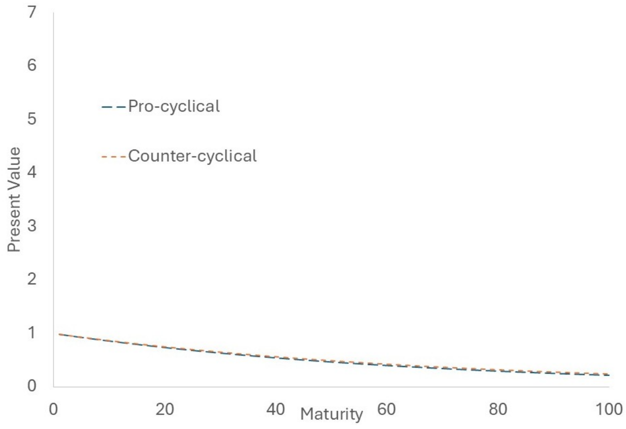

(a) No jump risk, no persistence in consumption growth

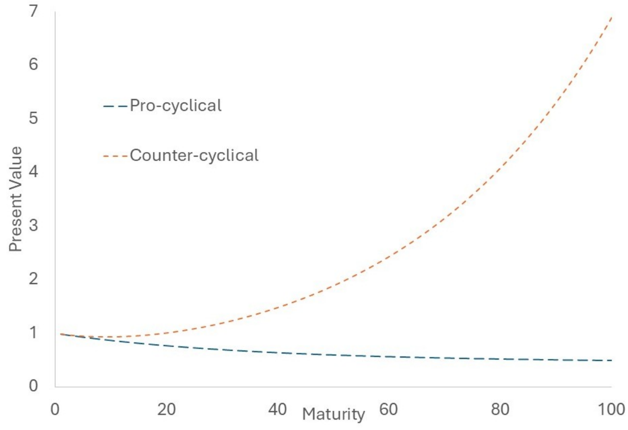

(b) No jump risk, full persistence in consumption growth

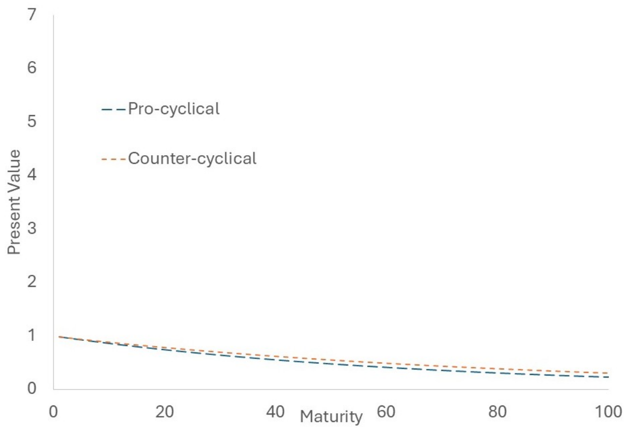

(c) No jump risk, partial persistence in consumption growth

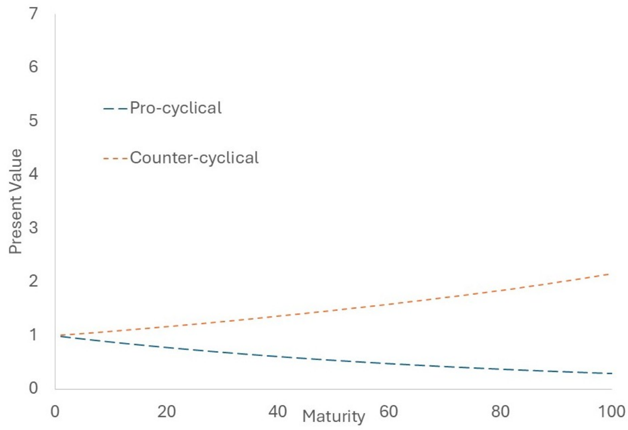

(d) One-period jump risk

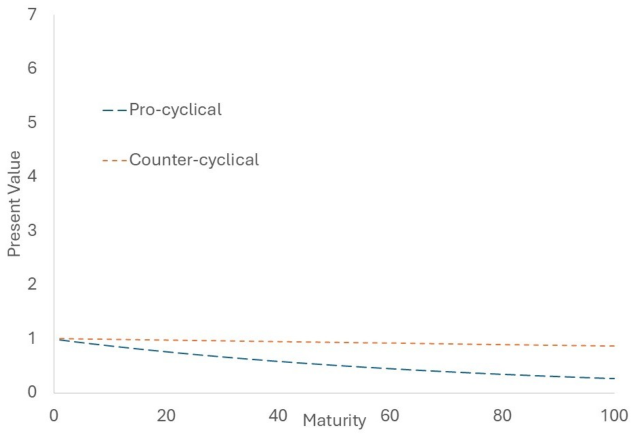

(e) One-period jump risk, with potential for recovery

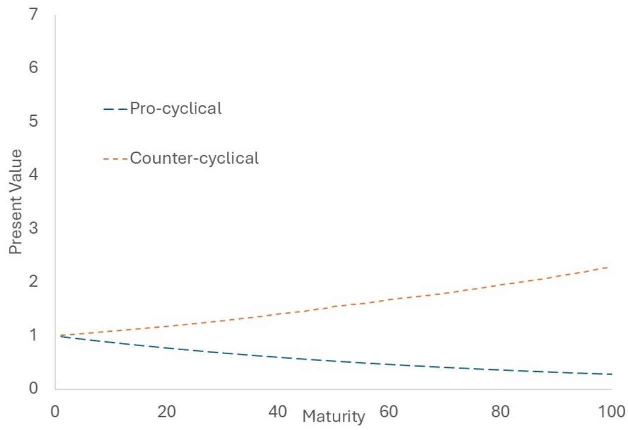

(f) Multi-period jump risk, with potential for recovery

Figure 12.3.: This shows the present values of the pro- and counter-cyclical projects for the six economies described in Figure 12.1.

this recommendation by trying to use the lower social insurance rate for projects that do not genuinely have this property. This recommendation should only be used in exceptional cases, and always with the explicit permission of HM Treasury beforehand.

R.15. Social insurance projects should have a discount rate below the average.

We recommend that any public project that delivers high social benefits when there is a significant negative shock to aggregate consumption should have a lower discount rate than is applied elsewhere through the Green Book. This, though, should be an exception, with explicit permission required form HM Treasury before this is actioned.

12.4.21. That rare catastrophe risk should be captured in the STPR model, lowering the risk-free rate while raising the average social risk premium, has been heavily influential in French guidance within an STPR framework; see France Stratégie (2023).

12.4.22. They apply a relatively complex model of rare disaster risk based on the Pareto model of Barro & Jin (2011). By choosing a relatively high value for $\mu = 2.478$ and other modelling choices, they derive an average risk premium of 2%. In our opinion, this is likely to be an upper estimate of the STPR risk premium given that we will be using a lower value of μ. Taking Panel (f) in Figure 12.1, for example, suggests a risk premium closer to 0.4% than 2%. We will calibrate this more precisely in Chapter 14.

12.4.23. We note that, taking 0.4% as a provisional estimate, plus another 0.1% for standard CCAPM effects, results in a risk premium still substantially below the 1.9% central estimate we have taken from an SOC approach in Chapter 11. The survey evidence from (Gollier et al. 2023) and our own survey sits towards the centre of these two values at c.1%, while the risk premium used in international guidance spans a wide range. As we emphasised in Chapter 6, we are not expecting that the STPR and SOC rates perfectly align — there are many reasons why this will not be the case. Yet, this analysis emphasises that the SOC-STPR divide is much narrower than an initial reading of the equity premium puzzle literature might suggest.

R.16. STPR discount rate models should be informed by rare disaster risk.

We recommend that HM Treasury extends the STPR Ramsey Rule model to incorporate rare disaster risk. This will result in a lowering of the risk-free rate, an increase in the average risk premium, and a case for applying negative risk premia in exceptional cases to projects that provide the greatest social insurance at times of economic collapse.

13. Declining Discount Rates: Empirical Evidence in in the UK

The UK Green Book has used declining discount rates for more than two decades, but the schedule has not been calibrated from UK consumption-growth data. Building on the work in Chapter 9, this chapter asks which extensions of the Ramsey rule can empirically justify that policy using updated GB+NI per capita consumption growth from 1830–2025. IID growth, AR(1) persistence, recurrent Markov regimes, GARCH volatility, and UK-calibrated disaster risk generate little long-run decline. Meaningful declining schedules arise only when historical growth variation is interpreted as unresolved uncertainty over persistent future states or low-frequency growth environments. Under $\delta = 0.5\%$ and $\mu = 1.3$, and normalising the candidate policy schedules to a 3.15% short forward rate, several plausible models decline more slowly than the current Green Book schedule. The Bai–Perron and Stock–Watson-style calibrations remain materially above the most aggressive long-run Green Book rates, while only stronger structural-uncertainty calibrations approach the 2003 schedule.

13.1. Introduction

13.1.1. The case for declining discount rates (DDRs) in public project appraisal is now a central part of the economics of long-run policy. The starting point remains the Ramsey rule, which links the consumption discount rate to pure time preference, expected consumption growth, and the elasticity of marginal utility (Ramsey 1928). Under certainty this delivers a constant discount rate. Under uncertainty, however, the certainty-equivalent discount rate can decline with the horizon because future states are not all valued symmetrically: payoffs in low-consumption states have high marginal social value, and long-horizon uncertainty changes the balance between expected growth and prudence. This is the central insight behind the modern DDR literature.

13.1.2. The literature gives several distinct reasons why long-run social discount rates may decline. One strand starts from uncertainty over future discount rates or interest rates, as in Weitzman (1998, 2001) and Newell & Pizer (2003). This is the tradition that most directly influenced early policy discussion of DDRs. A second strand derives DDRs from uncertainty in consumption growth, especially the work of Gollier (2002, 2008) and the synthesis in Gollier & Hammitt (2014). In that framework, the stochastic process for growth determines the term structure of certainty-equivalent discount rates. A third strand emphasises that the treatment of risk, ambiguity, and intertemporal risk aversion matters for the discount rate, as in Traeger (2014). A fourth strand focuses on normative disagreement and aggregation, including the nondogmatic social discounting approach of Millner (2020) and the broader review by Millner & Heal (2023). These contributions are complementary rather than substitutes: they show that DDRs can arise from market uncertainty, growth uncertainty, preference uncertainty, or disagreement about ethical primitives.

13.1.3. The policy literature has therefore moved from asking whether declining discounting is theoretically possible to asking which mechanism is empirically relevant for public CBA. Groom et al. (2005) review the early theoretical and empirical rationales for DDRs and emphasise the importance of distinguishing between them. Arrow et al. (2013) similarly argue that DDRs can be justified for long-term project analysis when future rates are uncertain and persistent. Cropper et al. (2014) focus on the practical problem of estimating DDR schedules for use by policy analysts. More recent reviews, including Freeman et al. (2018) and Millner & Heal (2023), stress that the choice of discount rate combines empirical evidence, market information, and normative judgement. This chapter contributes to that practical calibration problem for the UK.

13.1.4. The UK Green Book schedule is historically important in this respect, but its declining schedule was not calibrated from UK consumption data. It emerged from the wider DDR debate, especially the argument in Weitzman (2001) and the interest-rate evidence in Newell & Pizer (2003). That was a reasonable first policy step, but it leaves open the question of what an updated UK consumption-based calibration implies. The question matters directly for the current HM Treasury review of the discount rate. A UK-specific empirical exercise already existed in an earlier working paper by Groom et al. (2004), and comparative evidence across countries was subsequently reported by Hepburn et al. (2009). Yet the policy debate has continued to rely more on stylised or international arguments than on a transparent, updated UK calibration grounded in per capita consumption growth.

13.1.5. This chapter undertakes that calibration using updated GB+NI consumption and population data through 2025. It begins from a general stochastic discount factor and makes explicit that the implied term structure depends on the stochastic process for consumption growth. It then considers growth deviations around trend, IID Gaussian growth, AR(1) growth, recurrent Markov regimes, permanent parameter uncertainty, Bai–Perron structural breaks, Stock–Watson-style low-frequency blocks, and Barro-type disaster risk. The emphasis is deliberately comparative: each model corresponds to a different interpretation of the uncertainty that might justify DDRs.

13.1.6. The main empirical result is straightforward. Short-run persistence and recurrent regimes in UK per capita consumption growth do not, by themselves, support a strongly declining term structure. Much steeper decline appears only when the long-run mean and variance of growth are treated as structurally uncertain. The 2025 update also strengthens the need for careful interpretation: the recent sample contains the COVID-19 collapse and rebound, which can create a near-degenerate low-growth regime if mechanically interpreted as a permanent future state. The main results therefore focus on continuous truncated, Bai–Perron, and Stock–Watson-style calibrations of permanent parameter uncertainty, with the direct discrete regime-informed versions reported as appendix diagnostics and robustness.

13.2. A general stochastic discount factor

13.2.1. Let annual log growth in per capita consumption be $g_t = \Delta \ln c_t$ and cumulative growth over horizon T be

$$S_T = \sum_{j=1}^{T} g_{t+j}.$$

With isoelastic utility, pure utility discount rate δ, and elasticity of marginal utility μ (matching Green Book notation), the certainty-equivalent discount factor for a risk-free payoff at date T is

$$D_T = e^{-\delta T}\, \mathbb{E}\left[e^{-\mu S_T}\right].$$

The associated spot and one-period forward rates are

$$R_T = -\frac{1}{T}\ln D_T, \qquad F_T = -\ln\left(\frac{D_{T+1}}{D_T}\right).$$

The entire problem is therefore reduced to one question: what stochastic process should be used for growth? Different answers correspond to different economic interpretations of uncertainty, persistence, and tail risk.

13.2.2. Before turning to stochastic processes, it is useful to note Gollier's (2012) growth-deviations mechanism. In that setting the discount rate moves with the deviation of current growth from long-run trend. The point is conceptual: discount rates can vary because the economy is above or below trend, even in the absence of a fully stochastic long-run process. This provides a useful bridge from the deterministic Ramsey rule to the stochastic models that follow.

13.3. Model classes

13.3.1. This section separates the theoretical discounting mechanism from the empirical model used to estimate it. The starting point is the stochastic discount factor

$$D_T = e^{-\delta T}\mathbb{E}\left[e^{-\mu S_T}\right], \qquad S_T = \sum_{j=1}^{T} g_{t+j}.$$

If cumulative growth S_T is Gaussian with mean $\mathbb{E}[S_T]$ and variance $\mathrm{Var}(S_T)$, then

$$R_T = \delta + \frac{\mu}{T}\mathbb{E}[S_T] - \frac{\mu^2}{2T}\,\mathrm{Var}(S_T).$$

This expression is useful because it shows where the term structure comes from. The first term is pure time preference. The second is the expected-growth term. The third is the prudence term as discussed in Chapter 9. The way the prudence term changes with horizon depends on how the cumulative variance grows with T. In the IID case, cumulative variance grows linearly in T, so $\mathrm{Var}(S_T)/T$ is constant and the term structure is flat. In persistent-growth models, cumulative variance grows nonlinearly in T, so the prudence term varies with the horizon. In parameter-uncertainty models, uncertainty about permanent future growth states generates an additional long-horizon effect because the uncertainty is not resolved over time.

13.3.a. Growth-process models

13.3.2. The first class contains models in which the dynamics of growth itself determine the term structure. These models differ mainly in how cumulative expected growth and cumulative variance evolve with the horizon.

13.3.3. **Growth deviations around trend.** The simplest state-dependent model is the growth-deviations mechanism. Suppose current growth is x_0, long-run trend growth is $\bar{g}$, and growth mean-reverts deterministically according to

$$x_t = \bar{g} + \phi^t (x_0 - \bar{g}).$$

Then cumulative expected growth over horizon T is the finite sum of these deviations. Substituting this into the Ramsey Rule gives

$$R_T = \delta + \mu \left[\bar{g} + (x_0 - \bar{g}) \frac{1 - \phi^T}{T(1 - \phi)} \right].$$

This is a time-varying discount-rate model, but not a model of stochastic uncertainty.

13.3.4. IID Gaussian growth. The benchmark stochastic model assumes

$$g_t = m + \sigma \varepsilon_t, \qquad \varepsilon_t \sim N(0, 1),$$

with independent growth shocks. Then cumulative growth is Gaussian:

$$S_T \sim N(mT, \sigma^2 T).$$

The cumulative variance is linear in the horizon:

$$\mathrm{Var}(S_T) = \sigma^2 T.$$

Hence $\mathrm{Var}(S_T)/T = \sigma^2$, and the prudence term is constant. The certainty-equivalent discount factor is

$$D_T = \exp \left(-\delta T - \mu m T + \frac{1}{2} \mu^2 \sigma^2 T \right),$$

and the spot rate is

$$R_T = \delta + \mu m - \frac{1}{2} \mu^2 \sigma^2.$$

The term structure is flat.

13.3.5. AR(1) growth. The next model introduces persistence in growth:

$$g_t = m + \phi(g_{t-1} - m) + u_t, \qquad u_t \sim N(0, \sigma_u^2).$$

Under stationarity, cumulative growth remains Gaussian, so

$$R_T = \delta + \mu m - \frac{\mu^2}{2T} \mathrm{Var}(S_T),$$

where

$$\mathrm{Var}(S_T) = \gamma_0 \left[T + 2 \sum_{k=1}^{T-1} (T - k) \phi^k \right], \qquad \gamma_0 = \frac{\sigma_u^2}{1 - \phi^2}.$$

Unlike the IID case, $\mathrm{Var}(S_T)/T$ is not constant. Positive persistence raises long-horizon cumulative uncertainty and therefore strengthens the prudence correction.

13.3.6. Markov regime shifts. A more flexible recurrent-state model lets growth depend on a latent Markov state $s_t \in \{1, \ldots, K\}$:

$$g_t \mid s_t = i \sim N(m_i, \sigma_i^2), \qquad \Pr(s_t = j \mid s_{t-1} = i) = p_{ij}.$$

Let P denote the transition matrix and π the vector of stationary probabilities. Define the state-contingent pricing-kernel matrix

$$M = \mathrm{diag} \left(e^{-\mu m_i + \frac{1}{2} \mu^2 \sigma_i^2} \right).$$

Then

$$D_T = e^{-\delta T} \pi' M (PM)^{T-1} \mathbf{1}.$$

This model allows both the mean and variance of growth to vary across temporary recurrent states. But because the regimes are recurrent, a bad state today is expected eventually to be followed by other states.

13.3.b. Permanent parameter uncertainty

13.3.7. The second class follows the parameter-uncertainty logic of Gollier (2008). Here the uncertain objects are not temporary states but the permanent parameters of the future growth process. Let m denote the permanent mean and $v = \sigma^2$ the permanent variance. Conditional on (m, v), growth is Gaussian, but (m, v) itself is uncertain at date 0. The discount factor is

$$D_T = e^{-\delta T} \mathbb{E}_{m,v} \left[e^{-\mu m T + \frac{1}{2} \mu^2 v T} \right].$$

This is the theoretical model that generates the material DDRs. The empirical question is how to calibrate the distribution of (m, v) using historical UK data.

13.3.8. Regime-informed and continuous calibrations. A low-dimensional representation replaces the full distribution of (m, v) with a finite number of permanent states:

$$D_T = e^{-\delta T} \sum_i q_i \exp \left(-\mu m_i T + \frac{1}{2} \mu^2 v_i T \right), \qquad v_i = \sigma_i^2.$$

The preferred continuous specification is a truncated distribution over $x = (m, \log v)$. The support is truncated so that permanent mean growth is bounded below by zero in the main specification. The discount factor is evaluated by simulation:

$$D_T = e^{-\delta T} \mathbb{E}_x \left[e^{-\mu m T + \frac{1}{2} \mu^2 e^z T} \right].$$

The calibration uses recent UK regime estimates to discipline the location and spread of the distribution. Direct discrete regime-informed permanent-state calibrations are reported in the appendix because of the COVID-induced near-degenerate regime.

13.3.9. Bai–Perron and Stock–Watson-style low-frequency calibrations. The Bai–Perron calibration estimates historical changes in the mean of UK growth:

$$g_t = m_j + \varepsilon_t, \qquad t \in \mathcal{T}_j.$$

Each break segment has an estimated mean m_j and variance v_j. These segment-level estimates are then used as support points in the permanent parameter uncertainty model:

$$D_T = e^{-\delta T} \sum_j q_j \exp\left(-\mu m_j T + \frac{1}{2}\mu^2 v_j T\right).$$

Motivated by the Stock–Watson emphasis on low-frequency movements in macroeconomic growth, this chapter also considers a block-based calibration. The historical series is divided into long non-overlapping blocks, each providing an empirical estimate of a low-frequency growth environment.

13.3.c. Disaster-risk models

13.3.10. The final class is based on rare disasters. This is separate from both ordinary growth persistence and parameter uncertainty. The object of interest is left-tail risk.

13.3.11. Constant rare-disaster benchmark. The Barro model writes growth as

$$g_t = m + \sigma\varepsilon_t + J_t,$$

with

$$J_t = \begin{cases} \ln(1-b), & \text{with probability } p, \\ 0, & \text{with probability } 1-p. \end{cases}$$

If the normal component is recentered so that unconditional mean growth is unchanged, the disaster adjustment to the spot rate is

$$\Delta r_{\text{disaster}} = \mu p \ln(1-b) - \ln\left[(1-p) + p(1-b)^{-\mu}\right].$$

This is the central insight of Barro (2006): rare left-tail events can materially reduce the risk-free rate.

13.3.12. Time-varying disaster risk. If state i has normal-growth parameters (m_i, v_i) and disaster probability p_i, the state-contingent pricing kernel is

$$M_i = \exp\left(-\mu m_i + \frac{1}{2}\mu^2 v_i\right)\left[(1-p_i) + p_i(1-b)^{-\mu}\right].$$

The discount factor is then

$$D_T = e^{-\delta T}\pi' M (PM)^{T-1}\mathbf{1}.$$

This links the discounting problem to the broader macro-finance literature on time-varying disaster risk.

13.4. Data and empirical methods

13.4.1. The annual GB+NI consumption and population series run from 1830 to 2025, producing a per capita consumption-growth series for 1831–2025:

$$g_t = 100\,\Delta \ln(C_t/N_t),$$

as shown in Figure 14.1. Consumption is measured as a composite volume measure in 2022 reference-year prices, and population is measured for GB+NI. The main estimation window is the most recent 50 years, 1976–2025, with robustness demonstrated using the last 30 years and the full 1831–2025 growth series.[1]

13.4.2. All estimates are produced in Python using the accompanying replication code. Data handling and construction of growth rates use `pandas` and `numpy`. Linear Gaussian growth models are estimated using `statsmodels` as are the diagnostic tests and numerical calculations, including simulation of discount factors and low-frequency calibrations. All reported rates are annual percentage rates. The figures show forward rates, the tables show spot rates.

13.4.3. Growth-process models. For the IID Gaussian model, the mean m and variance $v = \sigma^2$ of annual per capita consumption growth are estimated directly from the relevant sample:

$$\widehat{m} = \frac{1}{n}\sum_t g_t, \qquad \widehat{v} = \frac{1}{n}\sum_t (g_t - \widehat{m})^2.$$

The IID discount factor is then evaluated analytically using the Gaussian moment-generating function.

13.4.4. The AR(1) model is estimated as

$$g_t = c + \phi g_{t-1} + u_t, \qquad u_t \sim N(0, \sigma_u^2).$$

In the code this is estimated by conditional Gaussian maximum likelihood, numerically equivalent to least squares on the AR(1) regression using `statsmodels`. The long-run mean is recovered as

$$\widehat{m} = \frac{\widehat{c}}{1-\widehat{\phi}},$$

[1] The historical series is checked for a possible territorial-boundary issue around the early 1920s, when the UK statistical boundary changed following the creation of the Irish Free State. In an earlier whole-UK series this appeared mainly as a population-denominator break around 1923. In the updated GB+NI series the 1923 population movement is much smaller and the large 1921 fall is driven by consumption rather than the population denominator. Appendix Table C.8 reports the diagnostic. Because the updated series appears harmonised, no baseline correction is applied; possible robustness corrections would include smoothing or omitting the 1922–1923 observation.

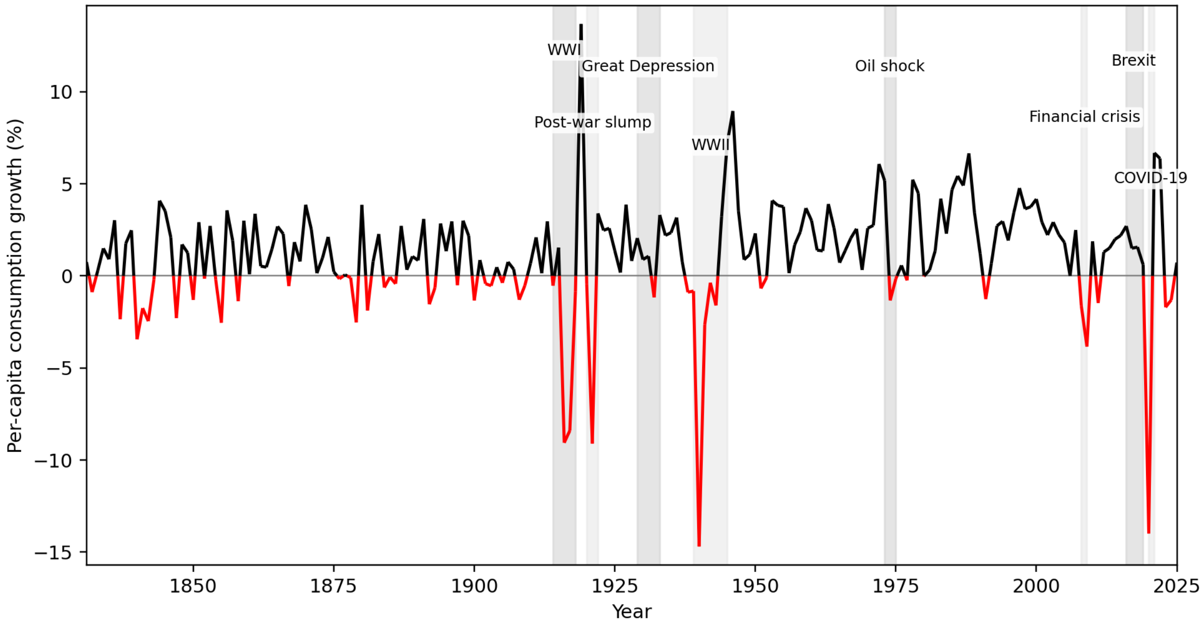

Figure 13.1.: UK per capita consumption growth, 1831–2025. Negative observations are shown in red, while shaded bands mark selected historical episodes.

and the innovation variance is estimated from the fitted residuals. The implied term structure is then computed from the analytical cumulative-variance formula for $S_T = \sum_{j=1}^{T} g_{t+j}$.

13.4.5. The growth-deviations model uses the estimated AR(1) persistence parameter and long-run mean, together with the terminal growth observation in 2025. It is therefore a current-state calculation rather than a stochastic parameter-uncertainty model.

13.4.6. The two-state and three-state Markov regime-switching models estimate

$$\begin{aligned} g_t \mid s_t &= i \sim N(m_i, v_i) \\ v_i &= \sigma_i^2, \\ \Pr(s_t = j \mid s_{t-1} = i) &= p_{ij} \end{aligned}$$

The models are estimated by maximum likelihood using `MarkovRegression` with switching variances and multiple starting values. The estimated states are then ordered by mean growth. The output used in the chapter includes the state means m_i, variances v_i, transition probabilities p_{ij}, and implied stationary probabilities π_i.

13.4.7. The updated 2025 sample creates one important estimation issue. In the raw recent-50-year regime-switching estimates, the likelihood isolates the 2020 pandemic contraction as a near-degenerate low-growth state. This is a reasonable statistical description of a temporary intervention-like observation, but it is not a credible permanent future growth environment. Consequently, the raw two- and three-state discrete permanent-state calibrations are not used as main policy schedules. For the illustrative discrete three-state comparison, we instead use a COVID-adjusted regime estimate that excludes 2020 from the regime-switching calibration. The raw estimates are reported in Appendix C as a diagnostic.

13.4.8. Permanent parameter uncertainty calibrations. The permanent parameter uncertainty models require an empirical distribution over permanent mean and variance pairs (m, v). The common discounting object is

$$D_T = e^{-\delta T} \mathbb{E}_{m,v} \left[e^{-\mu m T + \frac{1}{2}\mu^2 v T} \right],$$

or its discrete analogue,

$$D_T = e^{-\delta T} \sum_j q_j \exp\left(-\mu m_j T + \frac{1}{2}\mu^2 v_j T \right).$$

The different calibrations differ in how the support points (m_j, v_j) and weights q_j are obtained.

13.4.9. The regime-informed discrete calibration uses the estimated regime means, variances, and stationary probabilities as support points and weights. This is a deliberate reinterpretation: the Markov model estimates recurrent states, whereas the permanent-parameter model treats analogous mean–variance combinations as possible permanent future growth environments. Because the raw recent-sample regime estimates are distorted by the pandemic state, the main discrete three-state comparison uses the COVID-adjusted regime estimates described above.

13.4.10. The continuous truncated calibration uses recent regime estimates to discipline a continuous distribution over

$$x = (m, \log v).$$

The log-variance transformation ensures $v > 0$. The distribution is centred using the regime-weighted mean and covariance of $(m, \log v)$. The main specification truncates permanent mean growth at zero from below and bounds $\log v$ using the estimated range of regime variances. The discount factor is evaluated by simulation:

$$D_T = e^{-\delta T} \frac{1}{B} \sum_{b=1}^{B} \exp\left(-\mu m_b T + \frac{1}{2}\mu^2 v_b T\right),$$

where (m_b, v_b) are draws from the truncated distribution. This is the preferred permanent-uncertainty calibration because it preserves the low-growth tail without mechanically assigning permanent probability to a single extreme historical observation.

13.4.11. The Bai–Perron calibration estimates structural breaks in the historical UK growth series and treats the resulting segment-level mean–variance estimates as support points for possible permanent future growth environments. The break search is implemented directly in the replication code by dynamic programming over the residual sum of squares of piecewise-constant mean models, with a minimum segment length imposed and the number of breaks selected by BIC. For segment j, the estimated support point is

$$(\widehat{m}_j, \widehat{v}_j),$$

and the baseline weight is the segment-length weight,

$$q_j = N_j / N.$$

13.4.12. The Stock–Watson-style low-frequency calibration divides the historical series into long non-overlapping blocks. The baseline implementation uses 16-year blocks, with 10-year and 25-year blocks reported as robustness checks. For each block j, the block mean and variance are estimated directly:

$$\widehat{m}_j = \frac{1}{N_j} \sum_{t \in \mathcal{B}_j} g_t, \qquad \widehat{v}_j = \frac{1}{N_j} \sum_{t \in \mathcal{B}_j} (g_t - \widehat{m}_j)^2.$$

These block-level estimates are then interpreted as support points in the same permanent-parameter uncertainty model. This is a low-frequency calibration exercise rather than a claim that future growth follows fixed-length cycles.

13.4.13. **Disaster-risk and GARCH robustness.** For the Barro models, the constant rare-disaster benchmark uses the standard global calibration

$$p = 0.017, \qquad b = 0.29,$$

where p is the annual disaster probability and b is the proportional consumption loss in the disaster state. The constant-disaster model is evaluated analytically as a level adjustment to the risk-free rate.

13.4.14. The time-varying disaster specification uses a two-state Markov structure in which disaster probabilities differ across states. Disaster-like observations are identified as annual contractions exceeding 5 percent:

$$d_t = \mathbf{1}\{g_t < -5\%\}.$$

Let $\xi_{ti} = \Pr(s_t = i \mid \text{data})$ denote the smoothed probability that observation t belongs to state i. State-specific disaster probabilities are estimated using a shrinkage frequency:

$$\widehat{p}_i = \frac{\sum_t \xi_{ti} d_t + \kappa p_0}{\sum_t \xi_{ti} + \kappa},$$

where $p_0 = 0.017$ is the global Barro benchmark and κ is a pseudo-sample weight. The jump size b is kept fixed at the global Barro calibration. This specification is intended as a reduced-form UK illustration of time-varying disaster risk rather than a full structural estimate of disaster dynamics.

13.4.15. Appendix C also estimates an AR(1)-GARCH(1,1) model to test whether a separate conditional-volatility process materially alters the main conclusions. The model is

$$g_t = c + \phi g_{t-1} + \varepsilon_t, \qquad \varepsilon_t = \sqrt{h_t} z_t, \qquad z_t \sim N(0,1),$$

with

$$h_t = \omega + \alpha \varepsilon_{t-1}^2 + \beta h_{t-1}.$$

The model is estimated by Gaussian likelihood in the replication code, and the implied discount factors are computed numerically by simulation from

$$D_T = e^{-\delta T} \mathbb{E}\left[e^{-\mu \sum_{j=1}^{T} g_{t+j}}\right].$$

Finally, the appendix reports Ljung–Box tests for serial correlation, ARCH-LM tests for conditional heteroskedasticity, and AIC/BIC statistics for the Markov regime-switching models.

13.5. Results

13.5.1. Table 13.1 reports the 2025 results under $\delta = 0.5\%$ and $\mu = 1.3$. There are four key findings.

13.5.2. First, the pure growth-process models generate only limited decline. The IID Gaussian model is flat at 2.780%. The AR(1) model falls only from 2.812% at one year to 2.790% at 300 years, a decline of about 2 basis points. The two-state and three-state recurrent regime models are also essentially flat. This is an important result: estimating persistence in UK per capita consumption growth does not by itself generate a Green-Book-style declining term structure.

13.5.3. Second, meaningful DDRs arise from the permanent parameter uncertainty calibrations. The continuous truncated model falls from 2.637% to 1.310%, a decline of 1.33 percentage points. The Bai–Perron full-sample calibration gives a more moderate decline, from 1.960% to 1.287%. The Stock–Watson-style low-frequency block calibration falls from 1.959% to 0.898%. These calibrations differ in how they use historical data, but they share the same theoretical interpretation: historical mean–variance environments are treated as possible permanent future growth environments.

13.5.4. Third, the direct discrete regime-informed permanent-state calibrations are unstable in the 2025 update. The recent sample includes the COVID-19 collapse and rebound. The Markov regime model can isolate the 2020 observation as a near-degenerate low-growth state. If that state is mechanically treated as a permanent future growth environment, the resulting long-horizon rates become implausibly negative. For this reason the raw discrete regime-informed permanent-state schedules are reported in Appendix Table C.1, while the main text focuses on the continuous truncated and low-frequency calibrations.

13.5.5. Fourth, the Barro-type disaster models mostly affect the level rather than the slope. The constant rare-disaster benchmark gives a rate of 1.074%. The time-varying disaster specification is almost flat, moving from 2.569% to 2.573%. This confirms the earlier conclusion: UK-calibrated disaster probabilities do not generate a Green-Book-style declining term structure unless disaster-risk states are both severe and persistent.

13.5.6. Figure 13.2 shows the current Green Book comparison under the older parameterisation $\delta = 1.5\%$, $\mu = 1$. To compare slopes rather than levels, all schedules are shifted so that the initial one-year forward rate is 3.5 percent. The three-state calibration shown in this figure uses a COVID-adjusted re-estimation that excludes 2020, because the raw recent-sample regime model isolates the pandemic collapse as a near-degenerate one-year state. Figure 13.3 reports results for $\delta = 0.5\%$ and $\mu = 1.3$. For project appraisal, the comparison is normalised to start at 3.15 percent rather than the raw Ramsey benchmark of 2.45 percent. This reflects an illustrative adjustment for catastrophic risk and a positive project risk premium. The Green Book schedule is omitted from the new-recommendation figure so that the figure focuses on the candidate term-structure shapes.

13.6. Conclusion

13.6.1. The updated GB+NI calibration confirms the central result of the chapter: ordinary stochastic growth dynamics do not generate a strongly declining term structure of social discount rates. IID Gaussian growth is flat; AR(1) persistence produces only a very small decline; recurrent Markov regime-switching models are also essentially flat once states are treated as temporary rather than permanent. UK-calibrated disaster risk mainly changes the level of the risk-free rate rather than the slope of the term structure.

13.6.2. More substantial decline arises only when the long-run mean and variance of future growth are treated as structurally uncertain. The continuous truncated model, the Bai–Perron break-informed calibration and the Stock–Watson-style low-frequency block calibration therefore provide the central policy evidence. These models differ in how strongly they interpret historical low-growth and high-variance episodes as information about future permanent growth environments, and this judgement is what drives the resulting discount-rate schedule.

13.6.3. The updated 2025 results suggest that the current Green Book term structure is steeper than many plausible UK-based calibrations. When schedules are normalised to the current 3.5 percent short rate, several models decline more slowly than the Green Book schedule, falling to roughly 2.75–3.0 percent after around 100 years rather than to 2.5 percent, and then flattening out at higher long-run rates. Only more extreme forms of structural uncertainty, such as stronger permanent-state interpretations of low-frequency growth risk, come close to the decline embedded in the 2003 Green Book schedule.

13.6.4. With $\delta = 0.5\%$ and $\mu = 1.3$, the raw Ramsey benchmark with expected growth of 1.5 percent is 2.45 percent. The policy comparison in the chapter instead normalises the candidate schedules to 3.15 percent to allow for a lower catastrophe-adjusted risk-free rate and a positive project risk premium. On that basis, the Bai–Perron and Stock–Watson-style calibrations still imply more gradual decline than the most aggressive Green Book shape, while the continuous truncated model gives a stronger but still transparent structural-

Model	R_1	R_{30}	R_{75}	R_{125}	R_{300}	$R_1 - R_{300}$
Panel A. Growth-process models						
Gollier (2012) growth deviations: terminal vs estimated trend	1.402	2.755	2.787	2.795	2.802	-1.400
IID Gaussian growth	2.695	2.695	2.695	2.695	2.695	0.000
AR(1) growth	2.726	2.707	2.706	2.706	2.706	0.020
Two-state Markov regime model	2.857	2.967	2.969	2.970	2.971	-0.114
Three-state Markov regime model	2.638	2.660	2.660	2.659	2.659	-0.021
Panel B. Permanent parameter uncertainty: alternative UK calibrations						
Regime-informed continuous truncated permanent uncertainty	2.581	2.325	2.008	1.756	1.302	1.279
Bai–Perron break-informed permanent uncertainty	1.906	1.776	1.607	1.473	1.266	0.640
Stock–Watson-style low-frequency block calibration	1.906	1.689	1.428	1.234	0.903	1.003
Panel C. Disaster-risk models						
Constant rare-disaster benchmark, global calibration	1.063	1.063	1.063	1.063	1.063	0.000
Time-varying disaster probability	2.497	2.626	2.629	2.630	2.630	-0.133

Table 13.1.: UK social discount-rate term structures by model class (spot rates, percent). R_T denotes the continuously compounded spot discount rate to horizon T, reported in percent. Results use the updated GB+NI consumption and population data through 2025. The main calibration uses $\delta = 0.5\%$ and $\mu = 1.25$. Growth-process and continuous truncated models use the recent 50-year sample; Bai–Perron and Stock–Watson-style low-frequency calibrations use the full 1831–2025 growth series. The normalisation to 3 percent is applied only to the Figure 13.3 forward-rate schedules and is not applied to the unnormalised spot rates reported in this table.

Table 13.2.: Normalised old-parameter forward-rate term points (percent). All schedules are shifted to start at 3.5 percent. The three-state schedule uses a COVID-adjusted re-estimation excluding 2020.

Schedule	F_1	F_{30}	F_{75}	F_{125}	F_{300}	F_{350}
Green Book	3.500	3.500	3.000	2.500	1.500	1.000
Continuous truncated	3.500	3.182	2.811	2.548	2.179	2.134
COVID-adjusted three-state	3.500	2.597	1.902	1.669	1.570	1.569
Bai–Perron	3.500	3.335	3.131	2.989	2.853	2.847
Stock–Watson block	3.500	3.229	2.928	2.737	2.485	2.448

uncertainty case for DDRs. The policy conclusion is therefore not that DDRs lack empirical support, but that the degree of decline should be tied explicitly to the interpretation of long-run growth uncertainty.

13.6.5. As has been shown in this chapter, the precise speed with which the discount rate declines is highly model specific. The evidence strongly suggests that the current decline in discount rates is too steep, though. Having evaluated this evidence, we recommend that the forward discount rate declines by 0.5% at 31 years and then by a further 0.25% at 76 years, with it remaining flat thereafter. While in the very long-term the term structure may decline slightly more, our other recommendations are to treat extremely long term projects differently than simple discounted NPV analysis.

13.6.6. Finally, while there is also a literature that the risk premium should have a non-flat term structure (Gollier 2014), and examples of such effects were presented in Chapter 12, we feel it is not practical for HM Treasury to try to implement such a term structure at this stage.

R.17. The forward social discount rate should decline over time.

We recommend that HM Treasury retains a declining term-structure of the risk-free discount rate, with the forward rate dropping by 0.5% at 31 years and a further 0.25% at 76 years. We also recommend that the risk premium does not vary with maturity.

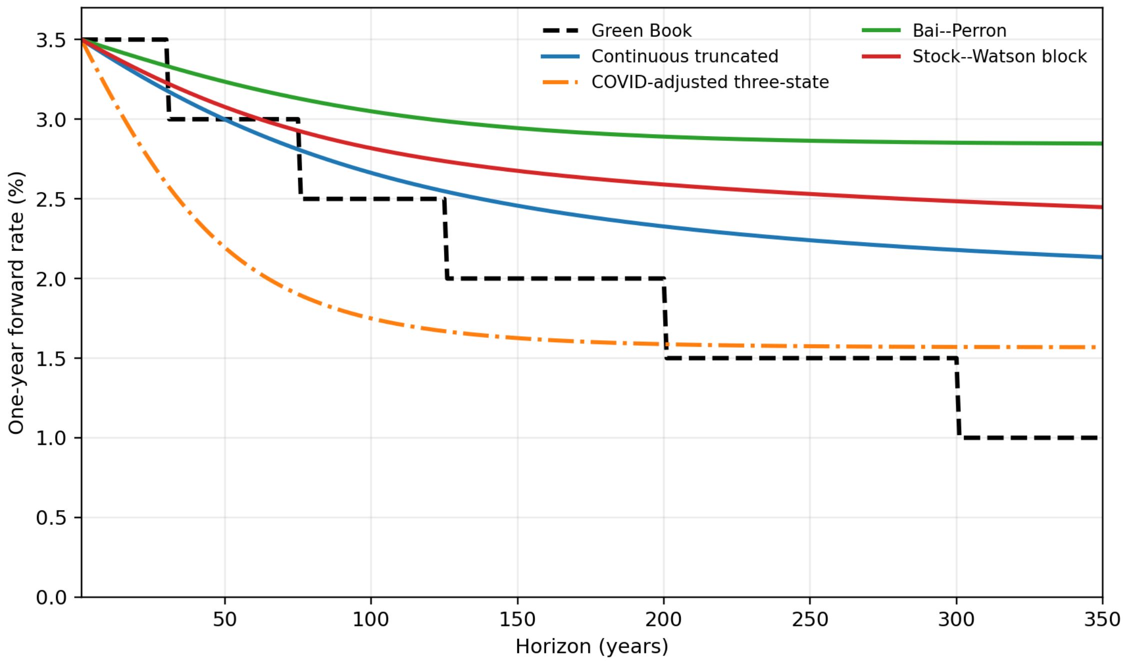

Figure 13.2.: Current Green Book comparison using one-year forward rates. All model schedules are vertically normalised to start at 3.5 percent. The Green Book schedule is shown at its current stepped values. The figure includes the continuous truncated model, the COVID-adjusted discrete three-state calibration, the Bai–Perron structural-break calibration, and the Stock–Watson low-frequency calibration. Model schedules use $\delta = 1.5\%$ and $\mu = 1$.

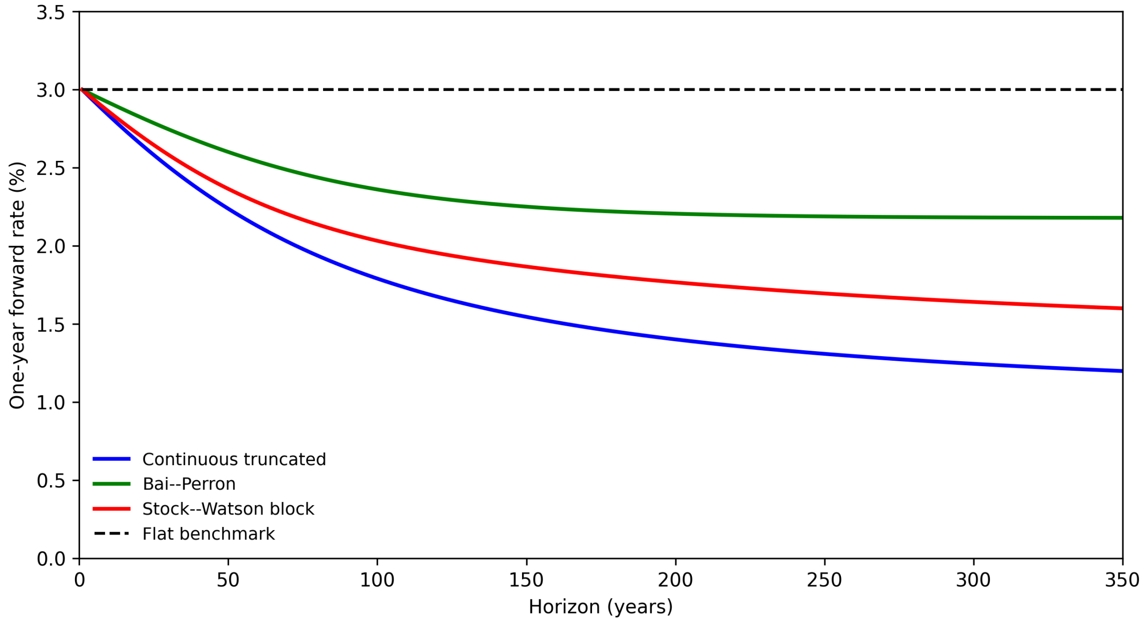

Figure 13.3.: New policy comparison using one-year forward rates. The Green Book schedule is omitted. All schedules are vertically normalised to start at 3.15 percent. This starting point is an illustrative project-appraisal benchmark that adjusts the raw Ramsey rate under $\delta = 0.5\%$ and $\mu = 1.3$ for catastrophic risk and a positive project risk premium. The figure shows the flat benchmark together with the continuous truncated model, the Bai–Perron structural-break calibration, and the Stock–Watson low-frequency calibration.

14. Catastrophic Risk, Risk-Free Discounting and Project-Specific Risk Premia

This chapter simulates the implications of catastrophic risk for public cost–benefit analysis using the updated GB+NI per capita consumption-growth series to 2025. The purpose is not to establish the signs of the effects, which follow directly from theory, but to quantify plausible magnitudes for UK appraisal. The analysis separates the catastrophe adjustment to the risk-free social discount rate from the systematic risk premium for project-specific payoffs. The main case uses $\mu = 1.25$, a no-disaster Ramsey rate $r_f^{ND} = 0.5 + \mu \times 1.5 = 2.375\%$, and project exposures $X = 20\%, 50\%, 80\%$. We compare a Barro (2006) inspired point-disaster benchmark with Pindyck & Wang (2013) inspired power-distribution calibrations using the full UK sample, the last 100 years and the last 50 years. Under $\mu = 1.25$, the Barro benchmark gives a risk-free correction of -1.63 percentage points and a pro-cyclical premium of 0.72 percentage points at $X = 80\%$. UK Pindyck–Wang calibrations give smaller risk premia, around 0.31–0.33 points at $X = 80\%$. A one percentage-point pro-cyclical premium requires either higher risk aversion, a more severe global disaster calibration, or nearly complete project loss in the disaster state.

14.1. Introduction

14.1.1. The purpose of this chapter is to test the possible quantitative implications of catastrophic risk for UK public cost–benefit analysis. The signs of the effects are well known; see Chapter 12. Safe payoffs become more valuable in low-consumption disaster states, so disaster risk lowers the risk-free social discount rate. Projects whose payoffs fall in the same states require positive systematic risk premia, while projects whose payoffs rise in disaster states may justify negative premia. The open policy question is not the direction of these effects, but their likely size under calibrations relevant to the UK and to the discounted-utilitarian framework used in the Green Book.

14.1.2. The analysis therefore distinguishes three objects. The first is the no-disaster risk-free rate, $r_f^{ND} = \delta + \mu g$, which is the ordinary Ramsey component. The second is the catastrophe adjustment to that risk-free rate, Δr^D. The third is the project-specific systematic risk premium, ρ_A. The project-specific discount rate is $r_A = r_f^{ND} + \Delta r^D + \rho_A$. This decomposition is useful because catastrophic risk can have a large effect on the risk-free component without implying an equally large premium for all projects. The risk premium depends on the project's exposure to disaster states.

14.1.3. The chapter compares two disaster-risk calibrations. The first is a Barro-inspired point-disaster benchmark using a fixed disaster probability and a fixed proportional loss (Barro 2006, Barro & Ursúa 2011). This is a useful comparator because it represents a severe global-disaster calibration. The second is a Pindyck–Wang-inspired power-distribution approach, in which disaster sizes are described by a fitted power distribution for the surviving fraction after a catastrophe (Pindyck & Wang 2013). We estimate the relevant event frequencies and size-distribution parameters from UK per capita consumption growth, using data from 1831 to 2025.

14.1.4. The main findings are numerical. With $\mu = 1.25$, the global Barro point-disaster benchmark implies a risk-free correction of -1.63 percentage points. The UK Pindyck–Wang calibrations imply smaller risk-free corrections, around -0.74 to -0.77 percentage points. The pro-cyclical project premia are smaller than the full risk-free correction because they are covariance premia rather than expected aggregate consumption losses. At $X = 80\%$, the Barro premium is about 0.72 percentage points, while the UK Pindyck–Wang premia are about 0.31–0.33 points. Under the Green Book-relevant value $\mu = 1.25$, even a pro-cyclical project that loses all of its payoff in disaster states does not reach a one percentage-point premium in the central calibrations. Higher risk aversion or a more severe disaster calibration is needed to reach that level.

14.2. Theory: risk-free disaster risk and project systematic risk

14.2.1. The discounted-utilitarian framework is the same as in the Green Book Ramsey-rule approach. Social welfare is based on expected discounted utility, with instantaneous utility $u(C_t) = C_t^{1-\mu}/(1-\mu)$ for $\mu \neq 1$, and the logarithmic limiting case when $\mu = 1$. Marginal utility is therefore $u'(C_t) = C_t^{-\mu}$. If aggregate consumption between today and the payoff date changes by the gross factor G, the social stochastic discount factor is proportional to $e^{-\delta t}G^{-\mu}$. A disaster that leaves the economy with surviving fraction ϕ raises marginal utility by the multiplier $\phi^{-\mu}$. This is the utility-theoretic source of both the risk-free disaster adjustment and the project-risk premium.

14.2.2. Let M denote the social stochastic discount factor and let A denote the project payoff factor relative to its normal-state payoff. The project-specific discount factor is

$$D_A = \frac{\mathbb{E}[MA]}{\mathbb{E}[A]}.$$

The risk-free discount factor is

$$D_f = \mathbb{E}[M].$$

The project systematic-risk premium is therefore

$$\rho_A = -100 \log\left(\frac{D_A}{D_f}\right) = -100 \log\left(\frac{\mathbb{E}[MA]}{\mathbb{E}[M]\mathbb{E}[A]}\right).$$

This expression is the central object in the project-risk calculation. It says that the premium depends on the covariance between the project payoff and the social pricing kernel. If the project pays less in states where marginal utility is high, the premium is positive. If the project pays more in those states, the premium is negative.

14.2.3. The project-specific discount rate is decomposed as

$$r_A = r_f^{ND} + \Delta r^D + \rho_A.$$

The first term is the no-disaster Ramsey rate,

$$r_f^{ND} = \delta + \mu g.$$

In the main case,

$$r_f^{ND} = 0.5 + 1.25 \times 1.5 = 2.375\%.$$

The second term, Δr^D, is the catastrophe adjustment to the risk-free social discount rate. The third term, ρ_A, is the project-specific systematic risk premium.

14.2.4. In a point-disaster model, a catastrophe occurs with probability p, and consumption falls by fraction b. Let

$$\phi = 1 - b$$

be the surviving fraction of consumption. The disaster-state stochastic discount factor multiplier is $\phi^{-\mu}$. The risk-free disaster correction is

$$\Delta r^D = 100\,[\mu p \ln \phi - \ln\{(1-p) + p\phi^{-\mu}\}]\,.$$

The first component, $100\mu p \ln \phi$, is the expected disaster-loss component. The second component, $-100 \ln\{(1-p) + p\phi^{-\mu}\}$, is the prudence component. The latter is the part most closely related to the maximum systematic project premium in a binary project-exposure model.

14.2.5. The project payoff is modelled as a binary event exposure. A pro-cyclical project has payoff $A_G = 1$ in the non-disaster state and $A_D = 1 - X$ in the disaster state. A counter-cyclical project has payoff $A_G = 1$ in the non-disaster state and $A_D = 1 + X$ in the disaster state. The parameter X is therefore the size of the project payoff change in the disaster state. This is not a market CAPM beta. It is a state-contingent catastrophe exposure.

14.2.6. More generally, suppose disaster events arrive with probability λ, and let ϕ be random conditional on a disaster. Then $\mathbb{E}[M] = (1-\lambda) + \lambda\mathbb{E}[\phi^{-\mu}]$. For a pro-cyclical project, $\mathbb{E}[A] = (1-\lambda)+\lambda(1-X)$ and $\mathbb{E}[MA] = (1-\lambda)+\lambda\mathbb{E}[\phi^{-\mu}](1-X)$. For a counter-cyclical project, $\mathbb{E}[A] = (1-\lambda) + \lambda(1+X)$ and $\mathbb{E}[MA] = (1-\lambda) + \lambda\mathbb{E}[\phi^{-\mu}](1+X)$. Substituting these expectations into the general premium formula above gives the reported pro- and counter-cyclical premia. Thus X affects the project-risk premium, but it does not affect the risk-free disaster adjustment.

14.3. Barro and Pindyck–Wang catastrophe specifications

14.3.1. The Barro-inspired benchmark treats the disaster as a point event. The central global calibration is

$$p = 0.017, \qquad b = 0.29.$$

This is not intended as an estimate from the UK series. It is a severe global-disaster comparator that allows the implied project premia to be compared with a familiar rare-disaster calibration.

14.3.2. The Pindyck–Wang-inspired model generalises the point-disaster benchmark by replacing the single disaster size with a distribution over catastrophe sizes. Pindyck & Wang (2013) model catastrophes as Poisson events with a power distribution for the surviving fraction of capital. In the reduced-form CBA implementation here, a disaster arrives with probability λ, and the surviving fraction of consumption is denoted ϕ. Conditional on an event exceeding threshold τ,

$$0 < \phi \le \phi_{\max}, \qquad \phi_{\max} = e^{-\tau/100}.$$

The fitted power distribution is

$$F(\phi \mid \phi \le \phi_{\max}) = \left(\frac{\phi}{\phi_{\max}}\right)^{\alpha}.$$

The corresponding density is

$$f(\phi \mid \phi \le \phi_{\max}) = \frac{\alpha\phi^{\alpha-1}}{\phi_{\max}^{\alpha}}, \qquad 0 < \phi \le \phi_{\max}.$$

14.3.3. The likelihood for observed disaster-state survival fractions $\phi_1, \ldots, \phi_n$ is therefore

$$L(\alpha) = \prod_{i=1}^{n} \frac{\alpha\phi_i^{\alpha-1}}{\phi_{\max}^{\alpha}}.$$

The log-likelihood is

$$\ell(\alpha) = n \ln \alpha + (\alpha - 1) \sum_{i=1}^{n} \ln \phi_i - n\alpha \ln \phi_{\max}.$$

Differentiating and setting the derivative equal to zero gives

$$\frac{n}{\alpha} + \sum_{i=1}^{n} \ln \phi_i - n \ln \phi_{\max} = 0,$$

so that

$$\widehat{\alpha} = -\frac{n}{\sum_{i=1}^{n} \ln(\phi_i/\phi_{\max})}.$$

The Pindyck–Wang-inspired model therefore uses a distribution of disaster sizes to estimate $\mathbb{E}[\phi^{-\mu}]$ and $\mathbb{E}[\ln \phi]$. In the project-risk calculation, the project loss remains binary conditional on the occurrence of a disaster event.

14.4. Data and empirical implementation

14.4.1. The empirical analysis uses the updated GB+NI per capita consumption-growth series, 1831–2025:

$$g_t = 100\,\Delta \ln(C_t/N_t).$$

Figure 14.1 shows the historical growth series. A disaster event is defined by

$$d_t = \mathbf{1}\{g_t < -\tau\}.$$

The central threshold is $\tau = 5\%$, with $\tau = 3\%$ and $\tau = 10\%$ reported in the output files as robustness checks. The main empirical estimates are reported for the full sample, the last 100 years and the last 50 years. This is deliberate: rare-event estimation requires a long time series, and short windows without major events provide little information about catastrophe risk.

14.4.2. For each sample and threshold, the annual event probability is estimated as

$$\widehat{\lambda} = \frac{1}{T} \sum_t d_t.$$

For each event year, the observed surviving fraction is

$$\phi_t = \exp(g_t/100),$$

and the observed proportional loss is

$$B_t = 1 - \phi_t.$$

The empirical event-size distribution is the observed set of ϕ_t's in disaster years. The Pindyck–Wang-inspired calibration instead estimates α from the observed ϕ_t's and evaluates expectations under the fitted power distribution.

14.4.3. The project-risk calculation simulates hypothetical pro- and counter-cyclical projects. For each exposure $X \in \{20\%, 50\%, 80\%\}$, a pro-cyclical project has payoff 1 in the non-disaster state and $1 - X$ in the disaster state, while a counter-cyclical project has payoff 1 in the non-disaster state and $1 + X$ in the disaster state. This ex post simulation estimates the systematic risk premium implied by the historical frequency and size distribution of UK disaster events and by the chosen disaster model.

14.5. Empirical disaster estimates

14.5.1. Table 14.1 summarises the empirical disaster estimates. The full sample contains five $g_t < -5\%$ events: 1916, 1917, 1921, 1940 and 2020. The last-100-year sample contains two such events, 1940 and 2020. The last-50-year sample contains only the 2020 event. The full-sample and last-100-year estimates are therefore important for judging whether the 2020 observation should be treated as representative of long-run disaster risk.

14.6. Main results for $\mu = 1.25$

14.6.1. Table 14.2 reports the risk-free disaster adjustment. The Barro point-disaster benchmark gives the largest risk-free adjustment because it assumes a 29 percent consumption loss. The Pindyck–Wang estimates based on UK data generate smaller but still material adjustments.

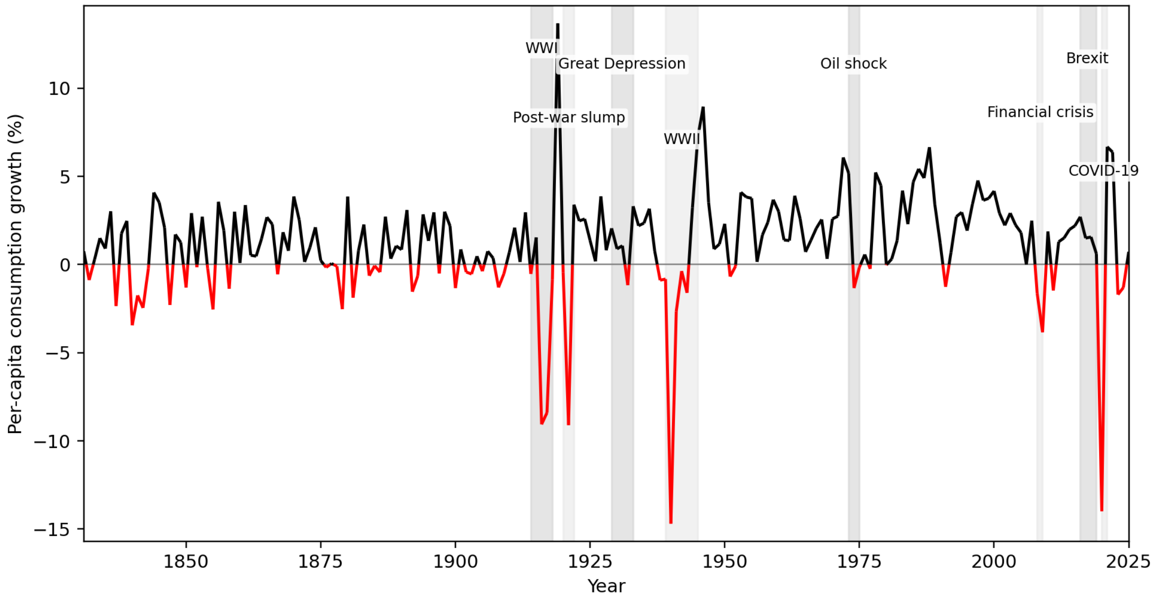

Figure 14.1.: UK per capita consumption growth, 1831–2025. This is the historical series used to identify disaster-event observations.

Table 14.1.: UK disaster-event estimates using the central threshold $g_t < -5\%$

Sample	N	Events	λ (%)	Years	Mean loss b (%)	α
Full sample	195	5	2.564	1916, 1917, 1921, 1940, 2020	10.429	16.527
Last 100 years	100	2	2.000	1940, 2020	13.360	10.705
Last 50 years	50	1	2.000	2020	13.054	11.126

Notes: Events are annual per-capita consumption-growth observations below −5%. Loss is $B_t = 1 - \exp(g_t/100)$. The power-distribution parameter α is estimated from the conditional survival fractions $\phi_t = \exp(g_t/100)$.

14.6.2. Table 14.3 separates the risk-free disaster correction into the mean-loss and prudence components and places these next to the project-risk premia. This is useful because the project premium is a covariance term, while the risk-free correction also includes the effect of expected disaster losses on aggregate consumption. The table therefore makes clear why the risk-free correction can be larger than the project premium even when project exposure is high.

14.6.3. Table 14.4 reports the systematic project-risk premia and combined project-specific discount rates. The Barro point-disaster benchmark generates the largest project premia and is a useful upper comparator. The UK Pindyck–Wang estimates are materially smaller. With $\mu = 1.25$, a pro-cyclical project with $X = 80\%$ has a premium of 0.72 percentage points under the Barro point-disaster benchmark, but only about 0.31–0.33 points under the UK Pindyck–Wang calibrations. The alternative counter-cyclical payoff normalisation considered in Appendix Tables D.1 and D.2 produces larger negative counter-cyclical premia, because the counter-cyclical project is then normalised to have a lower payoff outside the disaster state.

14.7. Exposure required for target premia

14.7.1. Table 14.5 reports the project exposure required to generate target pro-cyclical project premia. This table is useful because it shows when the project-risk premium can reach policy-relevant magnitudes. With $\mu = 1.25$, none of the central calibrations reaches a one percentage-point premium even when $X = 100\%$. Under the Barro benchmark, a 0.5 percentage-point premium requires $X \simeq 56\%$. A one percentage-point premium is reached under the Barro benchmark when $\mu = 1.5$ and $X \simeq 88\%$, and under the UK Pindyck–Wang calibrations only for much higher μ, around $\mu = 3$, with very high project exposure.

14.8. Conclusion

14.8.1. This chapter estimates the possible magnitude of catastrophic-risk adjustments in the UK Green Book discounted-utilitarian framework. The exercise is not designed to establish the signs of the effects: those follow from theory. Instead, it asks how large the risk-free adjustment and project-specific risk premium could be under alternative disaster-risk calibrations.

Table 14.2.: Risk-free disaster adjustments: Barro point-disaster comparator and Pindyck–Wang estimates, $\mu = 1.25$

Model	λ (%)	Mean loss b (%)	α	r_f^{ND}	Δr^D	r_f^D
Barro point disaster	1.700	29.000	–	2.375	-1.632	0.743
P–W Full sample	2.564	10.429	16.527	2.375	-0.742	1.633
P–W Last 100 years	2.000	13.360	10.705	2.375	-0.768	1.607
P–W Last 50 years	2.000	13.054	11.126	2.375	-0.747	1.628

Notes: The no-disaster risk-free rate is $r_f^{ND} = 0.5 + 1.25 \times 1.5 = 2.375\%$. Δr^D is the catastrophe adjustment to the risk-free social discount rate. The Barro row uses $p = 0.017$, $b = 0.29$. The Pindyck–Wang rows use UK event frequency and fitted power-distribution parameters at the $g_t < -5\%$ threshold.

Table 14.3.: Decomposing disaster-risk effects and project premia, $\mu = 1.25$

Model	λ (%)	Calibration	Mean-loss	Prudence	Δr^D	Pro $X = 50\%$	Pro $X = 80\%$	Counter $X = 50\%$	Counter $X = 80\%$
Global Barro	1.700	b=29.0%	-0.728	-0.904	-1.632	0.447	0.720	-0.438	-0.696
P-W Full sample	2.564	α=16.527	-0.354	-0.388	-0.742	0.191	0.309	-0.186	-0.295
P-W Last 100 years	2.000	α=10.705	-0.359	-0.410	-0.768	0.203	0.326	-0.198	-0.315
P-W Last 50 years	2.000	α=11.126	-0.350	-0.398	-0.747	0.197	0.317	-0.192	-0.306

Notes: All entries after the calibration column are percentage points. The risk-free correction is decomposed as $\Delta r^D = \Delta r^D_{\text{mean}} + \Delta r^D_{\text{prudence}}$, where $\Delta r^D_{\text{mean}} = 100\mu\lambda\mathbb{E}[\ln\phi]$ and $\Delta r^D_{\text{prudence}} = -100\ln\{(1-\lambda) + \lambda\mathbb{E}(\phi^{-\mu})\}$. The project-premium columns report systematic project-risk premia, not project discount rates.

14.8.2. The first numerical conclusion is that the risk-free disaster adjustment can be large. With $\mu = 1.25$, the Barro point-disaster benchmark lowers the risk-free rate by 1.63 percentage points. The UK Pindyck–Wang calibrations imply smaller adjustments, around 0.74–0.77 percentage points. These numbers can push the risk-free social discount rate close to zero under severe global-disaster assumptions, and under more extreme disaster probabilities or higher risk aversion could make the catastrophe-adjusted risk-free rate negative.

14.8.3. The second conclusion is that project premia are more limited unless project exposure is very high. With $\mu = 1.25$, a pro-cyclical project with $X = 80\%$ has a premium of 0.72 percentage points under the Barro point-disaster benchmark and about 0.31–0.33 points under the UK Pindyck–Wang calibrations. Even a pro-cyclical project losing all of its payoff in disaster states does not reach a one percentage-point premium under the $\mu = 1.25$ central calibrations. A premium of one percentage point requires either higher risk aversion, a more severe global disaster calibration, or extremely high project exposure.

14.8.4. The third conclusion is methodological. The risk-free catastrophe adjustment and the project-specific risk premium should be reported separately. The risk-free adjustment includes both expected disaster losses and prudence. The project premium is a covariance term. Combining them without decomposition can make the discount-rate implications difficult to interpret. For policy work, the Barro benchmark is a useful upper comparator for disaster-related risk premia, while the UK Pindyck–Wang estimates provide empirical lower-to-middle calibrations. The alternative counter-cyclical payoff normalisation in the appendix shows that the magnitude of negative counter-cyclical premia is sensitive to whether the counter-cyclical project is interpreted as having an above-normal disaster payoff or as reallocating payoff from normal states into disaster states.

14.8.5. From all the empirical work to this point, we are now in a position to draw conclusions on three things:

- Should the risk-free rate be adjusted for rare disaster risk? This analysis clearly implies that it should, and our best estimate is that the adjustment term should be minus 0.5%. Therefore, the STPR risk-free rate that we deduce from the Ramsey Rule with this adjustment is 0.5% + 1.25 × 1.5% - 0.5% = 1.875%. When triangulated against the SOC risk-free rate and survey data reported in Chapter 7, both at 2%, our best estimate of the risk-free component of the social discount rate for the UK is 2%. As concluded in Chapter 13, the forward rate should decline from this value by 0.5% at 31 years and a further 0.25% at 76 years.

- Should the average social project have a non-trivial value of L? Again, this analysis clearly states that it should. Considering the evidence above, we recommend an STPR risk adjustment of 0.5%. However, this requires triangulation with both the survey data and SOC estimates of the risk premium in Chapter 11. These respectively were around 1% and 1.9%. On balance, therefore, we recommend that the risk-adjustment term $L = 1\%$. As concluded in Chapter 11, this should not be varied by project except in exceptional

Table 14.4.: Project-risk premia and project-specific discount rates, $\mu = 1.25$

Model	X (%)	Δr^D	Pro premium	Counter premium	Pro rate	Counter rate
Barro point	20	-1.632	0.178	-0.176	0.921	0.567
Barro point	50	-1.632	0.447	-0.438	1.190	0.305
Barro point	80	-1.632	0.720	-0.696	1.463	0.047
P-W Full sample	20	-0.742	0.076	-0.075	1.709	1.558
P-W Full sample	50	-0.742	0.191	-0.186	1.824	1.447
P-W Full sample	80	-0.742	0.309	-0.295	1.941	1.338
P-W Last 100 years	20	-0.768	0.080	-0.080	1.687	1.527
P-W Last 100 years	50	-0.768	0.203	-0.198	1.809	1.409
P-W Last 100 years	80	-0.768	0.326	-0.315	1.933	1.292
P-W Last 50 years	20	-0.747	0.078	-0.077	1.706	1.550
P-W Last 50 years	50	-0.747	0.197	-0.192	1.824	1.435
P-W Last 50 years	80	-0.747	0.317	-0.306	1.944	1.322

Notes: For pro-cyclical projects, the payoff is $1 - X$ in the disaster state and 1 in the non-disaster state. For counter-cyclical projects, the payoff is $1 + X$ in the disaster state and 1 in the non-disaster state. X affects the project-risk premium but not the risk-free disaster adjustment.

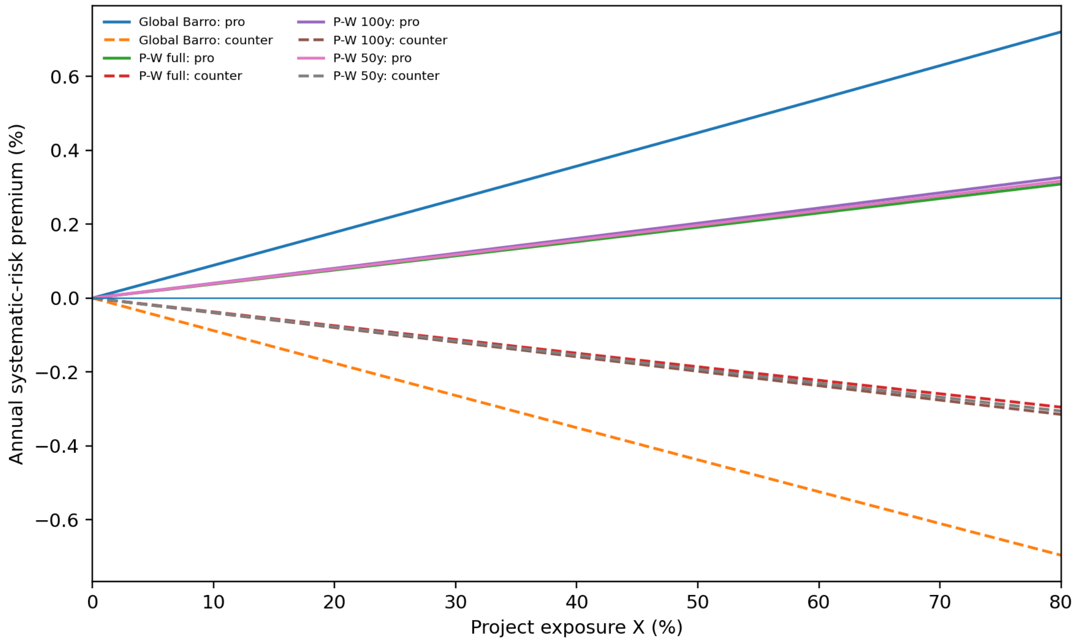

Figure 14.2.: Pro- and counter-cyclical systematic project premia by project exposure X, $\mu = 1.25$. Solid lines are pro-cyclical premia; dashed lines are counter-cyclical premia.

cases and, as concluded in Chapter 13, the term structure of this risk premium should be flat.

- Should social insurance projects have a different value for L? Again, it is clear from the analysis above that they should. We recommend that $L = -0.5\%$ for social insurance projects, again with a flat risk premium. These are projects that are characterised by high payoffs in the states of the world where aggregate consumption is very low and there is macroeconomic secular decline. Examples may include pandemic protection programmes and counter-terrorism measures that prevent the largest attacks to society.

R.18. The social discount rate should be 3.0% for standard projects and 1.5% for social insurance projects

For maturities of 0–30 years, the risk-free discount rate should be 2%, with a 1% risk premium for standard projects and a -0.5% risk premium for social insurance projects. At 31 years the forward rate declines by 0.5%, with a further decline of 0.75% at 76 years.

Table 14.5.: Exposure required to achieve target pro-cyclical project premia

μ	Model	Δr^D	Max prem. at X = 100%	X for 0.5%	X for 1.0%	X for 1.5%	X for 2.0%
1.00	Global Barro	-1.274	0.692	72.7	–	–	–
1.00	P-W Full sample	-0.588	0.305	–	–	–	–
1.00	P-W Last 100 years	-0.606	0.319	–	–	–	–
1.00	P-W Last 50 years	-0.589	0.310	–	–	–	–
1.25	Global Barro	-1.632	0.904	55.8	–	–	–
1.25	P-W Full sample	-0.742	0.388	–	–	–	–
1.25	P-W Last 100 years	-0.768	0.410	–	–	–	–
1.25	P-W Last 50 years	-0.747	0.398	–	–	–	–
1.50	Global Barro	-2.008	1.135	44.6	88.3	–	–
1.50	P-W Full sample	-0.899	0.474	–	–	–	–
1.50	P-W Last 100 years	-0.936	0.506	98.9	–	–	–
1.50	P-W Last 50 years	-0.910	0.490	–	–	–	–
2.00	Global Barro	-2.823	1.659	30.7	60.9	90.7	–
2.00	P-W Full sample	-1.224	0.658	76.6	–	–	–
2.00	P-W Last 100 years	-1.289	0.716	70.4	–	–	–
2.00	P-W Last 50 years	-1.252	0.692	72.7	–	–	–
3.00	Global Barro	-4.751	3.004	17.1	34.0	50.7	67.3
3.00	P-W Full sample	-1.920	1.070	47.5	93.7	–	–
3.00	P-W Last 100 years	-2.081	1.221	41.6	82.3	–	–
3.00	P-W Last 50 years	-2.014	1.175	43.2	85.5	–	–

Notes: Entries in the last four columns are the required pro-cyclical project exposure X, in percent, to generate the target annual systematic-risk premium. Dashes indicate that even a project losing 100 percent of its payoff in the disaster state does not generate the target premium under that calibration.

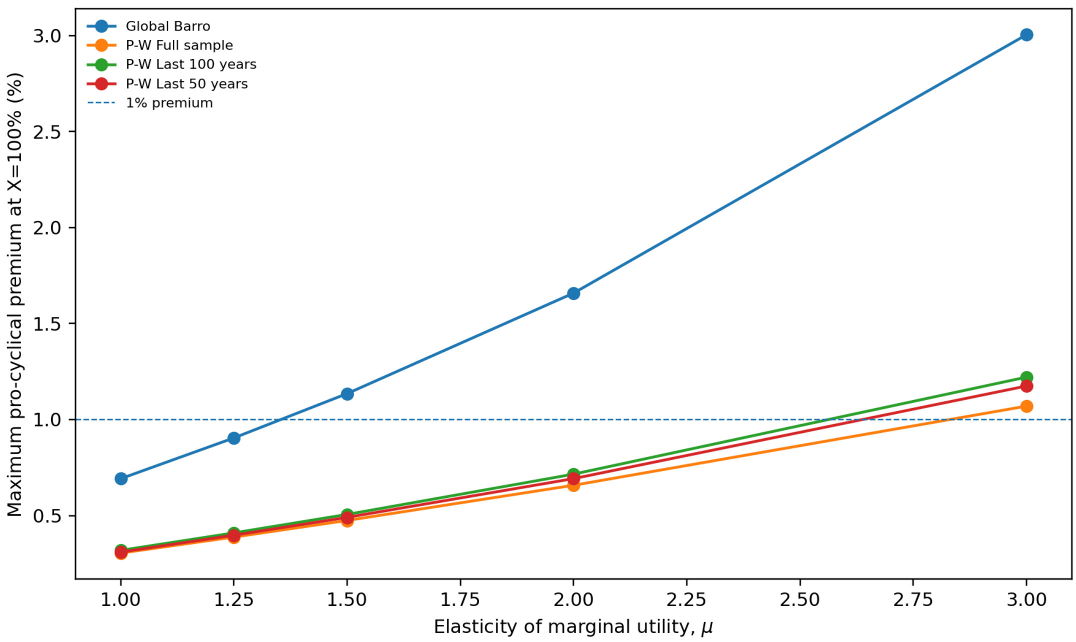

Figure 14.3.: Maximum pro-cyclical premium at X = 100% by μ.

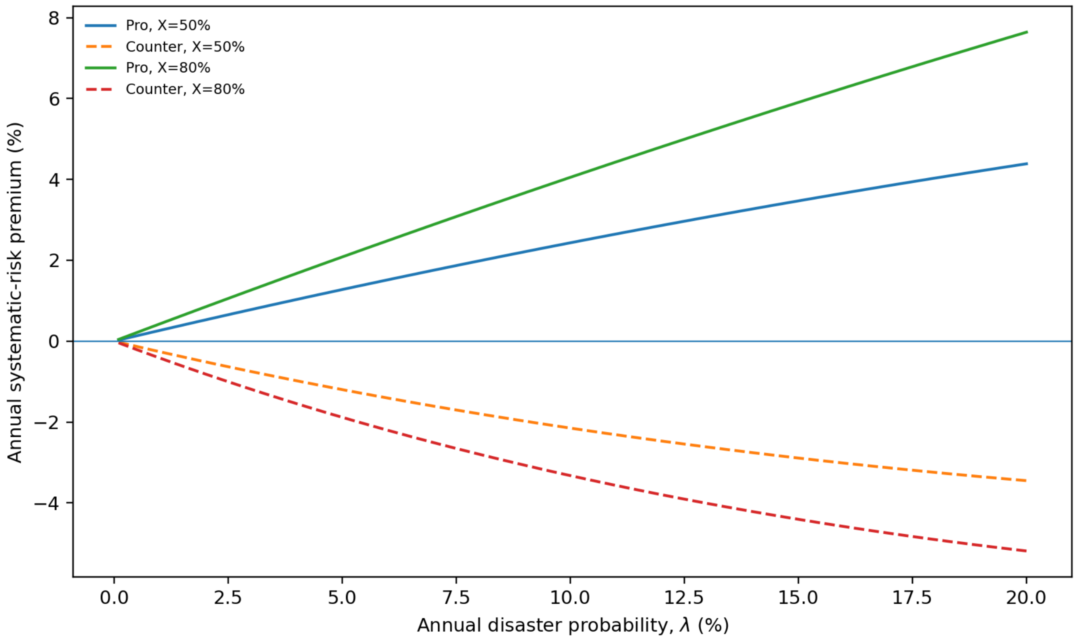

Figure 14.4.: Project-risk premia as the annual disaster probability varies in the Barro point-disaster benchmark.

15. Relative prices, commodity specific discount rates and sustainability

15.1. Introduction

15.1.1. In this chapter we review the basic theoretical ideas behind environmental discounting and the equivalent relative-price adjustment approach often proposed for CBA. The chapter illustrates that both operations lead to identical outcomes for CBA, but that in order to maintain the integrity of the social discount rate as means of handling just the maturity and risks (project level and secular) of the net-benefits, we recommend that the HMT Treasury uses a relative price approach in the Green Book.

15.1.2. There are two chief benefits of taking a relative pricing approach beyond maintaining the integrity of the discount rate across departments for a given time period and risk level. The first is that it emphasises the need for departments to undertake research to establish precisely how willingness to pay evolves with incomes and/or with scarcity for their particular sector. Departments are probably better placed than HMT, given their expertise, to undertake or commission this research. The second is that it reduces confusion among practitioners compared to a situation in which multiple discount rates are proposed. This is the present case with respect to health and air pollution in the Green Book. A relative price adjustment is also in line with the current guidance (section 6 HM Treasury 2026*c*), and so recommendations here can be seen as simply requesting that existing guidance on relative prices is simply applied more consistently and accurately. The recommendation would add more nuance to the Green Book's approach to different sectors while not entailing any additional principles to the guidance that were not there already.

15.1.3. The rest of this chapter spells out the theory and empirical work on the issue of discounting versus relative pricing. It turns out that much of this work has taken place in an often neglected sector of CBA, environment and ecosystem services. A general proposal here, in line with academic work in this area, is that the price of all ecosystem services should increase at the rate of growth of income (Drupp et al. 2024, 2025). This default mimics what is explicitly in place for the value of time saved in transport CBA and implicitly in the treatment of the value of QALYs in the evaluation of health interventions. While there is some evidence that this default is a reasonably accurate reflection of the elasticity of willingness to pay for ecosystem services in general, this recommendation may not be accurate for all ecosystem services. Our argument here is that this default is better than a default of zero given the evidence. Our recommendation is that HM Treasury should allow analyses that deploy the latest evidence on the evolution of relative prices in any given project, and insist on this for large projects.

15.1.4. The chapter first lays out the historical origins of the idea that different commodities should have their own specific discount rate. This is followed by an explanation of how this correction is in essence identical to a dynamic adjustment to relative prices for these commodities. We then discuss empirical evidence and some further issues that arise from the differential growth of environmental commodities and ecosystem services, such as the idea of *environmental drag*: where degradation affects growth in consumption (Zhu et al. 2019). We conclude by reiterating our recommendation that in the context of public CBA, relative price corrections are better pragmatically than specific discount rates.

15.2. Foundations

15.2.1. The intellectual origin of sector-specific discounting predates the modern environmental literature. Malinvaud (1953) is an important early precursor. His analysis of efficient capital accumulation in a multi-commodity economy supplied general intertemporal efficiency conditions in which commodity shadow prices need not move proportionately. When goods are not perfect substitutes, there is no reason for one physical unit of every commodity to share a common own-good discount rate. Later environmental work can be read as applying this general-equilibrium insight to a particularly important asymmetry: produced consumption may expand while environmental quality remains bounded or declines. Malinvaud did not present the modern environmental relative-price rule, but he supplied much of its multi-good valuation logic.

15.2.2. Guesnerie (2004) gave this logic a clear ecological interpretation. In a two-good model with a growing private good and a limited environmental good, he defined an ecological discount rate and investigated its long-run relation to the ordinary consumption rate. The analysis showed that ecological scarcity, preferences, and technological feasibility jointly determine the rate at which future environmental improvements should be discounted. Guéant et al. (2012) developed the argument further by analysing ecological and consumption discount rates along optimal growth paths and by clarifying how intergenerational concern, ecological concern, substitutability, and feasibility shape their asymptotic behaviour. These contributions were pioneering in two respects. First, they made the discount rate explicitly commodity-specific rather than assuming that there was one social discount rate. Second, it connected the ecological rate to the shadow value of an environmental commodity whose physical supply could not follow the growth of consumption. The core ecological intuition is already present in Guesnerie's work, but Guéant et al. (2012) also framed the issue as an evolving shadow price.

15.2.3. Among the earliest formal extensions, Weikard & Zhu (2005) derive distinct discount rates for environmental quality and consumption and ask directly when dual rates should be used. Their important contribution as far as this review is concerned is to show the equivalence between two procedures: discounting environmental quantities at an ecological rate, or allowing their relative price to change and discounting the resulting consumption-equivalent values at the consumption rate. They also explain why a lower ecological rate can be a pragmatic valuation device when a future environmental price path is not supplied separately. Hoel & Sterner (2007), Sterner & Persson (2008), Gollier (2010), and Traeger (2011) subsequently supplied alternative formulations, climate-policy applications, and transition dynamics. A formal analysis of the basic points now follows.

15.2.4. Our survey revealed broad acceptance of the view that substitutability and scarcity were issues that ought to be addressed in the Green Book. Views differed on how this should be done, reflecting the isomorphism of numerator (net-benefits) and denominator (discount rate) corrections to handle these concerns. Table 15.1 provides some selected qualitative evidence from the survey. All of these quotes agree that environmental scaricy is a problem, but there are differences on whether this leads to a discount rate adjustment or a relative price adjustment. The same issues arise in relation to other goods and services, not just the environment, but changing relative prices are not generally considered for the environment in government appraisal.

15.3. A two-good Ramsey framework

15.3.a. The one-good benchmark

15.3.1. In the standard normative Ramsey model (Ramsey 1928, Gollier 2013, Baumgärtner et al. 2015), instantaneous utility depends on an aggregate consumption good A_t and social welfare is

$$W = \int_0^\infty U(A_t)e^{-\delta t}\,dt,$$

where δ is the pure rate of time preference. Define the elasticity of marginal utility as

$$\mu_{AA,t} = -\frac{U_{AA,t}A_t}{U_{A,t}}.$$

With a deterministic consumption growth rate g_A, the consumption discount rate is the familiar Ramsey rule

$$SDR_{A,t} = \delta + \mu_{AA,t}g_A.$$

The second term is often called the wealth effect. If future people consume more, the marginal utility of an additional unit of consumption is lower. The same curvature parameter also encodes aversion to unequal consumption across dates in the discounted-utilitarian formulation. Reviews of the normative and positive approaches to this choice are provided by Drupp et al. (2018) and Groom et al. (2022).

15.3.b. Consumption and environmental quality

15.3.2. Disaggregate the composite good into market consumption C_t and environmental quality or ecosystem-service flow E_t:

$$W = \int_0^\infty U(C_t, E_t)e^{-\delta t}\,dt.$$

For $X, Y \in \{C, E\}$ define the normalized elasticity of the marginal utility of X with respect to Y by

$$\mu_{XY,t} = -\frac{U_{XY,t}Y_t}{U_{X,t}}.$$

The good-specific discount rates are then

$$SDR_{C,t} = \delta + \mu_{CC,t}g_C + \mu_{CE,t}g_E,$$
$$SDR_{E,t} = \delta + \mu_{EE,t}g_E + \mu_{EC,t}g_C.$$

Each rate depends on growth in both domains whenever utility is non-separable. The pure rate of time preference is the same in each case, so does not explain any difference between them. This is the kind of argument that illustrates the possibility of having good specific discount rates.

Table 15.1.: Selected qualitative comments on environmental discounting/relative price adjustments

For / against	Quote
For	"By not taking account of environmental scarcity and the lack of substitutability of nature. The current approach does not fully capture the costs and benefits of particular projects. As such there will be a misallocation of public funds, and for environmental scarcity and nature, then limited and often irreplaceable resources are being depleted."
For	"As environmental assets become scarcer over time, their marginal social value increases rather than declines, meaning that applying a standard consumption-based discount rate risks systematically undervaluing long-term environmental benefits."
Against	"Environmental scarcity and limited substitutability are better reflected through improved valuation methods, such as relative price adjustments and updated shadow prices, rather than through changes to the social discount rate. Adjusting the discount rate risks conflating time preference with valuation issues and reduces transparency and comparability across projects."
Against	"Values for environmental goods should be adjusted; if they become more scarce their value will go up. The discount rate is not the right tool for these adjustments."

Relative price interpretation

15.3.3. Let the current shadow price of environmental quality in units of consumption be

$$p_{E,t} = \frac{U_{E,t}}{U_{C,t}}.$$

Differentiating gives the environmental relative-price change

$$\begin{aligned} RPC_t &\equiv \frac{d\ln p_{E,t}}{dt} \\ &= SDR_{C,t} - SDR_{E,t} \\ &= (\mu_{CC,t} - \mu_{EC,t})g_C - (\mu_{EE,t} - \mu_{CE,t})g_E. \end{aligned}$$

The shadow price rises when environmental quality becomes scarce relative to consumption and the welfare loss cannot readily be offset by additional consumption. This expression is identical to the difference in discount rates in paragraph 15.3.2, making clear the equivalence of relative shadow price adjustments and adjustments to the discount rate. This result can be seen in Weikard & Zhu (2005).

15.3.4. **Aside: Hicksian substitution elasticity**: The Hicksian elasticity of substitution measures the proportional change in the relative quantities of consumption and environmental quality induced by a proportional change in their marginal rate of substitution, holding utility constant. It is defined locally as

$$\sigma_t \equiv \left.\frac{d\ln(C_t/E_t)}{d\ln(U_{E,t}/U_{C,t})}\right|_{dU_t=0}.$$

Using the general marginal-utility elasticities

$$\mu_{XY,t} \equiv -\frac{U_{XY,t}Y_t}{U_{X,t}}, \qquad X, Y \in \{C, E\},$$

the proportional change in the marginal rate of substitution is

$$d\ln\left(\frac{U_E}{U_C}\right) = (\mu_{CC} - \mu_{EC})\, d\ln C - (\mu_{EE} - \mu_{CE})\, d\ln E.$$

Along an indifference curve,

$$dU = U_C C\, d\ln C + U_E E\, d\ln E = 0.$$

Define the marginal-value shares

$$s_C \equiv \frac{U_C C}{U_C C + U_E E}, \qquad s_E \equiv \frac{U_E E}{U_C C + U_E E} = 1 - s_C.$$

The indifference-curve condition can then be written as

$$s_C\, d\ln C + s_E\, d\ln E = 0.$$

Combining the Hicks definition, the marginal-rate-of-substitution change, and the share-form indifference condition in paragraph 15.3.4 gives the general two-good Hicksian elasticity of substitution:

$$\frac{1}{\sigma_t} = s_{E,t}\left(\mu_{CC,t} - \mu_{EC,t}\right) + s_{C,t}\left(\mu_{EE,t} - \mu_{CE,t}\right).$$

15.3.5. The final display in paragraph 15.3.4 shows that gross substitutability is determined by a marginal-value-share-weighted combination of own- and cross-good marginal-utility effects. It is therefore distinct from either cross elasticity considered in isolation. In particular, the cross elasticities capture *net* substitutability or complementarity, whereas σ_t also incorporates the curvature of marginal utility within each domain.

15.3.6. For homothetic preferences $\mu_{CC,t} - \mu_{EC,t} = \mu_{EE,t} - \mu_{CE,t} = \frac{1}{\sigma_t}$. In this case marginal rate of substitution depends only on the quantity ratio C_t/E_t. In that special case the relative-price growth equation simplifies to

$$SDR_{C,t} - SDR_{E,t} = \frac{1}{\sigma_t}\left(g_C - g_E\right).$$

In the fully general case, the discount-rate difference should instead be retained as

$$SDR_{C,t} - SDR_{E,t} = \left(\mu_{CC,t} - \mu_{EC,t}\right) g_C - \left(\mu_{EE,t} - \mu_{CE,t}\right) g_E.$$

15.3.7. Defining the local Hicks elasticity of substitution σ_t along the relevant path gives the equivalent representation

$$RPC_t = \frac{g_C - g_E}{\sigma_t}.$$

The first display in paragraph 15.3.7 is an accounting identity once σ_t is defined by the response of the quantity ratio to the marginal rate of substitution. The expressions in paragraphs 15.3.3 and 15.3.7 establish the equivalence between dual discounting and relative-price adjustment. With constant rates,

$$p_{E,0}E_t e^{RPC\,t} e^{-SDR_C t} = p_{E,0}E_t e^{-SDR_E t}.$$

15.3.8. **Two options**: The previous analysis shows that it is possible to either: i) value a future environmental unit at its future shadow price and discount the resulting consumption-equivalent amount at SDR_C; or, ii) keep the current price fixed and discount the physical environmental unit at SDR_E. Similar adjustments and equivalence are possible under uncertainty. Weikard & Zhu (2005) emphasize that ii) can be a pragmatic device when future relative prices are unavailable. Hoel & Sterner (2007) and Gollier (2010) make the price-rate equivalence central to environmental appraisal. We now place mroe structure on the analysis to derive parametric expressions for the discount rate in this context to demonstrate dependence on the substitution terms.

15.3.c. CES–CIES preferences and transition dynamics

15.3.9. To neatly separate the substitution between contemporaneous goods from curvature over the aggregate consumption–environment composite, the CES - CIES approach is convenient and useful:

$$U(C_t, E_t) = \frac{1}{1-1/\gamma}\left[\alpha C_t^{(\chi-1)/\chi} + (1-\alpha)E_t^{(\chi-1)/\chi}\right]^{(1-1/\gamma)\chi/(\chi-1)},$$

where χ is the within-period elasticity of substitution, γ is the intertemporal elasticity of substitution for the composite, and λ_t is the value share of market consumption. The resulting rates can be written

$$SDR_{C,t} = \delta + \frac{1}{\gamma}[\lambda_t g_C + (1-\lambda_t)g_E] + \frac{1-\lambda_t}{\chi}(g_C - g_E),$$
$$SDR_{E,t} = \delta + \frac{1}{\gamma}[\lambda_t g_C + (1-\lambda_t)g_E] - \frac{\lambda_t}{\chi}(g_C - g_E).$$

The term scaled by $1/\gamma$ is common to both rates and captures overall growth and marginal-utility curvature. It is essentially the *wealth effect* of the composite. The terms scaled by $1/\chi$ redistribute this common component between the two goods according to relative scarcity. Their difference is

$$SDR_{C,t} - SDR_{E,t} = \frac{g_C - g_E}{\chi}.$$

Again, the pure rate of time preference is common to both rates in this CES-CIES context. Together with the composite wealth effect these determine the level of both rates. It is the contemporaneous substitutability that governs the wedge between these two rates. Because λ_t changes with quantities and prices, the two rates generally exhibit transition dynamics even when g_C, g_E, γ, and χ are constant. Nevertheless, the different between these rates is the same in structure as the relative-price expression in paragraph 15.3.7, albeit in a homothetic preference form. This decomposition and its implications for weak and strong sustainability are developed particularly clearly by Traeger (2011).

15.4. Substitutability, cross-partial effects, and environmental inequality aversion

15.4.a. Three different preference concepts

15.4.1. The environmental discounting debate often conflates three parameters. The first is own-domain curvature. In the environmental domain,

$$\mu_{EE} = -\frac{U_{EE}E}{U_E}$$

measures how rapidly the marginal utility of environmental quality rises as environmental quality falls. In an intertemporal social welfare interpretation, it expresses aversion to unequal environmental quality across dates or generations. Venmans & Groom (2021) call this environmental inequality aversion. It is analogous to consumption inequality aversion μ_{CC}.

15.4.2. The second concept is net, or Edgeworth, substitutability, represented by the cross partial. Under the sign convention introduced in paragraph 15.3.2,

$$U_{CE} < 0 \iff \mu_{CE} > 0 \text{ and } \mu_{EC} > 0,$$

so an increase in one good lowers the marginal utility of the other and the goods are net substitutes. If $U_{CE} > 0$, the goods are net complements and the normalized cross elasticities are negative. Although symmetry gives $U_{CE} = U_{EC}$, the normalized elasticities generally differ because their scale factors and marginal-utility denominators differ.

15.4.3. The third concept is the Hicks elasticity σ, a gross substitution measure. It records the total change in the quantity ratio induced by a change in the marginal rate of substitution, and therefore combines own- and cross-marginal-utility effects. A cross elasticity is not itself an elasticity of substitution. Additive separability sets $\mu_{CE} = \mu_{EC} = 0$, but it does not imply either perfect substitution or zero substitution. Conversely, a good may be a net substitute while gross substitution remains weak because own-good marginal utility changes rapidly.

15.4.4. This distinction is central to interpreting Venmans & Groom (2021). Their experiments estimate environmental inequality aversion rather than the CES between consumption and nature. They report substantially higher curvature in their intratemporal environmental allocation treatment than in intertemporal treatments, and lower estimates when environmental quality improves than when it declines. These results show that the ethical weighting of unequal environmental outcomes can be asymmetric and context dependent. They do not identify the rate at which consumption can compensate for environmental loss.

15.4.b. Limited and zero substitutability

15.4.5. If $g_C > g_E$ and $0 < \sigma < \infty$, the relative-price expression in paragraph 15.3.7 implies a positive environmental relative-price adjustment. The smaller the elasticity, the faster the environmental shadow price rises or equivalently the lower the environmental discount rate relative to the consumption rate. As σ approaches zero while the growth differential remains non-zero, the smooth-model relative-price adjustment becomes arbitrarily large.

15.4.6. At exactly zero substitutability, however, the issue is not merely that an unusually large finite uplift should be applied. In a Leontief or essential-input limit, a loss of the binding environmental component cannot generally be compensated by *any* finite increase in consumption. The marginal rate of substitution is undefined at a kink or becomes unbounded as an ecological minimum is approached. Conversion of environmental units into consumption-equivalent monetary units can therefore fail. Appraisal must then supplement monetary cost–benefit analysis with physical constraints, ecological thresholds, safe-minimum standards, or explicit non-compensation rules.

15.4.7. Baumgärtner et al. (2017) formalize this point by introducing a subsistence requirement for environmental services. Substitutability becomes state dependent. This has the effect of being a moderate adjustment to substitutability when far from the critical value but falling dramatically as environmental services approach the critical value. Drupp (2018) brings this mechanism directly into social discounting. His estimated initial mean elasticity is around two, but the effective elasticity of substitutability declines over time as environmental services approach subsistence, moving the system toward complementarity and generating rapidly increasing shadow prices. This is a different proposition from simply imposing a low constant elasticity, and warns against using today's empirical evidence to calibrate the longer future substitutability.

15.5. Empirical evidence

15.5.a. Indirect evidence from willingness to pay

15.5.1. Most empirical studies do not observe μ_{CE} and μ_{EC} directly. Instead, they estimate how marginal willingness to pay for an environmental service varies with income. Under homothetic CES preferences, holding environmental provision fixed, the income elasticity of marginal willingness to pay maps approximately into the inverse elasticity of substitution, $1/\sigma$. Combining that estimate with separate growth rates for consumption and ecosystem services then yields an implied relative-price path. The mapping is useful but conditional: income must proxy consumption, preferences must be sufficiently close to homothetic CES, and the underlying valuation studies must refer to comparable marginal changes.

15.5.2. Baumgärtner et al. (2015) provide an early operational implementation. They combine ecosystem-service growth measures with valuation evidence and estimate that, on average across their application, ecosystem services should be discounted at a rate about 0.9 ± 0.3 percentage points below manufactured consumption. Their contribution is not a universal number but a replicable decomposition of the rate differential into relative growth ($g_C - g_E$) and substitution ($1/\sigma$).

15.5.3. Heckenhahn & Drupp (2024) extend the approach to a wide range of German ecosystem services. Their estimates show substantial heterogeneity: services with weak or negative physical growth and limited substitutability can require relative-price adjustments of several percentage points per year, with particularly large adjustments for some regulating services. The result argues against a single generic "nature uplift" whenever service-specific evidence is available.

15.5.4. The global meta-analysis by Drupp et al. (2025) draws on 735 income–willingness-to-pay estimates from 396 studies. Its preferred income elasticity is around 0.6 and, when combined with global consumption and ecosystem-service growth, implies an aggregate relative-price increase of roughly 1–2 per cent per year, with a central estimate close to 1.7 per cent. The associated adjustments materially raise public natural-capital values. Drupp et al. (2024) translate this evidence into an accounting recommendation: values used in appraisal and natural-capital accounts should be updated to reflect the increasing benefits supplied by scarce ecosystems rather than treated as fixed real prices.

15.5.5. Specifically, Drupp et al. (2024) recommend that government guidelines on CBA should have ecosystem-service prices increasing at the rate of income growth as a default setting, rather than the typical current setting of zero growth in relative prices. This recommendation has been taken up by the German government in its recent guidelines.

15.5.b. Direct estimates of cross elasticities

15.5.6. A laboratory experiment by Disque et al. (2026) is designed to identify the four elasticities appearing in the good-specific discount-rate expressions in paragraph 15.3.2 by varying growth in both market consumption and environmental quality. Preliminary workshop results report

$$\widehat{\mu}_{CC} = 1.14, \qquad \widehat{\mu}_{EE} = 0.84,$$
$$\widehat{\mu}_{CE} = 0.17, \qquad \widehat{\mu}_{EC} = 0.34,$$

with an implied gross elasticity of substitution of approximately 1.31. The positive cross elasticities imply net substitutability, but their modest magnitude suggests that cross effects are not overwhelming. This study is important because it attempts to identify the cross elasticities directly rather than infer gross substitution from income–willingness-to-pay regressions. The estimates should presently be treated as preliminary because the work is not yet a peer-reviewed publication.

15.5.7. Taken together, the empirical evidence supports three cautious conclusions. First, relative-price growth is unlikely to be zero for many ecosystem services. Second, the adjustment differs across services because both physical growth and substitutability differ. Third, current evidence on gross substitution is much stronger than evidence on the two normalized cross elasticities. Appraisal guidance should therefore permit service-specific updating while making the maintained preference assumptions explicit.

15.6. Production, natural-capital stocks, and environmental drag

15.6.a. The production-side mechanism

15.6.1. The two-good utility model treats g_C and g_E as exogenous. Zhu et al. (2019) show why this can be misleading when ecosystem services are intermediate inputs into production. If environmental inputs grow slowly and are difficult to replace by produced inputs, they constrain output and consumption growth. The conventional Ramsey wealth effect, $\mu_{CC} g_C$, then falls endogenously. This production-side scarcity channel can be described as environmental drag: ecological scarcity reduces the growth rate of the consumption numeraire and thereby lowers the consumption discount rate itself.

15.6.2. Zhu et al. (2019) distinguish three cases: 1) If ecosystem services are readily substitutable in both production and utility, the economy converges toward the conventional Ramsey rate; 2) When substitution is easy in production but difficult in utility, consumption growth can remain strong and the preference-side relative-price effect is the main adjustment; 3) When substitution is difficult in production, environmental scarcity drags down consumption growth, so both sectoral discount rates decline. In that final case the long-run relative-price wedge may become less important because consumption growth is itself forced toward the low growth rate of ecosystem services.

15.6.3. *Environmental drag* is is a different effect to the relative-price effect. The latter values a given physical environmental flow more highly as it becomes relatively scarce. Environmental drag changes the endogenous path of market consumption and hence the common wealth-effect component of discounting. An applied model may contain both channels, and omitting either can bias appraisal.

15.6.b. Stocks and flows

15.6.4. A second extension distinguishes the stock of natural capital N_t from the flow of ecosystem services it produces:

$$\dot{N}_t = G(N_t) - D_t,$$
$$E_t = F(N_t).$$

A current shock can reduce N_t and thereby lower E_t in many subsequent periods. Stock persistence determines the physical duration of damage; the relative-price path determines the value attached to each lost service flow. These effects cumulate rather than substitute for one another.

15.6.5. Bastien-Olvera & Moore (2021) incorporate natural capital directly into an integrated assessment model. The stock generates market-production inputs, non-market ecosystem services, and non-use value. Climate damage to natural capital therefore propagates through multiple welfare flows rather than appearing only as a contemporaneous loss of aggregate output. Bastien-Olvera & Moore (2022) review this approach, while Bastien-Olvera et al. (2024) spatially disaggregate climate impacts on natural-capital values and show that losses in non-market ecosystem benefits can be large and highly unequal across countries.

15.6.6. Dasgupta (2021) places the same stock–flow distinction at the centre of comprehensive wealth accounting. Natural capital is an asset; ecosystem services are dividends from that asset; accounting or shadow prices measure the marginal contribution of the asset to social well-being. A rising accounting price can therefore coexist with a declining physical stock, with implications for sustainability.

15.6.7. Quaas & Bröcker (2016) examine another stock-based channel in a growth model containing produced capital, human capital, a non-renewable resource, and irreversibly accumulated greenhouse gases. Their contribution is to connect the social cost of carbon to substitution among dynamic stocks and irreversibility. They derive a rule of thumb under which the SCC grows approximately with economic growth divided by the relevant elasticity of substitution, with uncertainty capable of amplifying the result sharply.

15.6.8. **Summary**: These contributions suggest that there are limitations to the extent to which discounting or relative price adjustments can capture the full implications for sustainability associated with natural capital loss. The issue lies with the treatment of growth as exogenous in Cost Benefit Analysis, and treating the shadow prices of ecosystem services as responses to an exogenous change in income and physical scarcity. When consumption growth and GDP are endogenised and the relationship with the broad conception of wealth are included, then endogenous feedbacks reflect these dependencies. Environmental drag: consumption growth being hindered by a deteriorating environment, is a clear example of this. In principle the relative price effect is somewhat diminished if expected consumption growth is lower as a consequence of environmental degradation. However, the lower expected consumption growth would otherwise reduce the consumption wealth effect in the Ramsey Rule across all government project appraisals. Capturing these non-marginal secular macroeconomic interactions between consumption flows and environmental stocks (indeed other forms of wealth too, such as human and man-made capital) requires the Ramsey framework to be augmented. Accounting for stocks, flows and the dependency of production on inputs from nature and resources would provide a richer wealth orientation framework within which to appraise public projects and regulations. At the very least, projections of growth in the Ramsey Rule should be mindful of changes in wealth.

15.6.9. We return to the issue of sustainability below, and provide a framework within which a connection can be made between the relative accounting prices that we have discussed here, CBA and sustainability, noting that extra criteria are needed in order to establish whether a positive NPV for the flows of costs and benefits in CBA can be thought of as fitting a sustainable programme of investment and wealth generation for the long term. Before we do that, we discuss the importance of making the connection between rising ecosystem service values and the social cost of carbon, which is a key illustration of the importance of dynamic pricing in CBA.

15.7. Relative prices and the social cost of carbon

15.7.1. The quantitative importance of these mechanisms is especially visible in the social cost of carbon. In a simplified deterministic representation,

$$SCC_0 = \int_0^{\infty} e^{-\int_0^t SDR_{C,s}\,ds} \left[D_t^C + p_{E,t} D_t^E\right] dt,$$

where D_t^C is marginal damage measured directly in consumption units, D_t^E is marginal physical damage to environmental-service flows, and $p_{E,t}$ converts that environmental damage into consumption units. The formula separates physical damage, persistence through stocks and flows, and the shadow price attached to each environmental unit.

15.7.2. Sterner & Persson (2008) provide the pioneering integrated-assessment application of rising environmental relative prices. By modifying DICE to distinguish consumption from environmental amenities, they show that strong mitigation can be justified without changing pure time preference: if amenities become scarce and are imperfect substitutes for consumption, their rising relative price offsets part of conventional consumption discounting. Their main contribution is a preference-side relative-price correction, not an explicit natural-capital stock model.

15.7.3. Drupp & Hänsel (2021) examine the drivers and plausible magnitude of this correction in an updated DICE framework. In their central calibration, accounting for relative scarcity raises the SCC by more than

50 per cent; the policy effect is comparable to reducing pure time preference by about 0.6 percentage points. Their paper demonstrates that limited substitutability can be quantitatively first order even when the broad physical damage structure is held fixed.

15.7.4. The contributions of Frances Moore and co-authors identify complementary omissions. Moore & Diaz (2015) allow temperature to affect economic growth rather than only the level of output, so damages persist and compound. This is a damage-persistence mechanism, not a relative-price mechanism. Moore, Baldos, Hertel & Diaz (2017) replace older agricultural damage functions with crop-model and general-equilibrium evidence, improving the measurement of a sectoral damage flow and more than doubling the SCC in their application. Moore et al. (2024) synthesize 1,823 SCC estimates from 147 studies and show that structural model choices—including persistent growth effects, Earth-system representation, and distributional weighting—are major sources of variation.

15.7.5. The natural-capital work of Bastien-Olvera & Moore (2021) adds a further channel: climate change damages an asset that supports market, non-market, and non-use flows. The Quaas–Bröcker model adds substitution among productive stocks and irreversibility. The Baumgärtner–Drupp–Quaas and Drupp models add state-dependent essentiality as environmental services approach subsistence. These mechanisms should not be collapsed into one "low environmental rate." A complete SCC calculation can simultaneously require more persistent physical damages, a richer stock–flow structure, rising environmental prices, and altered consumption growth through environmental drag.

15.7.6. Strictly, once environmental effects are converted into consumption units, the appropriate numeraire rate is the consumption-good rate from the full two-good model, including any cross effect of environmental change on the marginal utility of consumption. The practical significance of this correction, and the design of long-run relative-price schedules, are considered in Section 15.10.

15.8. Relative-price appraisal and natural-capital sustainability tests

15.8.1. Ordinary environmental relative-price adjustments and sustainability analysis are closely related, but they are not equivalent. The relative-price approach developed above arises within a Ramsey–Koopmans welfare framework. It assigns an environmental service the contemporaneous consumption-equivalent price

$$p^R_{E,t} = \frac{U_{E,t}}{U_{C,t}},$$

where the superscript R denotes the ordinary Ramsey welfare framework without an additional sustainability constraint. This price reflects the demand for environmental services, their physical scarcity, their substitutability with consumption, and the production and ecological constraints represented in the model. Its growth is the environmental relative-price adjustment considered in the preceding sections.

15.8.2. Such a price can become very large when environmental services become scarce or difficult to replace. The subsistence formulation of Drupp (2018), for example, makes effective substitutability decline as environmental services approach a minimum requirement and causes their relative price to grow without bound. This provides a reduced-form representation of essentiality or proximity to a critical environmental level. It does not, by itself, impose a general requirement that intergenerational welfare, comprehensive wealth, or every natural-capital stock be maintained. Ordinary scarcity prices therefore do not make a positive-NPV project sustainable by construction.

15.8.3. The distinction is important for cost–benefit analysis. Ordinary CBA asks whether an intervention increases discounted intertemporal welfare relative to a baseline. A sustainability assessment asks whether the resulting economic programme maintains the relevant productive base, natural assets, ecological conditions, or intergenerational well-being. The Ramsey framework can supply the accounting prices needed for both calculations, but a positive NPV does not by itself answer the second question.

15.8.a. From environmental service prices to natural-capital prices

15.8.4. The relative price of an environmental service flow should be distinguished from the accounting price of the natural-capital stock that generates it. Let

$$\mathbf{K}_t = \left(K^P_t, K^H_t, N_t, \dots\right)$$

denote the vector of produced, human, natural, and other relevant capital stocks. Define ordinary Ramsey continuation welfare by the value function

$$V^R(\mathbf{K}_t, t) = \max_{a\in\mathcal{A}^R_t(\mathbf{K}_t)} \int_t^\infty U(C_s, E_s)e^{-\delta(s-t)}\,ds,$$

where a denotes a feasible allocation programme and $\mathcal{A}^R_t(\mathbf{K}_t)$ is the set of programmes permitted by the production technologies, ecological laws of motion, resource constraints, and institutions represented in the model. The superscript R denotes the ordinary Ramsey welfare problem without an additional sustainability constraint. For simplicity, the maximum is assumed to exist.

15.8.5. Let N_t denote the natural-capital component of $\mathbf{K}_t$, and let its current state affect future ecosystem-service flows according to

$$E_s = F_s(N_t, \ldots), \qquad s \geq t.$$

The Ramsey-consistent accounting price of the natural-capital stock, expressed in date-t consumption units, is

$$q^R_{N,t} = \frac{1}{U_{C,t}} \frac{\partial V^R(\mathbf{K}_t, t)}{\partial N_t}.$$

It measures the marginal contribution of an additional unit of the natural asset to current and future social welfare. Unlike the contemporaneous service-flow price $p^R_{E,t}$, it capitalises the entire future stream of effects generated by the stock.

15.8.6. Schematically, the natural-capital accounting price can be written as

$$q^R_{N,t} = \int_t^\infty D_C(t,s)\, p^R_{E,s} \frac{\partial E_s}{\partial N_t}\, ds$$
$$+ \text{ other production, resilience, non-use, and risk effects,}$$

where

$$D_C(t,s) = e^{-\delta(s-t)} \frac{U_{C,s}}{U_{C,t}}$$

is the consumption discount factor. The first term capitalises the future ecosystem-service flows generated by an additional unit of natural capital. The remaining terms represent welfare effects operating through production, ecological resilience, non-use values, risk, and other channels that are not already included in the service-flow term.

15.8.7. The environmental relative-price path is therefore an input into the accounting price of natural capital, but it is not itself the accounting price of the stock. Constructing $q^R_{N,t}$ also requires the ecological law of motion, the persistence and regeneration of the asset, the complete stream of services it generates, and its effects on production, resilience, non-use value, and risk. Rising environmental relative prices increase the accounting value of natural assets whose services are expected to become scarcer or more difficult to replace.

15.8.8. This stock–flow distinction is central to combining environmental relative-price appraisal with sustainability analysis. A discount schedule and a projected environmental unit price may be sufficient to value a specified stream of ecosystem-service changes. They are not sufficient to assess the natural-capital base unless the effects of the underlying stocks on all relevant future flows have also been represented. A project can be evaluated either by valuing the complete stream of service-flow changes or by valuing its change in the underlying natural-capital stock using $q^R_{N,t}$, provided that the chosen representation is comprehensive. The same welfare effects should not be included through both routes.

15.8.b. Ramsey accounting prices and genuine investment

15.8.9. The statement that the ordinary Ramsey framework does not impose sustainability should not be interpreted as implying that it cannot be used to evaluate sustainability. Under the conditions established in the comprehensive-wealth literature, the same Ramsey-based accounting prices can be used both for marginal project appraisal and for testing whether continuation welfare is non-declining (Dasgupta & Mäler 2000, Arrow et al. 2003, 2004).

15.8.10. For a marginal intervention whose effects can be represented as changes in the capital vector,

$$NPV^R_t(P) = \frac{\Delta V^R_t(P)}{U_{C,t}} \simeq \sum_i q^R_{i,t} \Delta K_{i,t}(P).$$

This expression compares two nearby economic programmes at a given date. A positive value means that the intervention raises discounted intertemporal welfare relative to the baseline.

15.8.11. A sustainability calculation applies the same accounting prices to changes in the productive base through time. Differentiating continuation welfare gives

$$\frac{\dot{V}^R_t}{U_{C,t}} = \sum_i q^R_{i,t} \dot{K}_{i,t} + \frac{1}{U_{C,t}} \frac{\partial V^R(\mathbf{K}_t, t)}{\partial t}.$$

The first term is genuine investment. The second captures explicit changes in technology, population, institutions, or other determinants of welfare not represented solely by the measured capital stocks. In a stationary setting,

$$GI^R_t \equiv \sum_i q^R_{i,t} \dot{K}_{i,t} = \frac{\dot{V}^R_t}{U_{C,t}}.$$

Non-negative genuine investment is therefore equivalent, under comprehensive and correctly priced accounting, to non-declining continuation welfare.

15.8.12. The same prices consequently perform two distinct operations:

$$\sum_i q^R_{i,t}\Delta K_{i,t} \quad \text{compares programmes for marginal CBA,}$$

$$\sum_i q^R_{i,t}\dot{K}_{i,t} \quad \text{tracks the productive base for sustainability assessment.}$$

Project appraisal differentiates welfare across programmes, whereas genuine investment differentiates continuation welfare through time along a programme.

15.8.13. Ramsey-based accounting prices can therefore be used to *test* sustainability even though the unconstrained Ramsey objective does not require the selected path to pass that test. A discounted-welfare-maximising programme need not have non-declining continuation welfare at every date. Similarly, a positive-NPV intervention can improve welfare relative to its baseline without making the resulting development path sustainable.

15.8.14. It is also important to specify what is being sustained. Non-negative comprehensive genuine investment permits substitution among produced, human, and natural assets to the extent represented in preferences, production, ecological dynamics, and the accounting prices. It does not require every individual natural asset to be maintained. Requiring a particular natural stock or ecological condition not to decline is a stronger, asset-specific requirement and may be appropriate for critical natural capital, protected habitats, irreversible losses, or assets close to ecological thresholds.

15.8.15. In this sense the accounting-price approach avoids imposing a mechanical weak-versus-strong sustainability dichotomy. Substitution possibilities are represented in the underlying welfare function, technologies, ecological dynamics, and feasible set, and are reflected in asset-specific accounting prices (Dasgupta 2021, Groom & Turk 2021). Where substitution is genuinely impossible, however, the appropriate representation is a physical or ecological constraint rather than merely a high finite price.

15.8.c. Ramsey CBA with a separate sustainability test

15.8.16. One coherent appraisal framework is to undertake ordinary Ramsey-consistent CBA and then subject the resulting economic programme to a separate sustainability test. Acceptance requires, for all relevant s

$$NPV^R_t(P) > 0 \qquad \text{and} \qquad \mathcal{S}_s(P) \geq 0,$$

where $\mathcal{S}_s(P)$ denotes the chosen sustainability criterion applied to the path that would result if the project were implemented.

15.8.17. If sustainability is defined as non-declining continuation welfare, the criterion concerns genuine investment or comprehensive wealth. If policy requires the preservation of a critical natural asset, the second test may instead concern a physical stock, ecological condition, resilience indicator, or legal environmental target. In both cases, the sustainability test must affect the decision rather than merely be reported alongside the NPV.

15.8.18. The resulting path, rather than only the incremental effect of the project, must be tested. A project can improve genuine investment relative to an unsustainable baseline without making the resulting path sustainable. A positive-NPV project can also reduce a particular natural-capital stock while increasing comprehensive wealth through investment in other assets. Whether that substitution is admissible depends on the sustainability criterion being applied.

15.8.19. Provided the Ramsey accounting prices are comprehensive and correctly constructed, no separate set of prices is required merely to undertake the genuine-investment test. The same prices $q^R_{i,t}$ can be used. What is added is a distinct, decision-relevant calculation concerning the evolution of the productive base or the relevant physical asset.

15.8.d. Sustainability-constrained accounting prices

15.8.20. A second coherent approach is to impose sustainability directly on the feasible set from which the accounting prices are derived. Suppose that admissible programmes must satisfy

$$\mathcal{S}_s(\mathbf{K}_s) \geq 0 \qquad \text{for all } s \geq t.$$

Let $V^S(\mathbf{K}_t, t)$ denote maximum discounted welfare subject to this requirement. The sustainability-constrained accounting price of asset i is

$$q^S_{i,t} = \frac{1}{U_{C,t}}\frac{\partial V^S(\mathbf{K}_t, t)}{\partial K_{i,t}}.$$

If the sustainability constraint is slack, the constrained and ordinary Ramsey prices coincide. If it binds, the constrained price can be represented schematically as

$$q^S_{i,t} = q^R_{i,t} + \Pi^S_{i,t},$$

where $\Pi_{i,t}^S$ is the present value of the relevant constraint-multiplier effects. The first term reflects ordinary welfare scarcity. The second reflects the opportunity cost of using scarce sustainability headroom.

15.8.21. For a marginal intervention that remains within the constrained feasible set,

$$NPV_t^S(P) = \sum_i q_{i,t}^S \Delta K_{i,t}(P) > 0$$

identifies a welfare improvement among sustainability-compatible programmes. Feasibility remains a separate requirement. A discrete project that crosses a hard ecological limit cannot be made admissible by a positive NPV calculated from local shadow prices.

15.8.22. The constrained-price method is especially useful where many interventions compete for a limited quantity of environmental headroom, such as a carbon budget, nutrient ceiling, habitat requirement, or legally binding emissions target. The value of that scarce headroom can then enter project appraisal through the associated multiplier. The separate-test method may be more transparent where the constraint is categorical, ecological non-convexities are important, or the multiplier cannot be estimated reliably.

15.8.e. Natural-capital values, revaluation, and double counting

15.8.23. Neither approach should equate sustainability with an increase in the reported monetary value of a natural asset. Consider the simple representation

$$A_{N,t} = q_{N,t}^R N_t.$$

Its change is approximately

$$\Delta A_{N,t} \simeq q_{N,t}^R \Delta N_t + N_t \Delta q_{N,t}^R.$$

The first term records investment or disinvestment in the natural asset at the prevailing accounting price. The second is a revaluation caused by a change in the accounting price.

15.8.24. A natural-capital asset value can therefore rise because the remaining asset has become scarcer and more valuable even while its physical extent, condition, or service capacity deteriorates. The revaluation term $N_t \Delta q_{N,t}^R$ should not be treated as investment in natural capital. A sustainability assessment should report physical extent, ecological condition, service capacity, resilience, and critical thresholds alongside monetary asset values.

15.8.25. The same distinction is needed to avoid double counting. Where CBA already includes the complete stream of ecosystem-service changes caused by a loss of natural capital, the capitalised value of those same service changes should not be added again as a separate project cost. Alternatively, the project can be evaluated through its change in the natural-capital stock valued at the full accounting price. Stock and flow representations should both enter the NPV only where they capture genuinely different welfare effects.

15.8.f. Two compatible appraisal procedures

15.8.26. Ordinary environmental relative-price appraisal and sustainability analysis are therefore compatible only when they are linked explicitly. The environmental service-price path should feed into the full accounting price of the affected natural-capital stock. CBA should use the relevant prices to measure the welfare effect of the intervention. Sustainability should then enter either as a separate, binding assessment of the resulting path or as a constraint whose multiplier is incorporated into the prices used in CBA. There are two procedures that could be built into guidance on CBA to serve this purpose.

15.8.27. The first procedure can be summarised as deploying ordinary Ramsey-consistent CBA together with a binding wealth sustainability test. The second can be summarised as CBA using sustainability-constrained accounting prices coupled with verification of feasibility. Neither ordinary relative-price uprating nor a positive NPV is, by itself, a sustainability test. In principle, the HMT Green Book HM Treasury (2026*c*) takes the first approach in recommending the valuation of flows of ecosystem services coupled with a screening via ENCA guidance of he impacts on natural capital. As the theory outlined above makes clear, this principle should be applied to all forms of wealth, not just natural capital to ensure sustainability.

15.8.28. The next section returns to the question of the prices used in valuation, how they are expected to evolve over time and how relative price effects are handled in international guidance on CBA.

15.8.g. Target-consistent accounting prices across sectors

15.8.29. The multiplier associated with a binding sustainability or policy constraint can sometimes be estimated through a target-and-cost calculation. This provides an operational link between the abstract sustainability-constrained accounting prices described above and the prices used in practical cost–benefit analysis. The approach is most familiar in carbon appraisal, but it can also be applied to biodiversity and, in principle, to binding targets concerning human capital, health, and other components of national wealth.

15.8.30. Let Z^j denote an outcome measured in an appropriate physical unit for sector j, and let $\overline{Z}^j$ be a target that government has committed to achieve. Consider the least-cost target-delivery problem

$$C_t^{j,*}\left(\overline{Z}^j\right) = \min_{a\in\mathcal{A}_t^j} \int_t^T D_C(t,s)c_s^j(a_s)\,ds$$

$$\text{subject to} \qquad Z_T^j(a) \geq \overline{Z}^j,$$

where a_s denotes the available interventions, $c_s^j(a_s)$ is their cost in consumption units, and

$$D_C(t,s) = e^{-\delta(s-t)}\frac{U_{C,s}}{U_{C,t}}$$

is the consumption discount factor. The target may alternatively be specified as a sequence of annual requirements or as a cumulative constraint; the terminal formulation is used only to simplify the exposition.

15.8.31. The target-consistent accounting price is

$$\tau_t^j = \frac{\partial C_t^{j,*}\left(\overline{Z}^j\right)}{\partial \overline{Z}^j} = \Lambda_t^j,$$

where Λ_t^j is the multiplier on the target constraint in the least-cost problem. It measures the additional consumption-equivalent cost of tightening the target by one unit. Equivalently, an intervention that increases the target outcome by $\Delta Z^j(P)$ avoids approximately

$$\Delta NPV_t^{T,j}(P) = \tau_t^j \Delta Z^j(P)$$

of target-delivery expenditure. An intervention that reduces the target outcome has $\Delta Z^j(P) < 0$ and therefore imposes an additional social cost because compensating action must be undertaken elsewhere if the target is still to be achieved.

15.8.32. This price is the practical counterpart of the constraint value in the sustainability-constrained welfare problem. For a single local target, the accounting price of asset i can be written schematically as

$$q_{i,t}^S = q_{i,t}^R + \tau_t^j \frac{\partial Z_t^j}{\partial K_{i,t}}.$$

The ordinary Ramsey price $q_{i,t}^R$ captures the asset's marginal contribution to discounted welfare. The second term captures its marginal contribution to satisfying the binding target. With dynamic targets, the corresponding expression contains the discounted sequence of current and future constraint multipliers.

15.8.33. A target-consistent price is conditional on the target rather than an estimate of the unconstrained willingness to pay for the outcome. It does not establish that the target itself is socially optimal. It identifies the opportunity cost of using scarce target headroom once the target has been adopted. More demanding targets, other things equal, generally imply higher shadow prices because progressively more costly interventions must be undertaken. If the target is slack, its multiplier and target-consistent price are zero.

15.8.34. **Carbon**. Let A^C denote cumulative greenhouse-gas abatement and let $\overline{A}^C$ be the abatement required to meet the emissions target. The target-consistent carbon price is the marginal abatement cost at $\overline{A}^C$:

$$\tau_t^C = \frac{\partial C_t^{C,*}\left(\overline{A}^C\right)}{\partial \overline{A}^C}.$$

If a project causes an additional $\Delta M_t > 0$ tonnes of emissions, government must undertake approximately the same quantity of additional abatement elsewhere. The project therefore incurs the target-compliance cost

$$\Delta NPV_t^C(P) = -\tau_t^C \Delta M_t.$$

Conversely, a project that reduces emissions releases target headroom and generates an avoided-abatement benefit. This is the logic of the target-consistent carbon values used in UK government appraisal (Department for Business, Energy & Industrial Strategy 2021).

15.8.35. **Biodiversity**. Groom et al. (2026) extend this logic to biodiversity through Target and Cost Analysis. Let B denote a specified biodiversity metric, such as species persistence, extinction risk, abundance, or species richness, and let $\overline{B}$ denote the adopted target. A marginal biodiversity recovery cost curve orders the available restoration measures according to their cost per unit of improvement. The accounting price

$$\tau_t^B = \frac{\partial C_t^{B,*}\left(\overline{B}\right)}{\partial \overline{B}}$$

is the marginal recovery cost at the target. A project-induced change $\Delta B(P)$ then enters CBA as

$$\Delta NPV_t^B(P) = \tau_t^B \Delta B(P).$$

A biodiversity loss receives a negative value because it increases the restoration effort needed elsewhere to meet the target. An improvement receives a positive value because it reduces that requirement.

15.8.36. The biodiversity target price should not be interpreted as the total value of all services supplied by the affected ecosystem. The Target and Cost Analysis price values the effect on the specified biodiversity target. Separately identifiable effects on carbon storage, water quality, recreation, flood regulation, health, and other ecosystem services should continue to be valued using their appropriate scarcity prices. Care is nevertheless required to ensure that the same ecological effect is not represented both in the biodiversity target price and in a separate service valuation.

15.8.h. Human and health capital: CBA and comprehensive wealth

15.8.37. Target-and-cost accounting prices are most natural where government has adopted a binding, quantitatively specified outcome constraint, as in the carbon and biodiversity examples above. They are not required for every component of national wealth. Education and health provide important examples in which a more complete application of ordinary cost–benefit analysis can be combined with the valuation of changes in comprehensive wealth. The relevant accounting prices are then the ordinary Ramsey welfare prices of human and health capital, unless an additional target or constraint has separately been imposed.

15.8.38. **Education and human capital**. Education can be represented as investment in a stock of knowledge, skills, and capabilities. Let K_t^L denote educational human capital and write its law of motion schematically as

$$\dot{K}_t^L = I_t^L - d_L K_t^L,$$

where I_t^L is gross investment through schooling, training, and other learning activities, while d_L captures depreciation through ageing, skill obsolescence, migration, and other losses. This treatment follows the human-capital tradition of Schultz (1961), Becker (1962) and the wealth-accounting approach of Jorgenson & Fraumeni (1992).

15.8.39. The Ramsey accounting price of educational human capital is

$$q_{L,t}^R = \frac{1}{U_{C,t}} \frac{\partial V^R(\mathbf{K}_t, t)}{\partial K_t^L}.$$

It measures the marginal contribution of an additional unit of knowledge and skills to current and future social welfare. In a comprehensive model, this contribution can include higher market productivity, increased future output, improved non-market production, greater capacity to benefit from technological change, and direct non-market effects of education.

15.8.40. A full cost–benefit analysis of an education programme begins with its resource costs, including teaching inputs, facilities, administration, and the opportunity cost of students' time. It then values the future flows generated by the resulting increase in human capital. Let $\Delta Z_{m,s}^L(P)$ denote the project's effect at date s on educational outcome or service m, and let $p_{m,s}^L$ denote its consumption-equivalent shadow price. A flow-based representation is

$$NPV_{L,t}^R(P) = -PV_t\,(\text{resource costs}) + \int_t^\infty D_C(t,s) \sum_m p_{m,s}^L \Delta Z_{m,s}^L(P)\, ds.$$

The relevant outcomes may include market production and non-market benefits that are not already captured by the productivity effect.

15.8.41. Where the accounting price $q_{L,t}^R$ capitalises the complete stream of these future effects, the same project can instead be represented locally as

$$NPV_{L,t}^R(P) \simeq q_{L,t}^R \Delta K_t^L(P) - PV_t\,(\text{resource costs}) + \text{direct effects not embodied in } K_t^L.$$

The flow and stock representations are alternative accounting routes. Future productivity gains should not be counted once through projected output or earnings and again through the capitalised human-capital price.

15.8.42. Observed lifetime earnings can provide evidence on the market component of human-capital value, as in the Jorgenson–Fraumeni approach. They are not, without further assumptions, a complete welfare accounting price. Wages need not equal social marginal product in the presence of taxes, labour-market distortions, externalities, or imperfect competition, while important non-market benefits may not appear in earnings. Similarly, additional tax receipts are primarily fiscal transfers and should not automatically be added to increased output as a separate social benefit.

15.8.43. The comprehensive-wealth calculation applies the same accounting price to the change in the human-capital stock:

$$GI_{L,t}^R = q_{L,t}^R \dot{K}_t^L.$$

Education CBA and human-capital accounting therefore use the same underlying valuation structure for different purposes. CBA asks whether a particular education intervention raises welfare relative to its baseline. Wealth accounting asks whether the stock of educational human capital is increasing or decreasing and how that change contributes to genuine investment.

15.8.44. **Health capital and productivity**. Health can be analysed in an analogous way. Following Grossman (1972), let K_t^H denote a stock of health capital that produces healthy life, directly contributes to utility, and affects the time and productivity available for market and non-market activities. Its law of motion may be written as

$$\dot{K}_t^H = I_t^H - d_H K_t^H,$$

where I_t^H includes prevention, treatment, public-health measures, and health-improving behaviour, while d_H represents biological deterioration and other adverse health processes.

15.8.45. The Ramsey accounting price of health capital is

$$q_{H,t}^R = \frac{1}{U_{C,t}} \frac{\partial V^R(\mathbf{K}_t, t)}{\partial K_t^H}.$$

It can include the direct welfare value of longer and healthier life, the effect of health on the enjoyment of consumption, changes in labour supply and productivity, effects on unpaid work and caring activity, and changes in future treatment and social-care requirements.

15.8.46. A health intervention can likewise be evaluated either through its complete stream of health and productivity effects or through the change it creates in the health-capital stock. Schematically,

$$NPV_{H,t}^R(P) \simeq q_{H,t}^R \Delta K_t^H(P)$$

$$-PV_t \text{ (intervention costs)}$$

$$+\text{direct effects not embodied in } K_t^H.$$

If life-years, QALYs, productivity effects, and future healthcare costs are valued separately as flows, the accounting price of health capital must not be added again unless it captures additional effects. The treatment of direct health benefits and productivity must also be internally consistent: a health value that already reflects the ability to work or enjoy consumption should not be supplemented by the same effect a second time.

15.8.47. The contribution of health investment to comprehensive wealth is

$$GI_{H,t}^R = q_{H,t}^R \dot{K}_t^H.$$

This places public-health and medical interventions within the same wealth framework as education, produced capital, and natural capital.

15.8.48. **Comprehensive wealth across sectors**. Combining these assets gives

$$GI_t^R = q_{P,t}^R \dot{K}_t^P + q_{L,t}^R \dot{K}_t^L + q_{H,t}^R \dot{K}_t^H + q_{N,t}^R \dot{N}_t + \cdots,$$

where K_t^P is produced capital, K_t^L is educational human capital, K_t^H is health capital, and N_t is natural capital. Under the usual qualifications concerning explicit changes in technology, population, and institutions, this genuine-investment measure indicates whether continuation welfare is increasing or decreasing (Arrow et al. 2012).

15.8.49. The example illustrates why the framework is not specifically environmental. Relative-price paths value the future services generated by each asset. Ecological scarcity affects the value of natural capital; skill scarcity, technological complementarity, and the future demand for capabilities affect the value of educational human capital; and income growth, health scarcity, and health–consumption complementarity affect the value of health capital. These service-price paths feed into the corresponding asset accounting prices.

15.8.50. No binding education or health target is required for this analysis. A target-consistent price could be introduced where government has adopted a specific binding requirement, but that would be an additional application. The primary point here is that ordinary Ramsey-consistent CBA and comprehensive wealth accounting can evaluate investment in education and health using the same system of welfare accounting prices, just as they can for natural and produced capital.

15.8.51. The next section returns to the question of the prices used in valuation, how they are expected to evolve over time and how relative price effects are handled in international guidance on CBA.

15.9. Government appraisal practice

15.9.a. The general principle beyond the environment: health

15.9.1. The same reasoning applies whenever social welfare depends on a non-consumption outcome whose relative scarcity changes. Replacing environmental quality by health H_t gives

$$SDR_{H,t} = \delta + \mu_{HH,t} g_H + \mu_{HC,t} g_C.$$

The cross term is important. If health and consumption are complements, $U_{HC} > 0$ and hence $\mu_{HC} < 0$. Growth in consumption then raises the marginal utility of health and lowers the health discount rate. The intuition is that consumption is more valuable when a person is healthy enough to enjoy it, while health is more valuable in a richer consumption environment. If the two goods are net substitutes, the sign reverses.

15.9.2. Gollier& Hammitt (2014) explain this equivalence in the context of long-run appraisal. A future health gain can be discounted at a health-specific rate, or its future monetary value can be increased relative to consumption and then discounted at the consumption rate. As with the environment, the appropriate relative-price path depends on income effects, complementarity, and the physical trajectory of the outcome.

15.9.3. Health also has its own stock–flow structure and can be considered as a component of national wealth. A medical intervention may alter a stock of health capital, producing future life-years, quality-adjusted life-years, productivity, and consumption-enjoyment benefits. Appraisal should distinguish those physical effects from the value attached to a unit of health.

15.9.4. These principles are general. The remainder of this section describes their treatment in government appraisal and natural-capital accounting before drawing together the recommendations for the Green Book.

15.9.b. The Netherlands

15.9.5. Dutch policy analysis has been an important application of differentiated environmental valuation. Koetse et al. (2018) recommend a 1 per cent annual increase in the relative price of non-substitutable nature and ecosystem services. Against the then standard 3 per cent discount rate, this was often presented as an effective environmental rate of 2 per cent. The recommendation was based on the expectation that relevant ecosystem services would become scarcer relative to consumption and could not be replaced one-for-one by manufactured goods.

15.9.6. The 2020 Discount Rate Working Group retained the 1 per cent nature adjustment while recommending a 2.25 per cent standard rate (Rijksoverheid 2020, 2025). If represented as a differentiated discount rate, the implied rate for qualifying nature effects was 1.25 per cent. The report nevertheless preferred to show the relative-price development separately. It also recognized that the appropriate uplift may be zero where substitution is easy or environmental provision keeps pace with consumption, and higher where services are essential or decline more rapidly.

15.9.7. The 2025 Working Group carries this separation further (Rijksoverheid 2025). Its rates, introduced from January 2026, are 2.8 per cent for horizons up to and including 35 years and 1.8 per cent thereafter. More importantly here, the report treats the discount rate as an adjustment for time and risk and treats changing scarcity through corrections to the values being discounted. It retains the existing 1 per cent annual correction for scarce ecosystem services pending further evidence, but reports expert concern that this correction is substantially too low and calls for generally accepted numerical price adjustments to be updated. Dutch practice has therefore evolved from an effective differentiated-rate presentation towards a more explicit relative-price framework.

15.9.c. The UK environmental review

15.9.8. The UK Treasury's Environmental Discount Rate Review reached a closely related conclusion but chose a different presentation. It rejected a generic lower environmental discount rate as an imprecise way to represent non-substitutability and increasing scarcity. Instead, it recommended improved environmental valuation and explicit uprating of environmental values over the appraisal horizon, an approach it described as consistent with the Dasgupta Review (HM Treasury 2026*c*, Dasgupta 2021). This recommendation also accords with the earlier natural-capital accounting report to the ONS by Freeman & Groom (2016*a*).

15.9.9. Economically, the Dutch effective-rate method and the Treasury relative-price method coincide under the assumptions of paragraph 15.3.7. Institutionally, however, explicit prices are preferable. They identify the source of the adjustment, allow different services to follow different paths, accommodate changing scarcity through time, and make it easier to detect double counting.

15.9.d. Natural-capital accounting at the ONS

15.9.10. The same structure appears in the UK natural-capital accounts. The Office for National Statistics reports physical ecosystem-service flows and annual monetary values separately from natural-asset values, which are calculated as the present value of expected future service flows. Its current quality and methods guide generally applies the Green Book schedule to projected monetary flows, while using the Green Book health schedule for air-pollution regulation, noise regulation, and recreation health benefits (Office for National Statistics 2025). This is a practical stock–flow framework: natural assets are stocks, ecosystem services are flows, and discounting is applied after those flows have been valued.

15.9.11. In two reports commissioned by the ONS, we recommended the relative-price approach explicitly. Freeman & Groom (2016*a*) recommend retaining the common Green Book discount schedule while increasing future natural-capital unit values to reflect scarcity and limited substitutability. Freeman et al. (2017) extend the recommendation across ONS valuations: for environmental and health assets, changing relative prices should be represented in the explicit pricing of costs and benefits rather than concealed in dual discount rates. The reports emphasize that the two methods are quantitatively equivalent under common assumptions, but that explicit prices separate scarcity and substitution from intertemporal welfare weighting.

15.9.12. The operational implication is that a natural-capital account should project physical service quantities and unit shadow prices separately. Their product is the monetary service flow to be discounted at the common consumption schedule. The ONS guide permits future asset values to reflect expected changes in physical quantities or unit prices, but where such evidence is unavailable it uses a transparent default projection. That fallback should not be interpreted as evidence that relative prices are constant. Service-specific price paths would bring the accounts more fully into line with the ONS-commissioned reports, the Green Book approach to relative prices, and the Treasury's environmental review.

15.9.13. The preceding sustainability analysis adds an important qualification. Monetary natural-capital asset values are not, by themselves, sustainability indicators. The ONS accounts should be interpreted alongside physical extent, condition, and ecosystem-service capacity. In particular, a rise in an asset value caused by scarcity-related revaluation should not be confused with positive genuine investment in the natural asset. CBA typically handles flows of goods and services. A separate analysis is required to assess stocks of wealth.

15.9.e. Health in the Green Book

15.9.14. Building on the Treasury evidence review by Freeman et al. (2018), current Green Book guidance (HM Treasury 2026*a*) decomposes the initial Social Time Preference Rate as

$$\begin{aligned} r_C &= \rho + \mu_C g_C = 3.5\%, \\ \rho &= \delta + L = 0.5\% + 1.0\% = 1.5\%, \\ \mu_C g_C &= 1.0 \times 2.0\% = 2.0\%. \end{aligned}$$

Health and life effects are discounted initially at the utility discount rate, or time-preference component, $\rho = 1.5$ per cent: the entire conventional wealth-effect term is removed (HM Treasury 2026*a*). This is not pure time preference alone, because ρ also contains the Green Book allowance L for catastrophic or systemic risk.

15.9.15. The health rule has an exact full relative-price interpretation:

$$\pi_H = \mu_C g_C = 2.0\%, \qquad r_H = r_C - \pi_H = \rho = 1.5\%.$$

It is equivalent to increasing the real shadow price of a fixed health unit by 2 per cent per year and discounting the resulting monetary value at the standard consumption rate. Current ENCA guidance states this equivalence explicitly (Department for Environment, Food and Rural Affairs 2026).

15.9.16. We recommend moving fully to this relative-price presentation. Health units such as life-years or QALYs should be assigned an explicit value path—initially one that reproduces the full wealth-effect deduction unless better evidence is adopted—and the resulting monetary values should be discounted using the standard consumption schedule. Under the current deterministic calibration, present values would be unchanged. The gain is conceptual and operational: physical health outcomes, their shadow prices, and intertemporal welfare weights would be visible separately. With this clear approach and the principle of one social discount rate across departments, fewer mistakes, e.g. double counting, may occur.

15.9.f. Transport and the value of time

15.9.17. UK transport appraisal already uses the explicit-price method. Department for Transport guidance allows work and non-work values of travel time to increase with real GDP per person, with an intertemporal income elasticity of approximately one for the principal values used in appraisal (Department for Transport 2026). The growing monetary value of an hour saved is then included in the annual benefit stream, which is discounted using the applicable Green Book schedule.

15.9.18. If ϵ_T is the income elasticity of the value of time and g_Y is real income growth per person, the relative-price increase is approximately

$$\pi_T = \epsilon_T g_Y,$$

and the effective rate applied to a fixed physical hour is $r_T = r_C - \pi_T$. Transport thus embodies the same principle as environmental and health appraisal: a future unit can warrant a lower effective discount rate because its relative monetary value rises, not because time preference differs by sector.

15.9.19. Freeman & Groom (2021*b*) emphasize that income growth is only one possible predictor of this relative price. In their stylized DeSerpa-type formulation, let necessary time for an activity X_t satisfy

$$T_t^w = b_t X_t,$$

where b_t is the time required per unit of the activity. Translating their curvature notation into the present

μ notation, the relative-price growth of time in the separable case is

$$\pi_{T,t} = \frac{d}{dt}\ln\left(\frac{U_{T,t}}{U_{C,t}}\right) = \mu_{CC,t}g_C - \mu_{TT,t}(g_b + g_X).$$

The term $g_b + g_X$ is the growth rate of necessary time. Time- or labour-saving technical change gives $g_b < 0$, while expansion of the time-using activity gives $g_X > 0$. The effective scarcity of discretionary time is therefore endogenous: innovation can relax the time requirement, while growth in time-using activities can tighten it.

15.9.20. The analogy with environmental relative pricing is useful, but the *supply side* differs. Environmental scarcity often reflects constrained natural regeneration, whereas the effective availability of time can be shifted by technology, work organization, and transport mode. The Department for Transport thinkpiece consequently recommends estimating the value-of-travel-time-saved uplift as a distinct empirical relative-price path, rather than assuming that it must rise one-for-one with income. Income growth is an appropriate uprating variable only to the extent that it predicts that path.

15.10. A recommendation for relative-price adjustments in CBA

15.10.1. We recommend the systematic use of explicit relative-price adjustments for non-consumption outcomes in government appraisal. The adjustments should reflect how the social values of environmental services, health, time, and other outcomes change as incomes, physical quantities, technologies, and substitution possibilities evolve. This approach is consistent with the existing Green Book framework, in which relative-price adjustments are already recognised, and is preferable to treating sector-specific discount rates as unexplained exceptions to a common appraisal rule.

15.10.2. The numerical adjustments should not be uniform across sectors. Environmental shadow prices depend on service-specific ecological trends, natural-capital stocks, thresholds, and substitution possibilities. Health shadow prices depend on income effects, health-capital dynamics, and health–consumption complementarity. The value of time depends on income, traveller and journey composition, the growth of time-using activities, and time-saving technological change. What is common is the accounting framework, not the numerical uplift.

15.10.a. A four-step appraisal framework

15.10.3. The recommended appraisal procedure has four steps. First, forecast the physical stocks affected by the intervention, the service flows those stocks generate, and the relevant production and ecological laws of motion. This stage should also identify legal constraints, critical thresholds, or sustainability requirements that may restrict the admissible set of options.

15.10.4. Second, estimate the shadow-price path for each material outcome. Where no separate sustainability constraint is imposed, these will be ordinary Ramsey-consistent scarcity prices reflecting preferences, physical scarcity, substitution possibilities, and so on. Where a sustainability or environmental target is incorporated directly into the feasible set, the prices should also include the value of the associated constraint.

15.10.5. Third, express the incremental physical outcomes in consumption-equivalent monetary units. As the theoretical section above shows, this can be done in either of two equivalent ways: using the affected natural-capital stocks and their complete accounting prices or the corresponding service-flow changes and their period-specific prices. The same effects should not be counted in both forms. In practice this will entail predictions of changes in service flows priced using a relative price adjustment reflecting scarcity.

15.10.6. Fourth, discount the resulting consumption-equivalent monetary flows using the appropriate STPR consumption discount schedule.

15.10.b. How sustainability enters the recommendation

15.10.7. The four-step procedure is a welfare-appraisal framework; it does not by itself guarantee sustainability. As established in Section 15.8, sustainability can enter in either of two ways.

15.10.8. Under the first, and currently more practicable, approach, appraisal uses ordinary Ramsey-consistent accounting prices and then applies a separate sustainability test. Provided the natural-capital and other asset prices are comprehensive, the same prices can be used to calculate genuine investment and changes in comprehensive wealth. Critical natural assets, ecological thresholds, and legal targets should additionally be tested.

15.10.9. The corresponding decision rule is that, for all relevant s,

$$NPV_t^R(P) > 0 \qquad \text{and} \qquad \mathcal{S}_s(P) \geq 0.$$

A positive NPV cannot compensate for failure of a binding sustainability test.

15.10.10. Under the second approach, sustainability is imposed directly on the feasible set and CBA uses the resulting sustainability-constrained accounting prices. These prices include the marginal value of target or constraint induced scarcity. This can be appropriate where interventions compete for a carbon budget, pollution ceiling, habitat requirement, or other quantifiable constraint. A positive constrained NPV is then meaningful only for a project that remains within the constrained feasible set.

15.10.11. The UK framework is most naturally interpreted as the first, two-step framework. Environmental effects are valued within welfare CBA, while the Green Book and ENCA require attention to natural-capital stocks, ecological condition, sustainable use, cumulative effects, tipping points, and legally binding environmental targets (HM Treasury 2026*c*, Department for Environment, Food and Rural Affairs 2026). This is theoretically coherent provided that the second-stage assessment is allowed to constrain the decision, uses a baseline consistent with the CBA, and does not duplicate benefits or costs already included in the monetary valuation.

15.10.12. This interpretation is consistent with the Government Economic Service review by Price & Durham (2009), which argued that appropriately specified shadow prices can allow social CBA to reflect the sustainability implications of marginal decisions while also recommending target-consistent prices, constraints on critical assets, and an additional asset check for non-marginal effects.

15.10.c. The consumption numeraire

15.10.13. The fourth step requires one qualification. Once a non-consumption outcome has been converted into consumption-equivalent units, it should in principle be discounted at the consumption-good rate implied by the relevant multi-good welfare model, rather than automatically at the simple one-good Ramsey rate. For environmental quality, this rate is

$$SDR_{C,t} = \delta + \mu_{CC,t}g_C + \mu_{CE,t}g_E,$$

rather than necessarily the approximation $\delta + \mu_{CC,t}g_C$. Environmental change can affect the marginal utility of consumption through the cross term $\mu_{CE,t}g_E$. Thus, even after environmental outcomes have been expressed in consumption units, the theoretically appropriate numeraire rate may itself depend on environmental change.

15.10.14. The preliminary estimates of Disque et al. (2026) suggest that this correction is small in their environmental application. They obtain $\widehat{\mu}_{CE} = 0.17$, which, under their illustrative environmental growth rate of $g_E = -0.135$ per cent, changes the consumption discount rate by approximately -0.02 percentage points. This is small relative to their estimated difference of approximately 1.55 percentage points between the consumption and environmental discount rates.

15.10.15. Their results therefore suggest that the principal error associated with the one-good Ramsey rule concerns the valuation of environmental outcomes rather than the consumption numeraire. Uprating environmental prices and then applying the simple consumption-based Ramsey schedule consequently appears to be a reasonable first approximation in this setting. It remains an approximation rather than a theoretical identity, and the corresponding cross effect should be examined separately for health, transport, and other sectors, where complementarity with consumption may be stronger.

15.10.d. Long-run discipline in relative-price schedules

15.10.16. For any non-consumption outcome X_t, a relative-price growth rate $\pi_{X,t}$ implies

$$p_{X,t} = p_{X,0} \exp\left(\int_0^t \pi_{X,s}\, ds \right).$$

A constant positive value of π_X therefore implies exponential growth in the shadow price. Although this may provide a reasonable approximation over a conventional appraisal horizon, there is little reason to assume that a single constant uplift will remain credible indefinitely.

15.10.17. The environmental case illustrates the issue. In a homothetic CES model, the unit environmental shadow price is not mechanically capped by an expenditure share. The corresponding environmental value, or expenditure-equivalent, share is nevertheless bounded:

$$s_{E,t} = \frac{p_{E,t}E_t}{C_t + p_{E,t}E_t} = 1 - \lambda_t.$$

When C_t/E_t rises and the elasticity of substitution satisfies $\chi < 1$, this share tends towards one. The result does not impose a finite upper bound on the unit shadow price. Rather, it shows that an indefinitely extrapolated relative-price path eventually assigns an overwhelmingly large share of total consumption-equivalent value to the environmental service.

15.10.18. Applied guidance should therefore not treat a default assumption—such as willingness to pay growing proportionately with income—as an immutable extreme-long-run law. Relative-price schedules should report their implied value shares at relevant appraisal horizons and should be supplemented by service-specific scarcity scenarios, periodic review points, and sensitivity analysis. Where appropriate, the schedule may need to taper, change with the state of the underlying asset, or be generated by a non-homothetic specification.

15.10.19. This flexibility is an important advantage of an explicit relative-price schedule over a permanent discount-rate deduction: the price path can be revised without changing the underlying rule for intertemporal welfare weighting.

15.10.e. The limits of monetary conversion

15.10.20. The relative-price method presupposes that the non-consumption outcome can be converted meaningfully into consumption-equivalent units. This is not always the case. Certain aspects of biodiversity, old-growth forests, endemic species, cultural assets, and essential health or ecological functions may be non-substitutable in the relevant range. Where an essential input is approaching a critical threshold, or where a loss cannot be compensated by any finite increase in consumption, Step 3 of the appraisal framework cannot be completed in the usual way.

15.10.21. In such cases appraisal guidance should require physical safeguards, constraints or safe-minimum standards alongside the usual CBA outputs of NPV and the benefit–cost ratio. The carbon values used in UK appraisal illustrate one version of this approach: they are linked to the marginal cost of achieving an emissions target rather than being derived solely from an unconstrained willingness-to-pay calculation. Similar target-consistent prices may be appropriate where biodiversity or other environmental constraints are sufficiently well specified. See Groom et al. (2026).

15.10.22. Relative prices reveal increasing scarcity, but they should not disguise cases in which the conditions for monetary commensuration break down. Where substitution is impossible, the correct response is not an arbitrarily large discount-rate adjustment but an explicit restriction on the admissible set of options.

15.11. Conclusion and final recommendations

15.11.1. The Ramsey framework has a larger role than the derivation of the conventional consumption discount rate. Ramsey-consistent accounting prices can be used both for marginal CBA and for testing whether continuation welfare is non-declining through genuine investment. Nevertheless, ordinary CBA does not ensure sustainability. Sustainability must additionally enter either as a separate, decision-binding wealth, natural-capital, or ecological test, or as a constraint whose shadow value is incorporated into the accounting prices used in appraisal.

15.11.2. The practical recommendation is as follows. Forecast physical stocks and service flows; estimate sector-specific shadow-price paths using the best available evidence; convert outcomes into consumption-equivalent values where such conversion is legitimate; and apply the consumption discount schedule. Accompany this appraisal either with a binding sustainability test or with accounting prices derived from a sustainability-constrained feasible set.

15.11.3. This approach is consistent with the direction chosen by the UK Treasury for environmental effects, recommended in the two ONS reports, and already used in transport appraisal. The Green Book's utility-rate treatment of health embodies this approach but should be presented as an explicit relative-price adjustment. The Netherlands' recent reviews likewise formalize the separation between discounting and relative scarcity.

15.11.4. Finally, zero substitutability marks a boundary to monetary appraisal. When an essential environmental or health input cannot be compensated by consumption, no finite discount-rate adjustment can make conversion valid. At that point public decision-making must impose physical safeguards, ecological thresholds, or other constraints. Relative prices are valuable because they reveal scarcity; they should not be used to disguise the cases in which substitution, and hence monetary commensuration, breaks down. Here, cost-based approaches may be eligible.

15.11.5. The final recommendations arising from this analysis for the review are:

R.19. Use explicit relative-price adjustments rather than sector-specific discount rates.

Non-consumption outcomes should be valued using explicit, sector-specific shadow-price paths and the resulting consumption-equivalent values should be discounted using the common consumption discount schedule. In the absence of more specific evidence, environmental unit values should increase in line with real income per person, corresponding to an income elasticity of willingness to pay of one. This provides a transparent default consistent with the existing treatment of health and the value of time in UK appraisal.

R.20. Further embed comprehensive-wealth accounting within the Green Book framework for sustainable growth.

The Green Book should further emphasise that appraisal should consider how interventions change the economy's productive asset base, including produced capital, educational human capital, health capital, and natural capital. A positive project NPV should therefore be accompanied by an assessment of whether the resulting path maintains or increases comprehensive wealth.

R.21. Consider *Target and Cost Analysis* to establish target-consistent biodiversity prices.

Non-monetisable objects such as biodiversity could be enter into CBA via a constraint on their depletion or a commitment for restoration. Carbon pricing currently follows this approach rather than having a specific discount rate. Binding government commitments, such as the Environment Act (2021) could be used to determine targets. (Groom et al. 2026).

16. Place-based objectives and welfare weights

16.1. Introduction

16.1.1. We now turn to place-based projects. We will argue that, from a discounting perspective, regionality itself is not a relevant considerations; the discount rate is not a function of geography. What matters for public economic appraisal is regional differences in consumption, and we will recommend that such differences should be treated via welfare weights in the numerator of the present value equation, as is current Green Book practice. This conclusion is supported by the respondents to our survey, as summarised in Table 16.1. There is a strong consensus across all cohorts, and particularly governmental respondents, that the discount rate should not be adjusted on this basis.

16.1.2. As one of the public sector respondents to our survey noted "One should not have a different discount rate according to place. This will likely make for overly complex economic cases. Instead, weights should be used on the benefits and costs to reflect a locations different socio-economic conditions" (paragraph A.2.10).

16.1.3. Another survey respondent "Adjusting the discount rate is not the correct way to handle it. Leaving aside all the practical difficulties of doing it well, if there are trade offs between efficiency and distributive impacts, those should be recognised explicitly in public choices" (paragraph A.2.10).

16.1.4. However, we stress from the outset that because, from a mathematical perspective, changes can be made in either the numerator or denominator of the PV equation to get the same present value, it is not factually incorrect to have a place-based discount rate. This is similar to the argument for environmental discount rates. It is not that place-based discount rates *can't* be done; rather it is just not the most appropriate way of doing it.

Regional discount rates

16.1.5. To demonstrate this equivalence, in Appendix F, we extend the Ramsey Rule to situations where there are regional representative agents with different consumption paths. This shows that the Ramsey Rule still holds subject to an adjustment term, which accounts for two things:

- The differences in consumption growth rates between different regions.
- The regions from which project funding comes, and the regions to which the project benefits are derived.

This collapses to the standard Ramsey Rule with no need for further adjustment if two conditions both hold:

- That all regional growth rates are the same. This is the standard condition needed for the aggregate representative agent.
- That each region pays for the project in proportion to the benefits that they receive from the project.

16.1.6. The first bullet loosely echoes a result in Emmerling et al. (2017). There, a relatively simple adjustment to the Ramsey Rule is suggested that accounts for the difference between the mean and the median rate of per capita growth across the population. The Ramsey Rule does not depend on regions being equally wealthy if the cost/benefit ratio for each region from the project are the same. Current regional inequality is no basis for changing the discount rate within this framework. Instead, the rate is adjusted because of differences in future growth rates, not levels.

16.1.7. The second bullet does relate to current consumption levels by region. The adjusted Ramsey Rule rate is lower than the standard Ramsey Rule if poorer regions have a higher benefit/cost ratio than wealthier regions from any given project. But this is, effectively, a simple wealth distributional benefit. Under isoelastic utility, social welfare is maximised when all members of the population have the same consumption level (because of decreasing marginal utility). Anything that therefore acts as a mechanism to distribute resources more evenly will have a perceived positive social value, which takes the form here of a lower discount rate.

16.1.8. But putting even such a stylised place-based discount rate into practice would be highly complex, and would violate a principle of this review that this variable should deal with maturity and risk only. We therefore turn instead to distributional weights in cost-benefit analysis as a problem of welfare-metric accounting.

	Count	Yes	No	Unsure
UK Academic	20	35%	50%	15%
Social Media	16	13%	63%	25%
SBCA	11	27%	55%	18%
Experts	11	36%	45%	18%
Total (non-government)	58	28%	53%	19%
Government	12	8%	75%	17%
Total (whole sample)	70	24%	57%	19%

Table 16.1.: Should the discount rate be adjusted for place-based objectives?

Welfare weights, discounting and reference points

16.1.9. Since the discount factor prices a unit change in a specified welfare metric, such as per capita or mean consumption, the *numerator* of the net present value calculation must be expressed in the same metric. Welfare weights can introduce accounting problems for CBA when they change the reference metric. The solution comes via some or all of three harmonising corrections: the initial level of the equity weight; the equity-weight reference unit when the NPV numerator and denominator use different reference paths; and a dynamic equity adjustment when affected groups evolve differently from the, typically national, reference path. These can be numerator corrections, akin to relative price adjustments, or equivalently folded into effective discount rates.

16.1.10. The Treasury Green Book approach illustrates the issue because the discount factor is based on mean real per capita consumption growth, through the Ramsey Rule, while its weights are national-median-referenced and use a different elasticity.

16.1.11. We recommend retaining the simple Ramsey Rule and making interpersonal, regional, place-based and dynamic distributional adjustments in the numerator. We also recommend harmonising the elasticities in interpersonal and intertemporal inequality aversion for similar accounting reasons.

16.2. Welfare weights

16.2.1. The use of welfare or equity weights in cost-benefit analysis (CBA) has seen renewed discussion in recent years. Two major recent government reviews of CBA have argued that growing concerns about regional inequalities, and the limited ability of existing policy frameworks to reduce them, justify a more central role for an explicit welfare account of the incidence of public interventions along the income distribution (OMB 2023*a*,*b*). At present, money-metric estimates of costs and benefits in CBA treat a pound to a rich region identically to a pound to a poor region and are blind to the possibility that society may place greater value on benefits and costs accruing to poorer people or places. Typically, the First and Second Welfare Theorems are invoked to separate efficiency from equity issues, with CBA focusing on the former. Arguments about the subjective nature of welfare weights are also raised as a barrier to their deployment. However, persistent concerns about inequality, alongside basic invocations of welfare theory and the more widespread use of welfare weights in CBA, suggest that inequality should become a more central component of public appraisal rather than an afterthought or an element of sensitivity analysis (Fleurbaey et al. forthcoming).

16.2.2. In this contribution we make two points. First, equity-weighted NPV should become the central result in project appraisal in the UK. This approach would not be a departure from current practice for two reasons: i) subjective or not, welfare weights are already used in public appraisal via the Ramsey Rule for social discounting, which stems from an explicit Social Welfare Function (SWF); ii) guidance on interpersonal equity weighting already exists in the Green Book (HM Treasury 2026*c*, Section 6), so default use of this guidance would be straightforward. Our second point is that, despite precedent and guidance, before equity-weighted NPV analysis becomes the default approach several adjustments are required because current guidance introduces a discrepancy in the units of account for costs and benefits in the numerator of the NPV calculation and the discount rate in the denominator.

16.2.3. The basic theorem is as follows. Suppose appraisal evaluates whether interventions increase or decrease the social objective reflected in the Social Welfare Function (SWF) $W(X)$ dependent on a welfare metric X. Accounting coherence requires that at any point in time costs and benefits be measured in terms of the welfare metric X and intertemporal comparisons are undertaken using a discount rate appropriate for the welfare metric X. If this accounting convention is broken then the measurement of welfare changes becomes problematic. Unfortunately, the use of welfare weights can introduce a disconnect between the *numerator* of the NPV calculation, where changes in the costs and benefits appear, and the *denominator* where the discount rate appears. Careful thought is required to ensure coherence between the two when welfare weights are used. Several problems can arise.

16.2.4. For instance, if appraisal uses the Ramsey Rule for discounting, the central welfare metric is mean per capita consumption and the intertemporal welfare weights compare per capita consumption at time t with the reference level at $t = 0$. In this context, we show that group-level intratemporal welfare weights should either use the same reference quantity or contain a correction to ensure compatibility of reference levels. Two kinds of time variation then have to be distinguished. First, national reference paths can evolve differently: for example, national mean consumption and national median income may grow at different rates. Second, the affected group, region or person may itself converge towards or diverge from the national reference path. Without accounting for these differences, the analysis compares intertemporal and intratemporal welfare units inconsistently. In sum, three corrections may be required to ensure welfare accounting integrity: i) a level correction for the initial welfare weight; ii) a national reference-path correction; and iii) a dynamic correction for the divergence or convergence of the affected group-level metric.

16.2.5. The UK Treasury Green Book guidance illustrates where accounting problems could arise in CBA. At present, the Green Book uses a discount rate suitable for real per capita consumption. However, its distributional weights are based on the ratio of national *median* equivalised income to the *median* equivalised income of the affected group. As we show, without appropriate corrections this approach is incorrect. The potential error is further exacerbated in the Green Book by the use of an elasticity of marginal utility of consumption of 1 in the Ramsey Rule discount rate, and an elasticity of 1.3 to capture welfare weights.

16.2.6. The main recommendation of the chapter is practical. The Green Book should retain the Ramsey denominator based on mean real per capita consumption unless it intends a larger change to the welfare framework. Interpersonal, regional, place-based and dynamic distributional corrections should then be placed in the numerator, so that project benefits are expressed in mean-consumption-reference units before being discounted. Equivalent discount-rate adjustments can be derived algebraically, as we show in Appendix F, but this is not our recommendation since group-specific or project-specific effective discount rates would be practically unwieldy and also potentially conceptually confusing.

16.2.7. The general principle here is that the social discount rate is determined by the social welfare function against which interventions are measured to establish social gains and losses. If the units in which money-metric benefits and costs are measured are different to what the discount rate prices intertemporally, then there is a welfare accounting problem that can lead to erroneous evaluations of welfare. Welfare weights can introduce this problem if not carefully accounted for, particularly if we expect inequality to increase or decrease over time. Once welfare weights are viewed through this lens, further complication can arise where welfare weights are established at the group level rather than the individual level. How well group-level approximations match the actual incidence of any given project determines the accuracy of the welfare weight correction, and hence the accuracy of the analysis. Group medians, group means, or other representative statistics perform differently depending on how the project benefits and costs are distributed within the group.

16.2.8. The following note provides a formal analysis of these points, starting with a *theorem* for welfare accounting in CBA, then describing the welfare weights and reference values in the Ramsey Rule and the ways in which ad hoc applications of welfare weighting can violate the theorem. We then argue that the Green Book currently breaches these welfare-accounting principles and characterise corrections that solve the problem. The recommendations are ultimately that welfare-weighted CBA should become the central approach to calculating NPV where place-based inequalities are material. However, further thought is needed on the use of median-referenced welfare weights and on the trajectories of different regions relative to the national average. Numerator corrections can handle these issues more clearly than adjustments to the discount rate.

16.3. A welfare-metric accounting theorem

16.3.1. Let F_t denote the distribution of consumption, income or consumption-equivalent welfare at time t, and let $c_t(\theta) > 0$ denote the welfare level of type θ. A welfare metric is a scalar representation of the period-t distribution whose unit changes are priced by the discount factor. Examples include mean consumption,

$$X_t = \bar{c}_t = \int_\Theta c_t(\theta)\ \mathrm{d}F_t(\theta),$$

Atkinson EDE consumption,

$$X_t = A_{\mu_I,t} = \left[\int_\Theta c_t(\theta)^{1-\mu_I}\ \mathrm{d}F_t(\theta)\right]^{1/(1-\mu_I)},$$

for $\mu_I \neq 1$, and the marginal-utility reference level,

$$X_t = Q_{\mu_I,t} = \left[\int_\Theta c_t(\theta)^{-\mu_I}\ \mathrm{d}F_t(\theta)\right]^{-1/\mu_I}.$$

These examples capture different aspects of what is considered a welfare-significant argument for a social welfare function. Mean per capita consumption is the typical argument underpinning the Ramsey Rule, EDE consumption has been proposed by Emmerling et al. (2017) among others as an argument for an alternative intertemporal welfare function, while Gollier (2013) proposes average marginal utility.

16.3.2. In general terms, suppose that intertemporal welfare is written as a function of welfare metric X,

$$W_X = \int_0^\infty e^{-\delta t} v_{\mu_T}(X_t)\ \mathrm{d}t,$$

where $v'_{\mu_T}(X) = X^{-\mu_T}$. The discount factor associated with a unit change in the metric X_t is

$$D_X(t) = \frac{e^{-\delta t} v'_{\mu_T}(X_t)}{v'_{\mu_T}(X_0)} = e^{-\delta t}\left(\frac{X_t}{X_0}\right)^{-\mu_T}.$$

16.3.3. For ease of exposition, we refer to this as being the *denominator* of the welfare calculation because the discount rate appears in the denominator of an NPV calculation, as opposed to estimates of the costs and benefits which appear in the *numerator*. The discount factor prices a unit change in X_t, perhaps caused by a public project or intervention.

16.3.4. Now suppose a project or intervention P generates a small money-metric benefit $b_t(\theta; P)$ for person- or group-type θ. In CBA this is typically a monetary measure of an individual welfare change. The social welfare change that this induces is determined by the social welfare function. Define the project-induced change in the welfare metric as

$$\Delta X_t(P) \approx \int_\Theta \frac{\partial X_t}{\partial c_t(\theta)} b_t(\theta; P)\ \mathrm{d}F_t(\theta).$$

16.3.5. The derivative $\partial X_t / \partial c_t(\theta)$ is the conversion term associated with the chosen metric that converts raw money-metric effects across all types θ into units of X_t. In many applications this conversion term is an equity or welfare weight. The welfare-consistent valuation of changes $b_t(\theta; P)$ is then

$$V_X(P) = \int_0^\infty D_X(t)\Delta X_t(P)\ \mathrm{d}t.$$

16.3.6. The theorem is therefore simple: if the denominator prices a unit change in the welfare metric X_t, the numerator must be the project-induced change measured in that same metric. Equivalently, each welfare metric has a specific discount factor and discount rate associated with it. In essence, we do not discount apples using the discount rate for oranges.

16.3.7. The theorem therefore describes how an appraisal can go wrong. Suppose the analyst uses $D_X(t)$ but discounts an unconverted raw benefit stream,

$$B_t(P) = \int_\Theta b_t(\theta; P)\ \mathrm{d}F_t(\theta),$$

yielding

$$\widetilde{V}(P) = \int_0^\infty D_X(t) B_t(P)\ \mathrm{d}t.$$

The resulting error is

$$V_X(P) - \widetilde{V}(P) = \int_0^\infty D_X(t)\left\{\Delta X_t(P) - B_t(P)\right\}\ \mathrm{d}t.$$

16.3.8. This error is zero only in special cases, such as when the project incidence happens to imply $\Delta X_t(P) = B_t(P)$, or when $B_t(P)$ has already been constructed in X_t-equivalent units. In welfare terms, the oranges have already been converted into apples. In general, applying a discount factor associated with one welfare metric to a numerator expressed in another metric amounts to a partial implementation of the underlying social welfare function.

16.4. The mean-consumption Ramsey denominator

16.4.1. The Green Book discounting framework is organised around expected growth in real per capita consumption. In the notation above, this means that the discount factor *denominator* prices a unit change in the metric $X_t = \bar{c}_t$. The associated social welfare representation is the familiar discounted utilitarian form

$$W^{\bar{c}} = \int_0^\infty e^{-\delta t} U_{\mu_T}(\bar{c}_t)\ \mathrm{d}t, \qquad U'_{\mu_T}(c) = c^{-\mu_T}.$$

The marginal welfare value of a small change in mean consumption at time t is

$$\frac{\partial W^{\bar{c}}}{\partial \bar{c}_t} = e^{-\delta t}\bar{c}_t^{-\mu_T}.$$

Thus the Green Book social welfare function, which as we show leads directly to the Ramsey Rule for the social discount rate, introduces welfare weights for the representative consumer at different points in time. A unit change in mean consumption at a later date is weighted by the product of pure time discounting and the marginal welfare value of mean consumption at that date.

16.4.2. Normalising by the reference marginal value of mean consumption at time zero gives

$$D_{\bar{c}}^{\mu_T}(t) = e^{-\rho t}\left(\frac{\bar{c}_t}{\bar{c}_0}\right)^{-\mu_T}.$$

This leads to the associated Ramsey Rule for the discount rate

$$r_{\bar{c}}(t) = -\frac{\mathrm{d}}{\mathrm{d}t}\log D_{\bar{c}}^{\mu_T}(t) = \rho + \mu_T g_{\bar{c}}(t), \qquad g_{\bar{c}}(t) = \frac{\dot{\bar{c}}_t}{\bar{c}_t}.$$

This is the appropriate social discount rate when mean real per capita consumption is the reference metric in the discounted utilitarian social welfare function.

16.4.3. The numerator measure of costs and benefits that is compatible with this denominator is a change in mean consumption. If F_t is normalised to integrate to one, then the value of project P described above is

$$\Delta\bar{c}_t(P) = \int_\Theta b_t(\theta; P)\ \mathrm{d}F_t(\theta).$$

The mean-consumption valuation is therefore

$$V_{\bar{c}}(P) = \int_0^\infty D_{\bar{c}}^{\mu_T}(t)\int_\Theta b_t(\theta; P)\ \mathrm{d}F_t(\theta)\ \mathrm{d}t.$$

16.4.4. From a welfare-accounting perspective this is internally coherent. However, as is well known, this approach is distributionally neutral within a period. All money-metric benefits occurring in the same period enter the numerator identically, regardless of to whom they accrue. For this reason interpersonal welfare weights are often proposed. As we now show, such weights must be chosen and applied in a way that preserves welfare-accounting coherence.

16.5. CBA with interpersonal welfare weights

16.5.1. This section provides the baseline welfare-accounting case by demonstrating the compatibility of the numerator and denominator in the standard Ramsey framework when interpersonal welfare weights are introduced. Alternative approaches can then be compared against this benchmark. In doing so, we demonstrate precisely what is meant by compatibility between the numerator and denominator according to the welfare-accounting theorem outlined above. We then investigate how this changes when intra- and intertemporal inequality aversion parameters differ, as is currently the case in the Green Book.

The baseline accounting case

16.5.2. In the first instance, assume that intratemporal and intertemporal inequality aversion are identical, $\mu_I = \mu_T = \mu$. The mean-reference welfare weight for type θ is

$$\omega_{\bar{c}}(t,\theta) = \left(\frac{\bar{c}_t}{c_t(\theta)}\right)^{\mu}.$$

This welfare weight converts type-specific consumption into common mean-consumption reference units. The welfare-weighted numerator of net benefits for project P is then

$$\widetilde{B}_{\bar{c},t}(P) = \int_\Theta \left(\frac{\bar{c}_t}{c_t(\theta)}\right)^{\mu} b_t(\theta; P)\ \mathrm{d}F_t(\theta).$$

The Ramsey discount factor converts mean consumption at time t into base-period mean-consumption units. Since both the welfare weights and the discount factor are normalised to mean consumption, this approach satisfies the welfare-metric theorem because the numerator and denominator are harmonised in terms of their reference units of account. With $b_t(\theta; P)$ interpreted as a money-metric net benefit, the corresponding weighted NPV is

$$V_{\bar{c},\mu}(P) = \int_0^\infty D_{\bar{c}}^{\mu}(t)\widetilde{B}_{\bar{c},t}(P)\ \mathrm{d}t.$$

16.5.3. Multiplying the mean-consumption discount factor and the mean-referenced equity weight gives

$$D_{\bar{c}}^{\mu}(t)\omega_{\bar{c}}(t,\theta) = e^{-\rho t}\left(\frac{\bar{c}_t}{\bar{c}_0}\right)^{-\mu}\left(\frac{\bar{c}_t}{c_t(\theta)}\right)^{\mu} = e^{-\rho t}\bar{c}_0^{\mu}c_t(\theta)^{-\mu}.$$

Thus, up to the fixed base-period unit $\bar{c}_0^{\mu}$, the combined intra- and intertemporal welfare weight for type θ is

$$M_{\bar{c},\mu}(t,\theta) \propto e^{-\rho t}c_t(\theta)^{-\mu}.$$

This is precisely the marginal social welfare value implied by the distributional extension of the same isoelastic Ramsey social welfare function,

$$W^D = \int_0^\infty e^{-\rho t} \int_\Theta u_\mu(c_t(\theta)) \, \mathrm{d}F_t(\theta) \, \mathrm{d}t.$$

16.5.4. In accounting terms, the discount factor prices a unit of mean-consumption-reference net benefit over time. The equity weight converts a type-specific money-metric benefit into that same mean-consumption-reference unit. Because both components use the same reference path and the same curvature parameter, the time-t reference path cancels. For this reason, the growth of the reference path does not enter the marginal welfare weight. Only the growth of type-specific consumption matters. No reference-path correction is required in this baseline case.

Intra and intertemporal inequality aversion differ

16.5.5. Now allow the intratemporal inequality-aversion parameter μ_I to differ from the intertemporal inequality-aversion parameter μ_T. This is the approach currently recommended in the Green Book and discussed theoretically in a number of places, including Emmerling et al. (2017). In this case the mean-reference welfare weights become

$$\omega_{\bar{c}}^{\mu_I}(t,\theta) = \left(\frac{\bar{c}_t}{c_t(\theta)}\right)^{\mu_I},$$

while the Ramsey discount factor, which forms the denominator of the appraisal calculation, remains

$$D_{\bar{c}}^{\mu_T}(t) = e^{-\rho t}\left(\frac{\bar{c}_t}{\bar{c}_0}\right)^{-\mu_T}.$$

When welfare weights are used, project P is therefore valued according to

$$V_{\bar{c},\mu_I,\mu_T}(P) = \int_0^\infty D_{\bar{c}}^{\mu_T}(t) \int_\Theta \left(\frac{\bar{c}_t}{c_t(\theta)}\right)^{\mu_I} b_t(\theta;P) \, \mathrm{d}F_t(\theta) \, \mathrm{d}t.$$

This collapses to the previous expression when $\mu_I = \mu_T$.

16.5.6. Again, up to a constant, the resulting marginal social welfare valuation becomes

$$M_{\bar{c},\mu_I,\mu_T}(t,\theta) \propto e^{-\rho t} \bar{c}_t^{\mu_I-\mu_T} c_t(\theta)^{-\mu_I}.$$

This expression is no longer the straightforward distributional extension of the social welfare function that motivates the Ramsey Rule. In fact, the social welfare function that implies this marginal valuation is

$$\mathcal{W}_{\bar{c}}^{\mu_I,\mu_T} = \int_0^\infty e^{-\rho t} \bar{c}_t^{\mu_I-\mu_T} \int_\Theta u_{\mu_I}(c_t(\theta)) \, \mathrm{d}F_t(\theta) \, \mathrm{d}t,$$

where $\bar{c}_t$ is treated as the exogenous reference path used for marginal project appraisal.

16.5.7. This two-parameter approach can be interpreted as a marginal welfare-accounting convention around the mean consumption reference path. It is not an extension of the typical single-curvature isoelastic social welfare function implied in the Green Book. The use of different intertemporal and interpersonal inequality aversion parameters is reasonable within this accounting framework, but this accounting convention (the underlying SWF) should be stated explicitly as the welfare objective.

16.5.8. There is, however, an important qualification. The integration of the marginal social welfare into an implied SWF has assumed that $\bar{c}_t$ is an exogenous reference path. In this case the two-parameter case provides a coherent marginal accounting convention for project appraisal. Since marginal analysis is the basis of CBA, this is not a problematic assumption necessarily. Yet if $\bar{c}_t$ is treated as the endogenous mean of the same distribution over which project benefits are being varied, the implied marginal weights need not be integrable to a single global social welfare function unless $\mu_I = \mu_T$. The reason is that in that case, the marginal social welfare for type θ depends both on the mean and changes in another person's consumption that also change the mean. The cross-effects need not be symmetric making integrability and representation as a SWF problematic. If these technical arguments are not enough, there exist normative arguments for treating inequality within time periods and across time periods in the same way.

16.5.9. A reasonable conclusion therefore, is that interpersonal and intertemporal inequality aversion might be better represented by the same parameter to avoid these accounting and theoretical complications. Further complications arise when reference levels between numerator and denominator diverge.

16.6. Reference units and dynamic equity corrections

16.6.1. If the welfare weights used in intra- and intertemporal analysis are normalised in different ways, ensuring compatibility between the numerator and denominator in CBA becomes more complicated. We explore this possibility because it already arises in some government guidance.

16.6.2. Let R_t denote the reference-path welfare metric used in the equity weights and let S_t denote the reference path priced by the discount-rate denominator. As discussed above, examples include mean consumption $\bar{c}_t$, national median equivalised income $m_{N,t}$, EDE consumption $A_{\mu_I,t}$, or the marginal-utility reference level $Q_{\mu_I,t}$. Different summary measures of the central tendency of the income or consumption distribution, such as the mean and median, are measured in the same consumption-equivalent monetary units. At a given date, changing the reference from the national median to the national mean changes the unit of account but not the underlying object being valued. Accounting problems emerge when reference paths evolve at different rates over time or when the numerator is expressed in the units of one reference path while the denominator prices another. Corrections are then required to maintain accounting integrity either in the numerator or, equivalently, in the denominator.

Reference path mismatch and possible corrections

16.6.3. Consider an R_t-reference welfare weight for group i, represented by a group statistic $z_{i,t}$,

$$\omega_{i|R,t}^{\mu_I} = \left(\frac{R_t}{z_{i,t}}\right)^{\mu_I}.$$

This weight expresses raw money-metric benefits in R_t-reference units. The corresponding discount factor is

$$D_R^{\mu_T}(t) = e^{-\rho t}\left(\frac{R_t}{R_0}\right)^{-\mu_T}.$$

Suppose instead that appraisal uses the discount factor

$$D_S^{\mu_T}(t) = e^{-\rho t}\left(\frac{S_t}{S_0}\right)^{-\mu_T},$$

which prices S_t-reference units. A correction is then required to preserve accounting consistency.

16.6.4. If the objective is to express the numerator in S_t-reference units while retaining the S_t-denominator, the appropriate adjustment is

$$\kappa_{S|R}(t) = \left(\frac{S_t}{R_t}\right)^{\mu_I}.$$

Multiplying the original welfare weight by this factor yields

$$\kappa_{S|R}(t)\omega_{i|R,t}^{\mu_I} = \left(\frac{S_t}{z_{i,t}}\right)^{\mu_I},$$

so that all benefits are now expressed directly in S_t-reference units. This is the natural correction when the analyst wishes to retain the discounting framework embodied in $D_S^{\mu_T}(t)$.

16.6.5. A different correction is required if the objective is to preserve the original R_t-reference valuation while continuing to discount using $D_S^{\mu_T}(t)$. In that case the relevant adjustment is the ratio of the two discount factors,

$$\chi_{R|S}(t) = \frac{D_R^{\mu_T}(t)}{D_S^{\mu_T}(t)} = \left[\frac{R_t/R_0}{S_t/S_0}\right]^{-\mu_T}.$$

This factor preserves the original R_t-reference valuation while expressing it through the S_t-reference denominator.

16.6.6. The two approaches are equivalent. The correction may be applied directly to the numerator, or folded into an effective discount rate. This mirrors the relative-price literature in environmental discounting, where changing relative prices can either be applied to the benefit stream or incorporated into an adjusted discount rate Hoel & Sterner (2007), Gollier (2010), Drupp et al. (2024).

16.6.7. In the present setting, the relative price is the conversion factor between two welfare-reference units. Since

$$r_R(t) - r_S(t) = \mu_T\{g_R(t) - g_S(t)\},$$

and

$$-\frac{\mathrm{d}}{\mathrm{d}t}\log\chi_{R|S}(t) = \mu_T\{g_R(t) - g_S(t)\},$$

the same correction can be implemented either through the numerator or through an equivalent adjustment to the discount rate.

16.6.8. If $\chi_{R|S}(t)$ is constant, changing reference paths merely alters the scale in which values are reported and leaves project rankings unchanged. If $\chi_{R|S}(t)$ varies over time, however, omitting the correction can alter NPVs, benefit-cost ratios and project rankings. The resulting valuation error is

$$V_R(P) - \widetilde{V}_{R|S}(P) = \int_0^\infty D_S^{\mu_T}(t)\{\chi_{R|S}(t) - 1\}\sum_i \omega_{i|R,t}^{\mu_I} B_{i,t}(P)\ \mathrm{d}t.$$

This is the formal sense in which a mismatch between reference paths can lead to an incorrect welfare assessment.

Changes in group outcomes relative to the reference unit

16.6.9. There is a further source of time variation when the affected group changes position relative to the reference path. It is useful first to consider the case in which the interpersonal and intertemporal curvature parameters are identical, $\mu_I = \mu_T = \mu$.

Let $z_{i,t}$ denote the representative income or consumption path of affected group i. The corresponding R-reference welfare weight is

$$\omega^{\mu}_{i|R,t} = \left(\frac{R_t}{z_{i,t}}\right)^{\mu}.$$

To separate the initial distributional position of the group from subsequent changes in its relative welfare position, write the welfare weight as the product of an initial equity level and a dynamic adjustment. The dynamic component is

$$\frac{\omega^{\mu}_{i|R,t}}{\omega^{\mu}_{i|R,0}} = \left[\frac{R_t/R_0}{z_{i,t}/z_{i,0}}\right]^{\mu}.$$

This term captures changes in the welfare position of group i relative to the reference path over time.

16.6.10. Suppose that the denominator prices S_t-reference units rather than R_t-reference units. As shown above, the corresponding reference-unit correction is

$$\chi^{\mu}_{R|S}(t) = \left[\frac{R_t/R_0}{S_t/S_0}\right]^{-\mu}.$$

Combining the dynamic equity adjustment with the reference-unit conversion gives the full time-varying numerator correction

$$\Pi^{\mu}_{i|R,S}(t) = \chi^{\mu}_{R|S}(t)\frac{\omega^{\mu}_{i|R,t}}{\omega^{\mu}_{i|R,0}} = \left[\frac{S_t/S_0}{z_{i,t}/z_{i,0}}\right]^{\mu}.$$

A notable feature of this expression is that the intermediate reference path R_t disappears entirely. Once the reference-unit conversion has been undertaken, the only remaining dynamic adjustment reflects the evolution of the affected group relative to the discount-rate reference path S_t.

16.6.11. The same correction can alternatively be incorporated into the denominator through an effective discount factor,

$$D^{\text{eff}}_{i|R,S}(t) = D^{\mu}_{S}(t)\Pi^{\mu}_{i|R,S}(t).$$

The associated effective discount rate is

$$r^{\text{eff}}_{i|R,S}(t) = r_S(t) - \frac{\mathrm{d}}{\mathrm{d}t}\log \Pi^{\mu}_{i|R,S}(t) = r_S(t) - \mu\{g_S(t) - g_i(t)\}.$$

Using the Ramsey expression $r_S(t) = \rho + \mu g_S(t)$, this simplifies to

$$r^{\text{eff}}_{i|R,S}(t) = \rho + \mu g_i(t).$$

Thus, once the initial equity level has been separated out, the remaining dynamic adjustment depends only on the growth path of the affected group. The reference path used in the original welfare weights no longer matters. The adjustment can be applied either in the numerator or in the denominator; these are simply two representations of the same welfare-accounting correction.

16.6.12. The situation changes when the interpersonal and intertemporal curvature parameters differ. In that case the full dynamic correction becomes

$$\Pi^{\mu_I,\mu_T}_{i|R,S}(t) = \left[\frac{R_t/R_0}{S_t/S_0}\right]^{-\mu_T}\left[\frac{R_t/R_0}{z_{i,t}/z_{i,0}}\right]^{\mu_I},$$

or equivalently,

$$\Pi^{\mu_I,\mu_T}_{i|R,S}(t) = \left[\frac{S_t/S_0}{z_{i,t}/z_{i,0}}\right]^{\mu_I}\left[\frac{R_t/R_0}{S_t/S_0}\right]^{\mu_I-\mu_T}.$$

Unlike the single-curvature case, the intermediate reference path R_t no longer cancels. The correction therefore depends not only on the relative evolution of the affected group and the discount-rate reference path, but also on the evolution of the welfare-weight reference path itself.

16.6.13. The corresponding effective discount rate is

$$r^{\text{eff}}_{i|R,S}(t) = r_S(t) + \mu_T\{g_R(t) - g_S(t)\} - \mu_I\{g_R(t) - g_i(t)\}.$$

Substituting $r_S(t) = \rho + \mu_T g_S(t)$ yields

$$r^{\text{eff}}_{i|R,S}(t) = \rho + (\mu_T - \mu_I)g_R(t) + \mu_I g_i(t).$$

This is a group-specific effective discount rate because it depends on the growth path of the affected group. When $\mu_I = \mu_T = \mu$, it collapses immediately to

$$r^{\text{eff}}_{i|R,S}(t) = \rho + \mu g_i(t),$$

recovering the single-curvature result.

16.6.14. The distinction between the two cases is important. With a single curvature parameter, reference-path conversion is purely a unit-of-account operation and the intermediate reference path disappears from both the numerator correction and the effective discount rate. With separate interpersonal and intertemporal curvature parameters, the reference path remains part of the correction. The resulting welfare weights should therefore be interpreted as a welfare-accounting convention around a particular reference path rather than as the direct implication of a single-curvature social welfare function. This provides a further argument for harmonising interpersonal and intertemporal curvature parameters in practical appraisal.

16.6.15. For policy guidance, however, the practical recommendation is unchanged. It is generally clearer to retain a common Ramsey denominator based on mean real per capita consumption and place initial equity levels, reference-unit conversions and dynamic equity corrections in the numerator, rather than introducing project-specific or group-specific discount rates.

16.7. Mean and median reference paths

16.7.1. In this section we explore the implications of the previous analysis looking at the current guidance in the HM Treasury Green Book (HM Treasury 2026*c*). Here, the distinction between the reference path for the discount rate, S_t, and the welfare weights, R_t, is different, the former being mean real per capita consumption and the latter national median equivalised income. Furthermore, the welfare weights define the group i outcome $z_{i,t}$ as median income. This allows us to define clearly the practical implications of our welfare-accounting theorem and isolate the possible corrections required.

Mean and median reference paths: the Green Book case

16.7.2. Let $\bar{c}_t$ denote mean real per capita consumption. The Ramsey discount factor used in the Green Book is interpreted as pricing a unit change in this mean-consumption reference path:

$$D_{\bar{c}}^{\mu_T}(t) = e^{-\rho t}\left(\frac{\bar{c}_t}{\bar{c}_0}\right)^{-\mu_T}.$$

If distributional benefits are to be expressed in the same mean-consumption reference units, a natural starting point for a compatible individual-level welfare weight can be based on the deviation of consumption for type θ from the national mean income:

$$\omega_{\bar{c}}^{\mu_I}(t,\theta) = \left(\frac{\bar{c}_t}{c_t(\theta)}\right)^{\mu_I}.$$

More generally, at group-level, if group i is represented by some consumption or income statistic $z_{i,t}$, the corresponding group approximation is

$$\omega_{i|\bar{c},t}^{\mu_I} = \left(\frac{\bar{c}_t}{z_{i,t}}\right)^{\mu_I}.$$

In each case the reference path is mean consumption, so the denominator prices a unit of mean consumption, and the numerator is converted into mean-consumption-reference units.

16.7.3. The Green Book distributional weight instead uses a national-median reference. Let $m_{N,t}$ denote national median equivalised income and $m_{i,t}$ the median equivalised income of affected group i. The Green Book-style group weight is

$$\omega_{i|N,t}^{\mu_I} = \left(\frac{m_{N,t}}{m_{i,t}}\right)^{\mu_I}.$$

It values a marginal pound to the median member of group i relative to a marginal pound to the national median person. However, if this is the welfare unit used in the numerator, the matching denominator is not $D_{\bar{c}}^{\mu_T}(t)$, but

$$D_N^{\mu_T}(t) = e^{-\delta t}\left(\frac{m_{N,t}}{m_{N,0}}\right)^{-\mu_T}.$$

Mean consumption and national median income are both monetary, consumption-equivalent reference paths. Thus, at a single date, using the median rather than the mean is mostly a change in the reference point. However, an accounting problem arises over time if $m_{N,t}$ and $\bar{c}_t$ grow at different rates. A national-median-reference numerator discounted with a mean-consumption denominator, as is currently the case in the Green Book, then requires a reference-path correction in order to ensure consistent welfare accounting. The reference-path correction is

$$\chi_{N|\bar{c}}^{\mu_T}(t) = \frac{D_N^{\mu_T}(t)}{D_{\bar{c}}^{\mu_T}(t)} = \left[\frac{m_{N,t}/m_{N,0}}{\bar{c}_t/\bar{c}_0}\right]^{-\mu_T}.$$

Consistent welfare evaluation using a national-median-reference valuation can be written either as

$$V_N(P) = \int_0^\infty D_N^{\mu_T}(t)\sum_i\left(\frac{m_{N,t}}{m_{i,t}}\right)^{\mu_I} B_{i,t}(P)\ \mathrm{d}t,$$

with $D_N^{\mu_T}(t)$ the appropriate discount factor for median consumption, or equivalently as

$$V_N(P) = \int_0^\infty D_{\bar{c}}^{\mu_T}(t)\chi_{N|\bar{c}}^{\mu_T}(t)\sum_i\left(\frac{m_{N,t}}{m_{i,t}}\right)^{\mu_I} B_{i,t}(P)\ \mathrm{d}t.$$

Here $D_{\bar{c}}^{\mu_T}(t)$ is the discount factor associated with the Ramsey Rule and the correction $\chi_{N|\bar{c}}^{\mu_T}(t)$ is a relative-price-type adjustment to the numerator. It converts national-median-reference weighted benefits into the units priced by the mean-consumption Ramsey denominator.

16.7.4. The same adjustment can be expressed in the denominator rather than in the numerator. If benefits are left in national-median-reference units, the compatible discount rate is the national-median-reference rate, $r_N(t)$, rather than the mean-consumption Ramsey rate, $r_{\bar{c}}(t)$. Since

$$D_{\bar{c}}^{\mu_T}(t)\chi_{N|\bar{c}}^{\mu_T}(t) = D_N^{\mu_T}(t),$$

the corresponding rate adjustment to the standard Ramsey discount rate is

$$r_N(t) = r_{\bar{c}}(t) - \frac{\mathrm{d}}{\mathrm{d}t}\log\chi_{N|\bar{c}}^{\mu_T}(t) = r_{\bar{c}}(t) + \mu_T\left\{g_N(t) - g_{\bar{c}}(t)\right\}.$$

Thus, if national median income grows more slowly than mean consumption, $g_N(t) < g_{\bar{c}}(t)$, the national-median-reference discount rate is lower than the mean-consumption Ramsey rate. This is the denominator version of the equivalent numerator correction that would be framed as a relative-price adjustment $\chi_{N|\bar{c}}^{\mu_T}(t)$. The value of net benefits for group i could alternatively be assumed to grow at the rate $-\mu_T\left\{g_N(t) - g_{\bar{c}}(t)\right\}$. Emmerling et al. (2017) provide evidence that median growth has lagged behind mean growth in the UK over the past 30 years or so. If this trend were expected to continue, this would imply that a lower discount rate or a positive relative-price correction would be appropriate for group i. Without these corrections, mean-based Ramsey discounting would therefore discount national-median-reference benefits too heavily.

16.7.5. As a precursor to our ultimate proposal for the Green Book, note that any welfare weight that has mean consumption as the reference quantity can be used to solve the reference-path issue. This means that both $\left(\frac{\bar{c}_t}{m_{i,t}}\right)^{\mu_I}$ and $\left(\frac{\bar{c}_t}{c_{i,t}}\right)^{\mu_I}$ are eligible corrections. Each of these changes the reference from the national median person to the mean-consumption reference person. The chapter's preferred recommendation is to start with the Ramsey denominator $D_{\bar{c}}^{\mu_T}(t)$ and then select welfare weights that ensure that the numerator is measured in mean-consumption-reference units. The Green Book median-to-median weight can be made compatible, but as we show below, this requires more complex corrections.

16.7.6. Figure 16.1 illustrates how the static mean- and median-referenced weights compare for different levels $\mu_I \in \{1, 1.3, 2\}$ using a lognormal distribution calibrated to a UK median disposable income of \$36,700 and a disposable-income Gini coefficient of 32.9%, values reported by the ONS for FYE 2024. The mean exceeds the median under the right-skewed calibration, so mean-reference weights lie above national-median-reference weights. Nevertheless, their relative ranking across recipients is unchanged.

Time varying welfare weights: A Green Book example

16.7.7. Beyond relative changes in the reference paths R_t and S_t, median and mean in the current example, changes over time in the welfare weights that arise due to the group-level outcome $z_{i,t}$ relative to the reference level S_t also need to be taken into account. Recall that $m_{N,t}$ is national median equivalised income, $\bar{c}_t$ is real per capita consumption, and $m_{i,t}$ is the median equivalised income of affected group i. In this case, the reference-unit correction from national-median units is the same as above. However, in addition to this, a correction is needed to account for the time path of the equity weight itself.

16.7.8. In the case of the Green Book, the welfare weight is $\omega_{i|N,t} = \left(\frac{m_{N,t}}{m_{i,t}}\right)^{\mu_I}$. Splitting this into an initial distributional level and a dynamic component gives

$$\omega_{i|N,t} = \omega_{i|N,0}\left[\frac{\omega_{i|N,t}}{\omega_{i|N,0}}\right],$$

where

$$\frac{\omega_{i|N,t}}{\omega_{i|N,0}} = \left[\frac{m_{N,t}/m_{N,0}}{m_{i,t}/m_{i,0}}\right]^{\mu_I}.$$

This term is greater than one when the affected group becomes poorer relative to the national median, and less than one when it becomes richer relative to the national median.

16.7.9. Define the combined reference-path and welfare-weight correction as

$$\Pi_{i|N,\bar{c}}^{\mu_I,\mu_T}(t) = \chi_{N|\bar{c}}(t)\frac{\omega_{i|N,t}}{\omega_{i|N,0}}.$$

Substituting the two components yields

$$\Pi_{i|N,\bar{c}}^{\mu_I,\mu_T}(t) = \left[\frac{\bar{c}_t/\bar{c}_0}{m_{i,t}/m_{i,0}}\right]^{\mu_I}\left[\frac{m_{N,t}/m_{N,0}}{\bar{c}_t/\bar{c}_0}\right]^{\mu_I-\mu_T}.$$

If $\mu_I = \mu_T = \mu$, the final term disappears and the correction collapses to

$$\Pi_{i|N,\bar{c}}^{\mu,\mu}(t) = \left[\frac{\bar{c}_t/\bar{c}_0}{m_{i,t}/m_{i,0}}\right]^{\mu}.$$

This is exactly the dynamic mean-reference welfare-weight correction

$$\Pi_{i|N,\bar{c}}^{\mu,\mu}(t) = \frac{\omega_{i|\bar{c},t}}{\omega_{i|\bar{c},0}}, \qquad \omega_{i|\bar{c},t} = \left(\frac{\bar{c}_t}{m_{i,t}}\right)^{\mu}.$$

Thus, when the same curvature parameter is used for interpersonal and intertemporal welfare comparisons, the national median reference path drops out once

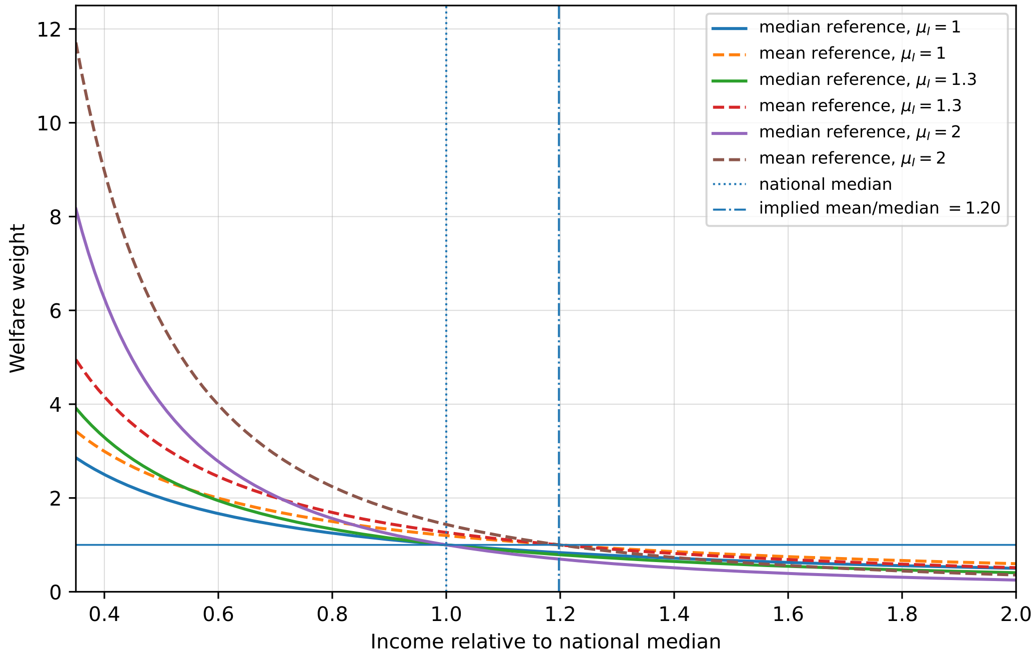

Figure 16.1.: Mean-reference and median-reference equity weights for different values of μ_I under an illustrative lognormal UK calibration. The horizontal axis is income relative to the national median. The mean-reference weights are higher by a constant factor within a given period because the mean exceeds the median.

the appropriate correction is made. Using national-median-reference weights plus the reference-unit correction is equivalent to using mean-reference dynamic weights directly. When $\mu_I \neq \mu_T$, this cancellation no longer occurs: the path of the national median relative to mean consumption remains in the correction. This finding informs the recommendations section below.

16.7.10. This result highlights an additional advantage of using a common curvature parameter for interpersonal and intertemporal welfare comparisons. When $\mu_I = \mu_T$, the national-median reference path disappears from the welfare-accounting correction entirely. The Green Book median-reference approach then becomes equivalent to a direct mean-reference approach once the appropriate reference-path correction has been applied. In contrast, when $\mu_I \neq \mu_T$, the path of the national median relative to mean consumption remains part of the correction. The use of a common curvature parameter therefore simplifies both the accounting and the interpretation of distributional appraisal.

16.7.11. Figure 16.2 illustrates this decomposition for the Green Book case, where the welfare-weight reference path is national median equivalised income, $R_t = m_{N,t}$, and the discount-rate reference path is mean real per capita consumption, $S_t = \bar{c}_t$. The affected group is represented by $m_{i,t}$. Figure 16.2 separates the reference-unit correction, $\chi_{N|\bar{c}}(t)$, from the dynamic national-median-reference welfare-weight correction, $\frac{\omega_{i|N,t}}{\omega_{i|N,0}}$. Their product is the full time-varying correction when national-median weights are used with the mean-consumption denominator: $\Pi^{\mu_I,\mu_T}_{i|N,\bar{c}}(t) = \chi_{N|\bar{c}}(t)\frac{\omega_{i|N,t}}{\omega_{i|N,0}}$. When $\mu_I = \mu_T = \mu$, this product collapses to the direct mean-reference dynamic correction $\frac{\omega_{i|\bar{c},t}}{\omega_{i|\bar{c},0}}$. The shaded wedge in Figure 16.2 shows the discrepancy that arises when $\mu_I \neq \mu_T$. To be clear, this wedge should not be interpreted as a failure of numerator–denominator equivalence: the same correction can still be placed in the numerator or folded into an effective discount rate. Rather, the shaded area shows that, once different curvature parameters are used, the national-median-reference route and the direct mean-reference route imply different, albeit more complex, welfare-accounting conventions.

16.8. From individual weights to group weights

16.8.1. The welfare-accounting theorem developed above is naturally defined at the individual level. In principle, if the incidence of a project across individuals is known, welfare weights should be applied directly to those individual impacts. In practice, however, appraisals rarely observe the distribution of benefits and costs at such a detailed level. Instead, analysts typically observe aggregate impacts on broad groups, such as regions, income classes, demographic groups or places. This creates a separate aggregation problem: how should an individual-level welfare weight

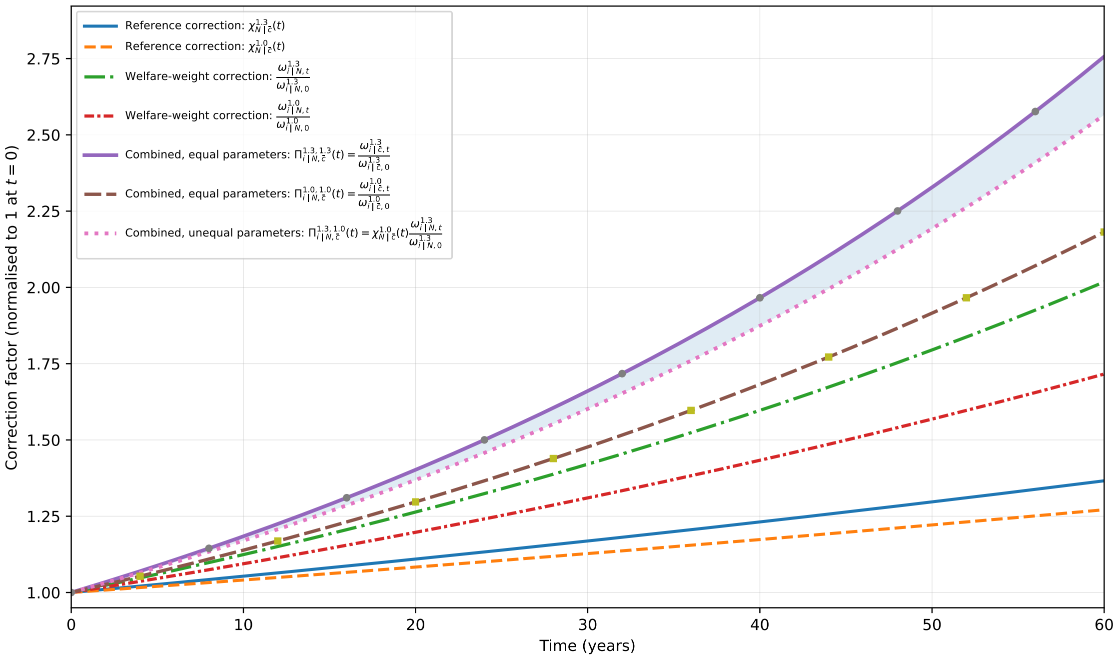

Figure 16.2.: Reference corrections, dynamic welfare-weight corrections and combined corrections in the Green Book mean–median example. The blue and orange lines show the reference correction from national-median-reference units to mean-consumption-reference units for $\mu_T = 1.3$ and $\mu_T = 1.0$, respectively. The welfare-weight correction lines show the change in the affected group's position relative to the national median. The combined lines multiply the reference correction by the dynamic welfare-weight correction. In the equal-parameter cases, $(\mu_I, \mu_T) = (1.3, 1.3)$ and $(1.0, 1.0)$, the combined correction is identical to the direct mean-reference dynamic correction. In the unequal parameter case, $(\mu_I, \mu_T) = (1.3, 1.0)$, the cancellation no longer holds. The shaded area shows the resulting divergence between the equal-parameter $(1.3, 1.3)$ correction and the unequal-parameter Green Book-style correction. The illustrative growth assumptions are $g_{\bar{c}} = 2.0\%$, $g_N = 1.6\%$, and $g_i = 0.7\%$.

be summarised as a single group-level weight?

16.8.2. The use of group-level welfare weights in the Green Book can be understood as a practical response to these data limitations. Rather than observing the income of every affected individual, the analyst is assumed to observe a representative statistic for the affected group. The use of median equivalised income is therefore best interpreted as an attempt to represent a typical member of the group. Whether the median is the most appropriate representative statistic depends on how project benefits and costs are distributed within the group. The choice between means, medians and other summary statistics is therefore not merely a statistical issue. It is also an implicit assumption about project incidence.

16.8.3. Let Θ_i denote affected group i. Suppose that the welfare framework implies individual-level weights $\omega(t, \theta)$. Under mean-reference CRRA weighting,

$$\omega(t, \theta) = \left(\frac{\bar{c}_t}{c_t(\theta)} \right)^{\mu_I}.$$

The exact weighted benefit accruing to group i is

$$\int_{\Theta_i} \omega(t, \theta) b_t(\theta; P) \; \mathrm{d}F_{i,t}(\theta).$$

If the full incidence $b_t(\theta; P)$ is known, there is no need to approximate the group weight. The analyst should use the individual weights. If only the aggregate group benefit

$$B_{i,t}(P) = \int_{\Theta_i} b_t(\theta; P) \; \mathrm{d}F_{i,t}(\theta)$$

is observed, the effective group weight that reproduces the individual-level calculation is

$$\omega^{\text{eff}}_{i,t}(P) = \frac{\int_{\Theta_i} \omega(t, \theta) b_t(\theta; P) \; \mathrm{d}F_{i,t}(\theta)}{\int_{\Theta_i} b_t(\theta; P) \; \mathrm{d}F_{i,t}(\theta)}.$$

Thus the appropriate group-level welfare weight is generally a benefit-incidence-weighted average of the individual weights within the group. It depends on how the project's benefits and costs are distributed within the group.

16.8.4. Several common approximations correspond to different implicit incidence assumptions. If benefits are distributed equally per person within the group, then

$$\omega^{\text{eff}}_{i,t} =_i [\omega(t,\theta)].$$

If benefits are concentrated around the group's median member, then a group-median approximation may be appropriate:

$$\omega^{\text{median}}_{i,t} = \left(\frac{R_t}{m_{i,t}}\right)^{\mu_I}.$$

If benefits are concentrated around the group's mean member, then a group-mean approximation may be appropriate:

$$\omega^{\text{mean}}_{i,t} = \left(\frac{R_t}{\bar{c}_{i,t}}\right)^{\mu_I}.$$

If benefits are proportional to income, the appropriate group weight is

$$\omega^{\text{income}}_{i,t} = \frac{{}_i[\omega(t,\theta)c_t(\theta)]}{{}_i[c_t(\theta)]}.$$

The group statistic used in the weight is therefore not innocuous. It is an assumption about incidence. A group-median weight does not generally reproduce the average of individual marginal welfare weights within the group. Nor does a group-mean weight. Each is an approximation that is appropriate only under particular assumptions about where project benefits fall inside the group.

A lognormal simulation of group-weight approximations

16.8.5. To illustrate the group-aggregation issue, consider a lognormal income distribution calibrated to the same UK median and Gini coefficient used above. The distribution is divided into income quintiles. For each quintile, Figure 16.3 and Table 16.2 compare four objects: the exact average individual weight under a mean-reference framework, the national-mean-over-quintile-mean approximation, the national-mean-over-quintile-median approximation, and the national-median-over-quintile-median approximation used in Green Book-style weights.

16.8.6. The comparison separates two issues. The difference between national mean and national median as the reference is a common within-period scale factor. The difference between a quintile mean, a quintile median and the exact average of individual weights is an approximation issue. In the UK-like lognormal calibration, the national-mean-over-quintile-mean weights are closer to the exact average individual weights than the median-over-median Green Book-style weights for much of the distribution. This does not prove that group means are always the right statistic. It shows that the best group statistic depends on the incidence of the project and on the shape of the within-group distribution.

16.8.7. The bottom quintile shows the largest difference. The exact average individual mean-reference weight is 4.080, while the national-median-over-quintile-median weight is 2.721. Part of this difference is the national median reference rather than the mean reference; the remainder is the use of the quintile median rather than the average of individual weights within the quintile. For broad per-person benefits in a skewed group, a group-median weight can understate the welfare value of benefits to poorer members of the group. This is why incidence matters for the construction of group-level weights.

16.8.8. More generally, the purpose of this exercise is not to establish that one particular group statistic should always be preferred. Rather, it illustrates that the choice of group statistic embodies assumptions about project incidence. The median may be a reasonable approximation when impacts are concentrated around the centre of the group distribution, but it need not reproduce the welfare implications of individual-level weighting when benefits are broadly distributed or concentrated among poorer members of the group. The appropriate group-level welfare weight therefore depends on both the within-group distribution and the incidence pattern of the project under consideration.

16.9. Diagnosis and recommendations for the Green Book

16.9.1. The preceding analysis leads to a simple recommendation: retain the Green Book's Ramsey denominator based on mean real per capita consumption, but make the distributional numerator compatible with that denominator. Distributional, regional and place-based concerns should be handled through welfare weights in the numerator, not through ad hoc changes to the discount rate.

Welfare accounting summary

16.9.2. The Green Book discounting guidance uses the social time preference expression

$$STPR = \rho + \mu g,$$

where g is expected growth in future real per capita consumption. The current short-term calculation uses $\mu_R = 1.0$, $g = 2.0\%$, and $\rho = 1.5\%$, giving an STPR of 3.5%. In the terminology of this chapter, this is a mean-consumption denominator: it prices a unit change in mean real per capita consumption.

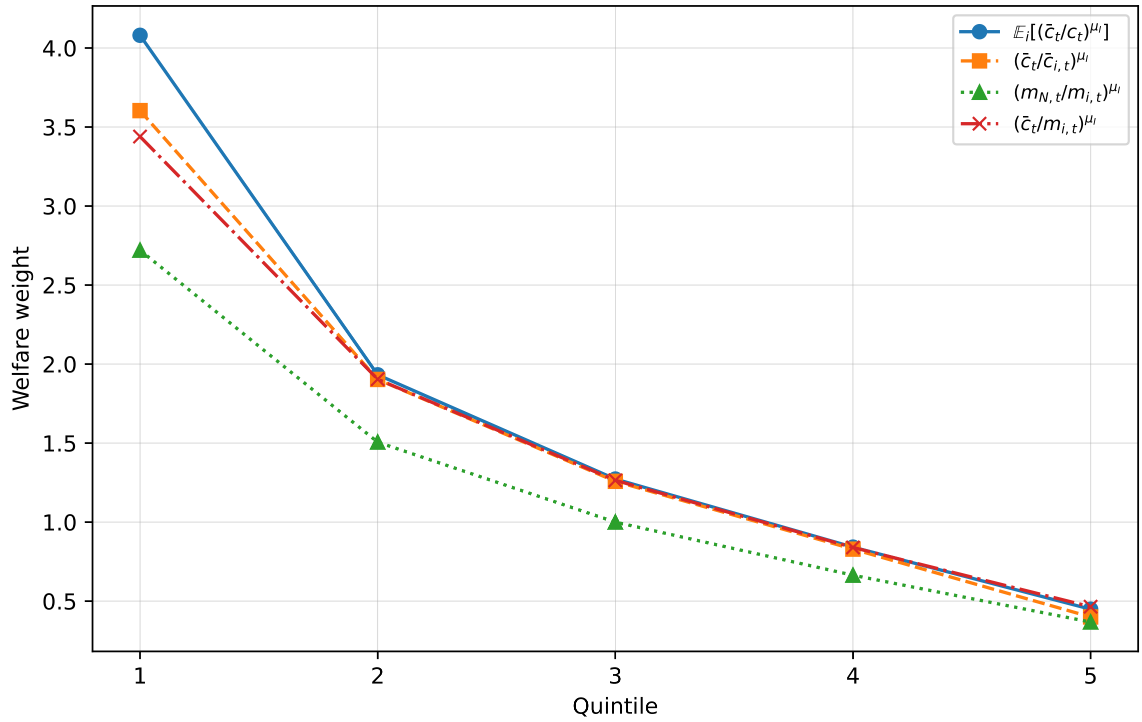

Figure 16.3.: Group-weight approximations under an illustrative UK-calibrated lognormal income distribution with $\mu_I = 1.3$. The exact benchmark is the average of individual mean-reference weights within each quintile. The group-level approximations use the national mean over the quintile mean, the national mean over the quintile median, and the national median over the quintile median.

Table 16.2.: Group-weight approximations under a UK-calibrated lognormal distribution

Quintile	$_i\left[\left(\frac{\bar{c}_t}{c_t}\right)^{\mu_I}\right]$	$\left(\frac{\bar{c}_t}{\bar{c}_{i,t}}\right)^{\mu_I}$	$\left(\frac{\bar{c}_t}{m_{i,t}}\right)^{\mu_I}$	$\left(\frac{m_{N,t}}{m_{i,t}}\right)^{\mu_I}$
1	4.080	3.604	3.440	2.721
2	1.932	1.903	1.904	1.506
3	1.273	1.258	1.264	1.000
4	0.842	0.829	0.839	0.664
5	0.449	0.400	0.465	0.368

16.9.3. The Green Book distributional weights are defined differently. They use the ratio of national median equivalised income to the median equivalised income of the affected group, raised to 1.3. Let $m_{N,t}$ denote national median equivalised income and let $m_{i,t}$ denote the median equivalised income of affected group i. The Green Book-style group weight is

$$\omega^{GB}_{i|N,t} = \left(\frac{m_{N,t}}{m_{i,t}}\right)^{\mu_I}, \qquad \mu_I = 1.3.$$

The current approach therefore combines a mean-consumption discount denominator with national-median-reference welfare weights. It also uses different curvature parameters in discounting and distributional weighting.

16.9.4. There are two possible accounting routes. One can start from the current Green Book welfare weights and then make the corrections needed to use them with the mean-consumption denominator. Alternatively, one can start from the Ramsey denominator itself and construct welfare weights directly in mean-consumption-reference units. We recommend the second route. It is simpler, more transparent, and better aligned with the welfare function that motivates the social discount rate.

16.9.5. Table 16.3 summarises the distinction. The first row starts from the Ramsey denominator and constructs compatible mean-reference weights. The second row starts from the existing Green Book median-to-median weights and shows the corrections required to make them compatible with a mean-reference denominator. The policy recommendation is to use the first row.

Table 16.3.: Two routes to mean-reference welfare accounting

Starting point	Reference path correction	Welfare weight correction	Overall correction, $\mu_I \neq \mu_T$	Overall correction, $\mu_I = \mu_T = \mu$	Dynamic numerator price adjustment
$S_t = \bar{c}_t, \quad R_t = \bar{c}_t$ $D_{\bar{c}}^{\mu_T}(t) = e^{-\rho t}\left(\frac{\bar{c}_t}{\bar{c}_0}\right)^{-\mu_T}$ $\omega_{i\mid\bar{c},0}^{\mu_I} = \left(\frac{\bar{c}_0}{\bar{c}_{i,0}}\right)^{\mu_I}$	$\chi_{\bar{c}\mid\bar{c}}^{\mu_T}(t) = 1$	$\frac{\omega_{i\mid\bar{c},t}^{\mu_I}}{\omega_{i\mid\bar{c},0}^{\mu_I}} = \left[\frac{\bar{c}_t/\bar{c}_0}{\bar{c}_{i,t}/\bar{c}_{i,0}}\right]^{\mu_I}$	$\Pi_{i\mid\bar{c},\bar{c}}^{\mu_I,\mu_T}(t) = \left[\frac{\bar{c}_t/\bar{c}_0}{\bar{c}_{i,t}/\bar{c}_{i,0}}\right]^{\mu_I}$	$\Pi_{i\mid\bar{c},\bar{c}}^{\mu,\mu}(t) = \left[\frac{\bar{c}_t/\bar{c}_0}{\bar{c}_{i,t}/\bar{c}_{i,0}}\right]^{\mu}$	$\frac{\mathrm{d}}{\mathrm{d}t}\log \Pi_{i\mid\bar{c},\bar{c}}^{\mu_I,\mu_T}(t)$ $= \mu_I\left\{g_{\bar{c}}(t) - g_{\bar{c}_i}(t)\right\}$
$S_t = \bar{c}_t, \quad R_t = m_{N,t}$ $D_{\bar{c}}^{\mu_T}(t) = e^{-\rho t}\left(\frac{\bar{c}_t}{\bar{c}_0}\right)^{-\mu_T}$ $\omega_{i\mid N,0}^{\mu_I} = \left(\frac{m_{N,0}}{m_{i,0}}\right)^{\mu_I}$	$\chi_{N\mid\bar{c}}^{\mu_T}(t) = \left[\frac{m_{N,t}/m_{N,0}}{\bar{c}_t/\bar{c}_0}\right]^{-\mu_T}$	$\frac{\omega_{i\mid N,t}^{\mu_I}}{\omega_{i\mid N,0}^{\mu_I}} = \left[\frac{m_{N,t}/m_{N,0}}{m_{i,t}/m_{i,0}}\right]^{\mu_I}$	$\Pi_{i\mid N,\bar{c}}^{\mu_I,\mu_T}(t) = \chi_{N\mid\bar{c}}^{\mu_T}(t)\frac{\omega_{i\mid N,t}^{\mu_I}}{\omega_{i\mid N,0}^{\mu_I}}$ $= \left[\frac{\bar{c}_t/\bar{c}_0}{m_{i,t}/m_{i,0}}\right]^{\mu_I}\left[\frac{m_{N,t}/m_{N,0}}{\bar{c}_t/\bar{c}_0}\right]^{\mu_I-\mu_T}$	$\Pi_{i\mid N,\bar{c}}^{\mu,\mu}(t) = \left[\frac{\bar{c}_t/\bar{c}_0}{m_{i,t}/m_{i,0}}\right]^{\mu}$ $= \frac{\omega_{i\mid\bar{c},t}^{\mu}}{\omega_{i\mid\bar{c},0}^{\mu}}$	$\frac{\mathrm{d}}{\mathrm{d}t}\log \Pi_{i\mid N,\bar{c}}^{\mu_I,\mu_T}(t)$ $= \mu_T\left\{g_{\bar{c}}(t) - g_N(t)\right\} + \mu_I\left\{g_N(t) - g_{m_i}(t)\right\}$ $\mu_I = \mu_T = \mu: \quad \mu\left\{g_{\bar{c}}(t) - g_{m_i}(t)\right\}$

Notes: The table keeps the mean-consumption Ramsey denominator, $D_{\bar{c}}^{\mu_T}(t)$, fixed. In both rows, $S_t = \bar{c}_t$ is the reference path priced by the discount factor. The table separates the base-period equity level from the time-varying correction. The full numerator multiplier is the base-period weight multiplied by the relevant $\Pi(t)$. Thus, in the first row, the full mean-reference group-mean weight is

$$\omega_{i\mid\bar{c},0}^{\mu_I}\Pi_{i\mid\bar{c},\bar{c}}^{\mu_I,\mu_T}(t) = \left(\frac{\bar{c}_t}{\bar{c}_{i,t}}\right)^{\mu_I}.$$

This row is the recommended route if the Green Book retains the Ramsey discount factor based on mean real per-capita consumption. The affected group is represented by its mean consumption or income, $\bar{c}_{i,t}$, so the dynamic numerator price adjustment depends on the growth gap $g_{\bar{c}}(t) - g_{\bar{c}_i}(t)$.
The second row starts from the Green Book-style national-median to group-median weight. Here $R_t = m_{N,t}$, where $m_{N,t}$ is national median equivalised income, and the affected group is represented by its median equivalised income $m_{i,t}$. The reference-path correction $\chi_{N\mid\bar{c}}^{\mu_T}(t)$ converts national-median-reference units into the units priced by the mean-consumption denominator. The welfare-weight correction captures the affected group moving relative to the national median. The product of these two terms is

$$\Pi_{i\mid N,\bar{c}}^{\mu_I,\mu_T}(t) = \chi_{N\mid\bar{c}}^{\mu_T}(t)\frac{\omega_{i\mid N,t}^{\mu_I}}{\omega_{i\mid N,0}^{\mu_I}}.$$

When $\mu_I = \mu_T = \mu$, the national median reference path cancels and the dynamic correction collapses to a mean-reference group-median correction, based on $g_{\bar{c}}(t) - g_{m_i}(t)$. This remains a group-median route, not a group-mean route. When $\mu_I \neq \mu_T$, the term involving $m_{N,t}/\bar{c}_t$ remains, so the national-median path continues to affect the welfare-accounting convention.
The final column gives the growth rate of the time-varying numerator multiplier, $\mathrm{d}\log\Pi(t)/\mathrm{d}t$. This is the relative-price-type adjustment kept in the numerator. If the same adjustment were instead folded into the denominator, the equivalent discount-rate adjustment would be its negative, $-\,\mathrm{d}\log\Pi(t)/\mathrm{d}t$. The paper recommends keeping these dynamic distributional adjustments in the numerator rather than replacing the common Ramsey discount rate with group-specific effective discount rates.

$$\omega^{\text{eff}}_{i,t}(P) = \frac{\int_{\Theta_i} \left(\frac{\bar{c}_t}{c_t(\theta)}\right)^{\mu} b_t(\theta;P)\ \mathrm{d}F_{i,t}(\theta)}{\int_{\Theta_i} b_t(\theta;P)\ \mathrm{d}F_{i,t}(\theta)}.$$

16.9.8. Thus the group statistic used in the weight should be chosen to approximate the incidence of the project. In the absence of better incidence information, using the affected group's mean income or consumption is the most natural group-level companion to a mean-consumption Ramsey denominator.

Specific recommendations

16.9.9. The Green Book should retain the current Ramsey structure for the social discount rate, but revise the treatment of distributional weights so that the weighted numerator and discount denominator are expressed in the same welfare units. We recommend the following.

1. **Present the unweighted NPV.** Appraisals should continue to present the unweighted NPV. This is the distributionally neutral valuation and provides continuity with current practice.

2. **Always present a distributionally weighted NPV.** For project appraisals and regulatory changes with material distributional or place-based consequences, the Green Book should also require a distributionally weighted NPV. This should not be treated as an optional sensitivity in cases where incidence is materially unequal. It should be a standard companion to the unweighted NPV.

3. **Harmonise the elasticity parameter.** The elasticity used in the Ramsey denominator and the elasticity used in the equity weights should be harmonised. Our recommendation is to use a single central value,

   $$\mu = 1.25,$$

   This avoids the ambiguity created by using one parameter for intertemporal welfare weights and another for interpersonal welfare weights. If different parameters are retained, the Green Book should state explicitly the two-parameter welfare-accounting convention being used.

4. **Use mean-reference equity weights.** The Green Book should replace national-median-reference weights with mean-reference weights. At the individual level, the appropriate weight is

   $$\omega_{\bar{c},t}(\theta) = \left(\frac{\bar{c}_t}{c_t(\theta)}\right)^{\mu}.$$

   At group-level, where only aggregate incidence is observed, the natural mean-reference approximation is

   $$\omega_{i|\bar{c},t} = \left(\frac{\bar{c}_t}{\bar{c}_{i,t}}\right)^{\mu},$$

   where $\bar{c}_{i,t}$ is the affected group's mean consumption or income. If the existing Green Book median-to-median weight is retained for presentational reasons, it should be recognised as a national-median-reference weight and converted before use with the mean-consumption denominator. The cleaner implementation is to construct the weight directly in mean-reference units.

5. **Consider time-varying equity weights for regional divergence.** Where a project has place-based effects and the affected region is expected to converge towards or diverge from the national mean, the equity weights should be allowed to vary over time. In the recommended mean-reference group-level implementation, this means using

   $$\omega_{i|\bar{c},t} = \left(\frac{\bar{c}_t}{\bar{c}_{i,t}}\right)^{\mu}$$

   rather than a fixed base-period weight. If the affected region grows more slowly than national mean consumption, the weight rises over time; if it grows more quickly, the weight falls. This is the numerator-side analogue of a relative-price correction. Developing a practical and proportionate method for forecasting these time-varying place-based weights would require further work, including guidance on regional income forecasts, convergence assumptions and the treatment of uncertainty.

16.9.10. The central recommendation is therefore not to alter the Ramsey discount rate project by project. The common discount denominator should remain the mean-consumption Ramsey denominator. Distributional, regional and place-based considerations should be handled in the numerator through welfare weights that are expressed in the same mean-consumption-reference units. This approach preserves the current intertemporal welfare framework, makes the equity-weight calculation compatible with it, and provides a clear basis for reporting both unweighted and distributionally weighted NPVs.

16.10. Conclusion

16.10.1. We recommend, in line with the principles of the review, that place-based, including regional disparities that can be captured by income differences, are accounted for via corrections to the welfare weights rather then group-wise, e.g. regional, discount rates. Differences in levels and trajectories can be effectively handled this way, largely in accordance with the current Green Book recommendations.

16.10.2. Yet, the welfare-accounting analysis above makes clear that to evaluate the NPV of a project the discount rate and the benefits and costs need to be measured in the same units of account particularly if the reference points evolve differently over time. The Green Book currently does not do this when it comes to using welfare weights since the reference point for welfare weights is national median income and the reference point for intertemporal welfare weights, via the discount rate, is mean consumption. Without correction, welfare evaluation could be erroneous. Mean and median incomes do not evolve together typically, so a correction term is strictly required to maintain accounting integrity.

16.10.3. There are two separate issues to handle: 1) secular changes in national reference points, such as the difference between national mean and median growth; and 2) changes in the affected group-level outcome relative to the national baseline. The importance of each depends on the welfare metric and reference path deployed in the CBA.

16.10.4. The simplest solution to this accounting problem is to proceed as follows: 1) ensure that inequality-aversion parameters are identical for intratemporal and intertemporal inequality aversion; and 2) ensure that the reference path for welfare weights is the same as for the intertemporal welfare weights, in this case mean real per capita consumption.

16.10.5. This proposal starts with the current intertemporal welfare function that defines the discount rate as the Ramsey Rule with a mean consumption reference path. Natural welfare weights for this case compare group mean to national mean consumption, not medians. Accounting corrections are made much simpler if the aversion to intra- and intertemporal inequality are made equal, which we also recommend.

16.10.6. A further recommendation to account for distributional effects is to consider moving away from static welfare weights and account for diverging or converging group performance over time. This would do a more accurate job of accounting for place-based distributional effects and would avoid the need to have confusing place-based discount rates.

16.10.7. In conclusion, this is a critical yet complicated area of welfare economics. Balancing issues of efficiency and equity in the context of overall place-based and regional disparities is not straightforward and depends on how we measure inequality and the factors that are driving inequality in the first place. The strategic objectives with regard to inequality and how public appraisal interfaces with that given current policy and institutional arrangements that already provide investment and regulatory frameworks that limit inequality, is an expansive topic that deserves further analysis.

17. High complexity

17.1. Introduction

17.1.1. In this chapter we relatively briefly discuss three of the most complex cases. Two of these are explicitly raised in the Terms of Reference — private finance models and transformational projects — and one is implicit in the question concerning declining discount rates — nuclear decommissioning and other ultra-long-term problems.

17.1.2. From Chapters 12 —16, it is clear that, even for moderately complex public projects, there are many welfare issues that need to be considered in a well-constructed economic appraisal. In the case of rare disasters and DDRs, we have accounted for these through adjustments to the discount rate. For place-based objectives and environmental sustainability, while we have recognised that the discount rate could, in principle, be adjusted, we have recommended that more emphasis be given to welfare weights and/or relative price adjustments in the numerator of the present value equation instead. We have provided some relatively detailed thoughts on how this might be done.

17.1.3. But some projects are so complex and significant to wider society that sophisticated adjustments should be made to the discount rate because of the risks that are placed on the public sector. Alternatively, it may just be infeasible to apply a standard discount rate in a way that can give the government confidence that it fully captures the welfare implications of the proposals. In the latter case, the precise economic value placed on the project will depend heavily on both the exact nature of the contract and more subtle considerations of whose welfare the government wishes to advance than δ and μ can account for alone.

17.1.4. This fits the principle of proportionality which has underpinned the whole of these Technical Annexes:

- The Ramsey Rule with a fixed risk premium, triangulated against SOC rates, survey responses, and other international guidance, is sufficient for many short-to-medium term public projects.
- Augmenting this with a DDR schedule is sufficient for many long-term projects up to 125 years in maturity.
- Reducing the risk premium for projects with social insurance characteristics will better reflect their true economic value.
- More consistently applying welfare weights and relative prices will help better deliver place-based projects and protect the environment.

17.1.5. In all cases, we are adding complexity beyond the standard Ramsey Rule in a proportionate way. But, at some point, the complexity gets too great for this to continue. Given the time and effort that public bodies put into the estimation of costs and benefits for large-scale projects, to then use a four-parameter model with DDRs to bring all that work back to the expectation of the change in social welfare self evidently fails the test of proportionality, in our opinion.

R.22. For highly complex projects, the answer is unlikely to lie in a discount rate.

For private finance models, the discount rate should reflect, in a proportional way, the risk that such contracts impose on the public sector. For transformational projects and ultra-long term projects such as nuclear decommissioning, HM Treasury should undertake more detailed welfare analysis and not rely on a discount rate. (HM Treasury 2021).

17.1.6. This is not a new suggestion. For example, such an opinion has been expressed by Professor Deborah Lucas of MIT in relation to very long-term environmental problems, as we quoted in paragraph 5.4.2.

17.1.7. In the remainder of this chapter we explain in more detail why we have reached this conclusion. We also make some more general observations on the way in which discounting is currently undertaken.

17.2. Private finance models

17.2.1. In this section, we discuss the economic appraisal of public projects that use a private finance model. This relates to projects covered under Annex A of HM Treasury (2026*c*). These come in wide variety of forms, the two main of which are Unitary Charge Models and User Pay models. We concentrate on the former throughout this section.

17.2.2. Specifically, we will consider Private Finance Initiatives (PFI) and its successor Private Finance 2 (PF2) models, for which there is greatest external transparency. We note, though, that there are a number of other types of Unitary Charge model that are not included in PFI and PF2, including NHS projects under the Local Improvement Finance Trust (LIFT) programme, the non-profit distributing (NPD) and hub models used in Scotland, and the Mutual Investment Model used in Wales.

17.2.a. The nature of the contract

17.2.3. Details on the operation of PFI and PF2 are given in National Audit Office (2018), where it is noted that the fundamentals of the financing structure and contract are the same across the two. In brief, as summarised by HM Treasury, these models are structured as follows:

> "PFI and PF2 are forms of Public Private Partnerships (PPPs). Public Private Partnerships (PPP) are long-term contractual arrangements between a public sector entity and a private sector provider.
>
> The private sector provider is engaged to design, build, finance, maintain and operate infrastructure assets and related services. The risks associated with construction delay, cost overrun and maintenance of the asset are transferred to the private sector partner.
>
> The public sector entity does not pay for the asset during construction, as the associated costs of construction are financed by the private sector. Once the asset is operational and services are being provided the public sector entity pays a monthly fee – sometimes referred to as a 'unitary charge' – to the private sector provider. This payment includes the costs of construction, financing costs, lifecycle replacement expenditure, maintenance and services.
>
> The payment is subject to performance, which means that payments are reduced if services are not delivered to the standards set out in the contract. This form of payment mechanism provides an incentive for the private sector provider to meet their performance obligations and underpins the transfer of risk to the private sector."

Soft services, such as cleaning and catering, are not generally involved in such contracts.

17.2.4. The government sees PFI and PF2 as a risk-reduction measure by giving reduced public exposure to the costs and delays from construction over-runs. There are also potential benefits to be gained from improved operational efficiencies and better maintained assets. However, there is evidence that the risks the private partner faces from construction delays may be incorporated into the payments they negotiate through the contract, as well as other cost premia for risk transfers. This, then, concerns risk reduction for the public sector rather than necessarily lower costs.

17.2.5. The National Audit Office (2018) also questioned whether there have, historically, been operational efficiencies from PFIs, with many departments finding these costs the same or higher under a PFI model. Maintenance spending under PFIs have, though, generally been higher than within the public sector.

17.2.6. PFIs have also helped departments, in the short term, keep within their spending budgets while increasing long term costs and compromising future flexibility (although PF2 looks to address the latter point). The obligations that arise from such arrangements are generally off balance sheet.

17.2.7. The National Audit Office (2018, #1.17) notes that:

> "Private finance procurement results in additional costs compared to publicly financed procurement, the most visible being the higher cost of finance. ... Data collected by IPA on PFI and PF2 deals entered into since 2013 show that debt and equity investors are forecast to receive a return of between 2% and 4% above government borrowing".

When combined with other costs from setting up a Special Purpose Vehicle, estimates suggest that the cash costs from PFI have historically been of the order of 40%-70% higher than the public funding; see, for example, National Audit Office (2018, Figure 4). Note, though, that under PF2, the government will often be a minority equity holder, allowing it to partially gain in any project upside.

17.2.8. Given this, the National Audit Office (2018, #1.24) states that

> "HM Treasury has noted that the higher cost of private financing means that the economic case for the model rests on achieving cost savings in the construction or operation of the project; or through the delivery of a qualitatively superior project. For PFI to

> offer value for money (VfM), these benefits must exceed the higher financing and other additional costs."

17.2.9. The National Audit Office (2018) report was written at a time of low Treasury yields and relatively high inflation, so the nominal STP rate of 3.5% + inflation = 6.09% was well above the rates that PFI providers could offer. This, under Green Book guidance, made PFI look attractive even though gilt yields were also low at the time, allowing for cheap public finance instead. This has ultimately led to the recommendation in HM Treasury (2026*c*, #A.10) that sensitivity analysis for the discount rate is run using the prevailing gilt rate.

17.2.10. The National Audit Office (2018, #2.6) makes clear what it is that the unitary charge pays for:

> "Data used in the Whole of Government Accounts (WGA) records that around half of current annual PFI charges relate to debt repayment and financing costs (interest and dividends). The balance is service charges – the costs of operating and maintaining the asset. The exact split of debt repayment, financing and service charges will vary over time, as debt is repaid, and from project to project. The service element of PFI payments increases each year in line with a retail price index (RPI) inflation measure. In the case of some health deals, the whole payment, not just the service charge, rises with inflation."

Interest rate swaps are commonly used to fix the annual debt repayment for the department over the life of the contract. However, as most PFI deals do not include break clauses, this can make such contracts look expensive if interest rates fall.

17.2.11. Although equity commonly makes up only around 10% of the financing of a standard deal, with the remaining 90% being senior debt, the National Audit Office (2018, #3.12) notes:

> "In some PFI deals, equity investors have been able to generate high investment returns, particularly when equity was sold after construction. For example our analysis of a recent equity sale in the M25 PFI contract showed that, over an eight-year period, equity holders have realised returns of around thirty-one per cent a year (Figure 15). This high return on the sale is likely to be because the new investor is willing to have a lower return as the project is in a lower risk operational phase, but may also mean that the project is more profitable than originally forecast. ... The Committee of Public Accounts has criticised the level of investor returns achieved on some projects and the NAO has previously concluded that inefficient pricing of equity has contributed to high returns."

17.2.12. The National Audit Office (2018) estimates that the average post-tax return to investors, combining debt and equity was around 4.5% - 5.0%; approximately double the gilt rate at the time.

17.2.b. Private finance discount rates

17.2.13. How private finance models should be discounted has been considered in detail in a utility regulation context by JRG (2011), with a more recent discussion in the context of financial regulation given by Oxera (2026). The former identifies four possible models.

- "Discount all costs (including financing costs as calculated based on a WACC) and benefits at the STPR. [The "Spackman approach"]
- Discount some costs and/or benefits at a WACC, and some at the STPR, depending on their likely systematic risk
- Discount all costs and benefits at a WACC
- Discount all costs and benefits at the STPR (excluding financing costs)" (JRG 2011, #1.11)

17.2.14. In general, JRG (2011) recommends the Spackman approach to regulators except when making adjustments to this is "proportionate", with these factors being based on systematic risk. Our understanding is that this is also the approach currently taken by the Green Book.

17.2.15. This emphasises that it is important to distinguish between two different rates of return. The first is the cost of capital (WACC) that is used to calculate what is essentially an annuity payment that the private partner receives from the public sector as a unitary charge. This compensates them for the initial construction costs plus a return on capital, together with the ongoing service costs. The second is the rate that is used to discount these payments in social economic appraisal, which is the STPR under the Spackman approach.

17.2.16. The complexity arises from the different treatment of risk between cost-benefit analysis in the public and private sectors and, to a lesser extent, the magnitude of the SOC-STPR divide for projects of equivalent maturity and risk. If private finance models were all risk-free, and if the Ramsey Rule rate equalled the gilt yield in the absence of taxes and other frictions, then the Green Book rate could be applied throughout. This, though, is not the case, with the public sector transferring the construction risk to the private sector and compensating them for taking this risk.

17.2.17. When it comes to the construction risks, which are borne by the private sector, these occur early in the life of the contract. To the private sector, the unitary charges that they will ultimately receive are close to being risk-free. But it is these unitary charges that are discounted at a risk-adjusted WACC, which does not reflect the risks of those payments. Many of the qualitative responses to our survey expressed opinions that reflected this point; e.g.:

- "PPP is a mechanism of providing almost riskless returns at higher than riskless rates. ... At the moment providers of capital get a very good deal from the UK government",
- "re the private entities actually taking on real risk in the public-private partnership, or are returns more guaranteed such that discounting based on generic private rates of return (which reflect higher degrees of risk as well as externalities etc.) would not be appropriate",
- "...the use of different discount rates between the public and private sector can create a "wedge" that overcompensates the private sector for the risk of a project, which is a cost ultimately borne by the taxpayer." (all paragraph A.2.9).
- "When we mix risk premia from the private sector, we end up with mismatches between the inherent risk of a project (from the perspective of government), which is very low and the valuation placed on that project. For PPP this becomes more prescient as that mismatch allows private sector investment to turn a paper difference into real returns. Which other countries do this?" (paragraph A.2.13).

Therefore, as another respondent to our survey noted: "Careful separation of risk allocation, transfer pricing, and financing structure is therefore essential, alongside transparency about which risks are borne by the public sector and which are genuinely transferred to private partners" (paragraph A.2.9).

17.2.18. Consider the following illustrative example. A build is expected to cost £100m, but with risk to the private sector. To compensate for that risk, it charges the public sector a WACC of 5% real over the 25-year life of the contract. This is equivalent to real annual unitary charge payments of £7.095m a year. This is discounted by the government under the Spackman approach at 3.5% a year currently using the Green Book STPR. This gives a total present value of unitary charge payments of £116.94m.

17.2.19. This is equivalent to the private partner quoting for a certainty-equivalent capital build cost of £116.94m to bake construction contingency risk into the initial price, but then accepting the STPR to calculate the unitary charge. In both cases, the annual payments would be the same but the economic interpretation of the way in which risk is being compensated for is different. In the former case, a risk-adjusted WACC is being paid over 25 years for cash flows that are almost risk-free. In the latter case, compensation for risk is being capitalised at the start, which is where the risk arises, and then the payments are discounted at a rate that reflects their very low levels of risk.

17.2.20. Once the construction risk has past, then, the unitary charges are over-compensating for the remaining risk. Therefore they can be sold at a premium. This is consistent with the finding reported in paragraph 17.2.11 that equity holders have historically received very large returns if selling their stake after the construction phase is completed, in part because "... the new investor is willing to have a lower return as the project is in a lower risk operational phase".

17.2.21. For the service contract, the risk to the private provider is ongoing. Therefore, using a risk-adjusted discount rate to calculate the annuity is, in this case, conceptually consistent with the time profile of the risk. However this rate should reflect the risk of the service contract itself and not a more generic WACC.

17.2.22. From the public sector perspective, we can draw a clear analogy with private firms. In the corporate world, a project is initially valued without consideration of finance. An asset beta (β_A) is used to estimate a cost of capital and a present value is calculated on that basis. Only once the NPV test has been passed will the firm concern themselves with questions of how it should be financed. This is mirrored in the public sector by the general Green Book approach.

17.2.23. Once finance is introduced, the debt and equity holders value their own individual claims to the project's future benefits using either a debt beta (β_D) or equity beta (β_E). As described earlier, $\beta_D < \beta_A < \beta_E$because of the seniority of debt over equity, so equity has the highest cost of capital.

17.2.24. In a private finance deal, the public sector's claims to the project resemble an equity claim. This claim is residual after the fixed unitary charges have been paid. Therefore, the risk to the public sector during the life of the project (after construction) is higher with PPP than without. In essence, as we understand these deals, the private finance component converts a reduction in short-term construction risk into a long-term residual benefits risk for the public sector. As one of the respondents to our survey noted, "It is essential to differentiate the risk born by the government from the actual risk-free benefits and costs being discounted" (paragraph A.2.9).

17.2.25. In theory, therefore, the Green Book should recommend that departments undertake two sets of present value calculations. One with the total social benefit (ignoring the private partner) discounted at the STPR, as is the general case under the Green Book. The other should be the residual claim to the public sector once the unitary charges have been deducted at a rate that reflects β_E rather than β_A. By using β_A = the STP rate in both cases under current practice, this must (all else being equal) overstate the public value of private finance for the average public project. In neither case should the unitary charges be discounted separately from considerations of the net benefits since these charges change the risk profile to the public sector.

17.2.26. The question, then, is whether this overstatement is materially significant to a social economic appraisal. We have previously recommended that the STPR should not be adjusted for specific project betas apart from in exceptional circumstances even though we know theoretically that this will overstate the value of highly pro-cyclical projects. We made this recommendation on the basis of practicality, proportionality, and the absence of a strong evidence base to make such an adjustment.

17.2.27. A similar argument applies here. If private finance does not significantly increase the risk to the public sector, and if the project is not very large in scale, then it may be proportionate to retain the use of the Spackman approach. When there are significant risk transfers, then the government should evaluate separately the risks to the private sector (to set the WACC that underlies the unitary charge calculation) and the residual risk to the public sector (to adjust the STP rate). These risks will, in general, be different.

17.2.28. But this will also depend on the precise nature of the contract and how this allocates risks and profits between the two parties. It is increasingly common in this type of arrangement to make the unitary charge more equity-like and the social residual benefit more debt-like through risk-sharing clauses. This will all affect the choices of WACC and β_E in a way that is difficult to generalise in a report such as this.

17.2.29. We also note that the Green Book recommends that sensitivity analysis be undertaken on the basis of the gilt rate as well as the STPR (HM Treasury 2026*c*, #A.10). As we noted in pargaraph 17.2.9, this recommendation came about when the gilt yield and private finance WACCs were substantially below the 3.5% STP rate. This is no longer the case, with the SOC-STPR divide being much narrower than it once was. If this divide widens again, then we have argued that this gives information to HM Treasury about the state of the economy that requires an early review of the STP rate and the impact of prevailing macroeconomic conditions on project appraisal. The 2% risk-free component of the STP rate that we are recommending in this review is extremely close to prevailing gilt rates.

17.2.30. Yet the overall STP rate of 3% that we are recommending is a risk-adjusted rate, with an average risk premium of 1% incorporated into it. This is conceptually distinct from a gilt rate that is effectively risk-free. We have argued throughout these Technical Annexes that SOCs and STPRs should only be compared on the basis of the same maturity and the same risk, which appears not to be the case here.

17.2.31. Our opinion is that, under current market conditions, if the social discount rate and WACC are set appropriately for the risk that each party bears, then the distinction that sensitivity analysis makes between the STPR and gilt rate will largely disappear. Whether such additional calculations remain useful for economic appraisal purposes remains a practical matter for HM Treasury's attention. In general, as we explained in Section 5.7, we think that if the discount rate is doing only two jobs — dealing with risk and maturity — then a decision should be taken on one rate without undertaking sensitivity analysis since the purpose of a present value calculation is to construct a single, expected change in social welfare that best reflects HM Treasury's view of what constitutes welfare. If HM Treasury views this more as a financing question than an economic appraisal one, then this matter falls outside the Terms of Reference for this review.

17.3. Ultra-long-term projects

17.3.1. The theory and empirics of social discount rates has developed tremendously over the last 20 years, particularly as a consequence of Stern (2007) on the Economics of Climate Change. Because carbon emissions stay in the atmosphere for very long periods of time, methods have been developed that allow for discount rates to be applied over periods of many centuries. This, in part, motivated the literature on declining discount rates.

17.3.2. As we noted in paragraph 12.3.11, this is not of direct relevance for HM Treasury because the Green Book calculates its Social Cost of Carbon based on a target consistent approach rather than discounted damages. However, there are a number of ultra-long-term projects in the public sector in the UK, most notably nuclear decomissioning. We also refer the reader to the work that the French government is doing on this (Cherbonnier et al. 2025) that we touched on in Chapter 12. As one of the respondents to our survey noted "Ensure that they are aware of the large range of areas where it is used and it can have perverse effects. For instance, the bulk of Ministry of Defence Annual Managed Expenditure is provisions for disposing of nuclear submarines. This went from a large deficit to a large surplus when the Treasury raised the discount rate" (paragraph A.2.13).

17.3.3. Once it comes to ultra-long-term costs and benefits, we think there is a difficulty with just continuing to use a discount rate exponentially. This has been more formally described by Lucas (2023, Abstract), which we cited in paragraph 5.4.2 but repeat here because of its centrality to the argument:

> "For policies with long-term impacts, intergenerational concerns become paramount, projections of cashflows and discount rates become highly uncertain, and present value calculations are an intrinsically unreliable measure of value. No approach to discount rate selection can overcome those problems; alternative decision criteria need to be established."

17.3.4. There are several reasons for this. We start with two statistical ones. First, we have emphasised in Chapter 6 that any estimate of the social discount rate has an error associated with it compared to the "true" discount rate. In Table 17.1 we consider the implications of this. Our estimate of the discount rate is 3%, but suppose instead that the "true" rate is 2.9% or 2.8% or 2.7% instead, within associated estimation errors of 0.1%, 0.2% and 0.3% respectively. The table shows the valuation error that arises in each case (as a percentage of the true value) at different maturities. At twenty years, the valuation error is less than 6% in all these cases, while at 200 years it is more than 17% in all these cases, rising to 44% when the estimation error of the discount rate is 0.3%. This is a simple effect of compounding. At some point, the errors introduced by continuing to raise one plus an estimated discount rate to the power of t becomes so large as to potentially undermine the economic appraisal.

17.3.5. Second, as we showed in Chapter 13, small changes to the way in which the economy is modelled leads to very large differences in the schedule of declining discount rates (DDRs). We made this point originally in Freeman & Groom (2016*b*). That discount rates decline over time is well established theoretically but estimating how fast that decline should be is exceptionally difficult. Therefore, there is a point where the estimated DDR schedule becomes driven by small changes in assumption rather than the "true" social discount rate.

17.3.6. Third, we have argued in detail in Chapter 8 that the case for having a utility discount rate of $\delta = 0\%$ is driven by considerations of intergenerational equity, where some resolution process may be put in place in parallel. Determining this variable just becomes much more complex in an ultra-long-term context.

17.3.7. There is a very loose parallel here with the place-based welfare analysis that we laid out in Chapter 16. We can view different generations as being temporally, rather than geographically, separated. This allows for more fundamental analysis to be undertaken on the transfers that are taking place.

17.3.8. We are recommending discount rates to 125 years and not beyond that point for these reasons. This is consistent with current Green Book practice (HM Treasury 2026*a*, Table 3.A) for $\delta = 0.5\%$. However, we believe this is the absolute maximum length that the standard discount rate to be applied before reverting to more formal welfare analysis, and in some cases it may be proportionate to make this transition earlier. This contrasts with much international guidance summarised in Table 4.1, where discount rates are prescribed over longer periods

17.4. Transformational projects

17.4.1. We were asked by HM Treasury to write this review following from a specific recommendation in the Findings and Actions of the Green Book Review 2025. This stated that:

> "HM Treasury will commission an independent review of the Green Book discount rate to make sure that the government is taking a fair view of the long term benefits that arise from transformational investments."

Yet it is only now, at the very end of these extended

Maturity	10	20	50	100	200	400
2.9%	-0.97%	-1.92%	-4.74%	-9.26%	-17.66%	-32.20%
2.8%	-1.92%	-3.81%	-9.26%	-17.66%	-32.21%	-54.04%
2.7%	-2.87%	-5.67%	-13.57%	-25.30%	-44.20%	-68.86%

Table 17.1.: Present value estimation errors when the estimate STPR=3% differs from the 'true' STPR which equalis 2.9%, 2.8% or 2.7% in the different rows.

Technical Annexes, that we explicitly turn to this issue.

17.4.2. The term "transformational change" is defined in HM Treasury (2026*c*, #4.28-#4.30). For the purposes of an economic appraisal, it is essential that this term is interpreted in an economic, as opposed to a political, sense. While not precisely the definition in the Green Book, we interpret this as the values of b_{it} being very large, often requiring co-ordination across a wide number of public and private bodies.

17.4.3. In our opinion, and as we will explain below, the entire analysis that we have laid out up to this point is of fundamental relevance to the matter. Before we get to that line of argument, though, we begin by addressing a technical question: do such transformational projects, with very large values of b_{it}, have a discount rate at all?

17.4.a. Marginality

17.4.4. Return to the very beginning of our discussion on the Foundations of Discounting in Chapter 5. Here we stated that, for present values to have the properties that we require of them, they must adhere to three mathematical assumptions: they must be scalable (A.1.), they must be additive (A.2.), and they must exist in a well defined probability space (A.3.) so that terms such as 'expectation' , 'variance' and 'correlation' all have a well-defined meaning.

17.4.5. In Chapter 7, we demonstrated that if the costs and benefits are marginal (small) compared to the aggregate consumption level then, in the derivation of the Ramsey Rule, that was sufficient to ensure that assumption A.1. holds.

17.4.6. However, transformational projects may not be marginal to the overall economy and therefore we can no longer be confident, a priori, that the core requirements of present values are met. In short, there may not be a discount rate, or a present value, for such projects at all. As one of the respondents to our survey noted: "In general, unless the project is non-marginal, the existing framework should be more than sufficient. If the project is truly non-marginal then the issue is not so much the discount rate as stepping outside marginal analysis and doing proper modelling in general equilibrium" (paragraph A.2.12).

17.4.7. This issue can be checked empirically through simple examples. Imagine a project that gives one benefit of b_{it} = \$1 per capita in 10 years time. Current per capita consumption is £25,000 and this is expected to grow at 2% per year so that, with certainty, in 10 years time real per capital consumption will be £30,535. The utility function sets $\delta = 0.5\%$, consistent with the current value in the Green Book, and we consider four possible values for μ: 0.5, 1, 1.5, or 2.

17.4.8. Because this is clearly a marginal project, we can derive its current value, p_{it}, either through the PV formula with the discount rate calculated via the Ramsey Rule, or directly from first welfare principles as described in Chapter 5 by setting $E[\Delta W] = 0$. These give virtually identical answers.

17.4.9. Now assume that the project increases in size, first by a factor of ten, then 100, then a 1,000, then 5,000 and finally 10,000. In this final case, the benefit of the project will contribute almost a third more to the overall consumption level in a decade's time than would otherwise have been the case. This is clearly a non-marginal project. In terms of the notation of Assumption A.1., we are picking $k \in \{10, 100, 1000, 5000, 10000\}$.

17.4.10. In each case, we can calculate precisely the exact maximum amount that the government should be prepared to pay for the project by setting $E[\Delta W] = 0$. This gives the values for $P_{k\times i}$. In Table 17.2 we present values for $P_{k\times i}/P_i$ for the four different values of μ and the five different values of k.

μ	0.5	1	1.5	2
k	$P_{k\times i}/P_i$			
10	10.0	10.0	10.0	10.0
100	99.8	99.7	99.5	99.4
1000	983.5	969.1	956.4	945.0
5000	4611.4	4312.1	4071.7	3872.4
10000	8549.8	7583.1	6873.0	6319.3

Table 17.2.: Non-marginal projects and Assumption A.1.

17.4.11. If Assumption A.1. holds precisely, all elements of this table should equal k, which they clearly do not once the project gets large. Because Assumption A.1. is violated, there is no theoretical basis to create a discount rate at all for society-changing projects. We need to go back to fundamental welfare analysis rather than the marginal analysis of PV calculations.

17.4.12. If we were, for practical purposes, to ignore this theoretical argument and just calculate the effective discount rate for the $k = 10,000$, $\mu = 2$ case, this is 9.09% compared to the Ramsey Rule value of 4.50%. This means that transformational projects should, all else being equal, have *higher* discount rates than marginal ones. This follows from the fact that we have assumed diminishing marginal utility (isoelastic utility functions have negative second derivative or, equivalently, are concave). The utility we get from very large benefits in future is much less than the loss in utility from having to pay a great deal for the project today.

17.4.13. In practice, the question arises as to at what point this non-marginality actually matters from a decision-making perspective. For example, suppose we define "material" from a practical perspective as being equivalent to the PV calculated from the Ramsey Rule being more that 10% greater than the 'true' PV calculated directly from setting $E[\Delta W] = 0$. In this case for the example above, the future benefit from the project divided by future consumption without the project must not be above 19.4% (when $\mu = 0.5$) or 5.6% (when $\mu = 2$). Dietz & Hepburn (2013, p.64) draw a similar conclusion saying that, when $\mu = 1$: "...if a project delivers a once-off benefit of 10% of current consumption, then conventional DCF analysis will overestimate the actual increase in utility by approximately 5%, simply because the marginal evaluation ignores curvature in the utility function".

17.4.14. But these are very large values in absolute terms. If we assume there are 45 million adults each consuming £30,535 in ten year's time, then the benefit needs to be around £77 trillion a year even at $\mu = 2$ for the Ramsey Rule to give a valuation error of more than 10% based on arguments around non-marginality.[1]

17.4.15. One of the expert members of our panel, with considerable experience of the economic appraisal of transformational projects, agreed that, from their experience, these are generally perfectly well dealt with by usual linear methods where the scalability assumption A.1. can be assumed to apply.

17.4.16. From this we conclude that even transformational projects are unlikely to violate the marginality assumptions that underpin the Ramsey Rule sufficiently seriously to make a material difference to most governmental decision-making. Therefore, we can think about the discount rate and present value of transformational projects as usual.

17.4.17. We make one other technical point. HM Treasury (2026*a*, #4.29), in its discussion of transformational change, says "The social value of such a group of projects and programmes is often greater than the sum of its parts". If taken at face value, that would violate Assumption A.2. in Chapter 5 — the present values are claimed not to be additive — and, again, the conditions for a present value or discount rate are violated. But this is not how we interpret what the Green Book is saying. By taking a departmental angle alone, potential synergies between public bodies are missed and this understates the present value; there is a need to value synergies as well as departmental gains. This is consistent with Assumption A.2., but suggests that the net benefits are being underestimated in the numerator of the present value equation. This certainly needs correcting for in an economic appraisal, but this is not the work of the discount rate.

17.4.b. Transformational project discount rates

17.4.18. In Table 17.3, we report the proportion of respondents to our survey who stated the opinion that the discount rate should be different for transformational projects compared to standard Green Book analysis. Among all responder types except those who came to our survey through links on Social Media, there was a strong opinion against, and this is reflected in the qualitative comments in paragraph A.2.12; for example, "Project valuation discounts future benefits using an adequate discount rate. The benefits are at the numerator and the discount rates at the denominator. A "transformational project" will simply be one with a big numerator, so there is no need to adjust the denominator (the discount rate)... ".

17.4.19. Our opinion is also that this is the correct approach. As discussions concerning the stochastic discount factor in Chapter 5 show, the risk-free rate can vary according to the maturity of the project but not its type irrespective of whether an STPR or SOC approach is taken. If we accept that transformational projects have discount rates, then there is no case to vary this argument here.

17.4.20. In terms of the risk premium, then it may be practical, proportionate, and evidence based to vary this on a transformational project by transformational project basis. We hope that Chapters 10 — 12 and 14

[1]Our consumption estimates here are somewhat conservative.

	Count	Yes	No	Unsure
UK Academic	21	24%	71%	5%
Social Media	19	42%	42%	16%
SBCA	11	27%	45%	27%
Experts	11	27%	55%	18%
Total (non-government)	62	31%	55%	15%
Government	12	8%	67%	25%
Total (whole sample)	74	27%	57%	16%

Table 17.3.: Should the discount rate be adjusted for transformational projects?

gives useful guidance to HM Treasury on the matters that would need to be considered to do this.

17.4.21. Therefore, HM Treasury *could* view the economic appraisal of transformational projects as a discounting problem and then apply many of the methods and arguments that we have laid out in these Technical Annexes to do so. If this is the route that it decides to take following from this review, then our opinion is that such evaluations should not deviate from the norms we have laid out here.

17.4.22. But this is not our central recommendation. Instead, our opinion is that transformational projects are just overall too complex to economically appraise by just taking the standard "off-the-shelf" discount rate of the Green Book. This, to us, makes the economic analysis disproportionately small compared to the scale of the resources at stake. As one of the respondents to our survey noted: "To take a fair view of long-term transformational benefits, HM Treasury should change how projects are valued and appraised, not change the discount rate applied to time itself" (paragraph A.2.12.)

17.4.c. Beyond discount rates

17.4.23. Throughout these Technical Annexes, we have been heavily influenced by the principle of proportionality (paragraph 3.4.9). We have largely remained within the existing Green Book setting of the Ramsey Rule (Chapter 7) with updated parameters (Chapter 8), with a fixed risk premium of 1% for almost all projects (Chapters 10 — 11) and 14. We have recommended one extension around rare disaster risk, which reduces the STPR risk-free discount rate but increases the average STPR risk premium (Chapters 12 and 14). This also introduces a specific, and much lower, discount rate for projects that provide clear social insurance. We also recommended that, under exceptional circumstances, the risk premium may be adjusted to be project-specific when there is robust evidence to do so and the scale of the economic case warrants it. We have continued to recommend a declining term structure of discount rates (Chapter 9), updating this schedule on the basis of up-to-date UK consumption data (Chapter 14). And we have argued that survey data, SOC rates, and international guidance all add important information to HM Treasury beyond the Ramsey Rule as it sets the Green Book discount rate (Chapter 6).

17.4.24. Yet we have also emphasised the many assumptions, conceptual choices, and calibration difficulties of operationalising this approach. For standard projects of medium maturity, this can be done in a way that can give government the confidence that economic appraisal is robust. Yet the rate that HM Treasury ultimately chooses must be seen as an estimated rate rather than the unobservable "true" rate.

17.4.25. Yet transformational projects embody many, if not all, of the economic characteristics that the standard discount rate most struggles to handle. Such projects are very likely to have Real Options characteristics where plans can be expanded, shrunk or terminated, and where lessons can be learned for future public investment (Section 5.4). For major technological investments, outcomes may be so difficult to predict that Knightian uncertainty comes into play (Section 5.4). Transformational projects may also be concentrated in certain regions (Chapter 16), where they may bring benefits across multiple generations (Section 17.3). Many of these transformational projects will also involve some types of private finance partnerships (Section 17.2). In short, *all* the complexities and difficulties of welfare analysis, which can often be simplified for more simple projects, come to the fore for transformational ones.

17.4.26. Many of the respondents to our survey (paragraph A.2.12) also argued that transformational projects may affect g itself, particularly on a regional basis. This may require a general equilibrium approach to handle rather than the use of discount rates.

17.4.27. Our recommendation, therefore, is that the welfare analysis that underpins the economic appraisal of transformational projects is enhanced beyond the usual Green Book discount rate standard.

R.23. Transformational projects require explicit welfare analysis.

We recommend that HM Treasury uses explicit welfare analysis rather than discount rates when undertaking the economic appraisal of transformational projects.

17.4.28. We understand that departments may view this as a requirement for additional resources and expertise that they do not currently possess. But we know, from speaking with the Observer Groups, the huge effort that goes into estimating costs and benefits for these largest projects. To then take unrealistic shortcuts when converting these to a present value just ensures that this information is not optimally used. Our recommendation here does not add complexity to the underlying economic problem. It instead aims to ensure that all the complexities are handled in the economic appraisal in a way that is practical, evidenced based, and above all, proportionate.

— ◇ — — ◇ — — ◇ — — ◇ —

Appendices

A. Survey responses

A.1. Characteristics of respondents

A.1.1. In this appendix, we report the characteristics of the respondents to our survey, and qualitative survey responses that we received. Recall from paragraph 3.1.3 that, while we asked all groups identical questions, we used different links to the survey to allow us to determine what route respondents came through. We use:

- “University" to denote those responses that came via the link we included in the emails that we sent to academic departments;
- “RES+BAFA" to the link distributed via the Royal Economic Society and the British Accounting & Finance Association;
- “Social Media" to the link that we put out on LinkedIn, BlueSky, and Substack;
- “SBCA" to the link sent via the Society of Benefit-Cost Analysis, which tends to be more international in membership;
- “Experts" to the link we sent to our academic Expert Panel members in advance of them speaking directly to us.
- “Government" to the link that HM Treasury sent to public sector bodies.

A.1.2. We begin in Table A.1 by noting the number of responses that we received to the survey via different routes. Over 400 people opened the link, of which just over 110 responded to at least one question, and just under 70 responded to at least 20 question:

	Opened		Answered	
		≥ 1 question	≥ 5 questions	≥ 20 questions
University	103	27	21	16
RES + BAFA	18	5	4	4
Social Media	119	28	22	16
SBCA	44	17	13	11
Experts	26	14	12	11
Total (non-government)	310	91	72	58
Government	94	21	13	11
Total (whole sample)	404	112	85	69

Table A.1.: Response rate to our survey through different routes.

A.1.3. Because the BAFA and RES responses were so low, we combine these responses with those received via the email to University departments. We call this combined group ‘UK Academic’ but note that this does not mean that UK academics did not also reply via other routes. As we will see from Table A.2, though, all respondents in this “UK Academic" group are either UK citizens and/or reside in the UK.

A.1.4. In Table A.2, we summarise whether our respondents are UK-based and/or have UK citizenship, as well noting whether their highest academic qualification was in economics or accounting & finance. Apart from via the social media and government channels, many of our respondents held doctorates, and economics was by far the dominant discipline. Fewer than 10% of our overall respondents came from outside economics, finance &

	% UK	% Doctorate	% Economics	% Finance & Accounting
UK Academic	100%	86%	70%	30%
Social Media	58%	32%	79%	5%
SBCA	18%	73%	73%	0%
Experts	45%	91%	100%	0%
Total (non-government)	63%	68%	78%	13%
Government	92%	38%	92%	8%
Total (whole sample)	68%	63%	81%	12%

Table A.2.: Academic and geographical background of our respondents.

accounting. The international mix of respondents varied greatly depending on the channel through which they arrived at the survey.

A.1.5. In Table A.3 we summarise the working experience of our respondents and their knowledge of discount rates. Via all routes, a large proportion of our respondents had more than 16 years working experience, although this was slightly less so for the Social Media and SBCA routes. For the other columns, we report the median response. Via all routes, the median respondent self-reported a high level of knowledge of discount rates and a good knowledge of the Green Book. Apart from the government sector, they also have solid knowledge of both international social discounting and private sector discounting. Note, though, that we report the median respondent per question, and so this is not necessarily the same individual across all columns.

	% More than 16 years experience	Median use of discount rates in daily work	Median Green Book knowledge	Median international social discounting	Median private sector discounting
UK Academic	83%	69%	51%	40%	60%
Social Media	42%	80%	70%	50%	35%
SBCA	45%	81%	75%	90%	50%
Experts	73%	82%	71%	60%	68%
Total (non-government)	63%	80%	70%	58%	60%
Government	54%	75%	82%	20%	26%
Total (whole sample)	61%	80%	70%	50%	57%

Table A.3.: The percentage of respondents with a doctorate, and then the self-reported discount rate knowledge of our median respondent for each question.

A.2. Qualitative comments

A.2.1. We now report the full set of qualitative comments that we received from our survey. Responses from the survey sent out by HM Treasury to public bodies are noted by "*[Government]*" at the end of the comment.

A.2.2. On STPR vs. SOC

- This is quite a hard one to answer as short term fluctuations should not affect long rates and the STPR is very much in that category. SOC is arguably better as it has some mechanism to adjust for changes in preference, but has its flaws. I think intellectually SOC is a better approach. I should point out, I very much believe in Arrow-Lind, whereby the social cost of risk should be near zero in the long run. Generally I think very long run discount rates are too high.
- There are a number of factor that produce a time-varying STPR. For example risk aversion changes through the business cycle and determines a time-varying risk free discount rate. SOC will better account for these time variations.
- I support retaining the STPR approach because public investment decisions should reflect society's time preferences and long-term welfare considerations, rather than being tied solely to market rates of return on capital. This is particularly important for projects with long-term social and environmental benefits.
- While I would support the SOC approach it would depend on how it was implemented given the major difficulties in measuring it.

- Public projects should take a different view than private, market-based projects.
- The government's resource envelope is not set by purely economic forces and used wisely there can be crowding-in of private investment, often hard to capture in CBA, as well as crowding-out. Also, there is no guarantee at all that market rates adequately reflect inter-generational considerations.
- The STPR is an intuitive approach to making decisions for policy makers and is now custom and practice in the Green Book. It also explicitly addresses both the marginal cost of public funds and social time preference.
- STPR seems to give smaller long term discount rates
- My main concerns with switch to rates of return on available capital is that (a) such rates might be more volatile and (b) they might be significantly higher in periods of elevated interest rates. Both factors would act to reduce societal investment, the benefits of of which extend beyond easily quantifiable financial returns.
- Social Opportunity Cost introduces possibly distortive, invariably subjective judgment variables that are anathema to effective central financial rate-setting.
- You frame the question as either/or. I would agree with the idea of making distinctly more use of market based measures of discount rates, but your paper does not really explore these fully, so I would not especially prefer SOC as defined in that paper
- I do not have enough knowledge of this to offer an answer.
- I dont understand the logic for the change- what we have had for years has generally worked. It is poor governmental policies that have lead to serious fluctuations.
- Whether we make welfare assessments using Kaldor-Hicks compensation tests, or by using a social welfare function, the correct way to trade off welfare between individuals is by using the social time preference rate. There isn't a well-founded welfare economics reason to use the SOC. If one is concerned with the opportunity cost of capital in terms of alternative private investments, one can incorporate the shadow price of capital additionally. But if there is very elastic supply for private sector capital (from international markets etc) then this is not necessary.
- Global risks are more likely to impact the UK and thus SOC is more realistic
- it is most appropriate since the SOC may be more biased
- STPR is more general than a market-based approach, and more applicable to a broad range of policy decisions
- A composite rate may be more appropriate (i.e., STPR & SOC).
- Implementing an STPR approach seems more transparent given that rates of return available to capital might reflect private returns rather than social returns
- An SOC based measure is less likely to under-estimate the discount rate in a future likely to feature lower and more variable growth rates than existed on average when the STPR approach was established.
- Does the job, and the use of the MUI (or its elasticity) makes the connection with simultaneous expected changes in values over time explicit and ideally consistent. That is, the net discount rate would probably reduce close to the rate of time preference, as increasing values over time and discounting cancel each other out
- The market assumptions underpinning SOC are highly unrealistic and I think the only defensible approach for tackling intergenerational questions is a normative one.
- one argument could be market failures.
- SOC answers the question "Could we do better by investing elsewhere?" STPR answers the question "Does this project increase social welfare?" It is the latter question that is relevant to cost-benefit analysis. The former can question can be answered in a framework in which the latter question is primary. The reverse is not the case.
- Since the SOC directly accounts for externalities and negative impacts on well-being caused by e.g. climate change damages and biodiversity loss
- Either approach has unique upsides but also flaws. I don't think that a pure form of one or the other approach is the optimal solution. If I had to decide, I would take the STPR approach as I see more fundamental flaws with the SOC approach.
- Given the decoupling of market based rates of return from the lived reality of the majority of people it makes no sense to adopt a SOC.

- Rates of return to capital is a very narrow frame for public projects with non-financial returns. Yes we can estimate the non-financial returns in monetary terms. But despite HMT guidance and acceptance, opponents of a project will always find reasons why non-financial (economic, social whatever you call it) monetary returns are not sufficiently robustly estimated and hence exclude them from consideration. Also while we all have impatience in common (albeit it at different individual levels and of course at a much lower social level) expected returns on capital would be more widely different...SOC for public projects also invites political discussions like is £1 return to MoD spending the same value as £1 return to DoE spending etc. If the valuation of return was sufficient broad and long term, we could differentiate them anyway but without that evidence (being taken seriously) SOC will collapse to the political opinion of the stronger stakeholders.
- Seems to have the strongest grounding in academic literature/evidence
- Government's remit is to maximise social welfare, which may or may not be the same as maximising returns on investment. SOC is a useful sensitivity where there is a narrative or logic in comparison between the two.
- SOC encapsulates the marginal cost of public funds in the discount rate. Better to handle this outside the discount rate using a shadow price for public funds which relates to the specific point in time and volume of public funds.
- The idea that the government should evaluate projects not based on the opportunity cost of the funds used for those projects is mad. A failure to do so means the government will fail to evaluate projects correctly. This can be demonstrated clearly during the period following the 2008 financial crisis when the risk-free rate was close to zero or even negative. Projects which in higher rate environments would be negative value suddenly looked like really good value and should have proceeded. Instead under the STPR approach there is no commensurate adjustment.
- we still consider the STPR to be an appropriate framework, whilst noting its following limitations:

  - not adequately reflecting societal risks in a strictly financial system. - Under-pricing of historical inputs (for example externalities of nature) - No consideration of long-term uncertainty/intergenerational fairness
- We would like to endorse the use of STPR instead of SOC to discount costs and benefits of proposed policy interventions in a Green Book setting.

  A SOC-based approach to discounting would be acceptable if we believed that individuals' decision-making process on present and future investment decisions can be extended to society as a whole. We would like to argue that government investment decisions should be made from a social planner perspective, which requires several considerations that are broader than just market returns. To this end, STPR allows a government to take ethical considerations into account, such as the trade-offs between present and future generations, levels of inequality aversion, and growth considerations that might not be accounted for in the market.

  Furthermore, for a SOC-based approach to hold we need to assume that actors in the market operate under perfect competition. The literature has shown and argued extensively that, when considering a world affected by climate change and nature degradation, market failures are very widespread. Therefore, failing to take this into account would lead to biased interest rates.

  In conclusion, we believe a SOC-based approach would undervalue the benefits of investing in climate action relative to STPR in multiple ways, hence we strongly favour the use of STPR in the Green Book.
- STPR provides a more appropriate framework than a purely financial returns-based approach for appraising projects with long-term societal and climate benefits (intergenerational benefits). But only if those benefits are properly identified, measured, and included in the appraisal in the first place.
- STPR is most relevant in many public decisions. However, there may be times that SOC is relevant. The 2024 New Zealand discount update allowed commercial and non-commercial proposals to be based on different discount rates, with mandatory sensitivity analysis using the other rate, eg STPR non-commercial proposal using SOC for sensitivity analysis.
- The ToR clarifies that the rate will apply to government, or public, initiatives, and not market investments. While market rates should apply to market investments, this does not look to be the focus of this study.
- Consistency is useful. I also think work by Solow and Farrow (JEEM, 2024) on expected discount rates in place of a single rate tends toward using a "lower rate" such as the STPR.
- Arrow et al. (REEP, 2014,) argue that with market imperfections, or if the economy is not already on the optimal consumption path, the STPR should be adopted in favour of the SOC.
- Rates of return to capital reflect multiple market distortions (externalities, market power, etc.) not appropriate to include in assessing the benefits and costs of regulations and public projects. Using the STPR in conjunction with the shadow price of capital is the theoretically correct way to discount for public projects, and with an open economy and a shadow price near one, the STPR should be used.

- The SOC approach bakes in the returns to average risk (beta), market power, unpriced externalities, etc.; it is far better to start with a STPR approach, and deal with issues that can be integrated into the discount rate (like systematic risk) separately.
- For many or most government actions, a STPR approach is probably more appropriate given those actions have effects on non-market outcomes (especially mortality). Based on current research, in cases where there is a strong effect on private capital, I favor adjustment through the shadow price of capital.
- Li and Pizer show that the risk-free consumption rate should be the central rate in benefit-cost analysis, though the social opportunity cost can be used in the short-run, though it should collapse to the consumption rate over a long enough time frame. "
- The rates of return to capital may represent the intertemporal MRS for a certain class of capital owners, but there are multiple issues with that line of thinking. First, discount rates will often be applied to future generations who don't yet exist or have a say in discount rates. Advocates of using the rate of return on capital argue that their approach is "descriptive" while the STPR approach is "normative", but the fact that we are applying current rates of return to capital to future generations for many policies (even just 20 years in the future there will be many people not born yet and many owners of capital will be dead) I think undermines this claim around rates of capital as being truly descriptive in this context. My view is that anytime one uses analysis that is intended to be considered as part of the policymaking process, this is inherently a normative exercise (even if one uses market rates of return on capital), and I believe the STPR approach has more coherent theoretical and normative underpinnings. Second, the STPR approach using Ramsey Discounting allows analysts to better align their discounting scheme to the expected future income growth scenario for a particular policy (because of the $\mu * g$ term). Third, having a clearly stated STPR with a clearly stated ""pure rate of time preference,"" allows analysts to coherently discount benefits and costs that have been income weighted. The Green Book allows income weighting, and in my view, that approach is superior for Benefit cost analysis compared to a Kaldor-Hicks approach, since accounting for diminishing marginal utility is something that reasonable Social Welfare Functions typically do. "
- I believe STPR is better aligned with the principles of social welfare that the Green Book adheres to.
- I believe that the risk free consumption rate is the appropriate central rate bSe on the work of Li and Pizer.
- STPR more consistent with utility theory
- The STPR approach is welfare-grounded, while the SOC approach is disconnected from the particular aspects of the project. For example, for a risk-free project, the SOC will systematically use a discount rate that is far too high.
- Each approach has its specific advantages. There may be a case for using both in some cases to examine whether there are major divergences which deserve further analysis
- It more closely reflects the actual oportunity costs of many projects which call on a time stream of public funds. Also it does not impose the particular social welfare function the is implied by a particular STPR. However, I would like to see a more nuanced approach where time stream of different types of cost and benefits are first valued then discounted using rates based on where costs and benefits fall. Using STPR not only implies a disputed welfare function it implicity makes very strong assumptions about capital markets, constraints on public expenditure and its allocation
- The STPR best represents the long term real relative social value of a unit of resource now versus in the future. This is the correct concept of the discount rate. Any scarcity of public resources (SOC>STP) should be handled via a MCPF or SOCEF.
- Both STPR and SOC have pros and cons. Maybe the main limitation of STPR originates from the difficulty to assess the time preference (being a behavioural variable to be estimated across diverse social groups), while the opportunity cost of capital in the private sector might be more tangible and easily detected. On the other hand, STPP provides an interesting perspective, based on individual preferences, and not only on private capital, potentially considering the preference of any type of citizen.
- If we think of the SOC as the risk-free social opportunity cost of capital, it is unclear to me why this should differ markedly from the STPR. Any difference either reflect who or what is being included.
- it applies to public investment and regulation. Opportunity cost could be considered but rather in relationship with the scarcity of public funding.
- Depends on context - SOC when allocating scarce public funds for near term projects with a potentially high BCR, STPR for policy issues with a longer time horizon and critically for climate change mitigation actions.
- Both approaches raise issues. STPR misses the point of risk and uncertainty. And SOC raises the issue of market short-termism.
- Many questions related to CBA concern welfare, so I would favor a STRP approach.

- The STPR approach allows to better distinguish between choosing normative parameters, such as the rate of pure time preference and the intertemporal inequality aversion, and relate this to expectations about future growth in consumption that are non-normative. This allows to better reflect the public's preferences and therefore argue to the public about which projects are conducted and which are not.
- The STPR is useful for conceptualizing some of the drivers of time preference but completely inappropriate for real-world evaluation of project cost in most instances. The SOC reflects the value that society places on cash flows with different risk, and that value is also relevant to governments acting on behalf of their populations. I would add the caveat that reference point for inferring market rates, using a STPR is a reasonable fallback. Also, to the extent that benefits are being evaluated and there is no natural for very long-term evaluations when market prices do not extend to the horizon of interest, then the STPR may be the best available alternative for discounting. But for very long-term costs and benefits, a discounted present value approach is too noisy to be a reliable and decision-relevant tool.
- As we look to optimise the net social value, we should therefore be using an equivalent form of discount rate. *[Government]*
- Very difficult to accurately predict what the Social Opportunity Cost is on the rates of return available to capital. Particularly when transport schemes can have a 60 year appraisal period. *[Government]*
- STPR has proven too rigid to reflect changing conditions. For example, assuming a 2% rate of real GDP growth seems too optimistic today. On top of this, market data doesn't need to justify its connection to agents' time preferences but, on the other hand, it will need to be cleaned from short term market noise. It can also better reflect varying levels of project size, risk and maturity. *[Government]*
- SOC provides a more intuitive view of comparing the cost and benefit. Currently, where projects have a BCR that hovers around 1, the social benefits may not be greater than the social costs as the cost of the finance is not included. Secondly, discounting is currently largely spoken of in terms of reflecting time preference, but this is a small feature of the STPR (with an over-estimate of growth). *[Government]*
- I think both approaches have been used in the past, e.g. when the discount rate was set at 6% and (I think) represented the opportunity cost of private capital back then. The STPR is probably more transparent and therefore if I had to chose, I would lean towards it. I have however responded 'unsure' due to limited knowledge of the pros and cons of either approach. Guessing others will be better placed to judge this. *[Government]*
- STPR is more universally applicable to different types of public sector projects, better capturing the value of future longer-term benefits *[Government]*
- Consistency is more important. *[Government]*
- I'm not clear why SOC is better than STPR and aware we have switched between them in the past. *[Government]*
- Insufficient evidence to conclude SOC is better than STPR and rates of return are much more variable *[Government]*
- The SOC approach uses higher private-sector-based rates, often unfairly undervaluing the long-term benefits of projects targeting household *[Government]*
- The STPR contains a reflection of the SOC in the growth rate. Wider economic growth - in economic theory - essentially refers to "all other goods". I am deeply skeptical that we are accurately appraising the benefits of any government policy. We do our best on a case by case basis to help make an individual decision. SOC seems to require us to have a full suite of policy proposals ranked and ready to go sufficient to actually produce a discount rate. I don't think it is technically feasible. Furthermore - discount rates are used widely in global statistics. If anything in National Statistical circles there has been pressure on the UK to use higher rates of discount that reflect global use. Our own - well developed internal discount rates protected us from that to some degree. Switching between SOC and STPR in different circumstances (e.g. international statistics) would be confusing and likely lead to mistakes. I may be wrong as I have not much experience of SOC but i am deeply skeptical that it is something we could make usefully operational across government. The decision on what the next best social project is is fundamentally a political one - for democratic discussion aided by economics. it is not one we are able to quantify. *[Government]*
- The UK government should retain its current STPR approach for the discount rate but should include a separate Social Cost of External Financing (SOCEF) applied to public sector costs rather than the discount rate to reflect the distortionary impact of taxes on the economy. Argument: - The use of the STPR is consistent with the Green Book approach which is concerned with how interventions impact on the social welfare over time of UK residents. The use of pure time preference and wealth effects in the STPR allows benefits and costs to be discounted over time in a way that is fully consistent with micro foundations of consumer theory. - The SOC is not appropriate for a social welfare based approach. It ignores the fact that 90% of government funding is through taxation. The small amount of public sector expenditure funded through debt is traded on international markets and does not reflect the preferences of UK residents. - However, in addition to STPR the Green Book should explicitly include a SOCEF measure to reflect the fact that money raised through taxation such as business taxes, income tax and national insurance contributions has a distortionary effect on the economy. This should be applied to public sector costs to reflect their true opportunity cost. This is discussed further in Spackman 2025 Paper on Alternative Approaches to Discounting. *[Government]*

A.2.3. On whether the current Ramsey Rule parameters remain broadly appropriate:

- No, but it is very hard to come up with something better.
- There is no rationale for discriminating against future individuals, so the "pure rate of time preference" so should be 0%. For elaboration, see Cowen and Parfit (1992).
- No. The assumed 2% long-run real per capita consumption growth rate appears optimistic in the current UK macroeconomic context, which would imply an STPR that is too high for discounting medium-term public projects. Updating the parameters would better reflect prevailing growth expectations.
- For a 10-year horizon, there is no declining discount rate issue.
- 2% growth in consumption seems on the high side given the experience of the last 15 years,
- Long-term discount rates should be lower.
- The expected annual rate of change has clearly decreased. The elasticity of marginal utility of income has always been more or less more supposition than rooted in sound evidence, we know more about the relationship now and that cross-sectional evidence differs from longitudinal impacts of general increases in income. But it remains a rather arbitrary number, and changing it significantly would make the Ramsey formula either more cumbersome to apply or dictate unrealistic parameters for the time preference and growth elements. So leave it at one. The same applies to a lesser extent about time preference. So, rather than 'angels on a pin-head' debates, simply use a lower and more realist expected growth rate of 1.5, still rather optimistic but guarded against reckless projects, for GDP per capita, to produce an overall STPR of 2 + 1= 3, So, 3 when we include the premium, which given prevailing market rates is also about the rate would use from an opportunity costs approach too. And certainly, public investment has lagged our G7 competitors, so a lower STPR for public investment seems warranted.
- The UK's real per capita consumption growth rate has been incredibly weak and remains low. While this is the case in many developed countries around the world, it is an acute problem in the UK. The empirical evidence for keeping this at 2% does not seem plausible, and the future growth trajectory of the UK at present shows no signs of material improvement.
- this elasticity is ideal. Moreover the Ramsey rule ignores great social/demographic changes that are occurring, i.e. world population peaking before the end of this century.
- Longer term discount rates should be higher, driven by increasing time preference rate and the elasticity of marginal utility of consumption that also should increase.
- The E of MU should be raised to 1.5
- Hard to be unequivocal on this. My stance is that it's important to encourage societal investment and that the costs of errors are asymmetric - more costly for society (particularly future generations) if the rate is set too high than too low. I think that, at 2%, the Expected average annual real per capita consumption growth is too high.
- Don't fix what isn't broken. There has been insufficient exploration to date of the under-anticipated consdequences of corrupting the present order.
- Where to begin? All three of the assumptions are contentious, and measured with significant error.
- That's probably a little bit low for elasticity of marginal utility of consumption and a little bit high for growth, but it's fine to get an overall 3.5 for 10 year horizon discounting.
- Consumption per capita growth rate seems too optimistic relative to recent and projected UK growth performance. PRPT and EMUC are broadly OK.
- 2% growth is no longer representative of realistic expectations for economic growth.
- Rate of pure time preference should be zero, or very low (at most 0.1%)
- The expected average annual real per capita consumption growth rate of 2% is too ambitious. There is nothing currently to suggest that productivity in the UK will improve enough to justify this rate.
- The elasticity of marginal utility of consumption should be revised higher
- A social discount rate should be a premium (comprising premia for both risk and social (non-market / externality) costs) above a non-constant (in both the maturity and time domains) risk-free rate derived from conventional and/or index-linked gilts. Gilts are, after all, the ultimate source of financing.
- The elasticity of consumption is too low and inconsistent with the latest empirical evidence and its use in other parts of the Green Book.

- Although I'm less concerned with 10-year discounting questions than longer ones, I have concerns about the wealth effect parameter of 2%. I think it is crucial to move towards an inclusive wealth measure here, rather than consumption or GDP. Ideally, we'd measure stocks not flows and the measure would include natural capital. I appreciate this is challenging in practice, but I still think some concession for climate change and nature loss is necessary; the current parameter states that we can discount future generations' welfare because they'll be wealthier. I'd ask 1) are we sure about this given the environmental harm we've stored up for later and 2) what's our measure of wealth/welfare/wellbeing?
- Considering the short time horizon, I agree, however the pure time preference could be adjusted to 0%
- I think 2. and 3. are very reasonable. For 1. it really depends - are we prescribing a value based on ethical arguments or should this reflect the (social) preferences of a broader public? In the latter case, I think the PRTP would- and should be significantly higher. "
- Arguably, both δ and g ought to be significantly lower:

  1. Based on the classical impartial Utilitarian philosophy, Pigou, Ramsey, Harrod, Stern, and others argue for a pure time preference of $\delta = 0\%$.

  2. It is safe to keep the current assumption of logarithmic utility ($\mu = 1$).

  3. It is overly optimistic to expect a 2% annual growth rate of future real per capita consumption. Even abstracting from cyclical shocks, the implied assumption of a structural productivity reversal is not empirically supported by the persistently low trend observed in recent decades. Using chained volume measures, the 2024 UK National Accounts show that real household final consumption expenditure (HHFCE) per capita (GBPOP) increased by less than 0.5% per year on average over both the past decade (0.45% since 2014) and the past two decades (0.46% since 2004)."
- 3 is too optimistic in a UK context where real GDP per cap will not grow
- Growth rate is and seems to remain unattainable
- I don't think risk adjustment should be through the discount rate. It should be through the quantity of costs and benefits. I am in general in favour of lower discount rates for environmental projects because the time preference should really not apply (given how dependent society is on nature) and growth rate is also not appropriate... yes we may be richer and have more ability to consume in teh future but what are we going to consume if we deplete nature. So I said NO as I'd like a lower discount rate. I wondered whether the relatively short period of 10 years warrant a higher rate....But no, because that could lead to even more cumulative longer term effects on the environment.
- g is more like 1%. elasticity of MU income more like 1.5. Pure time preference - don't have a strong view and willing to be swayed (perhaps normatively).
- Sadly no-one expects real per capita consumption growth of 2% p.a. anymore.
- Points 2. and 3. should be revised downwards.
- Growth is too high
- 1. 0.5 may be an understatement for some environmental matters where there are attributable long-term targets and lags between action and outcome (e.g., CCA and EA). No view on 2) or 3)
- Expected average annual real per capita consumption growth of 2% is beginning to look optimistic. The downturn looks like more than a blip. Rate of pure time preference of 0.5% may be on the high side."
- We can readily observe the real yield on UK government inflation linked bonds (approx. 1%). There is no need for the Ramsey Rule.
- Re 0.5% pure time preference. The pure time preference represents society's preference for consumption now rather than in future and is often the subject of debate in an intergenerational fairness context. We believe this debate is particularly relevant to transformational projects which are seeking to make changes to the present system which will impact future generations more than our own. We therefore recommend that this could be as low as zero for transformational projects.

  Re. expected average annual real per capita consumption growth rate of 2%: Increased uncertainty and the skewed nature of the outcomes distribution means there is now a very real plausibility of far lower growth outcomes than previously assumed, potentially exposing future generations to a far greater burden of costs than those faced by current generations. We therefore consider that a growth rate of 1.5% may be more appropriate.
- For a 10 year UK public project, the current Ramsey rule values in the Green Book is valid and appropriate for discounting a 10 year public project. Over a 10-year horizon the compounding effect that makes parameter choice so critical for 50 year plus projects, simply doesn't have time to materialise. (For example, even if g is slightly overstated at 2%, the distortion to present values over a decade is small enough to be within acceptable margins of error).
- It is worth reviewing the parameters, eg the growth rate may be lower.

- The pure rate of time preference might be expected to be a bit higher, at 0.75, based on the range of plausible values in the literature and in practice.

  The expected consumption rate over the next 10 years might be 1.5%, based on projected real GDP per capita growth rates by the OBR.

  The elasticity of marginal welfare wrt consumption might be a tad high at 1, but it's in the right ballpark. (Eg, something closer to 0.8 is probably more accurate.)

  The current values correspond to 2.5%, whereas the values I suggest here correspond to 2%
- Reasonable. The elasticity of 1 could be clarified as establishing the discount rate being for the median of the society given evidence on the non-constancy of this elasticity.
- Given the 2 choices, I have answered 'Yes'. However, I argue that one may wish to adopt a different rate of pure time preference (RPTP) for different goods - eg Parliament buildings may have a positive RPTP whereas the maintenance of democracy may have a (close to) zero rate.
- The risk-free STPR should be closer to 2.0%, as explained in the 2023 update to the U.S. Circular A-4, and as confirmed in multiple recent expert elicitations. I know of no valid reason why the U.K. STPR should be significantly different than the best estimate of the U.S. STPR (i.e., the estimate from the 2023 update to Circular A-4).
- Rather than estimating the STRP through a Ramsey method, I would encourage just setting the STRP equal to real yields (gilts) over the relevant time period (10 years). Right now that is roughly 2.5% (https://www.bankofengland.co.uk/statistics/yield-curves), making current values in the Ramsey approach fit. But the Ramsey rule is an incomplete description of the sources of the STRP (what if risk preferences are more Epstein-Zin than CRRA? Etc.). In addition, the values are likely off; the elasticity of the marginal utility of income (it's harder to measure consumption) is likely around 1.4 (see https://bidenwhitehouse.archives.gov/wp-content/uploads/2023/11/CircularA-4Explanation.pdf pages 50-52) and the expected average annual real per capital income (again, easier to measure) growth rate is likely closer to 1.5%. (See, e.g., IMF projections).
- I don't personally favor a Ramsey Rule approach to estimating the STPR in part because one could reasonably quibble with every one of these numbers. Of course the pure rate of time preference is hard/impossible to pin down (and to my taste, 0.5% is much too high). Estimates of the elasticity of marginal utility vary by orders of magnitude depending on the study, and even the conventional range is between 1 and 10. 1 is on the low side of conventional. The growth rate number seems optimistic.
- The 30-year average of per capita GDP growth and the long-run real return on the UK discount rate do not match these numbers. Real-rate of return is around 1.7% or so and the per capita growth rate is around 1.4% (according to World Bank data). So, I think a pure rate of time preference closer to 0.3% makes sense.
- I think the $\mu = 1$ term and the consumption growth rate=2% terms are good. I think the PRTP=0.5% term is high. For a STPR approach, I don't think there should be a reason to discriminate against future people's utility by virtue of them being future people. One could account for extinction risk as part of the PRTP (e.g., the Stern Report used 0.1% based on that motivation... although it did seem like that number was sort of pulled out of thin air). I would prefer a PRTP = 0%, or perhaps 0.1% (or somewhere in between) over the 0.5%.
- Explained in those terms, I can't see why not. But I may have answered "Not sure" if that option was available! As I would like to give it more thought
- The 30-year average of the real rate of return on a UK bond is somewhere around 1.5% to 1.8%. The average real growth rate of the UK is around 1.6% annually. So, if you are going to use the simple Ramsey rate, I think a pure rate of time preference around 0.2 to 0.4 can really be justified.
- The RPTP and EMUC are reasonable enough, but the mean growth rate of 2% seems a little high for the UK in recent history, which has averaged around 1% (https://data.worldbank.org/indicator/NY.GDP.PCAP.KD.ZG?locations=GB). Is there a reason to expect a reversal of this trend?
- It is simple and straightforward and any potential errors are not likely to be significant over a 20 year period.
- Pure time preference is fine for individual decision but not in social choices across individuals at different time points. It implies arbitrary normative values that are not consistent or supported. 2% real growth is silly (too high) Elasticity of 1 seems reasonable but I would cast this as inequality aversion for social choices and appeal to that evidence rather than individual preference or some supposed individual utility function
- My impression is that looking at the last fifteen to twenty years since the GFC, the actual real rate of GDP/cap is well below 2 per cent. I don't see strong evidence that is biassed downwards. Offsetting this, I think there is evidence that the MU of income elasticity is somewhat above unity. Overall, I think those two things roughly balance out, so the STPR should remain at around its current level.
- They do in general, although in 10 years many internal and external changes take place, so it might be necessary to update these values or combine this model with additional analyses with different methodology

- I'm not sure what 'remain valid' means. That is, is ok to discount at a constant rate? Or do I think the assumptions are valid in the first place? I think the pure rate of time preference and μ are a bit low. Whereas real income growth at the median (which I prefer to the mean for social questions) is below 1% for the UK.

  On the other hand, I'm not so worried about a constant risk free rate, especially when income growth is low.
- it is decreasing after 30 years (due to uncertainty about future values of its effects), not 10
- growth rate now expected to be considerably lower, for longer time horizons pure time preference should be lower, and small chances of bad outcomes, discount rate should be adjusted by probability and level of consumption in bad cases
- Elasticity too small. Substract a precautionary term.
- I would set the elasticity of marginal utility of consumption higher than 1 and the expected average annual real per capita consumption growth rate lower than 2%. I am a bit undecided when it comes to the rate of pure time preference for the 10 year horizon (for longer horizons, it should be lower).
- Since the choice of the first two parameters is ethical, I generally consider these values acceptable as they can reflect the preferences of society. The value of 0.5% for the rate of pure time preference is at the median of the survey by Drupp et al (Discounting Disentangled) so it is quite justifiable by this assessment. Inequality aversion or consumption smoothing motive at 1 also seems to be in line with what people prefer from my reading of the evidence. For assuming consumption growth at 2%, this value seems appropriate for a 10-year horizon as it is close to what was achieved prior to the pandemic. One may need to revise this for longer-term projects as growth at 2% could be too high. The UK could consider a lower rate of pure time preference as a sensitivity analysis to be more impartial to future generations.
- As noted in previous answer, none of these quantities reflect valuation in the real world. Rather, they are derived from estimates of parameters in highly stylzed models of the macroeconomy that are informative about basic mechanisms but not about how society evaluates the effects of time and risk on value.
- A regular (say every 5 to 10 years), pre-announced review of the expected average annual real per capita consumption growth rate should be undertaken. It is currently hard to justify this 2% based on recent and forecasted trends. *[Government]* "
- 1) Recent evidence shows that average per capita consumption does not grow at 2%. I think this assumption should be revised down. 2) I think pure time preference being 0.5% is a difficult assumption to make for all projects. There may be some longer term projects or different types of projects where this is not the case. *[Government]*
- All three are a bit outdated and per capita consumption growth, for example, is far below 2% *[Government]*
- The expected average annual real per capita consumption is no longer appropriate. This seems to be an over estimate. *[Government]*
- Points 1 and 2 are both fine to my knowledge (which is limited, admittedly). I guess the expected average annual real per capita consumption growth rate of 2% could be questioned these days based on more recent trends. It depends what the tipping point is however on this - longer terms it might still be appropriate. I guess only time will tell. *[Government]*
- I would argue that a growth in per capita consumption rate is too high. I think the last few years have shown that there also need to be more nuance to this - for example, in education we see that outcomes and pressures do not move smoothly with average consumption growth and this may be lead to us overstating how much (if at all) 'better off' future cohorts may be *[Government]*
- pure time preference could be less - due to growing climate concerns I think the public might now value the future, and future generations more highly than they used to *[Government]*
- I would have expected that consumption growth would have fallen, although recognise that may not be correct. *[Government]*
- I know to some extent growth rate here is a normative function, but seems to have been a structural break around GFC. Growth rate of 2% seems too high since GFC *[Government]*
- From Groom and Maddison (2018) it seems difficult to argue that the elasticity is anything other than 1.5, rather than 1. It's also clear from that paper that the per capita consumption growth rate has been on a clear downward trend and should be revised downwards accordingly. *[Government]*
- Largely because I have seen lots of discussions that play with these figures and have long debates over them. Only to find themselves falling back to where we are. I certainly would not want the rates to get any larger and crowd out necessary long term planning. Global uncertainty might edge pure time preference up a bit - but it is at best a guestimate and should probably be left alone. Growth rates are starting to ease back. Perhaps when/if the OBR changes its central assumptions for where long term growth falls back to we could adjust that and link it to that decision. Though who really knows how accurate that stuff is? Elasticity of marginal utility of consumption is a finger in the air following a lot of economic dancing and should probably stay the same. *[Government]*

- Values are an empirical question not ones that can be answered a priori. However, the current "risk-free" Ramsey Rule values are quite old and should be reviewed:: - For example is the Green Book long term real per capita growth rate of 2% still right to use? And what period should a long term growth rate be measured over?. - The elasticity of the MUC and the rate of pure time preference should be reviewed periodically (every 15-20 years?). *[Government]*
- It is likely that the consumption growth figure is somewhat on the high side, assuming consumption growth is aligned with GDP growth. *[Government]*

A.2.4. On what asset should be used to estimate the social discount rate should HM Treasury decide to go to an SOC approach

- A 10 year UK government bond, absolutely should not be a commercial discount rate.
- A government debt yield to maturity (YTM), because the Treasury always has the option to use the money to redeem / pay back government debt. Every pound used to pay back government debt will produce a return for HM Treasury equal to the Yield to Maturity. It terms of maturities, I suggest a weighted average of the Yield To Maturities of all outstanding public debt, weighted by the pound amount of outstanding debt per each maturity.
- A natural benchmark would be the real yield on medium-term UK government bonds (index-linked gilts), as they represent the closest observable proxy for the opportunity cost of low-risk capital in the UK economy. Using a 10-year real gilt yield provides a transparent, market-based reference consistent with a SOC framework.
- The real yield on 10-year UK index-linked gilts
- Rate on 10 year government bonds at the time of decision. Need to avoid switch in projects. Also need to take account of it being used for liabilities, such as nuclear decommissioning and pensions.
- I would not go for this.
- Don't know. It's very problematic, which is another reason why the STPR approach is better
- Long-term government bonds should be used as this reflects the cost of government borrowing.
- NB last question, the Ramsey rule seems to ignore yields on long term ILGs in the UK. These represent an important funding requirement.
- The average of UK Gilt yields, US Treasuries, German Bunds for 10-year maturity.
- 10-year gilt rate, although I'm concerned this may still be too high
- The US 10 year treasury bills rate x 125
- At last a question with a clear-cut answer! Implied zero coupon yield to maturity on a 10 year indexed bond
- Government bond market, as this reflects opportunity cost of borrowing more to fund this project.
- a public / private health-related project
- There is no financial asset that provides a good analogue for the UK's risk-free social discount rate over a 10-year horizon. The closest observable proxy is the real yield on UK index-linked government bonds at a 10-year maturity, but these yields are heavily affected by risk premia, institutional demand and market distortions.
- US Treasuries
- UK equity market real return
- Index of AAA government bonds
- conventional or index-linked gilts of 10 years to maturity
- government bonds
- maybe sovereignty bonds
- Any observed rate must be adjusted to reflect the externalities of the underlying investments. All standard options are private rates of return that completely ignore those external costs.
- A measure that is very sensitive to environmental externalities
- Bonds state
- The UK 10 Year Bond
- 10 year UK treasury bonds
- Using low-risk government bonds might serve as a proxy for the risk-free discount rate.

- there are no risk-free assets, projects or returns to capital in the real world. Let's be real. I know US bonds are used for example for risk-free example. But really?
- Private capital investment projects - perhaps tolls roads for transport (for example). Maybe this means looking at equity returns?
- Average return from investment in 10-year Treasury bonds.
- green bonds/gilts.
- Long-run real return on government bonds
- The ten-year UK inflation-linked bond rate (sure we can have some quibbles about whether this is risk-free or not, but surely much improved over the current approach).
- Under a SOC approach public spending displaces private investment? So, the benchmark is the real pre-tax return on marginal private sector capital (4-5% real per year).
- government bonds
- Risk free market rate would be the Treasury bond rate
- 10 year government bonds.
- Expected real return on a 10 year government bond
- A 10-year government bond.
- If a comparable public project can be found, that approach could work. Perhaps there are private powerplant projects that mirror public ones, and the financing could be used to back out a risk-inclusive discount rate. But in general, it is not possible. Benefits and costs in welfare units should be discounted at the real risk-free rate that market participants could borrow at, i.e., real gilts.
- Other than government bonds, I think this is a fool's errand.
- I think that this is difficult as it would include, as these returns include externalities, market power, and private risk premiums (which may match social risk premiums). Therefore, I do like the idea of calculating the social opportunity cost of capital directly from the consumption rate, following Newell et al. by adjusting for the tax wedge. This still leaves you with addressing the risk premium, but this is really project specific according to Gollier.
- Real interest rates on 10 Year U.K. Treasury Notes.
- Not sure what the alternatives would be
- 3%. This is roughly consistent with the pre-tax investment return consistent with the 2% implied by my recommended Ramsey parameters: RPTP=0.2%, EMUC=1, mu=1.5% => STPR = 2%. At a 33% tax rate, this implies a pre-tax risk-free ROI of 3%.
- In most cases this should be an investment in physical capital in the private sector, not financial assets. This would probably be larger than a typical STPR and thus treated with care for a public project.
- For project with time stream of public expenditure costs and benefits would use the goverment bond yields, in principle for each term of the liabilities. However, for the time stream of impacts on private capital would use estimates of (maybe even sector specific) opportunity cost of capital (eg pharma is about 11%). But then we would need to value the PV of these different time streams. We could use estimates of marginal cost public funds (a lot of assumptions about welfare and markets/ type taxation) or I would recommend using empirical estimates of marginal value of public funds (at least 15 for the NHS)
- They could be the government bonds, or the return on core infrastructure asset class, which are usually comparable to government bonds in terms of risk and return
- Maybe 20+ year inflation-linked gilts? I do like Giglio's Singapore paper too. Could also use US inflation adjusted treasuries. My hunch is you are going to end pretty close.
- government bond
- regulated public utilities such as transmission investment
- inflation-indexed 10-year sovereign bond
- Long-term government bonds for 10 years.
- Standard valuation principles should be applied, specifically, a fair value framework that determines value and discount rates for a project based on market-derived rates for similar projects or approximations thereof.
- Don't know, it is partly due to this I would not advise of this switch. *[Government]*

- I don't favour this approach. I think it's too difficult to forecast. *[Government]*
- public debt of similar maturity *[Government]*
- The marginal gilt rate used for the next spending should be used. Assuming additional spend comes from additional borrowing seems to be the most appropriate measure. An alternative would be to measure the social cost of additional taxation, but this is less clearcut. *[Government]*
- I'll admit that I don't know. I would guess it should be the average return to capital from private projects, assuming 100% crowding out. *[Government]*
- Bond yields for a government-backed private finance project. *[Government]*
- 10y gilt yields *[Government]*
- average of 10 year gilt over a recent period of time *[Government]*
- I thought SOC was supposed to be based on alternative social projects? I do not think it should be linked in any way to private returns it does not represent a realistic alternative to public spending. it is not its purpose. If it should look anywhere it should be at a key public good chosen by the cabinet. Defence, Education, Health or the Environment. If your intent is to reduce the impacts of climate change the next best way to spend the money is not on the current high returns on hyrdrocarbon exploration. *[Government]*
- The UK government should not switch to an SOC approach. The STPR is the appropriate discount rate to use as it is consistent with the approach put forward in the Green Book on maximising social welfare. 90% of funding for government expenditure comes through taxation rather than through borrowing. It is difficult to think what an appropriate SOC would be for taxation. However the distortionary impacts of taxation on the economy should be added to public sector costs through the use of a SOCEF it would need further work to determine the empirical value of this for the UK. *[Government]*

A.2.5. On whether the risk-free discount rate should vary with the time horizon

- There is no question we need to decay discount rates in the very long run.
- It seems unlikely that consumption growth will remain constant. Indeed, this is not what has happened historically (e.g., growth in the 20th century was much faster than growth in the 19th century). However, whether this can be reliably forecasted is another matter.
- Any term structure is usually the result of some sort of risk premium. Here we are talking about a risk free rate before risk.
- The discount rate should vary with the time horizon to reflect increasing uncertainty over long-term growth and interest rates. A time-declining structure avoids systematically undervaluing long-term public and environmental projects.
- A constant discount rate across all horizons is not theoretically and empirically well grounded. Uncertainty about long-run growth implies a declining certainty-equivalent discount rate as the horizon lengthens (Weitzman-type aggregation). Observable real yield curves (like index-linked gilts) are not flat. Using a single rate for all maturities implicitly misprices long-lived benefits and costs.
- Should reflect term structure of interest rates and increased risk at longer horizons,
- Need a downward-sloping term structure.
- Inter generational and uncertainty aspects become more important the longer the time period
- The discount rate should vary with declining discount rates being applied for very long-term projects and this is also compatible with a Ramsey set-up for the STPR.
- ..maybe they should vary with time as the YTM on the longest ILG changes..
- I think it should decline but that the decline is so small that it can be ignored
- I don't think it is helpful to revise the expected value of the project once the commitment has been made to invest (assuming appropriate due diligence and costing prior to project initiation).
- Counter-balanced to phase in the 8.3 year economic cycle: lower (possibly negative) immediately following recession, extra as end of cycle approaches.
- There is probably a logical and empirical case for it falling over long horizons, but any effect is almost certainly dwarfed by parameter uncertainty and general equilibrium considerations
- Weitzman 1998 uncertainty.

- While variation is more difficult to take into account it is also more realistic and will reflect the growing speed of changes
- We know from theory and applications of that theory that the certainty-equivalent risk-free rate should decrease with the horizon.
- Multiple reasons: growth rates vary; individual preferences are non-exponential; degree of commitment (which affects opportunity costs) decreases over time; uncertainty about future preferences (which is different than your definition of risk, as I understand it) goes up over time; etc
- If a composite rate is used (STPR + SOC) then it doesn't matter - the rate will adjust as parameters change.
- declining rate when as time horizons involve future generations
- The yield curve for government financing instruments is not constant across maturity, even in real terms.
- It is unlikely that economic / consumption growth will be constant over time, and particularly if there is correlation in consumption growth compelling evidence of a declining discount rate is provided by Arrow et al. (2014) and Weitzman.
- One issue I have is that the theory stems from infinite time-horizon utility maximisation, whereby you're essentially expanding the intertemporal behaviour of one representative agent to that of how society thinks about multiple generations' intertemporal importance. We as individuals discount our own future wellbeing but I consider that to be a very different phenomenon to discounting the welfare of others who do not yet exist. However, in both case, this points towards declining discount rates. In the former, evidence shows we tend to apply declining discounting wrt our own welfare. In the latter, I'd argue that our welfare is no more important than that of future generations, suggesting a pure time preference rate of 0 for intergenerational questions. Perhaps in practice this means a discontinuity from 0.5 to 0 after 15-20 years?
- One argument could be tail risks
- Ramsey discounting requires time-varying rates.
- Discounting is very sensitive to time horizons, therefore a varying discount rate is sensible
- Thinking in terms of the SOC approach, it is difficult to think of reference assets for certain long time horizons. Thinking of the STPR approach, other ethical judgements might apply for long time horizons than for short time horizons (i.e. discrimination of future generations through the PRTP) and if it should reflect population preferences, then these might change over time as well.
- decreasing over time to mimic non-exponential discounting long-term
- If δ is initially greater than 0%, it should approach 0% as the time horizon increases. Short-term positive δ can reflect impatience within a generation, while $\delta = 0$ embodies the ethical argument of intergenerational impartiality.
- Depends what you mean by vary, the discount rate should be updated to acknowledge expected changes in the relative scarcity or improvements in future resources
- the declining one that we currently have makes sense to me but can start from an even lower basis
- On the one hand evidence for DDR is compelling, but I would be keen see the relationship between that and the risk premium clarified in the specific case where unit benefit valuations grow with income
- The discount rate should increase as the project stretches over into the future.
- There are various technical explanation for the discount rate to decline over time, including uncertainty about the future discount rate. However, from a more values-based approach, a declining rate ensures that due account is taken of the welfare of future generations.
- The yield curve typically slopes upwards.
- we acknowledge the increased uncertainty at longer time horizons which has led to the time dependant rate, and note that climate change and other evolving uncertainties may increase this uncertainty further.
- Because climate change involves very long-term, uncertain and potentially irreversible impacts, using a constant discount rate would heavily undervalue distant damages. A declining rate better reflects growing uncertainty and gives appropriate weight to future generations.
- Factor in ethical and uncertainty considerations.
- There is uncertainty about what the correct discount rate is. There is a range of values it could take, and nobody is certain what it is. The certainty equivalent discount rate declines with time

- I think this gets into political economy where some possibly legitimate details, such as time varying discount rates, clash with explainability and consistency. Analysts can always use different discount rates in sensitivity analyses but in my view a common, constant, value assists in cross project comparisons and credibility to the general public/politicians. Don't let the ideal be the enemy of the useful.
- There are several convincing reasons for adopting declining discount rates. Reasons include: These include uncertainty about the rate of future consumption growth including potential non-linearities such as "tipping points" (Dasgupta, 2008; Gollier, 2012; Pindyck and Wang, 2013; Arrow et al., 2014), uncertainty about timing of payoffs (Dasgupta and Maskin, 2005), and uncertainty about the appropriate (risk-free) discount rate itself (Weitzman, 1998, 2001; Pearce et al., 2003).
- Due to uncertainty, the certainty-equivalent discount rate has a declining schedule. Economy-wide growth rates also could likely slow over time due to future impacts like climate change; limited substitutability of non-market goods will also increase uncertainty over growth rates over time; and future shocks to growth rate may be correlated over time, leading to a declining schedule.
- For the reasons described in Arrow et al. (2014): https://www.journals.uchicago.edu/doi/full/10.1093/reep/reu008.
- There are multiple good answers to this question that have now been published in the academic literature. I refer you to any of the many good papers by Gollier, Groom, or coauthors.
- This is consistent with the Ramsey rate.
- If one were to take the STPR approach, one should consistently use the same μ and ρ values. In my view, applying income weighting (that applies the μ across space and time) and using the ρ to account for the PRTP is the preferred approach. If one uses an alternative approach to STPR, one should consider using a declining discount rate, e.g., as was proposed by the U.S. 2023 Circular A-4 update.
- I find the arguements for declining rates convincing
- According to the standard Weitzman/Gollier reasoning, the term structure might rise or fall with the time horizon. But even setting that aside, the mean growth rate is likely to vary over time. The growth rate in the UK has been declining, which according to the Ramsey rule should lead to a declining term structure.
- Long-term infrastructure projects would always suffer if subjected to a discount rate based on the greater certainty of short-term projects but may need an allowance for the risk that could lead to larger long-term benefits
- Should value the liability or benefit using the opportunity cost faced for that term
- I don't see a compelling case for the numbers in the Ramsey Equation to vary with the time horizon.
- Probably it should vary to account for the changes that take place in the economy and society
- As projects get very long, then probably need to adjust with time if there is non-trivial income growth. That said, in practice, the effort might not be worth it.
- it needs to have a term structure, to account for uncertainty growing with the time horizon
- distant future has higher downside risks and lower growth prospects
- Strong scientific argument in favor of a decreasing term-structure: Positive serial correlation in consumption growth
- One good reason is uncertainty (Gollier style). Another is growth forecasts depending on the time horizon (flat, but horizon dependent). But it is also not clear to me that the pure time preference should be a constant rate. I find exponential utility discounting a bit restrictive.
- As shown by the Ramsey formula, the discount rate is time-varying if the annual growth of real consumption per capita varies with time. Thus, if long-term projections of consumption growth show a decline, the discount rate should decline over time as well. However, for the case of the UK, a constant rate is a good approximation as rapid industrialization with high consumption growth rates that converge over time is a past stage of development.
- The SOC would naturally take into account the term structure of the relevant rates. Incorporating an underlying term structure would modestly improve on the current approach.
- This should be based on the latest and best evidence. *[Government]*
- It should decrease over time to 0. *[Government]*
- As the rate curve shows, it should increase with the maturity of the project because rolling over short term rates means greater availability of funds *[Government]*
- Instinctively, vary with time horizons as people view time differently depending on how far away it is. However, this should be consistent with robust analysis. *[Government]*

- There are pros and cons here I think. It seems that there is some justification for variation over longer term time horizons (pure rate of time preference, Stern review issues etc). If fixed as a constant, I would suggest making varying it a required sensitivity test, particularly if likely to change the ranking of different projects with different appraisal periods. Same goes for sticking with a discount rate that varies over different time horizons - sensitivity testing should be a requirement I think (but should be standard practice really). *[Government]*
- Agree with the idea that as uncertainty increases as we look further into the future we should have lower discount rare to ensure we don't over penalise future generations *[Government]*
- Complexity is unhelpful. *[Government]*
- A constant time horizon puts too much weight on generations in the far future *[Government]*
- It has proven necessary for climate change work in the past. These figures appear to be objective but they are not and if they do not reflect they clear needs of humanity they are wrong. *[Government]*
- - There is a question about whether the wealth effect term will be constant in the long run. For example should the elasticity of the MUC be constant as income rises and does the growth rate remain constant?

  - There is also a question about whether the pure rate of time preference will reduce in the future as people live longer lives and have to manage consumption over longer periods.

  These are empirical issues which need further examination.

  Absent of contrary evidence though, different age cohorts should experience the same discount rate as they age. As new age cohorts enter UK society the assumption should be that the discount rate they apply to costs and benefits follows the same pattern as earlier cohorts. *[Government]*"
- The Green Book explains that the discount rate is reduced because of increased uncertainty about the parameters over a longer time horizon. It's not completely clear how a lower discount rate (which is effectively assuming values for those parameters anyway) addresses that. In addition, having step changes in the discount rate doesn't feel like a very elegant solution, plus the nature of discounting means it is compounding over time anyway. *[Government]*

A.2.6. On whether the 1% value for L remains broadly appropriate:

- See Arrow-Lind we are a big country.
- I don't see why 'systemic' or 'catastrophic' risk counts against investment. For example, if there is a high catastrophic risk, we may want investment to reduce this risk. More generally, this is not an easy "premium" to argue about since it is not microfounded.
- Systematic risk on the market is priced with a 5% risk premium. If we include catastrophic risk - and perhaps other forms of risk too - then the risk premium should be at least 6%. The current value of 1% is ridiculously small compared to the size of risk it should compensate for.
- If a single risk premium is retained, it should be set below the current 1% to avoid over-penalising long-term public investments. A lower premium better reflects the government's ability to pool and absorb systemic risk, while reducing the risk of systematically discounting socially valuable projects.
- 1% is too high if applied uniformly. 0.7% is less prone to systematic underinvestment in long-horizon public goods.
- It is a pure token amount. We have no idea about the probability or impact of the effects of climate change, nuclear war pandemics. But we would not want the rate to differ too much from the borrowing rate.
- It is lower than the market-based risk premium of 6 or 7% since it is macro risk that matters.
- An overall rate of 3% would allow more positive NPVs for public investments but guard against politicians topping-up with ancillary reasons to get vanity and popularity projects 'over the CBA' line.
- Keeping the single premium at 1% is not unreasonable but some decomposition of the sources of risk would be beneficial and provide greater insight. It would also allow for this to vary by project as the benefits and risks of different projects are materially different. However, this comment needs to also be linked/moderated by the previous observation of the persistent lack of consumption growth in the the UK economy that has been observed over the past 15+ years.
- Here is the conundrum, the long term risk has gone up (climate change) however to use a higher risk premium on long term climate projects means that FEWER climate mitigation strategies will be implemented, thus RAISING the risk. Different discount rates should be applied to projects that increase (high RP)/decrease (low or -ve RP) e.g. long term climate risk. These are the most important of public projects.
- I do not know why this figure is required along with the pure rate of time preference

- I'm assuming "systemic" and "catastrophic" risks exclude gross failures to plan and cost projects prior to approval. This is a heroic assumption but reflecting such risks in the discount rate fails to shine a light on the fundamental source of the problem.
- The Post-Dollar world environment and transition towards a new World Reserve Currency (probably a combination of *USD*,CAN, PRC Wan, possibly also some small stablecoin component introduces new systemic risks, which must be accomodated.
- An extremely rough and ready estimate of the market equity premium, also broadly consistent with available estimates of the return on housing
- Do we really think there's a 1% annual risk of annihilation of the UK economy? I don't.
- 1% seems to be a minimum. A much higher value may reflect the current (ST) reality better, but will also result in unintended consequences of project selections
- Within an STPR framework, the theoretically appropriate risk adjustment depends on covariance with aggregate consumption and is likely to be small for many public projects, which plausibly have low or at most modestly positive consumption betas. In practice, however, this may underestimate the appropriate risk premium, given uncertainty about tail risks and the well-known difficulties of reconciling observed asset returns with simple consumption-based models. I can therefore see a case for a pragmatic, ad hoc wedge on the risk premium, but it is difficult to justify a value as high as 1% if other sources of bias, such as optimism bias, are addressed elsewhere in the appraisal process. Rather than applying a single uniform risk premium to all public projects, HM Treasury could specify a small number of standard consumption betas by project type (e.g. resilience, state-independent public goods, cyclical infrastructure, demand-driven regeneration). Risk adjustments would then be derived from these betas and an STPR-consistent consumption risk premium, better reflecting systematic risk differences across projects while remaining operationally simple.
- Best estimate 0.5% for systemic risk and 0.1% for catastrophic risk
- It is modest enough to defend but it still matters.
- The effects of systemic and catastrophic risks are surely better measured by adjusting the numerator than the denominator, and by undertaking scenario or simulation analysis of the project under the different potential cash flow streams arising under catastrophe scenarios.
- Risk does not accumulate at a constant rate over time, which as a result penalises more certain benefits disproportionately. I don't think the discount rate is the correct vehicle for accommodating risk in appraisal. I would rather work with risk premiums (or apply the Arrow-Lind theorem if these risks can be diversified).
- I've left this blank because I don't feel I could give a response that was anything other than arbitrary.
- I apply an elasticity of 1 and a standard deviation of consumption of 10%. I double it to account for catastrophic risks.
- Considering the huge implications of an 1% increase due to risk, this value should be lower
- My knowledge on this aspect is not sufficient.
- We should ensure that the catastrophic risks of climate and biodiversity loss in this government assessment are taken account of: https://www.gov.uk/government/publications/nature-security-assessment-on-global-biodiversity-loss-ecosystem-collapse-and-national-security
- You mean of the SDR or the project costs/benefits? Of the latter 10%, of the DR shoudl be lower I think, by 1%
- Many public projects might mitigate "catastrophic" risk (see "Tail-Hedge Discounting and the Social Cost of Carbon" by Weitzman, 2013. DOI: 10.1257/jel.51.3.873).
- If we consider catastrophes to be more likely in the far future than in the near future then applying a premium may be useful conceptually but I don't think a one size fits all approach is going to improve welfare
- I don't know I am afraid
- 1% seems indefensible given academic evidence and CCAPM model?
- 1.5%=averaged over 1% for systemic and 2% for catastrophic risk, rexpectively.
- From a subjective perspective only, lower than 1% would seem to be unduly optimistic and above 1.0% erring on the pessimistic side. From a technical perspective, this seems like an element that requires deeper attention in research.
- There is no single risk premium for all public projects. To use a single rate despite varying risk is madness.
- Post COVID, and due to climate tipping points, rising geopolitical instability, emerging AI risks etc. perhaps suggest that the annual probability of civilisation-scale disruption is somewhat higher than the original 1% implied. However, 1.5% remains conservative enough to avoid overstating catastrophic risk in routine public project appraisal.

- In the 2024 New Zealand update we use a triangular distribution over the range 0% to 0.2%, with a mode of 0.05%, resulting in a mean estimate of 0.08%. Like Grimes (2023) we do not consider values of 1% per annum for cessation of all social welfare reasonable as it would imply too high a chance of extinction each 100 years.
- Catastrophic risk is perhaps expected to be 0.1%. Systemic risk is more project specific; some projects reduce risks whereas others increase it (eg, see Boardman et al on CBA)
- A risk premium should be added to reflect undiversifiable risk. 1% seems reasonable.
- The 2023 update to U.S. Circular A-4 estimated a risk premium of 2.5% and a correlation factor of 0.45. Many public projects and regulations can reduce systematic risk (e.g., investments in pandemic preparedness), so a generically high risk adjustment is inappropriate.
- Determining the average systematic risk of a public project is a borderline impossible task. Some projects, like investments in healthcare, may well have negative systemic risk. Others, like investments in commercial infrastructure, may well have a highly positive beta.

  In the absence of a clear way to determine this metric, I would defer to the rough correlation between utility equities and the stock market, multiplied by the leverage-adjusted equity risk premium, which is roughly around 1%. But I do not find this approach theoretically or practically advisable.
- A value around 1 seems appropriate based on the correlation between market risk and benefits from government projects.
- Based on adjusted rate in the 2023 Circular A-4, 1.1% was appropriate. They uses a market risk premium of 2.5%, which is low for private returns, but maybe right given the externalities and market power in the real world. They then adjust for leverage.
- I think this is defensible. Especially if you combine with a PRTP of 0
- I am unsure, but some arguments around inter-generational comparability have in the past made me question if this should approx. 0
- As a maximum. Some projects need to be allowed to take a riskier approach if potential benefits are to be realised.
- I can't see how a single risk premium can reflect both systemic and catastrophic risk. If a risk preium is included it should be strictly catastrophic for all projects and the economy with no possible recovery. So very small but higher than it was and the pa catastrophic risk maybe increasing over time (my presimistic view). Need to know annual catastrophic probability and the trend (up in my view) to get a term specific catastrophic risk premiums. Personally I would prefer to do this in the expected value of costs and benefits in each period over the time stream (including aspects that look like risk aversion). Then discount at the risk free rate for that term.
- I don't feel certain about this. I think a strand of your report should consider whether long term geopolitical risk (war, disease, climate change, financial market related) have changed.
- This would require specific research
- I'm not really sure. I think the risk premium should adjust to the project. The thing is many public projects are hedges against systematic risk in the macroeconomy. This makes it tricky. For things like standard capital investment, the 1% might ok (seems low). For things like health care or climate adaptation, these are hedges so they actual work to derisk the UK economy.
- the current 1% is too small to account for uncertainty. It should even be 3% for long term discounting
- Not the right way of treating risk - apply probabilities to possible outcomes, and then apply an optimism bias based on past project performances
- CCAPM: 2% observed aggregate risk premium
- I find 1% as an average a bit low. In the short run, I think there is good reason to think that the systematic part is higher. It is not clear to me how it should depend on the time horizon.
- single risk premium doesn't make sense. DIffernet projects have different risks. Use CAPM approach
- I have to admit that in answering this, I am referring to the practise of the Green Book. I find this question extremely hard to answer as it is hard to judge the average risk behind the plethora of projects that are undertaken by the UK government and how it correlates with systemic risks. If the choice of the risk premium relies on the fact that the social time preference rate without a risk-adjustment is too low so that the government takes on too risky projects in comparison to the return from long-term government bonds, it makes sense to choose 1 to move closer to the rate of return on the bonds, while maintaining the normative judgement of rather low pure time preference resulting in a lower STPR even with risk-adjustment.

- 1% is much lower than what is typically assumed to be the market risk premium. It biases decisions in favor of taking actions that do not generate value.
- This should be updated to reflect the most recent and most robust evidence. *[Government]*
- I think risk better sits elsewhere in the business case. Sensitivity tests can best assess risk. It should not be something that discounting accounts for. In some investments such as addressing climate change, you could make the argument that this risk premium should be negative, as acting now will reduce risk in the future. *[Government]*
- I don't think that a single value is good enough and the Green Book should include a table where it changes depending on size, risk and maturity of projects. Risk aversion increases as the sums involved become larger, more uncertain and for longer periods. *[Government]*
- I have stuck with the existing premium as I don't know the evidence base well enough to argue otherwise. If it worked before, and there's no compelling evidence for changing it, then it seems like it should be kept as is to me. I haven't seen anything to date that suggests that it should be changed (not saying that doesn't exist however). Once set, there doesn't seem to be a strong argument for changing it over time - the types of risks it addresses have and always will continue to exist, with no obvious evidence (at least from my standpoint) coming to light to suggest that they are more or less likely to materialise now than they have ever done. *[Government]*
- If done properly, I would argue that risk should be captured in the appraisal through risk analysis, scenario testing etc and due to the fact this effectively discounts long-term benefits (which occur more regularly that long term costs) on the basis they might not materialise (e.g., due to war or systemic shock), even though those are precisely the states of the world where welfare would be lowest and where resilience-building investments may matter most. *[Government]*
- Consistency appears appropriate. *[Government]*
- Inertia bias - I am aware of arguments for increasing and decreasing it in different circumstances. How do we reasonably weight those arguments? *[Government]*
- I am just guessing. The world is a bit riskier than it was. I don't want to put it up much though as I am cautious not to over react. *[Government]*
- At the moment what is contained in the risk premium is sketchy. Nor is the evidence behind the 1% estimate clearly set out.

  The discount rate should include something on catastrophic risk, as that's the part of risk that is cross-economy and essentially a temporal matter, but not include anything more / different on "systemic" risk. Other risks should be dealt with outside of the discount rate and in a consistent, but project by project way. *[Government]*

A.2.7. On whether the risk premium should vary depending on the risk of the project:

- Generally no, as any project will fall into the portfolio of projects and over a longer term the risk premium for this portfolio should be zero.
- I am not convinced by the need for this premium in the first place.
- The level of exposure to systematic risk of each project is different. Therefore, HM Treasury should first quantify the exposure to systematic risk of each project and then apply a risk premium proportional to the exposure to systematic risk of each project. The overall risk premium will be the same, but the level of exposure for each project will vary.
- Yes. Project-specific risk premia better reflect heterogeneity in uncertainty, exposure to systemic risk, and the degree of controllability across public projects. This approach improves appraisal accuracy and avoids inefficient one-size-fits-all adjustments.
- The present system does not really allow for the uncertainty about costs. This is a second best way of doing it. Could be done using the red, amber green rating of major projects.
- Some projects have a return that vary positively and other negatively with the business cycle, and this should be reflected in the risk premium.
- The costs and benefits of some projects are incredibly different when thinking about issues such as climate resilience. The societal benefit of climate resilience e.g., flood defences is very different to other projects e.g., high-speed rail. By having some choice in the discount rate and linking it more directly to the nature of the project, then there will be a stronger rationale for undertaking certain investments, which will have greater societal benefit.
- see the last question answer. The incremental "projects" that are under consideration are not independent of the (climate) risk outcome. Some will raise but many can lower the climate risk.
- While I can see value in this approach, it also introduces scope for bias (conscious or otherwise) and possibly political interference.

- Prone to subjective interpretation and politicisation.
- beta risk is clearly project specific
- This is the systematic, catastrophic risk again? I just don't really get its salience.
- Rather than applying a single uniform risk premium to all public projects, HM Treasury could specify a small number of standard consumption betas by project type (e.g. resilience, state-independent public goods, cyclical infrastructure, demand-driven regeneration). Risk adjustments would then be derived from these betas and an STPR-consistent consumption risk premium, better reflecting systematic risk differences across projects while remaining operationally simple.
- Apply to the costs or benefits directly. Both are unlikely to equally uncertain.
- Of course in reality it will vary, but this seems hard to quantify and hard / complicated to explain, including perhaps conveying too much certainty or precision than is warranted
- 0% risk for standard infrastructure projects and 0.5% - 1% for projects dealing with climate change mitigation, pandemic preparedness, etc.
- Variable discount rates across different government departments could create additional political debate that is probably best avoided. As previously mentioned, the variability across government departments is better finessed by rigorous modelling of the numerator (the cash flow stream) scenarios.
- See previous answer; what is the risk premium of a given project?
- I think this sounds good in principle but in practice it'd add a layer of bureaucracy to the business case process and would probably just result in 'gaming' the assumptions around what type of project it is. I also fear that it'd incentivise away from riskier projects with high expected NPV, such as nature-based solutions, in favour of more reliable 'grey' solutions which may, on average, result in lower benefits.
- Modern theory of finance
- For some projects the risk premium is definitely sensible. However, this is out of my expertises, therefore I can only advise to be cautious with extra increases in discount rates
- Risks differ in their variance and (expected) magnitude across domains so I would think so.
- Many public projects might mitigate "catastrophic" risk (see "Tail-Hedge Discounting and the Social Cost of Carbon" by Weitzman, 2013. DOI: 10.1257/jel.51.3.873).
- Projects differ in time and risk dimensions and this should be somehow accounted for
- public projects with non-financial returns should not be penalised because returns are estimated in biophysical or non-market economic (monetary) terms. The risks associated with, say, nature based solutions, compared to engineering solutiosn with supposedly more control should also not be penalised because of the perception of higher risks of the former. The same can be said in other public projects too — preventive health measures like green and social prescribing vs drug treatment for example.
- In theory, yes, but only if it will make enough difference to be worth the hassle and hours of pondering every time we do a project
- Because it would create a very expensive assessment system, with potential political incursions.
- Mega-projects like HS2, Trident and Hinkley C are large enough to have wider macro-fiscal consequences and this should be reflected in a higher risk premium.
- You can observe private investor behaviour. They do not use a single rate for all projects. Why would the government?
- Varying risk premiums makes it harder to compare alternative projects. The risks are generally better captured in the cash flow streams for the benefits and costs, and dealt with in sensitivity analysis. In the 2024 New Zealand update, the mandatory sensitivity analysis for commercial and non-commercial projects, could be an option.
- This might be too difficult and provide spurious accuracy. If a STPR is used, the consumption CAPM model would result in a very low risk premium, especially if a value of less than 1 is assumed for the elasticity of marginal valuation wrt consumption. That's assuming people can even identify the correlations of a project on national consumption per annum, which would be a small miracle. Most people struggle with basic CBA, let alone doing this sort of thing
- A risk premium based on the consumption capital asset price model should be added (e.g. a project with returns that vary inversely with the state of the economy - such as a mental health facility - should have a negative risk premium added to the 'standard' discount rate).
- Many regulations and public investments decrease systematic risk (such as investments in pandemic preparedness), in which case a positive risk premium is not appropriate.

- As would be obvious from some of my previous answers, I do not fully agree with Gollier's views on discounting. But I think he is essentially correct in Gollier (2026) that the welfare cost of ignoring the beta can be quite high: https://www.journals.uchicago.edu/doi/abs/10.1086/733779?journalCode=jpemi.

  The difficult part is accurately assessing a project's systemic risk. But better to state the aspiration and fall short than to entirely ignore the problem.
- In principle, yes. In practice, these are very hard to estimate for most settings. Balancing tractability and granularity, perhaps there could be a separate risk premium for things like climate policy where there is substantial research on the appropriate adjustment.
- Yes from the perspective of economic theory (i.e., the extended Ramsey Rate), though I get a bit confused if this is actually practical for government agencies in the real world. Moreover, I have concerns when these project specific risks are dependent on structural models, such as a climate beta, as the structure itself determines the beta. If anything is missing, then the beta can be wrong.
- I think this would complicate economic analysis, and may produce perverse incentives across different parts of the government with different interests.
- This could substantially undermine comparability, and project specific risk should be dealt with elsewhere in the appraisal
- As suggested in the previous answer. There are riskier projects that may nevertheless realise significant benefits that could be penalised by a blanket view of risk
- Personally I would want to value project risk outside discount rate. However if systemic risk is included then it will differ by project based on pro or counter cyclical risks
- I believe that the Govt should be somewhat risk averse and that therefore it should (a) represent project values as distributions and (b) should then report explicit certainty equivalents as the summary BCR values (or whatever). Actually I think this rather than the discount rate itself is the most significant area of weakness in appraisal practice.
- Yes because all projects have different risks, so called project-specific risk, and of different types (demand, supply, technical, etc)
- see prior answer
- Some projects provide an insurance value that should be accounted for through a lower discount rate of their net benefits. The reverse applies for those projects that enhance risk in the economy.
- as explained in previous question, but appraisal optimism is important bias to correct
- Absolutely! CCAPM. Survey Gollier, van der Ploeg, Zheng, JEEM
- Yes. But it is not clear to me a priori what the term structure should be (increasing or decreasing).
- That's what the CAPM / finance theory tells us should happen. Look at Biden administration proposed Ciruclar A-4 in the US
- It will be extremely valuable to assess the risks behind each project individually by estimating the beta of the project: how do the benefits of the project correlate with consumption and thus affect the overall discounting practise. In doing so, the government can argue more effectively to the wider public why some projects were taken on while others were not.
- Doing so would be consistent with the fundamental principles of financial economics and standard practice outside of the government. For the government to deviate from the private sector in its valuation practices, there should be a compelling reason. Using a standard approach increases the transparency and auditability of analyses.
- This is not for the discount rate to do. It could be that certain benefits and values could be discounted differently, like they already are for health, but this should not happen at the project level. This will encounter definitional issues about what different projects are. *[Government]*
- Could do, if as before, you are looking at investments that address climate change. I would prefer however, that the risk premium is simply removed and address elsewhere in the business case, through sensitivity tests. *[Government]*
- As per my previous comment, projects differ in size, risk and maturity. We can't assume that, compared with the size of the economy, a project a few thousand quid where government should be close to risk neutrality, is the same as preventing a disaster at a nuclear power plant. *[Government]*
- The current approach does not make sense to me. Understanding what is meant by these risks would be much more appropriate to enable it to be varied for projects. Understanding what these risks looks like for physical infrastructure projects feels very different to RDEL on youth clubs. Is this just GDP based or the metaphorical bridge falling down? *[Government]*

- I answer no as it would make cross project comparisons more difficult (unless I have misunderstood this point). I might be conflating issues a bit here too, but to my mind, this point is largely addressed via optimism bias allowances for different projects (so doesn't need to be controlled for again via the discount rate). *[Government]*
- The Green book discount rate should not in principle be applied to high-risk investments where we are looking to establish something more than just the time value of money between current and future periods. *[Government]*
- Risk will vary substantially across projects and regulations so it is much better to take this into account explicitly in the appraisal of different risk scenarios and how they affect outcomes which can be much more transparently set out for decision-makers *[Government]*
- because risk is relative *[Government]*
- Not in the discount rate (see answer to question 6).

  - But the project risk premium should be considered explicitly as part of the appraisal of benefits and costs and this should form a core part of the Green Book.

  - Project risk will vary from project to project. Identification of risks to benefits and costs and potential causes allows them to be explicitly managed as far as possible and handled properly using the avoid, reduce, transfer, share, accept framework. *[Government]*
- This could create confusion with respect to risk costs, which projects should be applying already. It is also arguably simpler for a decision-maker to see risk costs split out rather than risk being factored in via the discount rate. On a practical basis, including a project-specific element of risk in the discount rate would require very clear guidance to ensure it's applied consistently and accurately. *[Government]*

A.2.8. On how the discount rate should vary with both maturity and project risk:

- I know this is probably not a popular answer, but if the project is a "public project" then we need to treat it differently from the risk profile of private projects. Hence I don't know what a low, medium or high risk means in this case. If we are talking about investments in pools of liquid assets, then clearly this is different, but for infrastructure etc. the current choice of rates is vastly too high in sterling terms (clearly in dollar terms it will vary with the exchange rate).
- Risk free projects can be assumed to have no term structure. The 3% rate would be the weighted average of the Yields to Maturity of public debt, weighted by the pound amount of outstanding debt. There is a median 6% risk premium, in line with my answer to the previous question about risk premia. Risk premia increase along the risk dimension and along the maturity dimension. The 15% rate for high risk + 100 year reflects greater exposure to tail/catastrophic risk that may arise in the very long term.
- I apply a declining risk-free term structure to reflect greater long-run uncertainty and to avoid systematically under-valuing very long-lived public benefits. I then add simple, transparent risk premia by project risk quartiles, increasing the discount rate monotonically with risk while keeping increments modest to reduce double counting with other uncertainty adjustments.
- done relative to current 10 year gilt yield, so adjusted as that changes. Given the level of systemic uncertainty there should be a bias to shorter term projects.
- Low and high risk are not defined. If we know the beta, I would add beta times the risk premium.
- I have anchored the answers above on the Ramsey-Rule. The incremental increases in risk are arbitrary going from risk free to high-risk/upper quartile across each maturity. This is in part due to the subjectivity of what risk means in this setting, and I am of the view that risk in public projects is materially different from risk in markets. As such, the idea of the highest risk category being equivalent to something derived from the asset pricing literature e.g., the CAPM does not capture the correct meaning of risk. As well as this, the misapplication of discount rates on projects between the public and private sector can create a "wedge" that overcompensates the private sector for project delivery relative to the risk of the project. The other aspect is the use of declining discount rates moving from 10 to 100 year maturities. Again, the value chosen is arbitrary rather than based on some rigorous analysis, however, it is there to reflect the need for declining discount rates. The fact that the 10-year risk-free and 100-year high risk is purely coincidental.
- I would not have split them by maturity, I would have split them by their carbon impact in line with COP25 Paris alignment goals.
- I think that risk should be handled in a different way or ignored if the project is small
- Ultimately, I would be keen for these investments to occur, hence low rates in relative terms. This is especially true for low risk-short maturity projects
- Roughly geometric progression yield curve. Note: global tendencies has been towards shorter durations: you are missing a 5 year rate here.

- I adopt an STPR framework with a pure rate of time preference of 0.5% and an elasticity of marginal utility of 1. Short-run real consumption per capita growth is assumed to be around 0.5%, rising to around 0.75% in the long run, with the certainty-equivalent risk-free rate declining with the time horizon to reflect long-run uncertainty. Risk adjustments are treated separately and based on systematic risk through covariance with aggregate consumption; for many public projects these are small, but I allow for pragmatic upward adjustments to reflect tail risks and model uncertainty, rising to around 0.5% for high-risk projects rather than applying a single uniform premium.

- Not sure I fully understand how project risk relates to the earlier discussion of systemic and catastrophic risk, but I wouldn't add very much for risk - about 0.5%, and as noted previously I wouldn't try to distinguish different levels of risk too finely

- I don't have strong views on this question without more information as to what the risks are for a public project. Are these the risks that it may fail to meet its intended objectives, or risks of cost overrun, or risks of abandonment? Broadly, my view is that the discount rates used should be high enough to discourage wasteful expenditure, such that the first decision that seems to be suggesting itself for every public project is that it is not worth pursuing. This ensures greater scrutiny for every project and more efficient decision making. I have also suggested discount rates that rise as a function of both risk and maturity. While I stick by the former as valid relationship, the latter is more nuanced at such long horizons, where any component of the discount rates is just not observable. My reason for the positive relation to maturity is simply that ultra-long term projects run the risk of under-estimating the likelihood of obsolescence, and some premium for this is essential. In other words, whatever discount rate is used, it should not ignore the supply side of the economy.

- I don't have clear recommendations for this, apart that the same discount rate should apply across the levels of risk, but that it should be declining over time.

- I believe: -wealth effect should be set to 0 to acknowledge the upward bias in GDP due to its not accounting for environmental degradation, plus the irreversible future risks of climate and nature breakdown. -pure time preference = 0 after 20 years to address that discounting within one generation is not the same as across multiple generations -other parameters as they currently are -I don't personally believe that the DR is where project risk should be addressed. This is the remit of sensitivity analysis

- My knowledge on this aspect is not sufficient for a reasonable judgement.

- Using the Ramsey Rule, the current time preference of $\delta = 0.5$ is relevant for short-term discounting (e.g., 10-year maturity), but should be omitted ($\delta = 0$) for long-term discounting (e.g., projects with more than 30-year maturity).

  Assuming that total risk premiums are uniformly distributed from 0 to 1%, I infer risk premiums for low, median, and high risk projects to be 0.1 (0.125% on average for the lower quartile), 0.5, and 0.9 (0.875% to be precise) respectively.

- Hard to answer as can't convey a DDR in the response

- The real risk-free rate on UK inflation linked bonds is observable up to 30-years. This gives at end of January 2026 1.4% (approx.) for 10-year, 2.1% (approx.) for 30-year. The UK yield curve (at least my quick googling) seems to be downwards slopping towards higher maturities, so 50-year probably something like 1.7% (you should of course do this properly). 100-year is more difficult but sure you can think of something.

  For riskiness you should just use CAPM (as an improvement on current practice). I don't know the spread of projects that the government funds which means giving quartiles as difficult. But lets assume low is a beta of 0.2, median 0.5 and high 1. The numbers given are very approximate and have not been developed precisely.

- Climate change increases long run growth uncertainty and catastrophic risk, which strengthens the case for a declining discount rate over longer horizons. But I don't think all projects should automatically face higher risk premia. For example, climate adaptation infrastructure will deliver their highest benefits in bad states of the world (i.e., when climate damages are severe).

- The suggested values reflect where we landed in the 2024 New Zealand update. These might be relevant for UK considerations. A published working paper provides further information.

- First, I think a DDR should apply, but I cannot specify that, so what I've specified is the average constant discount rate that would apply to projects of given maturity. I read this from the black line in Figure 6 of Parker (2025) www.treasury.govt.nz/sites/default/files/2025-02/twp25-01.pdf

  Second, I concur with Boardman et al's CBA textbook that risk adjustments should happen in the cashflow profile, and not with the STPR. Risk should be measured as the covariance with the baseline measure of wellbeing; a risk-reducing project (such as a damn) would have a negative covariance (ie, providing benefits when consumption is low), but such a project does not conform to the ""Low, Median, High"" classification here. What I presume is meant here is whether it is likely or unlikely to realise benefits, which is not how risk is meant to be conceptualised in this context.

- The question does not explicitly allow for a declining discount rate; my suggested rates are for discount rates beyond the previous term (e.g. 2.0% for years 11-50 and 1% for years 51 to 100). My lower quartile risk premium reflects the assumption that many public projects have payoffs that are inversely related to the state of the economy so should have a negative risk premium. High risk projects should probably best be left to the private sector so I have specified a high risk premium for those.
- 40-year real gilts are currently less than 0.5%, so that pushes me towards a similar value for the 50-year maturity value (which is not traded). I use a simple model of declining discount rates to approximate the 100 year value. Because of a lack of awareness of data on systematic risk over time, I assume it is time invariant. I roughly assume that lower quartile projects have a beta of zero and upper quartile projects have a beta three times the median (reflecting a skew, as upper quartile projects should get closer to more commonly measured equity risk premia).
- I am basically basing the 10-year maturity risk-free rate on government borrowing costs, then using a simple declining rate that matches estimates of longer-horizon rates at 100 years. More sophisticated analysis could certainly be done. The median risk adjustment is based on my prior answer. The lower and higher risk adjustments are completely ad hoc.
- I cannot do low and high, as I am not really sure what I am doing here to be honest. I feel more comfortable leaving them empty.
- See past answers
- I find this difficult to answer without much more detailed analysis
- Risk free rate declines due to my pessimistic view of future growth. However risks accumulate with term (also my pessimistic view) which tends to off sets this. The rate for 100 years is probably much higher. My prior is 50/50 we are still around. But as before I think it's better to do this when estimating the EV and CE of a project at each point in the time stream.
- OK I might have misunderstood the previous question. I think project risk should be handled by risk adjusting the project benefit and cost flows, NOT the discount rate itself.
- I will look into this more carefully later, and provide my opinion during the roundtable
- This is hard for some many reasons and have near zero confidence in my answer. I don't know the portfolio of projects and the correlations well enough to be helpful here. The key is that lower bound must be negative.
- accounting for the risk if the project implies to also know their beta to derive a meaningful discount rate
- as explained I do not think risk should be included in discount rate, so the 10 yr projects are more like SOC
- Decreasing risk-free rate; increasing risk premium
- I feel unable to quantify all of these numbers on spot. But with an agreed upon method it should be possible. The risk premium should be project specific and time-dependent. "
- I assume I am assigning the risk-premium coming from systemic risk that is not already assessed on a project-basis. If not, then the projects should be assigned individual risk-premia according to their specific riskiness. I am starting from delta=0.5, eta=1 and g=2 percent for the risk-free rate for 10-year project (to have the current practise of the UK reflecting their preferences), and then let g decline towards 1.5 percent over the next 100 years. All projects receive a risk-premium of 1 percent for 10 years for systemic risks as argued before. The risk-premium is not adjusted for project-specific risk quartile because for evaluating systemic risk, the beta is important, so how are benefits from the project correlated with overall consumption growth. I personally cannot infer the beta from the riskiness of the project. This should be estimated on a project-specific basis. For 100 years, I also keep the risk-premium constant. This assessment is based on two countervailing effects. First, based on Arrow et al (2014), long-term risks that are systemic or catastrophic render public benefits from a project more valuable if consumption shocks are positively correlated (as has been observed). This would lead to a declining discount rate due to risk. However, one may also argue that in how far the benefits we estimate today are still beneficial in the future, if future technology renders them obsolete for example, a declining discount rate is not justified but the risk-premium should increase over time. The risk of unimaginable technological breakthroughs is hard to model and could therefore justify an increasing discount rate over time. Thus, I stick to a constant risk-premium over time and overall declining rate due to declining consumption growth. "
- I would recommend that the government entity overseeing the choice of rates should update these numbers annually for agencies to use for evaluating the cost of routine projects, The mandated rates should be updated to reflect current market conditions and fair value principles. For larger investments, a more careful choice of discount rates or other valuation approaches based on fair value principals should be made by valuation specialists in the government.
- Unable to provide. This should be ased on the most recent an robust empirical evidence we have. *[Government]*

- I don't have a clear view or evidence base for each option. I would split by type of risk rather than viewing projects as high or low risk. *[Government]*
- I wouldn't take these values on board as I might not have understood the question. I don't think the discount rate should vary according to project risks, so I have held them constant across the different types of project. I've dropped the discount rate for longer term projects as per the standard guidance, although not sure of the required schedule of reductions without checking. *[Government]*
- I don't think this is the right question. For risky projects, there should be a need to do a WACC based on comparable private sector projects. *[Government]*
- They shouldn't vary - better to be explicit in the analysis about risk. There are no risk free projects! *[Government]*
- The discount rate should include something on catastrophic risk, as that's the part of risk that is cross-economy and essentially a temporal matter, but not include anything more / different on "systemic" risk. Other risks should be dealt with outside of the discount rate and in a consistent, but project by project way.

  The exact value of catastrophic risk to include is an empirical question which needs to be explored scientifically. *[Government]*

A.2.9. On the social discount rate for public-private partnerships

- PPP is a mechanism of providing almost riskless returns at higher than riskless rates. In fact the problem here is the distortion, we need to set more aggressive public rates so that the value of private sector risk premia rise and the value of an investor choosing a portfolio of public and private capital for investment is less correlated. At the moment providers of capital get a very good deal from the UK government.
- The taxpayer does not care about rates that come from macroeconomic theory models. HM Treasury must show that these project create value after accounting for the opportunity cost of capital, and the clearest opportunity cost is simply to cut public debt.
- A central issue is the distinction between social discount rates used for welfare appraisal and private discount rates embedded in financial models, as these reflect fundamentally different objectives. In PPPs, there is a risk of inconsistency or double counting if private financing costs or risk premia are inappropriately incorporated into social discounting. Careful separation of risk allocation, transfer pricing, and financing structure is therefore essential, alongside transparency about which risks are borne by the public sector and which are genuinely transferred to private partners.
- Government can borrow more cheaply than the private sector. Treatment of risk is difficult. If the project is completed the SPVs can make a large profit in refinancing. But the probability of success may be larger. Government gets locked into an asset it may not need. Hospital in the wrong place.
- History suggests governments are poor at getting good PPF VFM
- As alluded to in Q9, the use of different discount rates between the public and private sector can create a "wedge" that overcompensates the private sector for the risk of a project, which is a cost ultimately borne by the taxpayer.
- The extent to which public projects increase or decrease global risks, climate and others.
- Potential volatility of market-based rates
- Don't know
- Recognition that many of the (most important) benefits may not be easily quantifiable in financial terms, meaning the hurdle rate should not be set too high (i.e., biased towards the rate that private finance may prefer).
- None. Extreme risk of subjective social objective being given false legititimacy by being quantified.
- Many! I would emphasise in particular the need for consistency with the way expected returns and hence discount rates feed into investment decisions in regulated utilities - all potentially part of the public sector. How would investment decisions change if they were renationalised? Does this make sense?
- I have not thought about this issue. However, I would think that appraisal should separate the valuation of social costs and benefits (using the social discount rate) from the assessment of financing costs, value for money, and risk transfer under private finance. I expect care is needed to avoid double-counting risk where private-sector pricing already reflects risks that are separately adjusted for in social CBA.
- What is the true social cost of private capital/finance? It may be higher than we think.
- The same issues as mentioned already. Ensuring robust scenario and simulation analysis is undertaken. Ensuring that contract restrictions and tenure are correctly valued. Ensuring that obsolescence is correctly valued. Ensuring that the consequences of delivering private monopolies, or near monopoly, or more centralized services, as seem to result from PPP, are fully represented in the calculations.

- Private finance models come to life in a different part of the appraisal methodology. In SCBA the question is about what a project is worth to society, not if the project can pay its loans back; or is a money making machine.
- No view-not an area I've worked closely on.
- Does not reflect our social values
- That relative price changes in non market goods might not be considered or underestimated
- Joint benefit private public mixed public goods
- My knowledge on this aspect is not sufficient for a reasonable judgement.
- To ensure comparability of appraisals, welfare-consistent social discounting can be maintained by treating private returns, financing costs, and risk transfer explicitly (e.g., expected values and probability-weighted costs), rather than embedding them implicitly in higher discount rates.
- It could be a deterrent to private sector involvement if they can get a better turn from a private-private partnership. But wouldn't the type and scale of return itself be a more important factor than the discount rate? I have no evidence either way. Just asking the question.
- How and whether it is valid to discount a flow of finance payments and call this part of social value. How to calculate the appropriate risk premium for the STPR.
- Creating (and updating) a basket including both public and private investment options.
- A public-private partnership is an implementation modality for a public project and, as such, should be assessed on the same terms as a public project implemented using a traditional procurement modality. The private partner will apply a different approach to assess whether participation is financially profitable and the public partner should also look at the project from the perspective of the private partner to enrich its negotiating position.
- The government will appraise these projects using the incorrect rate. The private party will be doing something more similar to the opportunity cost of capital approach.
- This is a complex issue, which might bring us back to what kind of discount rate we decide to use as discussed in question 1. The main consideration to account for is that the public sector and the private sector have different risk profiles when deciding to invest. It is likely that an individual company will not be able to diversify risk as effectively as the government. Therefore, the risk of failure for an individual company deciding to invest in the same project as the government will be much higher, which will impact the private sector's ability to commit to investing in risky projects. This is likely to impact in particular transformational projects for which benefits are uncertain, unknown, occur over a long time horizon and additional cost may arise unexpectedly.

  However, parameters in the standard Ramsey rule do not allow for systematic and project-specific risks to be accounted for in the discounting calculations. The only kind of risk included in the standard Ramsey rule is systemic risk. Systematic risks are usually assumed to be negligible, and project-specific risks are to be considered in the main cost benefit analysis. This is usually done by calculating an uncertainty-equivalent value of risks, and incorporating that value in the financial contingency of the project.

  Overall, if we believe that STPR is the best discount rate approach, serious consideration needs to be given to the extent to which appraisal practitioners effectively account for risks quantitatively in other parts of the cost benefit analysis. If this is not done appropriately for public-private partnerships, there will be serious risks not only for the private sector parties involved, but also for the public sector, which might lead to higher overall costs and lower benefits. To better assess whether this is currently an issue in government appraisal, we would encourage the reviewers to investigate how many HMT business cases accurately quantify uncertainty-equivalent values of risks, and include those in the contingency calculations. This could be done by requesting access to a sample of business cases across different policy domains and budget envelopes, focusing in particular on reviewing the economic and financial case components.

  Alternative approaches to account for private sector risks could be to modify the Ramsey equation, or to use a different type of discount rate that incorporates additional measures of risk. An alternative discount rate that is often mentioned in the literature for these purposes is WACC. However, this would raise a number of additional questions, such as what kind of discount rate to use if benefits occur to society as a whole, or whether different components of the cost benefit analysis should be discounted differently (a more comprehensive discussion can be found here: https://www.ofgem.gov.uk/sites/default/files/docs/2011/10/discounting-for-cost-benefit-analysis-involving-private-investment-but-public-benefit.pdf).
- The main conceptual problem is that social discounting reflects society's preferences, whilst private finance discounting reflects commercial return requirements, and these two frameworks sit uneasily together.
- A practical consideration is the consistency of discount rates used across various options, for comparisons. Sensitivity analysis can be useful, where these rates differ.

- First, set aside the procurement choice of government (which is a red herring), and consider more fundamentally is the project a market activity (championed and enabled primarily by the market) or a government activity (championed and enabled primarily by the govt). If it is primarily a market activity, then it should primarily be judged on a financial appraisal, which is very different and out of scope of this assessment. Many public initiatives can be commercially provided with funding from user revenues and market-priced debt (and may or may not require a return on equity capital), provided there is no step-up by parent governments (eg, ""revenue bonds""). A financial appraisal rather than social CBA suffices here, and therefore no need to comply with social discount rate policy.

  If it is primarily a government initiative (which means it should be mitigating market failure and/or in service of govt's redistribution or stabilisation/security functions), then it should be appraised with an STPR on the consumption equivalent cashflows.

  If a PPP procurement is being used for a government initiative (eg a highway), then it should still be appraised with a social CBA using a STPR on the cashflows. The cashflows will include the long-term repayments, which are determined by the supplier's cost of capital.

- It is important to understand the rationale for public involvement. If the project just displaces a pure private sector project then a private sector SOC rate is appropriate. If the project reflects a standard reason for public involvement (e.g. public good or merit good) then the public sector discount rate should be adopted. In practice, some projects - eg a new university - could be considered a mix of the two, in which case a practical solution may be to take an average of the two rates.

- Are the private entities actually taking on real risk in the public-private partnership, or are returns more guaranteed such that discounting based on generic private rates of return (which reflect higher degrees of risk as well as externalities etc.) would not be appropriate.

- It is essential to differentiate the risk born by the government from the actual risk-free benefits and costs being discounted.

- This is challenging. I have though about this. I think that a risk adjusted social rate is appropriate.

- Comparability across projects involving public funding (though to different extents)

- Analysing the relative risk to each partner and the risk that each partner fails to deliver to the initial agreement. The problems with most such partnerships is that there are regarded less as mutual partnerships and more as restrictive contracts. Any given risk is therefore viewed differently by each partner.

- The value of the public contribution would be valued appropriately. I would argue using empirical estimates of marginal value of public funds rather than estimates of marginal cost. In evaluating such projects all alternative should be consider to understand the true opportunity cost of the proposed project ie other ways to spend the same public money or raise more public money "

- The main issue here is who ultimately bears the risk. In situations where the ultimate risk of failure reverts to Government, the economic appraisal should be based on the STPR. In saying that, I am a supporter of the Ryrie Rules, but possibly with a recognition that if SOCEF>1, and if private capital is genuinely replacing/supplementing public capital, that might be represented in the appraisal.

- One issue is the diversity in the perspective: social preference vs private financing, as they might conflict. I think that understanding the financial viability (the industry viewpoint) is a necessary condition to avoid cost overruns and public budget issues. However, within the viable options, policymakers should assess what projects are mostly valued and needed by the society. The last problem is to determine the value of projects for the society. I am not sure time preference is enough as indicator, other indicators should be considered, also as depending on the different social groups

- Opportunity cost of capital and who holds the debt, who hold liabilities, and where do the returns go. Also, is the benefits general public benefits or fiscal benefits that could be transferred to the public.

- '-How to account for the difference between the 0% interest rate on public funding and the potentially large interest rates on the funds raised on financial markets? - How should all types of government subsidies be taken into account? (for instance, contracts for difference, CFDs) - How to compute a certainty equivalent of government subsidy?"

- Comparing PPP with public finance, the problems are higher WACC for PPP and arguably more costly or less careful supervision of public projects, and PPP suffers from danger of expedient bankruptcy so they take the upside and leave the downside with taxpayers. The key is to determine when the quality of management is self-evidently superior if provided by the private sector.

- As in the private sector, the public sector should adjust the discount rate to the riskiness of the project.

- That one does not outcompete private sector initiatives.

- The STPR considered here is 3.5%, which is lower than the average real return from financial market funds like the SP 500 (at roughly 7%). This leads to an inconsistency: the government could not invest in the public project and receive a higher return from their money investing in the private market. This could lead to public backlash to the investment made by the government as it there seem to be higher returns for the public possible. (I generally consider this line of thought correct but do not agree that one should increase the STPR due to this argument.)
- Here it is especially important to use the SOC/fair value approach in order to understand and communicate the differences between using a PPP and other alternatives.
- How a project is financed and owned should not come into this. We should be most concerned about how resources are used and looking to optimise outputs accordingly. *[Government]*
- No experience in this. *[Government]*
- the main issue would be where will market risk and the risk of the concession lie; who will buffer any deviation from the numbers assumed in the appraisal. *[Government]*
- There are two buckets of considerations depending on the question. Firstly, when doing cost effectiveness analysis assuming the project will go forward (which necessitates the inclusion of the public sector comparator). Given the cost to the economy is the same, there is a clear rationale for SOC to recognise the social cost more evenly given financing costs are included. Secondly, when considering through a lens of moving from short term spend to long term spend, there is a conceptual question of relative budget constraints. While affordability is not an acceptable reason for choosing private finance, PPP could enable choosing the private finance project in parallel with another project instead of consecutively. *[Government]*
- Depends on the rationale outlined for the social discount rate I would guess. I'm assuming the PFI or PPP options should always be compared against a public sector comparator to help evidence value for money? Not sure I'm on top of all the issues here in all honesty. *[Government]*
- Key question is whether government is taking commercial risk. If it isn't, as the exposure is limited, then it's appropriate to use STPR. If it is, then a private WACC should be used. *[Government]*
- No idea -sorry! *[Government]*
- For MHCLG household, business and public sector costs and benefits would be all be discounted by the STPR.

  The issue with public private partnerships is that they tend to be very complicated and are only justified if efficiency gains and handling of risk by the private sector offset the higher costs of finance that private sector providers face. But those should be handled through rigorous risk based assessment of costs and benefits not by altering the discount rate, as the decision still takes place within a social welfare framework. *[Government]*

A.2.10. On the extent to which the discount rate should be adjusted for place-based projects

- I think this reflects the fundamental difference in opinion that I have been what financing a public project looks like. I think the French model is vastly superior, where discount rates for public projects is generally at or neat the borrowing rate of the government.
- Different places have different exposures to systematic risk.
- No. Place-based objectives should be reflected through the identification and valuation of complementarities and spatial spillovers in benefits, rather than through adjustments to the social discount rate. Maintaining a common discount rate preserves consistency and comparability across projects.
- Projects should be assessed as a group, but standard procedures applied. Defining place may be difficult, e.g. in health or defence.
- The CBA test is about increasing the size of the national cake. It is merely one, albeit important, criterion. Regional and inequality considerations come after the Hicks-Kaldor test. It is perfectly legitimate to consider much more than the straight CBA NPV, such as additionality of places and distribution, but making these integral to the pure CBA risks the Green Book process becoming politicized and hence circular.
- I believe that there is a need to have more flexibility for projects to account for place-based differences. I am not sure what the quantum of adjustment should be, but to ensure an upgrading of local infrastructure, then a one-size-fits-all approach will lead to very skewed outcomes. One example would be rural broadband. Greater connectivity is critical for many things, but it is never economically efficient in the short-term and market forces dictate where the private sector cost-benefit (return) sits e.g., towns and cities. However, the lack of connectivity in rural areas is detrimental and imposes unnecessary frictions on economic activity. Removing these frictions is not in the interest of the private sector but is very much the business of government.
- to the extent that their (climate) impact differs

- This should be handled in a different way rather than adjusting the discount rate
- Non-financial returns will be much higher for certain place-based projects and so it is important this is reflected in the application of a lower discount rate
- No, should gravitate solely towards projects with the highest risk return, avoiding prospects of cross-subsidies.
- No idea what that means?
- In general I am wary that this will lead to putting political considerations ahead of economic considerations. To the extent that complementarities are relevant (and this could be the case), I struggle to see why the discount rate is the correct place to consider them, it should affect undiscounted benefits and costs.
- same as before - hard to get right, overly complex, conveys too much precision, opens the door to political wrangling and arguments
- I believe place-based projects invariably create systematic complementarities.
- As mentioned before, it is better to adjust the numerator and undertake scenario and simulation analysis of the project appraisal calculations. While "finance theory" might suggest that projects that somehow diversify risks might warrant a lower discount rate this could well be misunderstood and abused. Again, it would be better to bring project complementarities into the numerators.
- If we have a time based preference parameter and an income/consumption elasticity parameter, there is no reason for the discount rate to vary by context. The fact that health ignores the income effect is already a significant error in current practice; we should avoid reconsidering the question which was posed to the environmental discount rate a few years ago.
- I think that the discount rate serves a specific purpose and shouldn't be a parameter that is reduced in order to perform sensitivity analysis. In this instance, I would assume that some suggest it should be lowered in order to make the case for this type of 'systems' investment more favourable. I think that's great because CBA is weak on this front, but I don't think the DR is the parameter with which to do it.
- My knowledge on this aspect is not sufficient for a reasonable judgement.
- Place-based objectives should be reflected in what is counted and how benefits and risks are modeled, not in adjustments to the social discount rate.
- Discount rate should stay the same as a common factor across place so the effect of other place based factors can be identified more clearly. Those factors should include socio-economic factors like income distribution of course but that should be adjusted through income weights for example not through discount rate.
- There is no rationale for adjusting the discount rate in such cases.
- Is the discount rate on a bundle of projects in the same place different from the same bundle of projects but in different locations? Seems like it shouldn't be? Though perhaps you could make the argument that the discount rate should be higher for a place-based business case (all your eggs in one basket)? I think given real world evaluation you should just use the same discount rate (and I think the question is kind of missing the point, if there are complementarities between the projects its the benefits part of the consideration that should be adjusted not the discount rate).
- whilst not strictly related to place based discounting, please note the disclosure / guidance recommendations as per our paper.

  Rec 2 For all projects consider applying a sensitivity removing the wealth effect from appraisal calculations. This would ensure that decision makers for all projects understand how BCRs would be impacted if future growth in wealth is much lower than the central assumption assumes.

  Rec 4 Enhance the evidence base for discounting impacts by requiring users of the Green Book to calculate the duration of annual costs and benefits respectively for their projects and use this to assess and monitor any bias toward short termism on an ongoing basis. (See Appendix B for a definition and example calculation of duration.)

  Rec 6 Carve out what counts as transformational, using a top-down approach by defining key systems changes required and their implications, and setting out concrete examples to guide user

  Rec 7 Reduce reliance on quantitative measures when appraising transformational projects, providing clear guidance that, due to inherent uncertainties, the economic case, including cost benefit analysis, should form a smaller proportion of the decision making criteria than for non transformational projects.

  Rec 8 Embed assessment of co-benefits for all transformational projects via a succinct principles based framework. A useful starting point for this framework could be the 'Triple Dividend of Resilience' approach used to assess climate adaptation projects in Defra's supplementary accounting for climate change guidance.
- No, consistency of discount rates becomes even more critical in a place-based approach. Instead, consider critically inter-dependencies in the benefit and cost streams, and undertake sensitivity analysis.

- Not the value, but the application should be adjusted. Place-based activities are local public goods that raise local land values and the intensity of land use around them. This is the ideal taxbase for funding the share of benefits that are local public goods. This can be struck as a long-term recurring special tax to repay a special project loan that is supplied on market terms. Only the residual expenditure required from a general purpose Govt's general tax base needs to apply a CBA, and a STPR can apply to that public increment only.
- Easiest way in current framework is that the elasticity is likely to be different in low income compared with average or high income places.
- Place-based objectives are just one example of a myriad of reasons for government involvement in projects and have no special feature that distinguishes them from other valid reasons (e.g. population-based reasons such as a national mental health digital roll-out).
- I don't know enough about these business cases to assess whether they present meaningfully different risk profiles.
- I don't see why the discount rate would be different from such projects.
- Nothing in the description seems to implicate standard adjustments such as shadow price of capital or risk.
- It sounds reasonable in theory, but in practice, this sounds very complicated. More complication makes it more subjective and less defensible.
- Same as my answer to project-specific discount rates. Addding additional considerations such as this could introduce inconsistencies and arbitrariness (see, e.g., Sher 2024 AER)
- These issues are totally different and should be dealt with elsewhere in the CBA, not via discount rate
- Given that STPR will vary between places according to both different preferences and needs it seems appropriate that the discount rate should be adjusted to reflect this
- Best to value all the effects and thier interactions explicitly with empirical evidence or explicit assumtion and judgement that can be tested and updated. Using the discount rate to do this will inveitably be incoherent and unaccountable
- Adjusting the discount rate is not the correct way to handle it. Leaving aside all the practical difficulties of doing it well, if there are trade offs between efficiency and distributive impacts, those should be recognised explicitly in public choices.
- Yes because it is important to understand the synergies between different public projects when assessing their societal value
- If place is uncorrelated with the risk profile then no. If place is correlated with the risk profile then yes. If place is about equity and distribution, then that is a different decision and should be dealt with with some sort of equity weighted objective function.
- complementarity may affect the riskiness of the project(s)
- But all the synergies need to be included
- 1. It makes no sense to value interdependent prjects separately. 2. Different places may face different growth rates.
- I think these can be discussed in the qualitative parts of the CBA. But I would not add such clutter to the NPV calculation. Moreover, such concerns will feature in the public policy decision making setting that a CBA informs.
- I am not sure that I understand how the place-based business cases would be managed/regulated by the government, so my answer could be off. I would say that a risk-assessment should be done that is project specific (estimating the beta) and the discount rate should be adjusted for that as well as the time-horizon of the project. Other than that, the government could base its decisions on the STPR that is different than the interest rate used by the private businesses. Decisions by the government should then anticipate what the private partner will decide.
- To the extent that place has an effect on fair value rates, it would be automatically reflected if that approach were adopted. Incorporating an adjustment into rates using an alternative would be entirely ad hoc. Adjustments should be made to cash flows to reflect externalities, not to discount rates.
- One should not have a different discount rate according to place. This will likely make for overly complex economic cases. Instead, weights should be used on the benefits and costs to reflect a locations different socio-economic conditions. *[Government]*
- I don't follow how this would mean a different discount rate. However as indicated in previous answers, the principal of a much lower rate still holds and I would still believe it should diminish to 0 over time. *[Government]*

- National funding will need to take into account only national benefits. The discount rate should be the same everywhere in the country even though some redistribution coefficients could be used to reflect that some places are not as wealthy as others and the marginal utility of consumption will be different. This uplift shouldn't, however, change in time as this would mean assuming varying rates of growth across places. I don't think that we have enough data to robustly estimate convergence or divergence across regions and even less so across places *[Government]*
- I don't see a strong argument for varying it to be honest, so I have answered 'no' here. Some of the issues that you might think within the discount rate that could vary by place are probably better investigated via other metrics when it comes to project appraisals and selection (appraisal summary tables etc). Moreover, one rate for everywhere is easier to explain/justify and cuts down on the necessary guidance that would be required for varying rates by place (and the justification for that approach). *[Government]*
- Large regional variation in a number of the elements that make the discount rate *[Government]*
- Irrelevant. This is a distributional question. *[Government]*
- Shouldn't be necessary if we arrive upon an agreed upon rate. the point isn't to say 'this is the perfect rate for all circumstances', it needs to be the 'good enough' rate that can be taken as given in appraisal *[Government]*
- Separate guidance should be given on how to properly assess these so they are transparent taken account of - not hidden within the discount rate. Any amends to the costs and benefits should take account of the place based nature of the policy *[Government]*
- The discount rate isn't the problem. The issue is that transformative appraisal is not possible. If you are doing something to uplift overall productivity for an area and you keep using the GVA and population figures for the area before those shifts then you downplay the impact. I personally think this change is useful but it will be difficult to do - I actually think there should be an "ambition" based change in some of the underpinning economics for an area. Set the economic aims for all of the projects as a whole and appraise them on the basis that you will achieve that as well as not achieving it. By all means discuss the possibility of failure but don't assume it just because each individual change is marginal. Set the target and if the policies and projects suggested aren't expected to achieve that we know we should do more. *[Government]*
- The reason the discount rate should not be adjusted is because both the costs of place based interventions and the impacts of place based interventions are felt more widely than a particular place in which projects are prioritised and implemented.

  - Costs come from public sector funding through local government funding and specific government grants to local areas and are calls on the general taxpayer.

  - The impacts of place based initiatives will involve displacement and leakage beyond the area in which an initiative is taking place. A place based discount rate would effectively impose a rate on other areas being impacted by the policy. Displacement and leakage impacts are often very large.

  - Finally many major projects cross place boundaries eg HS2, strategic road projects and national policies. It is not clear what discount rate would be used for these interventions.

  A better approach is weight benefits and disbenefits to reflect the income group of the people being impacted on. *[Government]*
- Given the current components of the discount rate, it's not clear why the discount rate for place-based/transformational changes should take a different value. *[Government]*

A.2.11. On the extent to which the discount rate should be adjusted for environmental projects.

- What does "adjusted" mean in these cases. If we have an environmental risk, then typically this increases discount rates. But in many cases, the absence of investment is the most damaging issue. The current economic theories we have on this topic don't fit what is needed in the current scenario.
- There are very solid economic arguments that favour including some measure of the opportunity cost of environment depletion. Unfortunately there doesn't seem to be a consensus on how to estimate these costs, and therefore it is safer to ignore these. Alternatively, a flat rate of 1% could be considered, although this rate admittedly comes out of thin air, and therefore suffers from the same issues of social discount rates in general.
- No. Environmental scarcity and limited substitutability are better reflected through improved valuation methods, such as relative price adjustments and updated shadow prices, rather than through changes to the social discount rate. Adjusting the discount rate risks conflating time preference with valuation issues and reduces transparency and comparability across projects.
- Environmental scarcity and limited substitutability matter, but they should be accounted for in the valuation of environmental impacts over time, not by changing the discount rate applied to all future outcomes.

- The environmental elements should go into the costs or benefits or be treated as social values in awarding contracts
- The Treasury was correct to favour improved valuation for environmental impacts and updating these estimates to reflect latest evidence. But it is sensible to include sensitivity analysis for consideration where there is good reason to suspect that impacts are not being captured adequately.
- By not taking account of environmental scarcity and the lack of substitutability of nature. The current approach does not fully capture the costs and benefits of particular projects. As such there will be a misallocation of public funds, and for environmental scarcity and nature, then limited and often irreplaceable resources are being depleted.
- see all prior answers
- I believe that there are different ways of accounting for the growing scarcity of environmental goods other than through using a different discount rate
- Such risks/costs have been ignored for too long; Treasury decisions cannot continue to overlook these critical effects
- Criteria objectivity and statistical merit of such adjustments still face severe credibility issues.
- Possibly yes in the case of clear externalities; but far from clear whether this should be dealt with via discount rates, as opposed to the numerators in pv calculations.
- I agree with the current position, consistent with my answer to the previous question – deal with these issues through undiscounted benefit/cost estimation.
- Those costs are real but should be factored in via updated and comprehensive / holistic estimates of the actual impacts, not by using the discount rate
- Valuation techniques have improved a lot but many methodological and implementation issues remain. A declining discount rate would be preferable.
- I agree with HMTs approach.
- Values for environmental goods should be adjusted; if they become more scarce their value will go up. The discount rate is not the right tool for these adjustments.
- I think that the DR should be lower overall because of environmental considerations, and that WTP values for the environment should reflect future scarcity (I know Prof Groom has done good work on this question recently), but as per my previous points I don't think the DR should be a parameter that is 'flexed' like this. It makes it seem like it's an arbitrary parameter to start with-though I guess many people probably think it is :)
- Theoretically, each contribution to welfare has its own discount rate schedule. So the consumption discount rate should not be adjusted; rather, discounting should be applied to each determinant according to the social marginal rate of indifference between a unit of that determinant in adjacent time periods.
- Yes, absolutely. Environmental scarcity and limited substitutability is currently underestimated and overlooked. Very low, possibly even negative discount rates might be more applicable in some cases, due to present overuse and degradation of ecosystem and overshoot of planetary boundaries. Accounting for relative price changes, as well as biophysical limits, feedback loops and complexity is key on this context.
- Scarcity is a key driver of (relative) economic valuation in case perfect substitutability is not applicable. Nature plausibly becomes more scarce relative to the general economy and it seems highly unintuitive to think of the two as perfect substitutes.
- As is made clear in the government's own assessment (https://www.gov.uk/government/publications/nature-security-assessment-on-global-biodiversity-loss-ecosystem-collapse-and-national-security) there is catastrophic risk of continuing with our current accounting approach that does not properly consider long term environmental harms and benefits. This is evident to everyone with scientific understanding and should lead to a major reconsideration of our economic accounting approach.
- I agree with your assessment in Science (Drupp et al., 2024. DOI: 10.1126/science.adk2086): "an alternative to estimating future WTPs adjusted for relative price changes is to instead use different discount rates for ecosystem services and market goods. This, however, would also require changing the standard discount rate. The alternative that we propose here, which is mathematically equivalent (see SM), is to adjust future WTPs and use a single discount rate schedule. This proposal is simpler, more transparent, and often more compatible with how guidelines deal with other nonmarket goods."
- Discount rate should be lower. Two key reasons. The first is that nature has limited (an optimistic best) or no (more realistic) substitution. We need to understand what we mean by 'limited' - a single ecosystem service could be substituted but not substitution of a natural asset. I gave an answer to a previous question about the views on valuing environmental impacts...Those values, however robust, risk being ignored so nature does need the support of lower discount rates. The second reason is that we are assuming future generations will have more resources (financial,

technological etc.) to deal with negative impacts on nature. But we may pass tipping points by then, or changes may be irreversible even if not quite at tipping point. So the increased future resources may still be not sufficient. We cannot continue to allow economic analysis to, essentially, discriminate against nature or other non-financial returns, and against the poor now and the generations in the future.

- Uprate prices instead
- If the valuations of things increase with consumption growth then this is not required.
- The conclusion of the previous Treasury review is correct: better valuation of individual impacts is always going to be superior to the subsuming these in a 'blanket' coefficient, like the discount rate. The application of a declining discount rate for the distant future may also be seen as already capturing some of these concerns.
- Don't support a social discount rate (see answers to other questions).
- See planetary solvency papers https://actuaries.org.uk/news-and-media-releases/news-articles/2025/jan/16-jan-25-planetary-solvency-finding-our-balance-with-nature/ for discussion of environmental scarcity and policy responses needed
- We have answered yes to the above, but it should be interpreted as a "yes only under certain circumstances", as explained below.

  We believe the approach outlined in the question and in the relevant section of the Green Book refers to the methodology described in Drupp et al. (2024) (https://www.science.org/doi/10.1126/science.adk2086). While we agree with the problem outlined in the paper and the theoretical approach described to correct for the issue of environmental degradation in policy appraisal, often the data required for such adjustments is unavailable to policy makers and appraisal practitioners.

  One option to mitigate this problem could be to adopt an approach similar to what has been outlined by this French Government publication. That is, where information about relative price changes is available for the specific area of the policy intervention, an adjustment should be made using actual data. However, where that is not available, practitioners could take a scenario approach and show a range of results. For example, they could show how cost benefit analysis results change if relative prices were to increase by 1%, 2%, or 5%, explaining why those changes are reasonable. This would provide decision makers with more complete information about the value of the environment and the investment overall, even in absence of accurate data.

  However, while this approach is valid in the case of policy interventions specifically targeting reducing emissions or managing risks to the natural environment, it is not immediately straight forward how this would apply in cases when the intervention under consideration is not primarily climate related. A serious area of concern is whether proposals for non-climate related interventions effectively account for the impacts the policy will have on the emissions and therefore climate change, in particular climate costs and environmental degradation. Consider for example an investment into the expansion of an airport. This would affect the environment both in the short term (land use, construction emissions), but also in the long term (consistent higher emissions from increased air traffic). It is important that these environmental impacts are appropriately accounted for in a value for money analysis, in addition to the more direct costs and benefits of the investment.

  One way to solve this issue could be by changing the growth rate used in the Ramsey Rule, choosing a rate that better accounts for environmental impacts, as described in our answer to question 3. Otherwise, working with the current Green Book guidance, the only other alternative is to ensure that environmental costs are accounted for correctly in the appraisal. This is a complicated and concerning issue, because when future environmental costs of a policy are difficult to estimate, due to uncertainties or lack of data, the risk is that environmental costs will only be appraised qualitatively. This would lead to a biased cost benefit analysis where the value of benefits of the investment is high but costs are underestimated. For instance, in the airport expansion example, it might be difficult to estimate long term environmental costs. This is because the extent to which emissions will increase due to air traffic will not just depend on how many flights the airport will be able to accommodate, but also on future technological progress on sustainable fuels or aircraft efficiency, that cannot be straight-forwardly estimated in the present.

  We would encourage the reviewers to investigate to what degree environmental costs may be underestimated in government appraisal. Similarly to our suggestion in our answer to question 9, this could be done by requesting access to a sample of business cases across different policy domains and budget envelopes, focusing in particular on reviewing the value for money section in the economic case. If the review revealed that environmental costs are not appraised systematically and consistently across policy areas, we believe the next best alternative would be an adjustment to the discount rate as discussed above.
- The core social discount rate should remain consistent across appraisal. However, where environmental assets are critical and non-substitutable, relying solely on relative price adjustments may be insufficient, and lower effective discounting or constraint based approaches may be justified?
- No, environmental scarcity etc should instead be captured and considered in the benefit and cost streams, not in the discount rate.

- These impacts will vary by each project. If they were universally constant such that a discount rate could capture it, then it also means that CBA isn't required for each project, because the answer is known already.
- Similar to adjusting for time varying rates. I think the added detail would lose public credibility and also that the precision of such estimates may actually go down.
- If relative prices are changing due to scarcity then the returns to a project (that alleviates scarcity) will increase and that is the appropriate mechanism to reflect the scarcity in question. The standard discount rate should then be applied to these higher returns.
- If adjusting for relative price for every environmental good is (presumably) unrealistic in the near-term, a downward adjustment to the discount rate to reflect general limited substitutability of non-market goods (which are not just limited to environmental goods!) is appropriate as an actionable alternative. It should not necessarily be limited to environmental goods and should consider other non-market goods with limited substitutability.
- It is more accurate to attribute benefits and costs to the policy or project, than to indirectly approximate them using an adjustment to the discount rate.
- If I understand the previous statement correctly, there is already contemplation of adjusting for scarcity (which I think makes sense) just not through adjusting the discount rate. That seems sensible to me. Treat it like certainty equivalent calculations. The time path of environmental scarcity adjustment need not follow a geometric discount rate (as far as I recall).
- Recent work by Drupp shows that relative prices is relatively easy under certain simplifying assumptions.
- I am unsure of the latest literature on this subject, but I would consider it.
- As before, I'm yet unsure about the arguments, but could be still persuaded! At present, I see scope for improving consideration of environmental effects better elsewhere in the appraisal process, not via discount rate
- Environmental impacts are potentially measurable in the same way as for example value of time savings and effort should be devoted to improving such evaluations rather than adjusting discount rates. STPR should in theory include an element that expresses preferences relating to environment; change
- Better to value things explicitly using evidence and explicit judgement that can be accountable to reason, evolving evidence and social values
- Again, if society wants to place a 'high' value on landscape or biodiversity, it should do so directly, not by adjusting the discount rate.
- It might be important to include environmental parameters, although it might be also difficult to estimate these accurately. The risk is to distort an already complicated analysis. The advantage is that environmental issues would be considered
- Probably not, getting the valuation right and using certainty equivalents is probably better. Without good finance training people tend to get that approach faster.
- as the prior review, I favour improved valuation for environmental impacts and update of these estimates to reflect latest evidence, such as using relative price change adjustments.
- assuming that the correct time horizon in chosen - potentially long - and that relative price changes are included
- The cost-benefit analysis should value all impacts, each of them with its own flow of sociatal benefits and risk-adjusted discount rate.
- This is very important to account for and we have the tool. But adjust instead the net-benefits/ calculation prices.
- If increasing relative prices were taken into account when monetizing environmental impacts, this would change the overall composite consumption growth of the economy, which would change the discount rate. However, if inequality aversion over time is 1 and the elasticity of substitution between the environment and consumption is 1, then the discount rate should not be adjusted. I am not sure what the elasticity of substitution is for the UK, so I cannot answer this question.
- Here again, I definitely would adjust value for these factors but it obscures the analysis to incorporate these consideration via a discount factor.
- The above arguments still stand, unless we have good evidence which now contracts this. *[Government]*
- Some environmental benefits like, for example, those related to climate change, are quite uncertain and a specific discount rate can acknowledge this. If, in the future, we have better data about climate change then we could go back to a single discount rate but this is not likely to happen any time before the next revision to the Green Book *[Government]*

- I don't know the evidence base well enough here so have answered unsure. My gut response is 'no' however - as I don't see a reason to change - as is suggested by the prior review. It could potentially make appraisals more complex, reduce comparability across projects, and give rise to special pleading for different rates in order to game appraisals in my view. *[Government]*
- as environmental assets become scarcer over time, their marginal social value increases rather than declines, meaning that applying a standard consumption-based discount rate risks systematically undervaluing long-term environmental benefits. *[Government]*
- The above reasons still seem to hold to me. *[Government]*
- Shouldn't be necessary if we arrive upon an agreed upon rate. the point isn't to say 'this is the perfect rate for all circumstances', it needs to be the 'good enough' rate that can be taken as given in appraisal *[Government]*
- I strongly believe adjusting the discount rate in this way would be a mistake. There are a wide range of ecosystem services, each with a different value, which also differ in terms of quality, location and what environment valuation method was used to put a price on them. This means that they will each greatly differ in terms of how those benefits should evolve over time - each ecosystem service will be different. How these benefits change over time should be the subject of specific guidance, similar to how the value of GHG and time savings change over time. *[Government]*
- Fundamentally UK nature is quite substitutable (20 years plus as an environmental economist) where it isn't (e.g. rare species or ancient woodland) then the marginal processes of CBA won't be fixed by playing with the discount rate. Fundamentally there are some projects where this kind of appraisal doesn't work - e.g. for what price should we sell parliament to Vladimir Putin to knock it down and build some American Candy Stores? It is a fundamental violation of other moral considerations. CBA works where we are playing off marginal changes it doesn't work where we are eradicating a species. The discount rate isn't the problem. Moreover we have been over this many times in the last 40 years and always come back to the same conclusion, *[Government]*
- The review carried out last time was thorough. The recommendations from it are still valid. The focus of appraisal activity where environmental impacts are concerned should be on improving their measurement and understanding risks. There should be more focus on assessing specific environmental risks and interactions with other impacts and building that into the appraisal of costs and benefits eg through sensitivities, scenario analysis and distributional analysis. An environmental discount rate achieves very little.

  It would though be good to look at how far departments have got in terms of improving their assessments of environmental impacts and the degree to which there is a common treatment of environmental risks across departments eg through common scenarios. *[Government]*
- Defra economists provided several arguments in favour of adjusting the discount rate for environmental impacts in the previous review, which we believe still hold and should be taken into consideration.

  In short, there are good arguments for exclusion of the wealth effect (as is the case for health impacts) when discounting environmental projects: - diminishing marginal utility of consumption may not apply to the environment, similarly to health - irreversibility of loss (e.g. due to tipping points) and limited substitutability of natural assets compared to other assets

  We also note the ethical argument against applying the pure social time preference to cost-benefit analysis that involve significant and irreversible wealth transfers from the future to the present, as per the Stern Review and as referenced in the current Green Book guidance on discounting. Since a large portion of environmental projects are of this nature (and many environmental benefits are only realised decades into the future), this argument applies to the environment particularly acutely.

  While the current Green Book instructs practitioners to conduct sensitivity analysis with pure time preference element removed for these cases, we question to what extent this is applied consistently and effectively in practice across projects to sufficiently address the issue. We would therefore welcome this issue to be revisited and addressed in full, with clear and explicit guidance.

  At the same time, we acknowledge practical difficulties and arguments against applying a separate discount rate to environmental impacts. Having a unified discount rate ensures consistency across projects (despite the issues outlined above). A single discount rate also means the discussion around value for money can focus on the costs and benefits themselves, rather than the nuances of how those values were discounted. The argument regarding limited substitutability will also not apply to all natural assets and ecosystem services equally.

  Following the 2021 review, we updated the ENCA guidance accordingly (section 2.6), highlighting that issues specific to environment should be reflected in relative price adjustments, and key considerations around this. Even so, we believe that there are considerable practical difficulties in applying these recommendations for teams across the government.

  We would therefore welcome a considered revisiting of this question. Even if the review concludes against the adjusted discount rate for environmental impacts specifically, we think explicit discussion around this issue is required, with further practical guidance to relative price adjustments for environmental impacts included in the resulting Green Book guidance. This is particularly important now that the Green Book has been shortened and detailed discussions around complexities of environmental impacts have been removed from the main text. *[Government]*

A.2.12. On whether the discount rate should be changed for transformational projects

- See my previous answer. HM Treasury clearly does not believe in Arrow-Lind, I do. When the training is so heavily into the standard financial economics of risk premia and very little is on the discussion of what is the inherent difference between public vs private and PPP, we end up basing decisions on pure financial considerations under very high discount rates. This is in contrast to looking at what is needed and what can be afforded. Transformational projects should be financed and should not be a mechanism for making arbitrage premia from changing risk to effectively riskless projects that are often uncorrelated to the business cycle.
- Project valuation discounts future benefits using an adequate discount rate. The benefits are at the numerator and the discount rates at the denominator. A "transformational project" will simply be one with a big numerator, so there is no need to adjust the denominator (the discount rate). This sounds like a marketing/political soundbite.
- No. Transformational projects differ in the scale, persistence, and structure of their benefits, but this does not imply a different social time preference. Rather than adjusting the discount rate, transformational impacts should be reflected through declining discount rates, alternative growth scenarios, and explicit modelling of structural change, thresholds, and spillovers. This preserves consistency while allowing long-term benefits to be properly assessed.
- To take a fair view of long-term transformational benefits, HM Treasury should change how projects are valued and appraised, not change the discount rate applied to time itself.
- Transformation should go in the costs and benefits not the discount rate.
- It is an important matter, but one for sensitivity analysis rather than 'fixing' the NPV
- This answer is based on politics rather than economics. I have yet to see a transformational project that would merit such an adjustment. When projects like HS2 and Hinkley Point can be signed off. There is no need to adjust the discount rate, there will always be a government that sells "transformational" irrespective of whether it is actually a) transformational or b) needed at all. Creating a special category for the whims of the government of the day seems a slippery slope. I also see this position as being consistent with my place-based adjustments answer as those types of projects will be diverse and more likely than not cost considerably less and much more diverse, than what a government sees as transformational.
- Adjust the discount rate according to the carbon intensity of the project, low carbon projects will decrease risk (and future premia), hi carbon the opposite.
- Any transformative effects should be modelled via changes in the future cash flows, that are potentially generated via positive externalities of the project in question, and should be added to the valuation of the project, and not via changes in discount rate.
- I have never heard of this and I think it should be resisted.
- The UK has consistently underinvested in such projects over the last 50-60 years, particularly outside the South East. Adjustments could reflect the need for investment resulting to lack of prior-period investment (as a tangible way of actioning a levelling-up agenda).
- In corporate sphere, "transformational" tends to be synonymous with dynamic but mostly disasterous. Follow the private sector's lead and don't repeat/amplify their mistakes.
- This is again about the numerator rather than the denominator in pv calculations
- In general, unless the project is non-marginal, the existing framework should be more than sufficient. If the project is truly non-marginal then the issue is not so much the discount rate as stepping outside marginal analysis and doing proper modelling in general equilibrium.
- Hard to define, subject to political debate, might matter but shouldn't be shoehorned into discussions of the social discount rate
- It depends on the specific project. How are the benefits and costs of the project distributed?
- Shouldn't all public projects be transformational? If they're not, then then discount rate is not high enough.
- This relates to the endogeneity of the consumption growth parameter; if the project is large enough to affect national economic growth, than a differential discount rate by project might be justified. Evidence building, however, needs much more attention.
- As per previous answers, I favour this being dealt with differently. CBA underperforms in this area (it's essentially a question of non-linearities). Qualitative economic arguments perform well here in my view.
- My knowledge on this is limited but it seems intuitive. Otherwise it seems difficult to get out of certain undesired equilibria where change is more costly initially than mantaining the status quo.

- The transformative benefit of environmental action should be accounted for, of course including reduced emissions, but also the transformative benefits that arise from nature restoration such as increased tree cover, sustainable fisheries, etc.
- Transformational projects do not justify adjusting the social discount rate; their effects should be captured through explicit modelling of benefits and risk. Discounting should differ only if such projects materially change economy-wide growth or systemic risk, for example through declining or tail-hedge discounting applied consistently across projects.
- yes to lower discount rate in general. However, the key may be in ensuring costs and benefits of a transformational change are better defined, quantified and valued in monetary terms. In addition, the definition of transformational change needs to be sharper.
- Except to the extent that a transformational project is high-risk mega-project with potential systemic macro-fiscal risks, there does not seem to be a good reason for applying an adjusted social discount rate. Attention should be focused on identifying and valuing the transformation changes and the risks attaching to them. This sounds like an attempt to open up a backdoor route for marginal projects. Who determines what is and is not 'transformational'?
- Incredibly sceptical of this argument. I suspect its used to lower the discount rate when the benefits of the project cannot be adequately justified.
- see our paper:

  and in particular recommendation 3 Rec 3 For transformational projects consider removing the social time preference rate from baseline calculations, or, as a less preferred option, introduce its removal as a sensitivity calculation.
- We would advise against the use of adjusted discount rates for transformational projects. Typically, arguments in favour of adjusted discount rates for transformational projects would recommend using lower discount rates in these instances. This is based on the assumption that transformational investments will return significant long-term benefits to society which get discounted heavily if the discount rate employed is high. If that is the case, when compared to other projects, transformational projects get unfavoured unfairly.

  However, this line of argument might not consider two cost related issues. Firstly, non-transformational projects with spend happening in the short term will appear cheaper in present terms than transformational projects, which will lower the incentive to invest in transformational projects even further. Secondly, non-transformational projects might have long term costs on the environment and natural resources. However, if the discount rate applied to non-transformational projects is high, once again the present value of costs will be unjustifiably low.

  For these reasons, we would argue that it is preferrable to use a single, lower discount rate across the portfolio of government investments (in line with our answer to question 3). This would ensure that long term benefits of transformational projects are not overly discounted, while allowing a fair comparison across all projects.

  However, if our goal is ""...to make sure that the government is taking a fair view of the long-term benefits that arise from transformational investments"", we should also point out that regardless of the discount rate chosen, a quantitative cost benefit analysis alone is unlikely to provide a complete answer. When considering transformational project, a quantitative cost benefit analysis will almost always produce partial results, as benefits are uncertain and difficult to appraise quantitatively.

  In recent years there has been growing recognition among practitioners of the importance to incorporate a substantial qualitative component in any value for money analysis, to prevent decision making from solely being based on imperfect and incomplete benefit to cost ratios. The latest versions of the Green Book have started to be more explicit about the importance of qualitative approaches to value for money analysis, for example by referencing cost effectiveness analysis as an alternative tool to cost benefit analysis. However, the Green Book traditionally has mostly focused on quantitative techniques, providing little guidance to practitioners on what approaches HMT would endorse for qualitative value for money analysis. While this is probably beyond the scope of this review, we would encourage HMT to provide additional guidance on how to incorporate qualitative considerations in business case value for money analysis. This could be done in the form of supplementary Green Book materials, and ideally showing case studies and step by step examples for practitioners.
- Rather than a blanket adjusted rate for all transformational projects, HM Treasury should: 1.) Apply a lower long-term discount rate only where there is credible evidence of a permanent structural shift in the consumption growth path - requiring independent validation rather than departmental self-classification. 2.) Perhaps use scenario analysis across multiple discount rates rather than a single adjusted rate, reflecting genuine uncertainty about transformational outcomes.
- No, the discount rate is not the right place to capture transformation. Indeed, transformational projects would be better to have consistent discount rates, and then ensure that benefit-cost analysis is combined with other types of analysis eg to better understand the possible transformation paths.
- Yes, supposing that the change is significant enough to affect the growth rate g of consumption and/or the annihilation risk L assumed in the STPR. (Ie, like the Stern review did.) Otherwise, no, just continue to use the STPR on the risk adjusted or expected cash flows

- In the context of benefit-cost analysis, I think "transformational" projects raise large issues about baselines and discrete changes in the economy. Such issues, in my view, are better handled directly via scenarios, careful baseline work, and so in instead of adjusting discount rates.
- ’Transformational’ projects are normally just a political mirage. The standard policy tools should be applied.
- In a small number of cases, the discount rate is endogenous to the project being undertaken. In such cases, the project should be discounted using the discount rate that will obtain if the project is undertaken. (Consider, for example, massive investments that substantially raise real yields.) Otherwise projects will be discounted using a rate that does not actually obtain.

  I would not otherwise adjust the social discount rate for ""transformational"" projects, for reasons that equally apply to the environmental context just asked about.
- I frankly don’t understand what this means.
- Honestly, it is hard for me to think through. I can see where this could be useful, as governments made mistakes with wind and solar investments, which go very cheap from learning by doing. However, maybe better models of cost declines would be a better way forward for these projects.
- Again, I think this could become terribly arbitrary and perhaps subject to special or connected interests being unfairly advantaged. I think one should take a consistent ethical framework.
- As before, if a project is transformational, let’s capture those bigger impacts in the appraisal through better conceptualisation, measurement, and valuation methods
- Transformational change is difficult to measure but is clearly significant especially with major long-term projects. The main issue is that of risk with such projects but the danger is that the risk is assessed in such a way that the project becomes impossible to justify. The adjustment is about the response to risk rather than avoiding measuring the actual change more accurately
- Should justify and value the ‘transformation" with evidence reason and explicit judgement before these values are discounted. Adjusting the discount rate will embed judgements which may be unaccountable to reason, evidence and social values
- Firstly I am sceptical about such terminology and likely biasses in using it. Secondly and anyway, I think the impacts should be represented in the benefit and cost flows, not the SDR.
- As the society change, also the social discount rate should be adjusted to reflect these changes, namely the changed societal preferences. I can’t discuss here how this would be done technically, but I can provide some insights during the panel discussion.
- Short of changing the numeraire I don’t know how you identify such changes in an apolitical way. The machining machine led to transformation change, but we did not realize that until after the fact. This is too prone to political monkey business. Furthermore, sudden is still slow by most standards. AI is a good example, it is changing things, but is the pace of change really that much more rapid than that of the car, train, or internet?
- I not might fully get the concept of "transformational change", but according to my understanding I do not think that it affects the reasoning to define a suitable social discount rate.
- Again identifying and attaching probabilities is important as uncertainty is likely larger and there is the standard problem of valuing non-marginal changes that may change relative prices
- I am tempted to say yes. But I don’t have clear guidance. For projects that are clearly non-marginal one could imagine doing a total CBA instead of a marginal. Another possibility – to distinguish from standard projects – is to allow the analysis more freedom within agreed upon frameworks (for example in the calculation of systematic risk).
- These transformational changes seem to suggest that the change is not marginal. In such a case, a marginal concept like the discount rate would be inappropriate. A full analysis based on social welfare evaluation would be necessary. However, if the subjects to be transformed are of minor relevance to the wider economy, the STPR approach can still be used if the monetized benefits seem appropriate (i.e. the way of monetization is in line with the transformational change).
- The value of transformational change should be evaluated and taken into consideration but trying to do so via a discount rate doesn’t make any sense. The risk and cash flows associated with transformation change will differ widely across projects thus characterized and the use of a SOC would correctly reflect those differences.
- It is not clear why that would be the case from theory and whether empirical evidence would suggest this. Whilst a project may be transformational, it is not always the case they are transformational to all people. We will also have issues defining what we mean as transformational and which projects would have a different rate applied. *[Government]*

- As with place based business cases, I'm not sure a differing discount rate makes any sense, if before, my recommendations for a lower discount rate and that it tends to 0, this should be sufficient for longer term projects. *[Government]*
- There is too much "optimism" bias in appraisals of transformational projects and we shouldn't add a new one *[Government]*
- I have answered 'no' as I don't think we would be very well placed to identify truly transformational projects. Might be appraisal fatigue, but most projects are described as transformational, but few rarely are. And the ones that are, we could probably not have identified them beforehand anyway. Apologies for the cynicism. *[Government]*
- These projects are very hard to measure, and the discount rate is only one small part of the challenge of effective appraisal of such projects. I do not think that messing with the discount rate will make this any easier.

  As discussed above, there may be other factors (e.g. risks of failure, or additional risk premia) that might apply to high-risk projects. *[Government]*
- Projects that are politically important to the government of the day get a mandated boost to their economic appraisal. *[Government]*
- If the Green Book describes this as qualitative, including it within the discount rate (when applied to costs and benefits) will make it quantitative. Again, transformational change, as defined above, should be accounted for explcitily in the options appraisal. Perhaps using a framework like this which is already used for ODA programmes - https://www.gov.uk/government/publications/international-climate-finance-kpi-15-methodology-extent-to-which-icf-intervention-is-likely-to-lead-to-transformational-change/international-climate-finance-kpi-15-methodology-extent-to-which-icf-intervention-is-likely-to-lead-to-transformational-change *[Government]*
- As I explained - the block on transformational projects is that each individual project assumes the underlying economics of the area stays the same when the aim is the fundamentally change them across all projects. The time it takes to achieve that still matters. Assuming that growth or productivity in that area is higher in the future would better reflect the way the program as a whole will affect the area. It also amounts to the same effect. *[Government]*
- - Very few portfolios of interventions have impacts on national growth trajectories. Therefore the UK discount rate is unlikely to be impacted by major elements of spend. Even where there is a potential – such as the impact of AI – the nature of impact is so uncertain it would be extraordinarily difficult to make a reasonable judgement on what the adjustment should be.

  - A better approach is to better assess potential impacts for example by modelling - different future scenarios - to understand potential risks and opportunities (and how they should be handled if they begin to emerge).

  - Longer term if any evidence emerges that the long term UK growth rate has changed this can be built into the discount rate. *[Government]*
- Given the current components of the discount rate, it's not clear why the discount rate for place-based/transformational changes should take a different value. *[Government]*

A.2.13. Anything else that the respondent wished to add:

- I find the current version of the Green Book a depressing artefact of the 1980s. There is no doubt we need monetary discipline, but this should be focused on inflation and not intertwined with the necessary investment activities of government. When we mix risk premia from the private sector, we end up with mismatches between the inherent risk of a project (from the perspective of government), which is very low and the valuation placed on that project. For PPP this becomes more prescient as that mismatch allows private sector investment to turn a paper difference into real returns. Which other countries do this? Costs are already inflated because of the way outsourcing works and stripping the ability of central and local government to organise projects, then we allow the private sector further gains from the financing side.
- Thank you for reviewing your policies. It shows the UK is keen to remain at the forefront of modern policymaking.
- In addition to parameter choices, greater emphasis could be placed on transparency and guidance around implementation, including clearer separation between discounting, risk adjustment, and valuation. Regular review of key assumptions in light of macroeconomic evidence, alongside practical examples and sensitivity analysis, would help ensure consistent application across departments while preserving flexibility for long-term and complex projects.
- Ensure that they are aware of the large range of areas where it is used and it can have perverse effects. For instance, the bulk of Ministry of Defence Annual Managed Expenditure is provisions for disposing of nuclear submarines. This went from a large deficit to a large surplus when the Treasury raised the discount rate.
- It is imperative not to let pressure from politicians dilute the 'size of the national cake' Hicks-Kaldor CBA test. That risks a circular process of CBA according to political objectives. There is obviously much more to democracy than CBA, and the CBA can always be ignored if a politician wishes to do so, but it should be a pure Hicks-Kaldor test and. of course, such a test has limitations and is just one element of any consideration for policies projects and programmes.

- The Ramsey and other frameworks evolved in an era where investment was atomistic and did not affect the environment (exogenous). Now we live in a world where the scale of human investment has made the environment an endogeneity to the investment system, yet our investment appraisal process has not changed. Public investment is supporting high rate of return projects that are not climate friendly and which increase climate risk. The same high project hurdles should not be used for projects that reduce cc risk.
- It seems to me the key component is the E of MU and it is a pity that the current value is not consistent with the evidence.
- The UK has a very poor record over the 40-50 years of investing for future generations, and this is particularly true for the regions. Methods are needed that reverse this lack of investment.
- 5 year gilts rate. See in US as the component in Risk Free Rate has changed from 30 year bonds to ten years.
- Where to start? I have read the background paper, and while it deals carefully with some issues it fails to discuss some of major importance. Two in particular: 1. Many of the arguments for given models of the discount rate rely on general equilibrium, one example being the precautionary effect. But the paper only examines the implications for expected returns, rather than other macro variables, most notably saving. If an argument implies lower returns but higher savings, simply taking the first implication but ignoring the second is inconsistent and potentially very risky. 2. There is no discussion at all of the way expected returns and discount rates feed into investment decisions in the regulated sector. If these were renationalised, would this change investment decisions? Would this be desirable?
- '- Acknowledge the role of discounting and why something is worth more today than in the future (time preference + growth) - Don't use the discount rate as a tool to value the future more/less if and when unrelated to the above points.
- Thank you both for your work. As a govt. economist it's great to see academics so passionate about the key questions in public policy. Thanks for your service to such an important issue.

  My biggest overall view on the DR is that consumption and GDP are the wrong measures for the wealth effect, but I do appreciate that more theoretically defensible measures lack robust quantitative measures. Nonetheless, I think it's an area that could be improved."
- I'm curious whose welfare HM treasury is trying to maximise when choosing discount rates
- Keeping weitzman issues in check consumption growth rates could be Power in long run
- This survey has frontloaded a lot of technical economic questions, which may limit responses to those with economic training. This is disappointing as many of these questions should be discussed in light of expertise from other disciplines including environmental science, public health, and ethics. The Marmot review is a good example of relevant evidence that should be considered - the social determinants of health - but this survey is unlikely to collect that type of evidence from other disciplines.
- A further consideration is the treatment of distributional change over time. The Ramsey Rule relies on a representative-agent framework, implicitly assuming stable proportionality of consumption or abstracting from distributional effects. Over long horizons this is unlikely to hold, and it may therefore be appropriate to account explicitly for the higher marginal social value of consumption growth that reduces inequality relative to growth that increases it, for example through distributional weighting, where credible projections of the distribution of future real consumption growth are available.
- no thanks. Lower the better for environmental and other non-financial returns. Let's stop the false assumption that we can continue to substitute nature, or that one day when we are richer, we can fix nature. Higher discount rates essentially support this assumption.
- Consider whether it is valid to continue to rely on OBR forecasts for g.
- I would encourage the Treasury to spend less time on reviewing the discount rate, important though it is, and more time on improving the accuracy of cost estimates for projects, which is the real make-or-break factor for many projects.
- The discount rate changes over time. Just accept it.
- please note all the recommendations in our paper here:

  in particular we recommend better understanding of increased forward looking uncertainty relative to uncertainty underpinning analysis of past bond rates, and we acknowledge that the discount rate itself does not operate in isolation so a strengthening of guidance and the evidence base considered alongside BCR calculations (for example through the calculation of financial durations) will be key to translating economic analysis into better decision making.
- You could consider whether the commercial (SOC) vs non-commercial (STPR) discount rates with mandatory sensitivity analysis, that the 2024 New Zealand update landed on, is relevant in the UK context. All the underlying material is publicly available, including expert papers https://www.treasury.govt.nz/information-and-services/public-sector-leadership/guidance/reporting-financial/discount-rates .

- Perhaps focus more on what determines whether a project is fundamentally a market initiative versus a government initiative. Financial appraisals would apply for the former, and social CBA for the latter. Ensure that initiatives that are more market than government are subject to the same market disciplines as private firms (ie, no step-up, soft-budget constraints etc). This is a better way to account for the issues that SOC tries to account for.
- Good luck!
- I would highly encourage you to review the U.S. review of the evidence from 2023, on a variety of these questions: pages 57-94 at https://bidenwhitehouse.archives.gov/wp-content/uploads/2023/11/CircularA-4Explanation.pdf.
- It may want to consider including the precautionary effects, as uncertainty is an important issue, in addition to project specific risk.
- It would be useful to clearly identify and define the issues at hand (e.g. environmental effects, transformational projects, place-based, etc) and lay out what the alternatives are to improve their appraisal, and then asses how the discount rate may help and why. I remain unconvinced it is the right/best instrument to address those issues
- Please embed as little as possible in the discount rate. Instead offer guidance on how projects with different characteristics, including systemic and other project specific risks can be valued at each point in the time stream before discounting. This will lead to more coherent valuation but also add to transparency and accountability. Also found discounting on a more nuanced view of where the opportunity cost fall rather than adopting a discounting policy which implicitly assumes away the reality of constraints and implicitly imposes a particular and disputed view of social welfare. Therefore, a working principle to put the very minimum into the discount rate so that the evaluation of policy make the assumption and judgement more explicit would be good for evidence based policy that also reflects social values. It would also be good for democracy and parliamentary scrutiny through the Committees."
- Yes, given the relatively strong arguments for an STP discount rate and the obvious constraints on the public budget, I think HMT must ask itself whether its long term view is that SOCEF (or MCPF) > 1 across the entire public budget. I detect a considerable reluctance to engage with this question, but I don't think it is avoidable.
- To complement the discount rate analysis with analysis on specific projects, based on case study and qualitative methods, to capture the specificities of projects, the beneficial (and negative) spillovers each can have for the economy and society, especially to assess their performance over the years (in a dynamic perspective)
- The Biden administration put a lot of thought into this questions. There was good and not so good outcomes. Learn from that and don't repeat mistakes that were made. Ultimately, this is aways about opportunity cost where it is a Ramsey based social discount rate or opportunity cost of social capital approach. The difference really should come tricky bits of Jensen's inequality and approximations along the way.
- '- shall the interactions of discounting with equity weights and more broadly distributional impacts be considered? - Shall discounting account for the scarcity of public funding?"
- Actually applying SCBA would be an impressive step forward and considering a wider set of options before accepting what engineers with no skin in the game want
- Why not moving to a stochastic discount factor approach?
- I applaud the HM Treasury to take on such a thorough review to the practise. I hope that they take the opportunity to also communicate their careful analysis to the wider public.
- The distortions caused by governments making up discount rates rather than using private sector rates causes serious distortions in public sector investment policies and has caused enormous efficiency losses to society. Unfortunately, since one cannot see the counterfactual, those distortions are usually invisible. It is time for governments to drop that outdated and distortionary practice and adopt modern valuation practices. The benefits to that change could be enormous.
- Commit to regular (but not too often) independent reviews which are able to include the latest and most robust evidence to any potential changes. Ensure the discount rate is founded as objectively as possible, not subject to policy approaches or ways of doing things. *[Government]*
- The assumption that per capita increases in both consumption and GDP more generally do not appear to hold in recent data. *[Government]*
- No thanks. I do however appreciate the opportunity to offer views, so I am grateful for that (useful as a CPD brush up too, as not in a role that requires use of the discount rate that often these days). *[Government]* "
- I think the points raised in the questions about using private discount rates are important - if HMG is acting like a private investor, it should use private investor discount rates and find other ways to assess the benefits that come from investment.

  If it isn't, then consistency and continuity of rate seems most appropriate. Perhaps a small reduction to reflect growth trends. *[Government]*

- We seem to do this every 10 years or so? Can we have a moratorium on it for 20 years and focus on the other stuff? The work to better reflect regional work is important and difficult and doesn't require a discount rate change. *[Government]*
- 1. Should a periodic review be carried out of the parameters of the STPR "every 15 years? 2. The lifetime of assets is an area where the general guidance is light and is something that could be developed further.

  3. Given that the majority of funding is through taxation the Green Book should be adjusting public sector spending to allow for the distortionary impacts of taxation on the economy. (see Spackman, 2025, working-paper-419-Spackman.pdf). Costs would be factored up by this to reflect the distortionary impact. The STPR would still need to be applied to discount all costs and benefits though. Finalising a value for the marginal social cost of taxation would be a significant piece of work. *[Government]*

B. Appendix to Chapter 8

This Chapter assembles estimates of the elasticity of marginal utility from the literature.

Inequality aversion and the elasticity of marginal utility. Table B.1 summarises evidence on the elasticity of marginal utility of income or consumption, interpreted either as an inequality-aversion parameter or as the curvature term entering the Ramsey wealth effect. These estimates are closest in spirit to the social parameter used in distributional weighting and social discounting. They should nevertheless be distinguished from private risk aversion and from the inverse of the elasticity of intertemporal substitution.

Table B.1.: Inequality-aversion and elasticity-of-marginal-utility estimates

Source / context	Estimate	Object being estimated	Interpretation for $\mu = 1.25$
Domestic revealed-ethics evidence, surveyed by Del Campo et al. (2024)	$h \simeq 1 - -2$; broader domestic support roughly $1 - -4$	Resource-space inequality aversion, often inferred from domestic tax or income-distribution choices	$\mu = 1.25$ sits in the lower-middle of the domestic range.
Foreign-aid and international-transfer evidence (Tol 2010, Del Campo et al. 2024)	Tol-style refined range $h = 0.385 - -0.554$; direct high-leakage foreign-aid estimate $h \simeq 0.7$	International inequality aversion inferred from costly or leaky transfers between countries	Below $\mu = 1.25$; useful evidence for separating domestic from foreign/global redistribution.
Axiomatic benchmarks in Del Campo et al. (2024)	$h > 0$, $h \geq 1$, or $h > 2$, depending on the equity principle	Normative restrictions implied by Pigou–Dalton, proportional-transfer, solidarity, or no-envy principles	Shows that higher values can be defended normatively, but this evidence does not identify a single empirical value.
UK Green Book, STPR discounting (?)	$\mu = 1.0$	Elasticity of marginal utility of consumption in the Ramsey wealth effect, $r = \rho + \mu g$	Slightly below $\mu = 1.25$.
UK Green Book, distributional welfare weights (HM Treasury 2026*c*)	1.3	Elasticity of marginal utility of income used in distributional weights	Very close to $\mu = 1.25$.
US OMB Circular A-4 (OMB 2023*a*)	1.4	Absolute value of the income elasticity of marginal utility for income-weighted benefit-cost analysis	Close to, and slightly above, $\mu = 1.25$.
Layard et al. (2008)	Range $1.19 - -1.34$; combined estimate 1.26	Marginal utility of income estimated using subjective wellbeing data	Almost exactly matches $\mu = 1.25$.
Acland & Greenberg (2023)	Central value 1.6; sensitivity range $1.2 - -2.0$	Meta-analytic estimate of EMUI for distributional weighting and social discounting	$\mu = 1.25$ is within the lower part of their recommended range.
Drupp et al. (2018)	Mean $\eta = 1.35$; median and mode 1.0	Expert-elicited wealth-effect parameter for the risk-free social discount rate	Places $\mu = 1.25$ near the centre of expert views.
Groom & Maddison (2019)	Fixed-effects pooled estimate 1.528 (95% CI : $1.443 - -1.613$); random-effects estimate 1.594 (95% CI : $1.362 - -1.827$)	UK elasticity of marginal utility from equal-sacrifice, Euler-equation, insurance, additive-preference, and wellbeing evidence	Higher than $\mu = 1.25$, but in the same broad policy range.
French / Gollier-style benchmark discussed in Freeman et al. (2018)	2.0	Normative or guideline benchmark used in social-discounting discussions	A useful high benchmark; materially above $\mu = 1.25$.

Notes: h, η, and ε are used differently across the literature. The common feature is that they describe the elasticity of marginal social value with respect to income or consumption. These estimates are therefore closer to the social curvature parameter used in welfare weighting and Ramsey discounting than the private risk-aversion estimates in Table B.3.

Elasticity of intertemporal substitution. Table B.2 reports estimates of the elasticity of intertemporal substitution, denoted ψ. The inverse $1/\psi$ is reported because, under CRRA and time-separable expected utility, $1/\psi$ corresponds to the curvature parameter. This mapping is useful for comparison but should not be treated as definitional: with Epstein–Zin preferences, habits, liquidity constraints, or nonseparabilities, risk aversion, intertemporal substitution, and social inequality aversion need not coincide.

Table B.2.: Estimates of the elasticity of intertemporal substitution

Study	Data / method	EIS, ψ	Implied $1/\psi$	Interpretation for $\mu = 1.25$
Blundell et al. (1994)	UK Family Expenditure Survey; repeated cross-sections and pseudo-panel Euler equations	Baseline absolute ISE about $0.75 - -0.77$; broader specifications roughly $0.55 - -1.21$	Baseline about 1.3; broader range $0.83 - -1.82$	Classic UK micro evidence is broadly consistent with μ around one to one-and-a-half.
Groom & Maddison (2019), Euler-equation approach	UK ONS and Bank of England quarterly data; updated UK Euler-equation estimate	OLS coefficient $0.631 - -0.632$; IV estimate 0.593	About $1.58 - -1.69$	A more recent UK time-series estimate implies somewhat more curvature than $\mu = 1.25$.
Best et al. (2020)	UK mortgage interest-rate notches; bunching around loan-to-value thresholds	Around 0.1	Around 10	Much higher implied curvature, but the identifying variation is specific to mortgage debt and refinancing margins. It should not be read as a general social μ.
Havránek (2015)	Meta-analysis of 2,735 EIS estimates from 169 published studies	Corrected micro estimate for asset holders around $0.3 - -0.4$; corrected macro mean near zero	About $2.5 - -3.3$ for the micro asset-holder estimates	Suggests caution about high EIS calibrations; the implied curvature is above $\mu = 1.25$.
Crump et al. (2022)	US subjective expectations; Survey of Consumer Expectations	About $0.5 - -0.8$, depending on specification	About $1.25 - -2.0$	Places $\mu = 1.25$ at the lower end of the implied curvature range.
Barsky et al. (1997)	Health and Retirement Study hypothetical choices over lifetime income paths	Mean EIS about 0.2	About 5	Useful mainly as evidence that EIS and risk tolerance are empirically distinct; less direct for social discounting.

Notes: The sign convention in some Euler-equation papers reports the intertemporal substitution elasticity as negative; the table reports absolute magnitudes for comparability. The column $1/\psi$ is not a separate estimate; it is included only to compare EIS evidence with a CRRA curvature parameter.

Risk-aversion evidence. Table B.3 reports published evidence on private risk aversion. These studies are useful for disciplining plausible curvature values, but they should not be treated as direct estimates of the social discounting parameter. They generally identify individual preferences over lotteries, labour supply, insurance choices, or hypothetical income risk, rather than society's attitude toward uncertain aggregate consumption.

Taken together, the tables suggest that $\mu = 1.25$ is a moderate choice. It is close to the UK distributional-weighting value of 1.3, the US A-4 value of 1.4, the Layard et al. estimate of 1.26, and the centre of expert-survey evidence. It is above most foreign-aid estimates and some private experimental risk-aversion estimates, but those are not the closest objects for setting a social intertemporal welfare parameter. Conversely, some EIS-based estimates imply substantially larger $1/\psi$, but those estimates identify private intertemporal substitution and need not be read as social inequality aversion or social risk aversion.

Table B.3.: Published evidence on private risk aversion

Study	Data / risk domain	Estimate or finding	Interpretation for social discounting
Holt & Laury (2002)	Laboratory lottery choices with real monetary incentives	CRRA estimates centred around $\gamma \simeq 0.3 - -0.5$	Good evidence against risk neutrality, but small-stakes lab risk is weak evidence for the social curvature of aggregate consumption.
Andersen et al. (2008)	Representative Danish field experiment jointly eliciting risk and time preferences	CRRA estimate about $\gamma \simeq 0.74$	More externally credible than student-lab evidence, but still identifies private monetary risk preferences.
Chetty (2006)	Labour-supply behaviour and the link between consumption and labour	Central benchmark around $\gamma \simeq 1$; values above 2 are hard to reconcile with labour-supply evidence	Strong discipline against very high private curvature, but still not a direct estimate of social risk aversion.
Barsky et al. (1997)	Health and Retirement Study hypothetical gambles over lifetime income	Mean risk tolerance about 0.25, implying rough relative risk aversion around 4	Closer to large-stakes income risk, but hypothetical and highly heterogeneous. Also shows that risk tolerance and EIS are essentially uncorrelated.
Dohmen et al. (2011)	Representative German survey and paid lottery validation	No single CRRA parameter; strong heterogeneity in willingness to take risks	Useful for external validity and heterogeneity, but not a scalar social discounting parameter.
Cohen & Einav (2007)	Auto-insurance deductible choices	Large and skewed heterogeneity; some implied relative-risk-aversion estimates are double-digit	Poor as a direct social-discounting input because deductible choices mix risk preferences, beliefs, liquidity, and insurance-market frictions.
Barseghyan et al. (2013)	Auto and home insurance choices	Probability distortions explain much of apparent insurance risk aversion	Important caution: some estimated “risk aversion” is not utility curvature at all.
Einav et al. (2012)	Employer-provided insurance and 401(k) choices across domains	Evidence of a domain-general component of risk preferences, but not a universal scalar parameter	Supports some stability in private risk attitudes, while also showing that risk preferences remain context-specific.

Notes: γ denotes private relative risk aversion. These estimates should be separated from the social parameter μ. In a CRRA expected-utility model, $\gamma = \mu = 1/\psi$, but that equality is an imposed structure rather than an empirical fact. For the social discount rate under uncertainty, the relevant object is society’s attitude toward uncertain aggregate consumption, together with the covariance of project benefits with aggregate consumption.

C. Appendix to Chapter 13

This appendix provides robustness checks and additional estimation detail for the results in Chapter 13. It includes raw recent-50 permanent-state calibrations, 30-year and full-sample robustness, diagnostics, regime estimates, data-boundary checks, Bai–Perron and Stock–Watson-style low-frequency robustness, and AR(1)-GARCH robustness.

C.1. Raw recent-sample permanent-state calibrations

C.1.1. Table C.1 reports the raw recent-sample permanent-state calibrations. These results are included as diagnostics because the unconstrained recent-sample regime calibration treats the 2020 pandemic contraction as a near-degenerate state.

C.2. Robustness: alternative time samples

C.2.1. Tables C.2 and C.3 report the corresponding discount-rate results for the last-30-year and full-sample windows. The tables follow the same model classification as the main text.

C.3. Diagnostics and model selection

C.3.1. Table C.4 reports serial-correlation, ARCH and information-criterion diagnostics for the main estimated models.

C.4. Regime-switching estimates

C.4.1. The regime-switching estimates are central to the paper because they are used in two different ways. First, they are estimated as recurrent growth-process models, in which states are temporary and the economy can move between them. Second, selected state estimates are used as empirical inputs to calibrate the permanent-parameter uncertainty model. This second use is deliberately more demanding: a temporary low-growth state in the Markov model becomes a possible permanent future growth environment in the parameter-uncertainty model.

C.4.2. The updated 2025 data introduce an important diagnostic issue. In the raw recent-50-year sample, the likelihood isolates the 2020 pandemic contraction as a near-degenerate low-growth regime. This is statistically understandable, but it is not economically appropriate to reinterpret that single-year pandemic state as a permanent future growth state. For this reason, Tables C.5 and C.6 report stable regime estimates for the recent-50-year and last-30-year windows excluding 2020, alongside the full-sample estimates. Table C.7 reports the raw recent-sample estimates to show the source of the problem.

Table C.7.: Raw recent-50-year regime estimates: diagnostic showing the pandemic state

Model	State	m (%)	σ (%)	v (%^{2})	π
2-state	1	-13.988	0.000	0.000	0.025
2-state	2	2.377	2.284	5.218	0.975
3-state	1	-13.988	0.000	0.000	0.023
3-state	2	0.904	1.704	2.903	0.619
3-state	3	4.306	1.404	1.972	0.358

Notes: This diagnostic table reports the raw regime estimates for the recent-50-year sample including 2020. The low-growth state has mean growth close to -14% and near-zero variance, indicating that the likelihood isolates the 2020 pandemic contraction as a near-degenerate state. This is why the raw discrete permanent-state calibration is not used as a main policy schedule.

Model	R_1	R_{30}	R_{75}	R_{125}	R_{300}	$R_1 - R_{300}$
Panel A. Growth-process models						
Gollier (2012) growth deviations: terminal vs estimated trend	1.438	2.845	2.878	2.887	2.894	-1.456
IID Gaussian growth	2.780	2.780	2.780	2.780	2.780	0.000
AR(1) growth	2.812	2.791	2.790	2.790	2.790	0.022
Two-state Markov regime model	2.947	2.951	2.951	2.951	2.951	-0.003
Three-state Markov regime model	2.720	2.684	2.682	2.682	2.682	0.039
Panel B. Permanent parameter uncertainty: alternative UK calibrations						
Regime-informed two-state permanent uncertainty	2.947	-5.650	-12.786	-14.745	-16.460	19.407
Regime-informed continuous truncated permanent uncertainty	2.637	2.366	2.033	1.773	1.310	1.327
Regime-informed three-state permanent uncertainty	2.720	-5.382	-12.643	-14.659	-16.424	19.144
Bai–Perron break-informed permanent uncertainty	2.780	2.780	2.780	2.780	2.780	0.000
Stock–Watson-style low-frequency block calibration	2.779	2.406	1.882	1.505	1.053	1.726
Panel C. Disaster-risk models						
Constant rare-disaster benchmark, global calibration	1.074	1.074	1.074	1.074	1.074	0.000
Time-varying disaster probability	2.569	2.573	2.573	2.573	2.573	-0.005

Table C.1.: Raw recent-50 results including direct discrete regime-informed permanent-state calibrations (spot rates, percent).

Model	R_1	R_{30}	R_{75}	R_{125}	R_{300}	$R_1 - R_{300}$
Panel A. Growth-process models						
Gollier (2012) growth deviations: terminal vs estimated trend	1.438	2.095	2.109	2.112	2.115	-0.677
IID Gaussian growth	2.099	2.099	2.099	2.099	2.099	0.000
AR(1) growth	2.003	2.006	2.006	2.006	2.006	-0.003
Two-state Markov regime model	2.082	1.991	1.983	1.981	1.979	0.102
Three-state Markov regime model	2.443	2.181	2.131	2.117	2.104	0.339
Panel B. Permanent parameter uncertainty: alternative UK calibrations						
Regime-informed two-state permanent uncertainty	2.082	1.505	0.801	0.388	-0.044	2.126
Regime-informed continuous truncated permanent uncertainty	2.575	2.365	2.077	1.823	1.320	1.255
Regime-informed three-state permanent uncertainty	2.443	1.608	0.768	0.332	-0.115	2.558
Bai–Perron break-informed permanent uncertainty	2.099	2.099	2.099	2.099	2.099	0.000
Stock–Watson-style low-frequency block calibration	2.099	2.009	1.876	1.752	1.510	0.588
Panel C. Disaster-risk models						
Constant rare-disaster benchmark, global calibration	0.393	0.393	0.393	0.393	0.393	0.000
Time-varying disaster probability	0.841	0.645	0.629	0.624	0.620	0.221

Table C.2.: Appendix results: last 30 years, 1996–2025 (spot rates, percent).

C.5. Data-boundary diagnostic

C.5.1. The updated GB+NI data are checked around the post-WWI period. The large 1921 fall in per capita consumption is driven by the consumption numerator rather than by a population-denominator discontinuity. The 1923 observation is no longer the stark population break seen in earlier whole-UK data.

Table C.8.: Data-boundary diagnostic around the post-WWI period

Year	Consumption	Population	$\Delta \ln C$	$\Delta \ln N$	$\Delta \ln(C/N)$
1916	187822	43.45	-8.669	0.392	-9.061
1917	173123	43.56	-8.149	0.253	-8.402
1918	171490	43.52	-0.948	-0.092	-0.856
1919	196270	43.44	13.497	-0.184	13.681
1920	196721	43.73	0.230	0.665	-0.436
1921	180543	43.96	-8.582	0.525	-9.106
1922	188602	44.39	4.367	0.973	3.394
1923	194375	44.63	3.015	0.539	2.476
1924	200810	44.93	3.257	0.670	2.587

C.6. Bai–Perron and low-frequency block robustness

C.6.1. Tables C.9, C.10 and C.11 report the structural-break and low-frequency block calibrations used to construct alternative permanent-parameter uncertainty distributions.

Model	R_1	R_{30}	R_{75}	R_{125}	R_{300}	$R_1 - R_{300}$
Panel A. Growth-process models						
Gollier (2012) growth deviations: terminal vs estimated trend	1.438	2.016	2.031	2.035	2.038	-0.600
IID Gaussian growth	1.960	1.960	1.960	1.960	1.960	0.000
AR(1) growth	1.962	1.927	1.926	1.926	1.926	0.037
Two-state Markov regime model	1.961	1.938	1.935	1.934	1.934	0.027
Three-state Markov regime model	1.951	1.893	1.888	1.886	1.885	0.066
Panel B. Permanent parameter uncertainty: alternative UK calibrations						
Regime-informed two-state permanent uncertainty	1.961	1.817	1.475	1.015	0.126	1.834
Regime-informed continuous truncated permanent uncertainty	2.076	1.941	1.758	1.592	1.228	0.848
Regime-informed three-state permanent uncertainty	1.951	1.650	1.234	0.788	-0.080	2.032
Bai–Perron break-informed permanent uncertainty	1.960	1.819	1.639	1.498	1.287	0.673
Stock–Watson-style low-frequency block calibration	1.959	1.724	1.447	1.243	0.898	1.061
Panel C. Disaster-risk models						
Constant rare-disaster benchmark, global calibration	0.254	0.254	0.254	0.254	0.254	0.000
Time-varying disaster probability	1.084	0.746	0.702	0.690	0.679	0.405

Table C.3.: Appendix results: full sample, 1831–2025 (spot rates, percent).

Table C.4.: Diagnostics and model-selection statistics

Sample	IID LB(5)	AR1 LB(5)	IID ARCH	AR1 ARCH	2-state AIC	2-state BIC	3-state AIC	3-state BIC
Recent 50 years	0.733	0.799	0.997	0.966	187.890	199.362	–	–
Last 30 years	0.823	0.837	0.996	0.996	153.513	161.921	151.652	168.466
Full sample	0.095	0.415	0.251	0.252	906.566	926.204	902.731	942.007

Notes: LB(5) denotes the Ljung–Box test for serial correlation up to lag 5. ARCH denotes the ARCH-LM test for conditional heteroskedasticity. The information criteria are reported for the two-state and three-state Markov regime-switching models. Dashes indicate cases where the likelihood did not converge reliably.

Table C.9.: Bai–Perron selected break regimes

Sample	Segment	Years	N	m (%)	v (%^{2})	Weight
Recent 50 years	1	1976–2025	50	1.820	10.239	1.000
Last 30 years	1	1996–2025	30	1.317	13.271	1.000
Full sample	1	1831–1943	113	0.524	8.943	0.579
Full sample	2	1944–2025	82	2.091	8.262	0.421

Table C.10.: Bai–Perron model selection

Sample	Breaks	Segments	SSR	BIC
Recent 50 years	0	1	511.956	120.223
Recent 50 years	1	2	442.827	120.794
Recent 50 years	2	3	441.797	128.501
Last 30 years	0	1	398.133	80.969
Last 30 years	1	2	345.518	83.519
Last 30 years	2	3	345.439	90.315
Full sample	0	1	1804.792	439.187
Full sample	1	2	1688.062	436.695
Full sample	2	3	1610.604	438.081
Full sample	3	4	1589.032	445.998
Full sample	4	5	1573.347	454.609
Full sample	5	6	1565.107	464.131
Full sample	6	7	1556.829	473.643

Table C.11.: Stock–Watson-style low-frequency block robustness, full sample (spot rates, percent)

Block length	R_1	R_{30}	R_{75}	R_{125}	R_{300}	$R_1 - R_{300}$
10-year blocks	1.959	1.695	1.350	1.061	0.507	1.452
16-year blocks	1.959	1.724	1.447	1.243	0.898	1.061
25-year blocks	1.959	1.787	1.591	1.447	1.206	0.754

Table C.5.: Regime-switching estimates: state means and variances

Parameter	Recent 50 years excl. 2020	Last 30 years excl. 2020	Full sample
Panel A. Two-state model			
Low-growth state: m_1 (%)	0.981	1.723	-0.516
s.e.	(0.498)	(0.647)	(1.733)
Low-growth state: v_1 (%^{2})	2.983	7.923	53.771
s.e.	(0.912)	(2.720)	(19.104)
High-growth state: m_2 (%)	4.409	2.083	1.383
s.e.	(0.759)	(0.224)	(0.154)
High-growth state: v_2 (%^{2})	1.887	0.281	3.634
s.e.	(1.043)	(0.167)	(0.436)
Panel B. Three-state model			
Low-growth state: m_1 (%)	0.124	0.810	-0.673
s.e.	(0.504)	(0.827)	(1.622)
Low-growth state: v_1 (%^{2})	2.542	8.728	54.454
s.e.	(1.000)	(3.578)	(18.773)
Middle-growth state: m_2 (%)	2.212	2.086	1.022
s.e.	(0.224)	(0.203)	(0.050)
Middle-growth state: v_2 (%^{2})	0.317	0.273	2.867
s.e.	(0.191)	(0.146)	(0.389)
High-growth state: m_3 (%)	4.592	3.862	3.997
s.e.	(0.412)	(0.270)	
High-growth state: v_3 (%^{2})	1.571	0.295	1.475
s.e.	(0.713)	(0.209)	(0.372)

Notes: States are ordered by estimated mean growth. The table reports the state mean m_i and variance $v_i = \sigma_i^2$ of annual per-capita consumption growth, in percent and percentage-squared units respectively. Standard errors are reported in parentheses. The recent-50-year and last-30-year columns exclude 2020 to avoid treating the pandemic collapse as a separate near-degenerate regime.

C.7. AR(1)-GARCH robustness

C.7.1. The AR(1)-GARCH robustness exercise separates persistence in expected growth from persistence in conditional volatility. The results do not overturn the main conclusion that ordinary historical persistence generates little DDR decline under the new parameterisation.

Table C.6.: Regime-switching transition probabilities and stationary probabilities

Probability	Recent 50 years excl. 2020	Last 30 years excl. 2020	Full sample
Panel A. Two-state model			
p_{11}	0.865	0.881	0.784
s.e.	(0.080)	(0.093)	(0.125)
p_{12}	0.135	0.119	0.216
s.e.	(0.080)	(0.093)	(0.125)
p_{21}	0.285	0.269	0.025
s.e.	(0.145)	(0.164)	(0.016)
p_{22}	0.715	0.731	0.975
s.e.	(0.145)	(0.164)	(0.016)
Stationary probability π_1	0.679	0.693	0.105
Stationary probability π_2	0.321	0.307	0.895
Panel B. Three-state model			
p_{11}	0.699	0.854	0.788
s.e.	(0.142)	(0.115)	
p_{12}	0.137	0.097	0.211
s.e.	(0.115)	(0.099)	
p_{13}	0.164	0.048	0.000
s.e.	(0.183)	(0.152)	
p_{21}	0.233	0.247	0.027
s.e.	(0.160)	(0.001)	(0.035)
p_{22}	0.690	0.753	0.933
s.e.	(0.154)	(0.002)	(0.028)
p_{23}	0.077	0.000	0.040
s.e.	(0.222)	(0.002)	(0.021)
p_{31}	0.216	0.000	0.000
s.e.	(0.116)		
p_{32}	0.070	0.068	0.285
s.e.	(0.081)	(0.096)	
p_{33}	0.714	0.932	0.715
s.e.	(0.141)	(0.096)	
Stationary probability π_1	0.426	0.434	0.102
Stationary probability π_2	0.259	0.256	0.788
Stationary probability π_3	0.314	0.309	0.110

Notes: Transition probabilities are $p_{ij} = \Pr(s_t = j \mid s_{t-1} = i)$. Approximate standard errors are reported in parentheses where available. For the final probability in each transition row, standard errors are approximated from the reported transition-probability standard errors because the probabilities in each row sum to one. Stationary probabilities are implied by the estimated transition matrix.

Table C.12.: AR(1)-GARCH(1,1) robustness under $\delta = 0.5\%$, $\mu = 1.3$

Sample	N	ϕ	ω	α	β	$\alpha + \beta$	R_1	R_{300}	$R_1 - R_{300}$
Recent 50 years	50	0.318	8.649	0.140	0.000	0.140	2.256	2.557	-0.301
Last 30 years	30	0.203	11.877	0.126	0.000	0.126	1.599	1.636	-0.037
Full sample	195	0.228	3.702	0.139	0.446	0.584	1.889	2.010	-0.122

D. Appendix to Chapter 14

The main text in Chapter 14 treats a counter-cyclical project as having payoff 1 in the non-disaster state and $1 + X$ in the disaster state. This appendix reports the alternative normalisation in which the counter-cyclical project has payoff $1 - X$ in the non-disaster state and 1 in the disaster state. The pro-cyclical payoff is unchanged. This alternative normalisation corresponds to a project that reallocates payoff from normal states into disaster states rather than adding extra payoff in disaster states.

Table D.1.: Appendix: alternative payoff normalisation, project-risk premia and rates, $\mu = 1.25$

Model	X (%)	Δr^D	Pro premium	Counter premium	Pro rate	Counter rate
Barro point	20	-1.632	0.178	-0.220	0.921	0.523
Barro point	50	-1.632	0.447	-0.866	1.190	-0.123
Barro point	80	-1.632	0.720	-3.261	1.463	-2.518
P-W Full sample	20	-0.742	0.076	-0.094	1.709	1.539
P-W Full sample	50	-0.742	0.191	-0.367	1.824	1.266
P-W Full sample	80	-0.742	0.309	-1.359	1.941	0.273
P-W Last 100 years	20	-0.768	0.080	-0.100	1.687	1.507
P-W Last 100 years	50	-0.768	0.203	-0.392	1.809	1.215
P-W Last 100 years	80	-0.768	0.326	-1.473	1.933	0.134
P-W Last 50 years	20	-0.747	0.078	-0.097	1.706	1.531
P-W Last 50 years	50	-0.747	0.197	-0.381	1.824	1.247
P-W Last 50 years	80	-0.747	0.317	-1.430	1.944	0.197

Notes: This appendix table reports the alternative payoff normalisation in which a counter-cyclical project has payoff $1 - X$ in the non-disaster state and 1 in the disaster state. The pro-cyclical payoff is unchanged.

Table D.2.: Appendix: alternative payoff normalisation, decomposition and project premia, $\mu = 1.25$

Model	λ (%)	Calibration	Mean-loss	Prudence	Δr^D	Pro $X = 50\%$	Pro $X = 80\%$	Counter $X = 50\%$	Counter $X = 80\%$
Global Barro	1.700	b=29.0%	-0.728	-0.904	-1.632	0.447	0.720	-0.866	-3.261
P-W Full sample	2.564	α=16.527	-0.354	-0.388	-0.742	0.191	0.309	-0.367	-1.359
P-W Last 100 years	2.000	α=10.705	-0.359	-0.410	-0.768	0.203	0.326	-0.392	-1.473
P-W Last 50 years	2.000	α=11.126	-0.350	-0.398	-0.747	0.197	0.317	-0.381	-1.430

E. Environmental Discount Rate Review

In this appendix, we provide the informal report to HM Treasury that underpinned the conclusions of the 2022 Environmental Discount Rate Review. This was written by Ben Groom but has not previously been published. We have edited this document to be in keeping with the style of the rest of this document and to maintain anonymity. It is a summary of meetings with experts organised by HM Treasury and chaired by Professor Groom.

E.1. Background

E.1.1. It should be noted that there is scope in the Green Book and in most CBA guidelines for consideration of relative price changes in the appraisal of projects. The Green Book (HMT 2020) notes this in passing on p 50 (para 5.14), the Appendix, and the term "relative price effect" is also in the glossary. So in a sense the Green Book allows for relative price effects already, and so there is a precedent for thinking about such effects in relation to the Environment. **Greater emphasis may be needed in the guidelines for the importance of such effects in the case of Environmental quality and environmental amenities.**

E.1.2. For the environment/environmental quality the arguments are organised around three key issues:

- **Non-marketed good**: The non-marketed nature of environmental goods and services, hence the need for a careful and separate analysis of the shadow price and how it evolves over time. It is not always possible to look at market prices and see the historical trend, as recommended in the Green Book para 5.14;
- **Environmental scarcity**: There is a structural reason why the shadow price may vary over time due to environmental degradation and associated physical scarcity of natural resources;
- **Non-substitutability**: Environmental resources may not have close substitutes. As Krutilla (1967) puts it: "While the supply of fabricated goods and commercial services may be capable of continuous expansion for a given resource base by reason of scientific discovery and mastery of technique, the supply of natural phenomena is virtually inelastic." The effect of environmental scarcity and non-substitutability will be exacerbated with incomes if demandsincrease with incomes and income is growing.

E.1.3. However, there is an anchoring on the 1.5% discount rate in the positioning paper which I think is misplaced. The reason, according to the Green Book for ignoring the wealth effect in relation to health is that the main measure of "health" is QALYs. It is widely understood that QALYs are measured in units of *utility* (on a scale from zero to 1) rather than in terms of *consumption.* For this reason using a consumption discount rate is inappropriate. Utility is linearly related to welfare in the utility function, whereas consumption is non-linearly related to welfare via the conversion to utility.

E.1.4. Paragraph A6.16 of the Green Book (HMT 2020) states:

> "*The recommended rate for risk to health and life values is 1.5%. This is because the 'wealth effect' or real per capital consumption growth element of the discount rate is excluded. As set out in Annex 2 [sic, actually Annex 1], health and life effects are expressed using welfare or utility values such as Quality Adjusted Life Years (QALYs), as opposed to monetary values.*"

E.1.5. There is no perfectly offsetting wealth effect such that values of health (in consumption terms) increase precisely at the rate of income growth and are not subject to the wealth effect. It is purely that the units of measurement are utils not consumption.

E.1.6. Nevertheless, in relation to health, even this argumentation is theoretically questionable particularly at the point at which QALYs are converted into cash values (£30k is the typical value these days I understand, but that is another issue altogether).

E.1.7. In relation to the environment, it is possible, I suppose, that a similar argument could be used. However, I find it unlikely since typical environmental valuation methods measure marginal values in terms of consumption and monetary units. Converting the 3 rationales for relative prices changes into an operational measure is a matter of measurement and measurement of preferences and scarcity, rather than estimating changes in utility directly. This is what various papers have attempted to articulate (e.g. Venmans and Groom, 2019; Baumgartner et al. 2015; Drupp 2016).

E.1.8. To see that the health context is not necessarily appropriate to the environmental one, consider the following formal analysis and the odd set of preferences that are required to obtain the perfectly offsetting effect on the wealth effect.

E.2. Relative prices and uprating: Brief formal analysis

E.2.1. Hoel and Sterner (2007) show that when we treat C and E separately in the utility function: $U(C_t, E_t)$ the rates of change of the respective shadow prices, i.e. SDRs, for each argument are given respectively by:

$$\begin{aligned} SDR_C(t) &= \rho + \gamma_{CC} g_C + \gamma_{EC} g_E \\ SDR_E(t) &= \rho + \gamma_{EE} g_E + \gamma_{CE} g_C \end{aligned}$$

Where $\gamma_{ij} = -jU_{ji}/U_i$ for $i = C, E$; that is these are the equivalent elasticities and cross-elasticities of marginal utility, as opposed to the typical framework in which E (amenity value) is assumed to be priced in terms of consumption, placed within C and is not treated differently from any other consumption good: $SDR = \rho + \gamma g$. Each SDR measures the rate of change in the shadow price of the respective commodity.

E.2.2. **Note that** the commodity E is a kind of index of all environmental amenities, in the same way that C is an index of all possible items in the consumption bundle that makes up aggregate consumption. With consumption the metric that allows use to add all things together is money. With E we are in principle thinking about physical units, and so the metric that allows adding together of air quality with non-use values (apples and oranges) is not entirely clear. In the end, some valuation is required in a common metric, usually cash or, as in Venmans and Groom (2020), the preference parameters are measured directly and the metric is essentially utility. In any event, the treatment of E as a composite index of the "environment" is quite heroic. This index/aggregation of environmental quality issue will be returned to below.

E.2.3. Returning to the relative price issue: importantly, when values are placed in terms of the common numeraire for CBA, consumption C (cash), the marginal willingness to pay for a unit of the environment E (its price) is $\frac{U_E}{U_C}$. Denoting the rate of change of the price of E compared to the numeraire C as $\left(\frac{d}{dt}\frac{U_E}{U_C}\right)/\frac{U_E}{U_C}$, simple algebra shows that this relative price change, $\Delta RP_{EC}(t)$, is given by:

$$\Delta RP_{EC}(t) = (\gamma_{CC} g_C + \gamma_{EC} g_E) - (\gamma_{EE} g_E + \gamma_{CE} g_C)$$

E.2.4. But this is just the difference between the two discount rates in paragraph E.2.1 (Weikard and Zhu, 2005). Dual discounting and relative price adjustments are equivalent.

E.2.5. The offset equation in paragraph E.2.3 tells us what the relative price effect for "environment" depends on in principle, and how it can be estimated. As discussed it depends on: (i) non-marketability; (ii) increasing scarcity: reflected by g_E; and, (iii) substitutability: reflected by the social preference parameters γ_{ij}.

E.2.6. To understand the intuition behind the expression in paragraph E.2.3, we present some simple examples at the end of this appendix chapter.

E.3. How to view the Green Book Health Discount Rate

E.3.a. Social preferences for health

E.3.1. Suppose that health benefits are not measured in terms of QALYs but in terms of a particular type of benefit such as reduced sick days. We can then apply the framework above and understand how, meaning what form of social preferences, one arrives at just using the pure rate of time preference as the discount rate for health.

E.3.2. If utility takes the form: $U(C_t, E_t) = U(C_t) + \theta E_t$ is linear in Health (e.g. sick days)[1] then the SDR for consumption is the typical Ramsey Rule: $SDR_C(t) = \rho + \gamma_{CC} g_C$, the SDR for health: $SDR_E(t) = \rho$. In terms of relative prices, one would make sure that the relative prices of health increased at the rate $\Delta RP_{EC}(t) = \gamma_{CC} g_C$ (the difference between the two discount rates) in the numerator, leaving an effective discount rate of $SDR_C(t) - \Delta RP_{EC}(t) = \rho$ for health related benefits and costs.

E.3.3. With this interpretation, the substitution of health for consumption is not perfect, the marginal utility of consumption decreases with consumption, and the relative price increases and perfectly offsets the wealth effect. Whether or not the relative price growth term is in reality equal to $\gamma_{CC} g_C$ is an empirical question.

E.3.4. It is not immediately obvious that this is an appropriate form of social preference for the environment, in fact, Venmans and Groom (2020) provide evidence that this is not the case for a variety of different environmental amenities. See also Drupp et al for the idea that there are subsistence limits to the levels of amenity that society should have, just as there are subsistence limits to certain elements of consumption in principle (Drupp 2018).

E.3.b. A Normative position on health?

E.3.5. Another approach is to simply argue that health ought not to be treated differently across different income groups. Looked at via a social welfare function, the normative stance is to posit a welfare function of the form $U(C_t, E_t) = U(C_t) + \theta E_t$, as above, which is linear in health. On reflection this seems like a peculiar normative position since it suggests that in relation to health there is no inequality aversion: we do not prefer to give a unit of health to someone with low health compared to someone with high health. Not only is it a peculiar normative stance, it is also refuted by empirical evidence (Dolan and Tsuchiya 2011). One issue that also arises here is that intra- and inter-temporal inequality aversion are not necessarily the same. High inequality aversion intra-temporally leads to a high SDR, yet there is evidence that people feel differently about inequality in these different dimensions (Venmans and Groom 2020; Emmerling et al. 2017)

E.3.6. So ultimately, for health we are left with the Green Book rationale, which is that QALYs are measured in utils not consumption, and so should be discounted with the utility discount rate not the consumption discount rate. There is no QALY for environment, so this argument does not hold.

E.4. *E* in aggregate versus its components

E.4.1. By now it is probably clear that there is an important separation to be made between the individual components of amenity value from the environment, and their aggregate E that appears in the discount rate. The analogy with consumption as a whole and its individual components should also be clear. Making this separation shows that there is a potential role for dual discounting, but only via the aggregate effect. All individual components should be treated in the numerator as a pricing effect to maintain the clean separation of discounting and price effects.

E.4.2. The overall purpose of the discount rate is to take a position on how society as a whole will value a unit of consumption at some point in the future. The SDR tells use how that value or shadow price changes over time (the rate of change). In addition to the pure rate of time preference (impatience), the SDR does this by capturing our best guess of what the future state of the world will be in the simple terms of whether we are richer or poorer. This is converted into a marginal utility effect using the elasticity of marginal utility with respect to aggregate consumption, which is the numeraire.

E.4.3. The relative price effect recommended by the Green Book picks out particular aspects of aggregate consumption (travel time, apples, particular goods and services) and treats them differently in terms of their price relative to the numeraire. The value of addition consumption (as a whole) changes at rate X, a particular component Z becomes relatively more valuable in terms of the numeraire at a rate Y, so we makes sure that the price of Z increases at the rate Y and the *effective* discount rate on Z is $X - Y$.

E.4.4. The same logic should apply to the "environment" E. E is an index of all amenities. In the same spirit of reflecting the future state of the world, the consumption discount rate should reflect the future state of the

[1] This is one possibility, but there are other ways in which the terms of the SDR for consumption could be perfectly offset by the relative price effect. E.g. if in the equation in paragraph E.2.3, it was the case that $(\gamma_{EE} g_E + \gamma_{CE} g_C) = 0$. There are other possibilities. The point is that the terms cancel in some way.

world both in terms of future aggregate consumption, C, and future aggregate amenity, E. If our valuation of consumption is affected by the level of aggregate *amenity*: C and E could be aggregate complements or substitutes for instance, then this should be reflected in the SDR for the numeraire, which is aggregate consumption. That is, the SDR for consumption should be:

$$SDR_C(t) = \rho + \gamma_{CC} g_C + \gamma_{EC} g_E$$

E.4.5. This is a legitimate adjustment to the denominator of the Net Present Value calculation: the discount rate for consumption, if we think that aggregate amenity E can be treated separately in utility to aggregate consumption. This would be true if we expect the state of the world to evolve differently: $g_C \neq g_E$, and/or the elasticity of amenity differs from that of consumption: $\gamma_{CC} \neq \gamma_{EC}$.

E.4.6. Then, if, in parallel to the relative price effects for consumption, an *individual component* of E has an increasing relative price compared to the numeraire, C, that is not captured by the aggregate adjustment due to E, there is a further argument for adjustment of the numerator: a relative price adjustment.

E.4.7. The details of all this need to be worked out properly going forward. E.g. How sensible is it, and what are the conditions for aggregate E to be separated out from C? But for now the discussion in the previous paragraphs is supposed to illustrate the distinction between the job that the SDR is doing for the numeraire aggregate consumption: reflecting the aggregate state of the world in the future, and what relative price adjustments are supposed to capture, which is reflect scarcity in individual components of the aggregate economy (components of E and C) that are relevant to a particular project.

E.4.8. A good example of this is the department of transport's treatment of the value of time. They have an "uprate" to this individual component in their appraisal. It is not an augmentation of the discount rate though (even though it is sometimes represented as this).

E.4.a. Catastrophic risk and the environment

E.4.9. The way in which catastrophic risk is treated, and whether this is a numerator (pricing/benefit) or a denominator (discounting) effect in CBA, is an interesting question. Perhaps some principles are:

1) Catastrophic risk should be embodied in the discount rate if:
    a. The catastrophic risk refers to aggregate C, or aggregate E, or indeed the exogenous existential risk, e.g. societal collapse (which would remove the people who obtain utility, and hence benefits altogether).
    b. The risks are correlated with growth of the numeraire: aggregate consumption.
2) Catastrophic risk is a pricing issue if:
    a. There are identifiable elements of E that may be subject to tipping points and irreversibilities;
    b. There are specific subsistence amenity requirements discussed in Drupp (2018).

Overall, the separation between aggregates and components is important.

E.4.10. The treatment of catastrophic risk in the SDR in the current Green Book is not very clear and quite inconsistent. In the 1% catastrophic risk element that appears in the components of the pure rate of time preference, appear several elements: societal collapse, obsolescence, mortality risk death, each of which deserve a separate treatment and some of which do not reflect the existential risk element that can be the only justification for adjusting the pure rate of time preference this way.

E.4.11. I see that one of the experts has provided extensive comments on the recent work in this area, in relation to risk and the environment, and in relation to the subsistence elements of the environment. I happen to have read that contribution and I agree with those points. Another expert also makes some points about risk which are repeated here.

E.4.b. Conclusion

E.4.12. The Treasury initiative to investigate dual environmental discounting is an important attempt to internalise the influential academic literature on dual discounting into the CBA guidelines. The problem with having specific discount rates for the "environment" is that the "environment" is not well defined. Different amenities have different relative price trajectories and hence are probably not captured well by a single "environmental" discount rate. However, there are certainly some special aspects of environmental amenities and environment per se which should be specifically addressed, with a view to the changes over time in the relative values / shadow prices of these aspects. Issues of subsistence, tipping points and irreversibility should be considered in the shadow pricing of particular projects. A good outcome of the report will be if these **intertemporal** aspects of environmental CBA are better addressed. More studies are required of environmental values, and how they change with the economic fundamentals. In the same way as in transport there are numerous studies of the value of time (a key ingredient to the CBA of transport projects) the same is required for different dimensions of the environment. Building up a social demand system for environmental amenities, in the same way as we have a demand system for consumption goods, would be a key input into the appropriate social valuation of environmental amenity as a whole, and its individual components.

E.4.13. This work should dovetail with the *policies* or intentions that are in place for the environment in the UK. For instance, the environment bill and the 2011 white paper all talk of no net loss or net gain for biodiversity. In the framework above, if E is biodiversity, then that means that $g_E = 0$ and the consumption discount rate becomes:

$$SDR_C(t) = \rho + \gamma_{CC} g_C$$

nd the relative price effect becomes:

$$\Delta RP_{EC}(t) = (\gamma_{CC} g_C) - (\gamma_{CE} g_C)$$

which means that the effective discount rate for the environment becomes:

$$SDR_C(t) = \rho + \gamma_{CC} g_C - \Delta RP_{EC}(t) = \rho + \gamma_{CE} g_C$$

which cancels the wealth effect, but adds $\gamma_{CE} g_C$. This makes sense here because there is a national target for an aggregated "environment" target: biodiversity. Nevertheless, defining this aggregated index, and its social preference parameter γ_{CE} remains a challenge as discussed in Venmans and Groom (2020).

E.4.14. The other policy that any action on environmental discounting needs to be mindful of is the treatment of carbon pricing in the Green Book and the Climate Act. Here a net zero target is reflected by target compatible prices for carbon. The relative price effects here are already built into the analysis. Whether there are any further residual catastrophic risks that need to be built into the analysis is a good question. Like deep recession in general, it is quite likely that catastrophic risk remains an issue in general since it depends on global action, not just that associated with the UK. These are all issues that are in need of further thought. Another good thing that could come out of the review is a return to the definitions contained in the Green Book SDR around existential risk, obsolescence and project risk. These aspects are treated inconsistently in the Green Book at present and of course relate strongly to the environment and climate disasters.

E.5. Simple Examples of Relative Price Effects

E.5.a. Example 1: Cross elasticities are zero.

E.5.1. In this case the utility function is additively separable $U(C_t, E_t) = U(C_t) + V(E_t)$, hence the marginal utility of C does not depend on E or vice versa. Changing relative prices are now given by:

$$\Delta RP_{EC}(t) = \gamma_{CC} g_C - \gamma_{EE} g_E$$

which is the difference between two "wealth effects" and depends on the relative growth of each commodity and how growth affects marginal utility. Practically speaking, to discount E we would ue the **suitable** consumption rate of discount and ensure that the relative prices increased at the rate $\gamma_{CC} g_C - \gamma_{EE} g_E$ in the numerator. This would leave an effective discount rate of $SDR_E(t) = \rho + \gamma_{EE} g_E$ for E. Calibration of this discount rate requires estimating the growth of E and γ_{EE}; e.g., as in Venmans and Groom (2020).

E.5.b. Example 2: CES Utility.

E.5.2. If

$$U(C_t, E_t) = \frac{1}{1-\gamma}\left[\alpha C^{1-\frac{1}{\varphi}} + (1-\alpha) E^{1-\frac{1}{\varphi}}\right]^{\frac{\varphi(1-\gamma)}{\varphi-1}}$$

then the change in relative prices becomes the following:

$$\Delta RP_{EC}(t) = \frac{1}{\varphi}(g_C - g_E)$$

where φ is the elasticity of substitution between C and E, measuring how easy it is to compensate a loss in E with a gain in C. This illustrates clearly the importance of substitutability. If φ is small (large) then, for a given difference in growth rates, relative prices of E will diverge quickly (slowly) reflecting rapidly (moderately) increasing scarcity values. If E is perfectly substitutable then there will be no relative price effect since E is not really economically scarce.

F. A Place-based Ramsey Rule

F.0.1. Let us assume that there are M people in the population of the country, with each person denoted by $m \in \{1, \ldots, M\}$. The government is interested in the total welfare across all of its population.

$$W = \sum_{t=0}^{T} \sum_{m=1}^{M} \exp(-\delta t) U(c_{mt})$$

where both U, δ are the same across all members of the population. Since we are extending the Ramsey Rule framework for regionality only, we will assume that is no stochasticity and that the model is entirely deterministic. The extension to the case of uncertainty is relatively straightforward conceptually, but less analytically tractable and also less directly comparable with the Green Book.

F.0.2. For the purposes of regional discounting, we assume that the country can be broken down into N regions. In each region there are $M_n = M\theta_n$ people, with $\sum_N \theta_n = 1$; θ_n is the proportion of the total population in region N. If there is a representative agent in each region that consumes c_{nt} at each time, t, then:

$$W = M \sum_{t=0}^{T} \sum_{n=1}^{N} \theta_n \exp(-\delta t) U(c_{nt})$$

F.0.3. Now the government is considering the maximum price, Mp_{it} that it will pay at $t = 0$ to derive the single non-stochastic benefit Mb_{it} at time t (and nothing else), so that p_{it}, b_{it} represent, as usual, the real per capita PV and benefit from that one single benefit from the project.

F.0.4. However, because of regional investing and the differential of taxation between regions, the price and benefit are both divided unequally across the country on a per capita basis. Let $w_n M p_{it}$ represent the total amount that region n contributes towards achieving that single benefit, with $\sum_n w_n = 1$. Similarly, let $\varpi_n M b_{it}$ be the total benefit that region n receives from that single benefit, with $\sum_n \varpi_n = 1$. Then, per capita, the cost to each person in region n is $w_n M p_{it}/M_n = w_n p_{it}/\theta_n$. Similarly, the benefit to each person in the region at time t is $\varpi_n b_{it}/\theta_n$. Therefore, if the government decides to pursue the project, reducing consumption by the investment amount this period but increasing it by the benefit at time t, then the change in social welfare, ΔW, is:

$$\Delta W = M \sum_{n=1}^{N} \theta_n [U(c_{n0} - w_n p_{it}/\theta_n) - U(c_{0n}) + \exp(-\delta t)[U(c_{nt} + \varpi_n b_{it}/\theta_n) - U(c_{nt})]]$$

F.0.5. Taking a first order Taylor's series expansion this under the assumption that the cost and benefit of the project are marginal to per capital consumption in each region, and setting $\Delta W = 0$ to derive the maximum price the government should pay, as usual:

$$p_{it} \sum_{n=1}^{N} w_n U'(c_{n0}) = b_{it} \sum_{n=1}^{N} \exp(-\delta t) \varpi_n U'(c_{nt})$$

F.0.6. Notice here that θ_n and M have cancelled in the numerator and denominator — the maximum price for the project does not depend on the population level in each region. Taking isoelastic utility, the present value of the project is given by

$$p_{it} \sum_{n=1}^{N} w_n c_{n0}^{-\mu} = b_{it} \sum_{n=1}^{N} \exp(-\delta t) \varpi_n c_{n1}^{-\mu}$$

F.0.7. Taking the standard expression for the continuous social discount rate, and noting this is a risk-free rate, r_{ft}:

$$r_{ft} = \frac{1}{t}\ln\left(\frac{b_{it}}{p_{it}}\right) = \frac{1}{t}\ln\left(\frac{\sum_{n=1}^{N} w_n c_{n0}^{-\mu}}{\sum_{n=1}^{N} \exp(-\delta t)\varpi_n c_{nt}^{-\mu}}\right) = \delta + \frac{1}{t}\ln\left(\frac{\sum_{n=1}^{N} w_n c_{n0}^{-\mu}}{\sum_{n=1}^{N} \varpi_n c_{nt}^{-\mu}}\right)$$

F.0.8. Now consider the regional annual growth rate, $g_{nt} = (1/t)/\ln(c_{nt}/c_{n0})$. Then $c_{nt}^{-\mu} = c_{n0}^{-\mu}\exp(-\mu g_{nt} t)$. Now, suppose the project is particularly important for the region $n = 1$. Then we can rewrite the growth rate in each other region in terms of the average per capita national growth rate:

$$c_{nt}^{-\mu} = c_{n0}^{-\mu}\exp(-\mu t g_t)\exp(-\mu t(g_{nt} - g_t))$$

and therefore, letting $\kappa_i = \varpi_i c_{i0}^{-\mu} / \sum_n w_n c_{n0}^{-\mu}$:

$$r_{ft} = \delta + \mu g_t - \frac{1}{t}\ln\left(\sum_{n=1}^{N} \kappa_n \exp(-\mu t(g_{nt} - g_t))\right)$$

F.0.9. The first two terms on the right-hand side are just the Ramsey Rule. The third term adjusts for the fact that regional growth rates differ from the national average. It also makes an adjustment for the fact that costs and benefits are not equally spread across regions.

F.0.10. This collapses to the Ramsey Rule if two assumptions hold. First, g_{nt} must be the same for all n — there must be no difference in growth across regions, all of which must equal g_t. In this case, the exponential terms all become equal to one. Note that the Ramsey Rule does not require the current income level to be the same in each region, only that future growth will be the same. The second condition is that $\sum_n \kappa_n = 1$. A sufficient condition for this to hold is if $w_n = \varpi_n$ for all n; regions pay for the project in proportion to the benefits that they receive from it.

F.0.11. In Table F.1 we present an example of this regional discounting model in practice. We assume that $\delta = 0.5\%$, $\mu = 1$ consistent with current Green Book guidance. The population is assumed to stand at 45 million (this does not affect the final results) and that we are considering a net benefit of £22.5 billion that arises in 10 years time. This is a marginal project in the sense that the approximation errors from using a first-order Taylors series expansion are insignificant.

F.0.12. It is assumed that the country has three regions. Region 1 is the largest in terms of population size, has the lowest current consumption level and the lowest growth rate of consumption over the next two decades. Region 3 is the smallest, wealthiest, and fastest growing. Region 2 sits in the middle. The project is going to primarily benefit the less prosperous regions of the country, with $\varpi_1 = 60\%$. However, because of the progressive taxation system, Region 3 primarily funds the project, with $w_3 = 60\%$. The cost per person is almost ten times higher for Region 3 than Region 1 (because there are fewer people in Region 1) although the benefits are three times higher per capita in Region 1 than Region 3.

F.0.13. From this, we can derive the discount rate. The aggregate growth rate of consumption in this economy is 1.826%, and so the Ramsey Rule rate is 2.326%. The adjustment term is -1.507%. The total social discount rate is then 0.891%, leading to a PV of £19,101,059. Had the Ramsey Rule been used instead, the discount rate would have been estimated at £14,129,336. This is clearly a material difference from a practical perspective.

	c_{n0}	g_{nt}	c_{nt}	θ_n	w_n	ϖ_n	Cost/person	Benefit/person
Region 1	20,000	1.50%	26,997	44.4%	15.0%	60.0%	0.1433	0.6750
Region 2	25,000	1.75%	35,477	33.3%	25.0%	30.0%	0.3184	0.4500
Region 3	35,000	2.25%	54,891	22.2%	60.0%	10.0%	1.1461	0.2250
Aggregate	25,000	1.83%	36,022	100%	100%	100%	0.4245	0.5000

Table F.1.: An example of applying a regional discounting model

Bibliography

Acland, D. J. & Greenberg, D. H. (2023), 'Distributional weighting and welfare/equity tradeoffs: A new approach', *Journal of Benefit-Cost Analysis* **14**(1), 68–92.

Addicott, E. T., Fenichel, E. P. & Kotchen, M. J. (2020), 'Even the representative agent must die: Using demographics to inform long-term social discount rates', *Journal of the Association of Environmental and Resource Economists* **7**(2), 379–415.

Adler, M. D. (2012), *Well-Being and Fair Distribution: Beyond Cost-Benefit Analysis*, Oxford University Press, New York.

Andersen, S., Harrison, G. W., Lau, M. I. & Rutström, E. E. (2008), 'Eliciting risk and time preferences', *Econometrica* **76**(3), 583–618.

Andersen, S., Harrison, G. W., Lau, M. I. & Rutström, E. E. (2014), 'Discounting behavior: A reconsideration', *European Economic Review* **71**, 15–33.

Andreoni, J. & Sprenger, C. (2015), 'Risk preferences are not time preferences', *American Economic Review* **102**(7), 3357–3376.

Arrow, K., Cropper, M., Gollier, C., Groom, B., Heal, G., Newell, R., Nordhaus, W., Pindyck, R., Pizer, W., Portney, P. et al. (2013), 'Determining benefits and costs for future generations', *Science* **341**(6144), 349–350.

Arrow, K. J. (1999), Discounting, morality, and gaming, *in* P. R. Portney & J. P. Weyant, eds, 'Discounting and Intergenerational Equity', RFF Press, Washington, DC, pp. 13–21.

Arrow, K. J., Dasgupta, P., Goulder, L. H., Daily, G. C., Ehrlich, P. R., Heal, G. M., Levin, S. A., Mäler, K.-G., Schneider, S. H., Starrett, D. A. & Walker, B. (2004), 'Are we consuming too much?', *Journal of Economic Perspectives* **18**(3), 147–172.

Arrow, K. J., Dasgupta, P., Goulder, L. H., Mumford, K. J. & Oleson, K. (2012), 'Sustainability and the measurement of wealth', *Environment and Development Economics* **17**(3), 317–353.

Arrow, K. J., Dasgupta, P. & Mäler, K.-G. (2003), 'Evaluating projects and assessing sustainable development in imperfect economies', *Environmental and Resource Economics* **26**(4), 647–685.

Arrow, K. J. & Lind, R. C. (1970), 'Uncertainty and the evaluation of public investment decisions', *American Economic Review* **60**, 364–378.

Asheim, G. B. & Nesje, F. (2016), 'Destructive intergenerational altruism', *Journal of the Association of Environmental and Resource Economists* **3**(4), 957–984.

Asheim, G. B. & Zuber, S. (2014), 'Escaping the repugnant conclusion: Rank-discounted utilitarianism with variable population', *Theoretical Economics* **9**(3), 629–650.

Australian Centre for Evalauation (2025), 'Cost-benefit analysis (cba)'. The Treasury: Australian Government.

Backus, D. K., Routledge, B. R. & Zin, S. E. (2004), 'Exotic preferences for macroeconomists', *NBER Macroeconomics Annual* **19**, 319–390.

Bailey, A. J., Bridges, J., Harrison, R., Jones, J. & Mankodi, A. (2020), 'The central bank balance sheet as a policy tool: Past, present and future'.

Barberis, N. C. (2013), 'Thirty years of prospect theory in economics: A review and assessment', *Journal of Economic Perspectives* **27**(1), 173–196.

Barrage, L. (2018), 'Be careful what you calibrate for: Social discounting in general equilibrium', *Journal of Public Economics* **160**, 33–49.

Barrage, L. & Nordhaus, W. D. (2024), 'Policies, projections, and the social cost of carbon: Results from the dice-2023 model', *Proceedings of the National Academy of Sciences* **121**(13), e2312030121.

Barro, R. J. (2006), 'Rare disasters and asset markets in the Twentieth Century', *The Quarterly Journal of Economics* **121**(3), 823–866.

Barro, R. J. (2009), 'Rare disasters, asset prices, and welfare costs', *American Economic Review* **99**(1), 243–264.

Barro, R. J. & Jin, T. (2011), 'On the size distribution of macroeconomic disasters', *Econometrica* **79**(5), 1567–1589.

Barro, R. J. & Ursúa, J. F. (2011), 'Rare macroeconomic disasters', *Annual Review of Economics* **3**, 83–109.

Barseghyan, L., Molinari, F., O'Donoghue, T. & Teitelbaum, J. C. (2013), 'The nature of risk preferences: Evidence from insurance choices', *American Economic Review* **103**(6), 2499–2529.

Barsky, R. B., Juster, F. T., Kimball, M. S. & Shapiro, M. D. (1997), 'Preference parameters and behavioral heterogeneity: An experimental approach in the health and retirement study', *The Quarterly Journal of Economics* **112**(2), 537–579.

Bastien-Olvera, B. A., Conte, M. N., Dong, X., Briceno, T., Batker, D., Emmerling, J., Tavoni, M., Granella, F. & Moore, F. C. (2024), 'Unequal climate impacts on global values of natural capital', *Nature* **625**, 722–727.

Bastien-Olvera, B. A. & Moore, F. C. (2021), 'Use and non-use value of nature and the social cost of carbon', *Nature Sustainability* **4**(2), 101–108.

Bastien-Olvera, B. A. & Moore, F. C. (2022), 'Climate impacts on natural capital: Consequences for the social cost of carbon', *Annual Review of Resource Economics* **14**, 515–532.

Baumgärtner, S., Drupp, M. A. & Quaas, M. F. (2017), 'Subsistence, substitutability and sustainability in consumption', *Environmental and Resource Economics* **67**(1), 47–66.

Baumgärtner, S., Klein, A. M., Thiel, D. & Winkler, K. (2015), 'Ramsey discounting of ecosystem services', *Environmental and Resource Economics* **61**(2), 273–296.

Baumstark, L. & Gollier, C. (2014), 'The relevance and the limits of the Arrow-Lind Theorem', *Journal of Natural Resources Policy Research* **6**(1), 45–49.

Becker, G. S. (1962), 'Investment in human capital: A theoretical analysis', *Journal of Political Economy* **70**(5, Part 2), 9–49.

Best, M. C., Cloyne, J., Ilzetzki, E. & Kleven, H. J. (2020), 'Estimating the elasticity of intertemporal substitution using mortgage notches', *The Review of Economic Studies* **87**(2), 656–690.

Blundell, R., Browning, M. & Meghir, C. (1994), 'Consumer demand and the life-cycle allocation of household expenditures', *The Review of Economic Studies* **61**(1), 57–80.

Boadway, R. W. & Bruce, N. (1984), *Welfare Economics*, Basil Blackwell, Oxford.

Broome, J. (2004*a*), Same-lifetime aggregation, *in* 'Weighing Lives', Oxford University Press, Oxford, pp. 215–232.

Broome, J. (2004*b*), *Weighing Lives*, Oxford University Press, Oxford.

Buchholz, W. & Schumacher, J. (2010), 'Discounting and welfare analysis over time: Choosing the ', *European Journal of Political Economy* **26**(3), 372–385. Special Issue on Ethics and Economics.

Bureau of the Budget (1969), 'Circular A-94'. Congressional Record, Proceedings and Debates of the 91st Congress, First Session, pages 18277–18279.

Cavatorta, E. & Groom, B. (2025), 'Definitely maybe: Inter-temporal discounting when future payments are in doubt', SSRN Working Paper. Posted 13 June 2025.

Chamberlain, G. (1983), 'Funds, factors, and diversification in arbitrage pricing models', *Econometrica* **51**, 1305–1323.

Chamberlain, G. & Rothschild, M. (1983), 'Arbitrage, factor structure, and mean–variance analysis on large asset markets', *Econometrica* **51**, 1281–1304.

Chen, H., Dou, W. W. & Kogan, L. (2024), 'Measuring "dark matter" in asset pricing models', *Journal of Finance* **79**(2), 843–902.

Cherbonnier, F., Gollier, C. & Pommeret, A. (2025), 'Stress discounting', *Journal of Risk and Uncertainty* pp. 1–25.

Chetty, R. (2006), 'A new method of estimating risk aversion', *American Economic Review* **96**(5), 1821–1834.

Cochrane, J. H. (2009), *Asset pricing: Revised edition*, Princeton University Press.

Cochrane, J. H. (2017), 'Macro-finance', *Review of Finance* **21**(3), 945–985.

Cohen, A. & Einav, L. (2007), 'Estimating risk preferences from deductible choice', *American Economic Review* **97**(3), 745–788.

Cohen, J., Ericson, K. M., Laibson, D. & White, J. M. (2020), 'Measuring time preferences', *Journal of Economic Literature* **58**(2), 299–347.

Constantinides, G. M. & Duffie, D. (1996), 'Asset pricing with heterogeneous consumers', *Journal of Political Economy* **104**(2), 219–240.

Cropper, M. L., Freeman, M. C., Groom, B. & Pizer, W. A. (2014), 'Declining discount rates', *American Economic Review* **104**(5), 538–543.

Crump, R. K., Eusepi, S., Tambalotti, A. & Topa, G. (2022), 'Subjective intertemporal substitution', *Journal of Monetary Economics* **126**, 118–133.

Damodaran, A. (2026), 'Equity Risk Premiums (ERP): Determinants, estimates and implications — The 2026 edition'. Available at SSRN: 6361419.

Dasgupta, P. (2005), Three conceptions of intergenerational justice, *in* 'Ramsey's Legacy', Oxford University Press.

Dasgupta, P. (2008), 'Discounting climate change', *Journal of Risk and Uncertainty* **37**(2–3), 141–169.

Dasgupta, P. (2021), *The Economics of Biodiversity: The Dasgupta Review*, HM Treasury, London.

Dasgupta, P. & Heal, G. (1974), 'The optimal depletion of exhaustible resources', *The Review of Economic Studies* **41**(5), 3–28.

Dasgupta, P. & Mäler, K.-G. (2000), 'Net national product, wealth, and social well-being', *Environment and Development Economics* **5**(1), 69–93.

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

Dasgupta, P. S. & Heal, G. M. (1979), *Economic Theory and Exhaustible Resources*, Cambridge University Press, Cambridge.

Del Campo, S., Anthoff, D. & Kornek, U. (2024), 'Inequality aversion for climate policy', *Review of Environmental Economics and Policy* **18**(1), 96–115.

Department for Business, Energy & Industrial Strategy (2021), Valuation of greenhouse gas emissions: For policy appraisal and evaluation, Technical report, UK Government, London.

Department for Environment, Food and Rural Affairs (2026), Enabling a natural capital approach guidance, Technical report, Department for Environment, Food and Rural Affairs, London.

Department for Transport (2026), TAG unit A1.3: User and provider impacts, Technical report, Department for Transport, London.

DESNZ (2023), 'Valuation of energy use and greenhouse gas (GHG) emissions: Supplementary guidance to the HM Treasury Green Book on Appraisal and Evaluation in Central Government'. London: Department for Energy Security & Net Zero.

Diamond, P. A. (1965), 'The evaluation of infinite utility streams', *Econometrica* **33**(1), 170–177.

Dietz, S. (2011), 'High impact, low probability? An empirical analysis of risk in the economics of climate change', *Climatic Change* **108**(3), 519–541.

Dietz, S. & Hepburn, C. (2013), 'Benefit–cost analysis of non-marginal climate and energy projects', *Energy Economics* **40**, 61–71.

Dietz, S., Hepburn, C. & Stern, N. (2009), Economics, ethics and climate change, *in* K. Basu & R. Kanbur, eds, 'Arguments for a Better World: Essays in Honor of Amartya Sen, Volume 2: Society, Institutions, and Development', Oxford University Press, Oxford, pp. 365–386.

Dietz, S., Rising, J., Stoerk, T. & Wagner, G. (2021), 'Economic impacts of tipping points in the climate system', *Proceedings of the National Academy of Sciences* **118**(34), e2103081118.

Disque, S., Drupp, M. A. & Groom, B. (2026), An experimental test of dual discounting for consumption and the environment. Unpublished manuscript and workshop presentation; preliminary results.

Dohmen, T., Falk, A., Huffman, D., Sunde, U., Schupp, J. & Wagner, G. G. (2011), 'Individual risk attitudes: Measurement, determinants, and behavioral consequences', *Journal of the European Economic Association* **9**(3), 522–550.

Drupp, M. A. (2018), 'Limits to substitution between ecosystem services and manufactured goods and implications for social discounting', *Environmental and Resource Economics* **69**(1), 135–158.

Drupp, M. A., Freeman, M. C., Groom, B. & Nesje, F. (2018), 'Discounting disentangled', *American Economic Journal: Economic Policy* **10**(4), 109–134.

Drupp, M. A. & Hänsel, M. C. (2021), 'Relative prices and climate policy: How the scarcity of nonmarket goods drives policy evaluation', *American Economic Journal: Economic Policy* **13**(1), 168–201.

Drupp, M. A., Turk, Z. M., Groom, B. & Heckenhahn, J. (2025), 'Global evidence on the income elasticity of willingness to pay, relative price changes and public natural capital values', *Environmental and Resource Economics* **88**(12), 3765–3804.

Drupp, M., Hänsel, M. C., Fenichel, E., Freeman, M., Gollier, C., Groom, B., Heal, G., Howard, P., Millner, A., Moore, F. et al. (2024), 'Accounting for the increasing benefits from scarce ecosystems', *Science* **383**(6687), 1062–1064.

Edge, P. (2011), Stochastic discount factors and real options, *in* 'Real Options Theory Meets Practice 15th Annual International Conference', pp. 15–18.

Einav, L., Finkelstein, A., Pascu, I. & Cullen, M. R. (2012), 'How general are risk preferences? choices under uncertainty in different domains', *American Economic Review* **102**(6), 2606–2638.

Emmerling, J., Groom, B. & van der Ploeg, F. (2025), The equity-adjusted social cost of carbon. Revised version of CESifo Working Paper No. 9777; revise and resubmit, Journal of the Association of Environmental and Resource Economists.

Emmerling, J., Groom, B. & Wettingfeld, T. (2017), 'Discounting and the representative median agent', *Economics Letters* **161**, 78–81.

Enke, B., Graeber, T. & Oprea, R. (2023), Complexity and hyperbolic discounting, Working Paper 31047, National Bureau of Economic Research.

Epper, T. F. & Fehr-Duda, H. (2023), 'Risk in time: The intertwined nature of risk taking and time discounting', *Journal of the European Economic Association* **22**(1), 310–354.

Evans, D. & Sezer, H. (2002), 'A time preference measure of the social discount rate for the uk', *Applied Economics* **34**(15), 1925–1934.

Fama, E. F. & French, K. R. (1992), 'The cross-section of expected stock returns', *Journal of Finance* **47**(2), 427–465.

Fama, E. F. & French, K. R. (2015), 'A five-factor asset pricing model', *Journal of Financial Economics* **116**(1), 1–22.

FCA (2024), 'Statement of policy on cost benefit analyses'. London: Financial Conduct Authority.

Fernandez, P., Garcia de la Garza, D. & Fernández Acín, L. (2025), 'Survey: Market risk premium and risk-free rate used for 54 countries in 2025'. Available at SSRN 5260463.

Fletcher, J., Marshall, A. & O'Connell, M. (2024), 'Model scan of factors in UK stock returns', *The European Journal of Finance* **30**(13), 1548–1561.

Fleurbaey, M. et al. (forthcoming), Weighting distributive impacts in BCA: The new mainstream? Unpublished manuscript; introduction to the special issue on distributional weights in benefit–cost analysis, Journal of Benefit-Cost Analysis.

Foldes, L. & Rees, R. (1977), 'A note on the Arrow-Lind Theorem', *American Economic Review* **67**(2), 188–193.

France Stratégie (2023), 'Guide de l'socioéconomique des investissements publics'. Paris.

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

Freeman, M. C., Clacher, I., Claxton, K. & Groom, B. (2017), Reviewing discount rates in ONS valuations, Technical report, Office for National Statistics, London. Version 2; released by the Office for National Statistics in July 2018.

Freeman, M. C. & Groom, B. (2015), 'Positively gamma discounting: Combining the opinions of experts on the social discount rate', *Economic Journal* **125**(585), 1015–1024.

Freeman, M. C. & Groom, B. (2016*a*), Discounting for environmental accounts, Technical report, Office for National Statistics, London.

Freeman, M. C. & Groom, B. (2016*b*), 'How certain are we about the certainty-equivalent long term social discount rate?', *Journal of Environmental Economics and Management* **79**, 152–168.

Freeman, M. C. & Groom, B. (2021*a*), 'Appraisal accounting think-piece'. Report for the Department for Transport.

Freeman, M. C. & Groom, B. (2021*b*), Appraisal accounting think-piece, Technical report, Department for Transport, London. Version 2, 24 March 2021.

Freeman, M. C., Groom, B. & Spackman, M. (2018), 'Social discount rates for cost-benefit analysis: A report for HM Treasury'. London: HM Treasury.

Gabaix, X. (2012), 'Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance', *Quarterly Journal of Economics* **127**(2), 645–700.

Gollier, C. (2002), 'Discounting an uncertain future', *Journal of Public Economics* **85**(2), 149–166.

Gollier, C. (2008), 'Discounting with uncertain growth', *Journal of Risk and Uncertainty* **37**, 171–186.

Gollier, C. (2010), 'Ecological discounting', *Journal of Economic Theory* **145**(2), 812–829.

Gollier, C. (2013), *Pricing the planet's future: The economics of discounting in an uncertain world*, Princeton University Press.

Gollier, C. (2014), ‘Discounting and growth’, *American Economic Review* **104**(5), 534–537.

Gollier, C. (2026), ‘The welfare cost of ignoring the beta’, *Journal of Political Economy Microeconomics* **4**(1), 000–000.

Gollier, C. & Hammitt, J. K. (2014), ‘The long-run discount rate controversy’, *Annual Review of Resource Economics* **6**, 273–295.

Gollier, C., van der Ploeg, F. & Zheng, J. (2023), ‘The discounting premium puzzle: Survey evidence from professional economists’, *Journal of Environmental Economics and Management* **122**, 102882.

Goodin, R. E. (1995), Government house utilitarianism, *in* ‘Utilitarianism as a Public Philosophy’, Cambridge University Press, Cambridge, pp. 60–78.

Graham, J. R. (2022), ‘Presidential address: Corporate finance and reality’, *The Journal of Finance* **77**(4), 1975–2049.

Greenstone, M., Kopits, E. & Wolverton, A. (2013), ‘Developing a social cost of carbon for US regulatory analysis: A methodology and interpretation’, *Review of Environmental Economics and Policy* .

Grijalva, T. C., Lusk, J. L., Rong, R. & Shaw, W. D. (2018), ‘Convex time budgets and individual discount rates in the long run’, *Environmental and Resource Economics* **71**(1), 259–277.

Grijalva, T., Lusk, J. & Shaw, W. D. (2014), ‘Discounting the distant future: An experimental investigation’, *Environmental and Resource Economics* **59**(1), 39–63.

Groom, B., Drupp, M. A., Freeman, M. C. & Nesje, F. (2022), ‘The future, now: A review of social discounting’, *Annual Review of Resource Economics* **14**, 467–491.

Groom, B. & Hepburn, C. (2017), ‘Looking back at social discounting policy: The influence of papers, presentations, political preconditions, and personalities’, *Review of Environmental Economics and Policy* **11**(2), 336–356.

Groom, B., Hepburn, C. J., Koundouri, P. & Pearce, D. W. (2005), ‘Declining discount rates: The long and the short of it’, *Environmental and Resource Economics* **32**(4), 445–493.

Groom, B., Koundouri, P., Panopoulou, E. & Pantelidis, T. (2004), Declining discount rates: The long and the short of it, Working paper, University College London, Department of Economics.

Groom, B., Lowe, J., zu Ermgassen, S., Milner-Gulland, E. J., Atkins, T., Balmford, B., Binner, A., Butler, A., Day, B., Duffus, N., Hails, R., Maier-Peveling, H., Mancini, M., Meier, S., Nicholas, H., Rinaldo, D., Smale, R., Snowdon, P., Venmans, F. & Bateman, I. J. (2026), ‘Bringing the economics of biodiversity into policy and decision-making: A target and cost-based approach to pricing biodiversity’.

Groom, B. & Maddison, D. (2019), ‘New estimates of the elasticity of marginal utility for the uk’, *Environmental and Resource Economics* **72**(4), 1155–1182.

Groom, B. & Turk, Z. (2021), ‘Reflections on the dasgupta review on the economics of biodiversity’, *Environmental and Resource Economics* **79**, 1–23.

Grossman, M. (1972), ‘On the concept of health capital and the demand for health’, *Journal of Political Economy* **80**(2), 223–255.

Guéant, O., Guesnerie, R. & Lasry, J.-M. (2012), ‘Ecological intuition versus economic “reason”’, *Journal of Public Economic Theory* **14**(2), 245–272.

Guesnerie, R. (2004), ‘Calcul économique et développement durable’, *Revue économique* **55**(3), 363–382.

Guthrie, G. A. (2009), *Real Options in Theory and Practice*, Oxford University Press.

Harrison, J. M. & Kreps, D. M. (1979), ‘Martingales and arbitrage in multiperiod securities markets’, *Journal of Economic Theory* **20**(3), 381–408.

Harrison, M. (2010), ‘Valuing the future: The social discount rate in cost-benefit analysis’, *Available at SSRN 1599963* .

Harrod, R. F. (1948), *Towards a Dynamic Economics: Some Recent Developments of Economic Theory and Their Application to Policy*, Macmillan and Co., London.

Harsanyi, J. C. (1955), ‘Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility’, *Journal of Political Economy* **63**(4), 309–321.

Harsanyi, J. C. (1982), Morality and the theory of rational behaviour, *in* A. Sen & B. Williams, eds, ‘Utilitarianism and Beyond’, Cambridge University Press, Cambridge, pp. 39–62.

Hartwick, J. M. (1977), ‘Intergenerational equity and the investing of rents from exhaustible resources’, *American Economic Review* **67**(5), 972–974.

Havránek, T. (2015), ‘Measuring intertemporal substitution: The importance of method choices and selective reporting', *Journal of the European Economic Association* **13**(6), 1180–1204.

Heal, G. M. (1998), *Valuing the Future: Economic Theory and Sustainability*, Columbia University Press, New York.

Heal, G. M. & Millner, A. (2014*a*), 'Agreeing to disagree on climate policy', *Proceedings of the National Academy of Sciences* **111**(10), 3695–3698.

Heal, G. M. & Millner, A. (2014*b*), 'Agreeing to disagree on climate policy', *Proceedings of the National Academy of Sciences* **111**(10), 3695–3698.

Heal, G. & Millner, A. (2013), Discounting under disagreement, NBER Working Paper 18999, National Bureau of Economic Research.

Heckenhahn, J. & Drupp, M. A. (2024), 'Relative price changes of ecosystem services: Evidence from germany', *Environmental and Resource Economics* **87**(3), 833–880.

Heimer, R. Z., Myrseth, K. O. R. & Schoenle, R. S. (2019), 'YOLO: Mortality beliefs and household finance puzzles', *The Journal of Finance* **74**(6), 2957–2996.

Henderson, N. & Bateman, I. (1995), 'Empirical and public choice evidence for hyperbolic social discount rates and the implications for intergenerational discounting', *Environmental and Resource Economics* **5**(4), 413–423.

Hepburn, C., Duncan, S. & Papachristodoulou, A. (2010), 'Behavioural economics, hyperbolic discounting and environmental policy', *Environmental and Resource Economics* **46**(2), 189–206.

Hepburn, C., Koundouri, P., Panopoulou, E. & Pantelidis, T. (2009), 'Social discounting under uncertainty: A cross-country comparison', *Journal of Environmental Economics and Management* **57**(2), 140–150.

HM Treasury (1973), 'Use of Discounted Cash Flow and the Test Discount Rate in the public sector'. London: HM Treasury.

HM Treasury (1980), 'Investment appraisal and discounting techniques and the use of the Test Discount Rate in the public sector'. London: HM Treasury.

HM Treasury (1982), 'Investment appraisal in the public sector'. London: HM Treasury.

HM Treasury (1984), 'Investment appraisal in the public sector: A technical guide for government departments'. London: HM Treasury.

HM Treasury (1991), 'Appraisal and evaluation in central government'. London: HM Treasury.

HM Treasury (1997), 'Appraisal and evaluation in central government'. London: HM Treasury.

HM Treasury (2003), 'The Green Book: Appraisal and evaluation in central government'. London: HM Treasury.

HM Treasury (2018), 'The Green Book: Central government guidance on appraisal and evaluation'. London: HM Treasury.

HM Treasury (2021), 'Environmental discount rate review: Conclusion'. London: HM Treasury.

HM Treasury (2023), 'Public service pensions: Consultation on the discount rate methodology — Government response'. London: HM Treasury.

HM Treasury (2025), 'The orange book management of risk — principles and concepts'. London: HM Treasury.

HM Treasury (2026*a*), 'Discounting: Green Book supplementary guidance'. London: HM Treasury.

HM Treasury (2026*b*), 'The government financial reporting manual: 2026-27'. London: HM Treasury.

HM Treasury (2026*c*), 'The Green Book: UK government guidance on appraisal'. London: HM Treasury.

Hoel, M. & Sterner, T. (2007), 'Discounting and relative prices', *Climatic Change* **84**(3–4), 265–280.

Holt, C. A. & Laury, S. K. (2002), 'Risk aversion and incentive effects', *American Economic Review* **92**(5), 1644–1655.

Hoogendoorn, S. & Romijn, G. (2020), 'Discontovoet voor de Nederlandse economie'. Centraal Planbureau.

Hou, K., Xue, C. & Zhang, L. (2015), 'Digesting anomalies: An investment approach', *Review of Financial Studies* **28**(3), 650–705.

House of Lords (2021), 'Quantitative easing: a dangerous addiction?'.

Howard, G. (2013), 'Discounting for personal and social payments: Patience for others, impatience for ourselves', *Journal of Environmental Economics and Management* **66**(3), 583–597.

Iverson, T. (2013), 'Minimax regret discounting', *Journal of Environmental Economics and Management* **66**(3), 598–608.

Jenkins, G. & Kuo, C. (2007), The economic opportunity cost of capital for Canada —an empirical update, Technical report, Queen's Economics Department Working Paper.

Johnson, E., Johnson, S. B., Stark, G., Reed, H., Lee, R. & Johnson, M. (2025), 'What support is there in the UK for renationalisation of public utilities? Evidence on drivers and fluidity of support via adversarial narrative co-production', *Next Research* p. 100615.

Jorgenson, D. W. & Fraumeni, B. M. (1992), 'Investment in education and U.S. economic growth', *The Scandinavian Journal of Economics* **94**(Supplement), S51–S70.

Joslin, S., Walford, W., Clark, M., Goh, S., Jarrett, L. & Chauhan, C. (2026), 'Green Book reform An actuarial perspective on transformational change and the Green Book discount rate'. London: Institute and Faculty of Actuaries.

Jouini, E. & Napp, C. (2014), 'How to aggregate experts' discount rates: An equilibrium approach', *Economic Modelling* **36**, 235–243.

JRG (2011), 'Discounting for CBAs involving private investment, but public benefit'. Joint Regulators Group.

Karp, L. (2005), 'Global warming and hyperbolic discounting', *Journal of Public Economics* **89**(2-3), 261–282.

Kelleher, J. P. (2025*a*), A foundation for (discounted) utilitarian social welfare functions, *in* 'The Social Cost of Carbon: Ethics and the Limits of Climate Change Economics', Oxford University Press, Oxford, pp. 89–121.

Kelleher, J. P. (2025*b*), *The Social Cost of Carbon: Ethics and the Limits of Climate Change Economics*, Oxford University Press, Oxford.

Kimball, M. S. (1990*a*), 'Precautionary saving and the marginal propensity to consume'. National Bureau of Economic Research Cambridge, Mass., USA.

Kimball, M. S. (1990*b*), 'Precautionary saving in the small and in the large', *Econometrica* **58**, 53–73.

Koetse, M. J., Renes, G., Ruijs, A. & de Zeeuw, A. (2018), Relative price increase for nature and ecosystem services in cost-benefit analysis, Technical Report PBL Publication 3214, PBL Netherlands Environmental Assessment Agency, The Hague.

Koopmans, T. C. (1960), 'Stationary ordinal utility and impatience', *Econometrica* **28**(2), 287–309.

Koopmans, T. C. (1972), Representation of preference orderings over time, *in* C. B. McGuire & R. Radner, eds, 'Decision and Organization'.

Krasker, W. S. (1980), 'The 'peso problem'in testing the efficiency of forward exchange markets', *Journal of Monetary Economics* **6**(2), 269–276.

Kreps, D. M. & Porteus, E. L. (1978), 'Temporal resolution of uncertainty and dynamic choice theory', *Econometrica* **46**(1), 185–200.

Kula, E. (1985), 'An empirical investigation on the social time-preference rate for the united kingdom', *Environment and Planning A: Economy and Space* **17**(2), 199–212.

Kula, E. (1987), 'Social interest rate for public sector appraisal in the United Kingdom, the United States and Canada', *Project Appraisal* **2**(3), 169–174.

Layard, R., Mayraz, G. & Nickell, S. (2008), 'The marginal utility of income', *Journal of Public Economics* **92**(8–9), 1846–1857.

LeRoy, S. F. & Werner, J. (2014), *Principles of Financial Economics*, Cambridge University Press.

Li, Q. & Pizer, W. A. (2021), 'Use of the consumption discount rate for public policy over the distant future', *Journal of Environmental Economics and Management* **107**, 102428.

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**(1), 13–37.

Little, I. M. D. & Mirrlees, J. A. (1974), *Project Appraisal and Planning for Developing Countries*, Heinemann Educational Books, London.

Lowe, J. (2008), 'Intergenerational wealth transfers and social discounting: Supplementary Green Book guidance'. London: HM Treasury.

Lucas, D. (2014), 'Rebutting Arrow and Lind: Why governments should use market rates for discounting', *Journal of Natural Resources Policy Research* **6**(1), 85–91.

Lucas, D. (2023), 'Reflections on what financial economics can and cannot teach us about the social discount rate', *Annual Review of Financial Economics* **15**(1), 185–195.

Malinvaud, E. (1953), 'Capital accumulation and efficient allocation of resources', *Econometrica* **21**(2), 233–268.

Markowitz, H. (1952), ‘Modern portfolio theory’, *Journal of Finance* **7**(11), 77–91.

Martin, I. W. & Pindyck, R. S. (2015), ‘Averting catastrophes: The strange economics of Scylla and Charybdis’, *American Economic Review* **105**(10), 2947–2985.

Mehra, R. (2007), ‘The equity premium puzzle: A review’, *Foundations and Trends in Accounting* **2**(1), 1–81.

Mehra, R. & Prescott, E. C. (1985), ‘The equity premium: A puzzle’, *Journal of Monetary Economics* **15**(2), 145–161.

Merton, R. C. (1973), ‘An intertemporal capital asset pricing model’, *Econometrica: Journal of the Econometric Society* pp. 867–887.

Millner, A. (2013), ‘On welfare frameworks and catastrophic climate risks’, *Journal of Environmental Economics and Management* **65**(2), 310–325.

Millner, A. (2020), ‘Nondogmatic social discounting’, *American Economic Review* **110**(3), 760–775.

Millner, A. & Heal, G. (2023), ‘Choosing the future: Markets, ethics, and rapprochement in social discounting’, *Journal of Economic Literature* **61**(3), 1037–1087.

Moore, F. C., Baldos, U. L. C., Hertel, T. & Diaz, D. (2017), ‘New science of climate change impacts on agriculture implies higher social cost of carbon’, *Nature Communications* **8**, 1607.

Moore, F. C. & Diaz, D. B. (2015), ‘Temperature impacts on economic growth warrant stringent mitigation policy’, *Nature Climate Change* **5**, 127–131.

Moore, F. C., Drupp, M. A., Rising, J., Dietz, S., Rudik, I. & Wagner, G. (2024), ‘Synthesis of evidence yields high social cost of carbon due to structural model variation and uncertainties’, *Proceedings of the National Academy of Sciences* **121**(52), e2410733121.

Moore, M. A., Boardman, A. E. & Vining, A. R. (2017), ‘Risk in public sector project appraisal: It mostly does not matter!’, *Public Works Management & Policy* **22**(4), 301–321.

National Audit Office (2018), ‘PFI and PF2’. London: HM Treasury; Report by the Comptroller and Auditor General.

Nesje, F. (2024), ‘Cross-dynastic intergenerational altruism’, *Journal of Environmental Economics and Management* **128**, 103059.

Nesje, F., Drupp, M. A., Freeman, M. C. & Groom, B. (2023), ‘Philosophers and economists agree on climate policy paths but for different reasons’, *Nature Climate Change* **13**(6), 515–522.

Nesje, F., Drupp, M. A., Freeman, M. C. & Groom, B. (2024), ‘Sense and sensitivity: An argument against reporting multiple net present values’. Available at SSRN 4663124.

Nesje, F. & Lund, D. (2018), ‘Risikojustering av kalkulasjonsrenta i samfunnsøkonomiske analysa’, *Samfunnsøkonomen* **132**(4), 34–42.

Newbery, D. (1992), ‘Long term discount rates for the forest enterprise’. Department of Applied Economics, Cambridge University, for the UK Forestry Commission, Edinburgh.

Newell, R. G. & Pizer, W. A. (2003), ‘Discounting the distant future: How much do uncertain rates increase valuations?’, *Journal of Environmental Economics and Management* **46**(1), 52–71.

Newell, R. G., Pizer, W. A. & Prest, B. C. (2022), ‘A discounting rule for the social cost of carbon’, *Journal of the Association of Environmental and Resource Economists* **9**(5), 1017–1046.

Nordhaus, W. D. (2007), ‘A review of the stern review on the economics of climate change’, *Journal of Economic Literature* **45**(3), 686–702.

Norwegian Ministry of Finance (2012), ‘Cost-benefit analysis’. Official Norwegian Reports NOU 2012: 16.

Norwegian Ministry of Finance (2021), ‘Rundskriv r: Prinsipper og krav ved utarbeidelse av samfunnsøkonomiske analyser’. Det Kongelige Finansdepartement, Norway.

Obstfeld, M. (2023), Natural and neutral real interest rates: Past and future. National Bureau of Economic Research.

Office for National Statistics (2025), UK natural capital accounts quality and methods guide, Technical report, Office for National Statistics, Newport.

Office of Best Practice Regulation (2014), ‘Research report: Environmental valuation and uncertainty’. The Treasury: Australian Government.

Office of Impact Analysis (2023), ‘Cost benefit analysis’. The Treasury: Australian Government.

Olson, M. & Bailey, M. J. (1981), ‘Positive time preference’, *Journal of Political Economy* **89**(1), 1–25.

OMB (1972), 'Circular A-94: Revised'. Executive Office of the President: Office of Management & Budget; Washington D.C.

OMB (1992), 'Circular A-94: Guidelines and discount rates for Benefit-Cost Analysis of Federal Programs'. Executive Office of the President: Office of Management & Budget; Washington D.C.

OMB (1996), 'Economic analysis of Federal regulations under Executive Order 12866'. Executive Office of the President: Office of Management & Budget; Washington D.C.

OMB (2003), 'Circular A-4'. Executive Office of the President: Office of Management & Budget; Washington D.C.

OMB (2023*a*), 'Circular A-4'. Executive Office of the President: Office of Management & Budget; Washington D.C.

OMB (2023*b*), 'Circular A-94'. Executive Office of the President: Office of Management & Budget; Washington D.C.

Oxera (2002), 'A social time preference rate for use in long-term discounting'. Report to the Office of the Deputy Prime Minister and the Department of the Environment, Food and Rural Affairs.

Oxera (2026), 'Discount rates and appraisal periods for the cost–benefit analysis of financial services regulation'. Prepared for the Financial Conduct Authority.

Parfit, D. (1986), Overpopulation and the quality of life, *in* P. Singer, ed., 'Applied Ethics', Oxford University Press, Oxford, pp. 145–164.

Pearce, D. & Ulph, D. (1995), 'A social discount rate for the United Kingdom'. CSERGE Working Paper GEC 95-01.

Pigou, A. C. (1930), *The functions of economic analysis*, number 12, Oxford University Press, H. Milford.

Pindyck, R. S. (2013), 'The climate policy dilemma', *Review of Environmental Economics and Policy* **7**(2), 219–237.

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

Price, R. & Durham, C. (2009), Review of the economics of sustainable development: Interim report, Technical report, Government Economic Service and Department for Environment, Food and Rural Affairs, London. Revised November 2009.

Quaas, M. F. & Bröcker, J. (2016), Substitutability and the social cost of carbon in a solvable growth model with irreversible climate change, Economics Working Paper 2016-09, Department of Economics, Christian-Albrechts-Universität zu Kiel, Kiel.

Ramsey, F. P. (1928), 'A mathematical theory of saving', *The Economic Journal* **38**(152), 543–559.

Rao, M., Freeman, M., Thijssen, J., Bland, S., Sculpher, M., Claxton, K. & Revill, P. (2026), 'Insights from financial economics to value healthcare investments that reduce system-level risks: Example of disease elimination and eradication', *Applied Health Economics and Health Policy* pp. 1–10.

Rawls, J. (1971), *A Theory of Justice: Original Edition*, Harvard University Press.

Rees, M. J. (2003), *Our final hour: A scientist's warning: how terror, error, and environmental disaster threaten humankind's future in this century–on earth and beyond*, Basic Books (AZ).

Rennert, K., Errickson, F., Prest, B. C., Rennels, L., Newell, R. G., Pizer, W., Kingdon, C., Wingenroth, J., Cooke, R., Parthum, B. et al. (2022), 'Comprehensive evidence implies a higher social cost of CO2', *Nature* **610**(7933), 687–692.

Rietz, T. A. (1988), 'The equity risk premium a solution', *Journal of Monetary Economics* **22**(1), 117–131.

Rijksoverheid (2020), 'Rapport werkgroep discontovoet 2020'. Government of the Netherlands.

Rijksoverheid (2025), 'Rapport werkgroep discontovoet 2025'. Government of the Netherlands.

Robson, A. J. & Szentes, B. (2014), 'A biological theory of social discounting', *American Economic Review* **104**(11), 3481–3497.

Samuelson, P. A. (1937), 'A note on measurement of utility', *The Review of Economic Studies* **4**(2), 155–161.

Samuelson, P. A. (1956), 'Social indifference curves', *The Quarterly Journal of Economics* **70**(1), 1–22.

Sandmo, A. & Drèze, J. H. (1971), 'Discount rates for public investment in closed and open economies', *Economica* **38**(152), 395–412.

Schmelzing, P. (2020), Eight centuries of global real interest rates, R-G, and the 'suprasecular'decline 1311–2018. Bank of England Staff Working Paper, No. 845.

Schultz, T. W. (1961), ‘Investment in human capital’, *American Economic Review* **51**(1), 1–17.

Sen, A. K. (1967), ‘Isolation, assurance and the social rate of discount’, *The Quarterly Journal of Economics* **81**(1), 112–124.

Sen, A. & Williams, B. (1982), Introduction: Utilitarianism and beyond, *in* A. Sen & B. Williams, eds, ‘Utilitarianism and Beyond’, Cambridge University Press, Cambridge, pp. 1–22.

Sharpe, W. F. (1964), ‘Capital asset prices: A theory of market equilibrium under conditions of risk’, *Journal of Finance* **19**(3), 425–442.

Solow, R. M. (1974*a*), ‘The economics of resources or the resources of economics’, *The American Economic Review* **64**(2), 1–14. Richard T. Ely Lecture, delivered at the Eighty-Sixth Annual Meeting of the American Economic Association, New York, December 1973.

Solow, R. M. (1974*b*), ‘Intergenerational equity and exhaustible resources’, *The Review of Economic Studies* **41**(5), 29–45.

Sozou, P. (1998), ‘On hyperbolic discounting and uncertain hazard rates’, *Proceedings of the Royal Society B* **265**, 215–220.

Spackman, M. (2013), ‘Government time discounting and required rates of return: UK history and current issues’, *Economic Affairs* **33**(2), 190–206.

Spackman, M. (2020), ‘Social discounting and the cost of public funds: A practitioner’s perspective’, *Journal of Benefit-Cost Analysis* **11**(2), 244–271.

Spackman, M. (2021), ‘Social discounting and the equity premium’. Grantham Research Institute on Climate Change and the Environment, London.

Spackman, M. (2024), ‘The social discount rate and the cost of public funds: A search for more consistency and better practice’, *Journal of Environmental Economics and Policy* **13**(2), 228–242.

Starmer, C. (2000), ‘Developments in non-expected utility theory: The hunt for a descriptive theory of choice under risk’, *Journal of Economic Literature* **38**(2), 332–382.

Stefanski, R. & Trew, A. (2024), ‘Selection, patience, and the interest rate’, *Journal of Political Economy Macroeconomics* **2**(1), 149–181.

Stern, N. (2007), *Stern Review: The Economics of Climate Change*, Cambridge University Press.

Stern, N. H. (1977), Welfare weights and the elasticity of the marginal valuation of income, *in* M. Artis & A. R. Nobay, eds, ‘Studies in Modern Economic Analysis’, Basil Blackwell, Oxford.

Sterner, T. & Persson, U. M. (2008), ‘An even sterner review: Introducing relative prices into the discounting debate’, *Review of Environmental Economics and Policy* **2**(1), 61–76.

Stiglitz, J. E. (1974), ‘Growth with exhaustible natural resources: Efficient and optimal growth paths’, *The Review of Economic Studies* **41**(5), 123–137.

Sunstein, C. R. (2014), ‘On not revisiting official discount rates: Institutional inertia and the social cost of carbon’, *American Economic Review* **104**(5), 547–551.

Tharyan, R., Gregory, A. & Chen, B. (2024), ‘An investigation of multi-factor asset pricing models in the UK’, *Available at SSRN 4933529* .

Tol, R. S. J. (2010), ‘International inequity aversion and the social cost of carbon’, *Climate Change Economics* **1**(1), 21–32.

Traeger, C. P. (2011), ‘Sustainability, limited substitutability, and non-constant social discount rates’, *Journal of Environmental Economics and Management* **62**(2), 215–228.

Traeger, C. P. (2014), ‘Why uncertainty matters: discounting under intertemporal risk aversion and ambiguity’, *Economic Theory* **56**(3), 627–664.

Treasury Board of Canada (2022), ‘Canada’s cost-benefit analysis guide for regulatory proposals’. Published by Treasury Board of Canada, Secretariat.

Trigeorgis, L. (1996), *Real options: Managerial flexibility and strategy in resource allocation*, MIT press.

UKRN (2023), ‘UKRN guidance for regulators on the methodology for setting the cost of capital’. UK Regulators Network.

USEPA (2023), ‘Report on the social cost of greenhouse gases: Estimates incorporating recent scientific advances: Final guidance’. US Environmental Protection Agency, Washington DC.

Van Eck, C. W., Messling, L. & Hayhoe, K. (2024), ‘Challenging the neutrality myth in climate science and activism’, *npj Climate Action* **3**(1), 81.

Venmans, F. & Groom, B. (2021), ‘Social discounting, inequality aversion, and the environment’, *Journal of Environmental Economics and Management* p. 102479.

Wachter, J. A. (2013), ‘Can time-varying risk of rare disasters explain aggregate stock market volatility?’, *Journal of Finance* **68**(3), 987–1035.

Weikard, H.-P. & Zhu, X. (2005), ‘Discounting and environmental quality: When should dual rates be used?’, *Economic Modelling* **22**(5), 868–878.

Weil, P. (1989), ‘The equity premium puzzle and the risk-free rate puzzle’, *Journal of Monetary Economics* **24**(3), 401–421.

Weitzman, M. L. (1998), ‘Why the far-distant future should be discounted at its lowest possible rate’, *Journal of Environmental Economics and Management* **36**(3), 201–208.

Weitzman, M. L. (2001), ‘Gamma discounting’, *American Economic Review* **91**(1), 260–271.

Weitzman, M. L. (2007), ‘A review of the Stern Review on the economics of climate change’, *Journal of Economic Literature* **45**(3), 703–724.

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

Zeckhauser, R. J. & Viscusi, W. K. (2008), ‘Discounting dilemmas: Editors’ introduction’, *Journal of Risk and Uncertainty* **37**(2), 95–106.

Zhu, X., Smulders, S. & de Zeeuw, A. (2019), ‘Discounting in the presence of scarce ecosystem services’, *Journal of Environmental Economics and Management* **98**, 102272.

Zuber, S. & Asheim, G. B. (2012), ‘Justifying social discounting: The rank-discounted utilitarian approach’, *Journal of Economic Theory* **147**(4), 1572–1601.