

A formula for success: reviewing the social discount rate





Contact
Matthew Shepherd
Principal

The social discount rate (SDR) is a crucial component of the UK government’s approach to project and policy appraisal. Government guidance has stipulated the same SDR (3.5%) since 2003. In the first of two *Agenda in focus* articles, we show that recent research can reasonably support an SDR of under 2.5%; and that this SDR declines considerably more slowly as the appraisal period is extended than the one that is currently used

SDRs enable us to place a present value on the future costs and benefits of projects and policies that are intended to provide a societal benefit: they specify the degree to which we ‘discount the future’. SDRs are largely unrelated to discount rates used to appraise private projects, such as the weighted average cost of capital (WACC).

Identifying the ‘correct’ SDR is one of the most important aspects of cost–benefit analysis, as changing the discount rate by even a fraction of a percentage point can affect whether a policy is approved. In this article, we show that a case can be made for reducing the current UK SDR by over 1%.

It is often debated whether an SDR ought to reflect the rate that society **should** discount at (prescriptivism), or the rate that it **does** discount at (descriptivism). Here, we do not take a position on which is more appropriate, as this is more of a question for moral philosophers than for economists. Prescriptivist arguments for changing the SDR would state that the current SDR does not represent how society should trade off consumption across time, while descriptivist arguments might claim that empirical facts have changed or become better understood over time, and that the SDR should therefore be adjusted to reflect this.

Since 2003, the UK Treasury’s Green Book (which provides guidance on appraisal and evaluation to inform government decision-making) has stipulated an SDR of 3.5% in real terms, based on a formula that represents the socially optimal discount rate, r :¹

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

Parameter	Current value	Description
δ	0.5%	Rate of pure time preference , representing the extent to which the current population’s welfare is prioritised or deprioritised relative to the future population’s welfare
L	1%	Risk parameter —a measurement of risk, which according to some economists covers only the probability of societal collapse,* while in the Green Book it also includes systematic and other risks
μ	1	Elasticity of marginal utility with respect to consumption ,** defined (approximately) as the percentage change in marginal utility,** divided by the percentage change in consumption
g	2%	Expected growth rate of future per-capita consumption

Table 1 Elements of the SDR

Note: * Stern, N. (2006), ‘The Economics of Climate Change: The Stern Review’, 30 October, p. 45. ** HM Treasury (2018), ‘The Green Book’, p. 102 incorrectly calls this the ‘marginal utility of consumption’, which is a measurement of the rate at which utility changes with different levels of consumption. *** The additional utility or happiness that an individual receives from an infinitesimally small change in consumption, at a given level of consumption.

Source: HM Treasury (2018), ‘The Green Book’.

The above formula shows how the SDR is calculated for projects that are evaluated over 30 years or less. In the case of longer-term projects, a long-term SDR is calculated. This results in a formula for the SDR that differs from the equation

above by making the SDR a function of time. Consequently, after 30 years, the SDR in the Green Book declines from 3.5% in accordance with the term structure shown in Figure 1. This term structure was described in a 2002 report by Oxera that




Figure 1 The term structure of SDRs in the UK

Source: Lowe, J. (2008), ‘Intergenerational wealth transfers and social discounting: Supplementary Green Book Guidance’.

summarised the research available on the subject at the time.²

The remainder of this article discusses the parameters of the SDR equation. The rate at which the SDR should decline, whether it should be adjusted for inequality, and the implications for UK policymaking, will be covered in a future *Agenda in focus* article.

The rate of pure time preference

The rate of pure time preference (δ) represents the relative weight that is placed on the welfare of people alive now versus people alive in the future.³

A δ of 0 implies that equal weight should be placed on current and future people’s

wellbeing, while $\delta > 0$ implies that the current population's welfare counts for more, and $\delta < 0$ implies that it counts for less.

Currently, the UK government sets δ at 0.5%,⁴ based on the median result from a survey of experts.⁵ While the value $\delta = 0.5\%$ may seem very close to $\delta = 0$, the former implies that the welfare of people alive in 139 years is worth half that of people alive today. This is potentially inconsistent with many people's moral intuitions.

While several arguments have been made for $\delta > 0$, one of the most famous is by economist Kenneth Arrow who showed that, when $\delta = 0$ (and under other additional assumptions), we are led to claim that current generations should save over two-thirds of their income for the benefit of future generations.⁶ Arrow contended that this was implausibly high, and that therefore δ should be greater than 0.

There are, however, (at least) two reasons why a δ of 0 could be reasonable.

- The assumption that $\delta = 0$ reflects the view that the value of a person's wellbeing does not depend on when that person is born. This view was operative in the Stern Review;⁷ and according to the Intergovernmental Panel on Climate Change, it is the majority view among economists.⁸
- Since 2003, society has become increasingly concerned for the welfare of future generations—exhibited, for example, by the current focus on climate change and plastic pollution. In Wales, the government even passed a 'Well-being of Future Generations Act', requiring the government to consider long-term impacts in policymaking.⁹

Therefore, even if δ should not be reduced all the way to 0, it could perhaps be brought closer to this value. It should also be noted that the results of the expert survey referred to above may be consistent with a belief that $\delta = 0$, because the survey did not separately identify δ and L , meaning that experts' estimates may have covered $\delta + L$ rather than δ alone.

Does the current model treat risk correctly?

The risk parameter L (measured at 1% in the Green Book) captures unforeseeable factors that could significantly reduce or completely eliminate the benefits associated with a project.¹⁰ It also includes a small premium for 'systematic' risk, to account for the fact that the net benefits of a project are typically positively correlated with real GDP growth.¹¹

Consequently, L captures a mixture of different types of risk, instead of considering each of them in different ways. This is described by Freeman, Groom and Spackman (2018) as a 'practical shortcut' for taking risk into account.¹²

In reality, projects face three types of risk that can affect their net benefits:¹³

- **project-specific risk** (unsystematic risk)—for example, the probability that a project will be managed well or badly, leading to over- or underperformance;
- **factors outside the project's control** (systematic risk)—for example, the risk that the level of benefit it provides will be correlated with changes in the macroeconomic environment, such as economic growth. This risk cannot be eliminated through diversification (i.e. having lots of different projects);
- **catastrophic risk**, reflecting the probability that society will 'fail to survive' a given year. Catastrophic risks come from existential threats such as meteorites as well as significant political events such as large-scale anarchy or nuclear war.

With each of the above, an increase in risk should intuitively raise the SDR because the future benefits of a project are more likely to be reduced, or—in the event of catastrophic risk—destroyed, making future investment less attractive.

While there does not appear to be consensus on how the different types of risk should be treated, we outline below an approach that *could* be taken.

Unsystematic risk could be dealt with through explicit adjustments to the profile of benefits and costs, as is already suggested in the Green Book. Methods for doing this include decision trees and real options techniques, which directly adjust the streams of benefits and costs rather than the discount rate.

Systematic risk can be dealt with in two ways: analysts can either convert net benefits into certainty equivalents¹⁴ and then discount them at an SDR where L represents only catastrophic risk (sometimes called a 'risk-free' SDR), or they can keep net benefits as they are and add a term for systematic risk to the SDR.¹⁵ In the context of policy appraisal, economists differ in opinion on the importance of systematic risk, based mainly on whether they take a normative or a descriptivist approach to discounting. In the former case, systematic risk is often quantified at below 0.1%,¹⁶ while in the latter case it is considerably larger, with one estimate putting it at 2%.¹⁷ Some economists who take the former approach consider that systematic risk can be largely ignored for

government projects.¹⁸ The question of which is the correct treatment is clearly material, but also highly complex and controversial.

Catastrophic risk could be captured by a parameter reflecting the probability of societal collapse. As this risk is independent of other macroeconomic events (unlike systematic risk), it would enter L directly and linearly. This would essentially result in L being set to equal catastrophic risk, as in the Stern Review,¹⁹ in which catastrophic risk was estimated at 0.1%, implying a probability of society's survival after 100 years of approximately 90%.²⁰ Toby Ord, a philosopher at Oxford University's Future of Humanity Institute, places the probability of the complete destruction of humanity at approximately 17% across the next century, implying an estimate for L of around 0.2%.²¹

If systematic and non-systematic risk are removed from L then its role in discounting will change to focus exclusively on catastrophic risk.

The elasticity of marginal utility

The elasticity of marginal utility with respect to consumption (μ) reflects three conceptually distinct features of a social welfare function:

- **the level of intratemporal inequality aversion**—the aversion that a society has to (consumption) inequality at a given point in time;²²
- **the level of intertemporal inequality aversion**—the aversion that a society has to (consumption) inequality across time;
- **the level of risk aversion**—the extent to which a society is concerned about not just the central estimate of a project's net benefits, but also its variance.

As mentioned above, μ measures the percentage change in marginal utility (i.e. the rate at which an individual's wellbeing changes as their consumption increases) caused by a 1% increase in their consumption. Higher levels of μ reflect a society that derives less utility from higher levels of consumption than would be the case if μ were lower.²³ Such a society therefore benefits less from shifting resources to the future (assuming that the society is expected to become wealthier in the future). Two issues arise with the estimation of μ .

First, some economists take the view that the elasticity of marginal utility depends on how society *should* value consumption, whereas empirical analysis uses real-world data and so is based on the *actual*

behaviour of individuals. Disagreement on whether μ should be empirically estimated can have material implications for its magnitude, because normative and empirical economists will use different methodologies. For example, empirical economists will be happier using real-world data, while normative economists may test the implications of different values of μ for various ethical thought experiments. Empirical methods were used to derive the value of $\mu = 1$ used in the Green Book, while some academics who prefer normative methods have found higher values, such as between 1.5 and 3.²⁴

Second, to the extent that empirical methods are appropriate, there is debate over which methods should be used and the values that they suggest. Empirical estimation is particularly difficult in the case of μ because, as discussed above, it measures three conceptually distinct aspects of social welfare functions, and methods based on one of the three aspects may not produce results that are consistent with methods based on another.

Notwithstanding this, the Green Book currently sets a value of $\mu = 1$, based on results presented in Oxera (2002), which drew on available research at the time to suggest a value of between 0.5 and 1.2.²⁵ Since 2002, more work has been done on the estimation of μ —most recently in Groom and Maddison (2019), whose ‘central’ estimate was 1.5.²⁶ While any values will be dependent on the particular estimation method adopted, and the estimation method may or may not be disputed, Figure 2 shows that all of the μ estimates are higher than 1 (the current value), indicating that an increase may be appropriate.

The growth rate of consumption

The growth rate of consumption, g , attempts to capture the expected percentage change in real per-capita consumption. This is

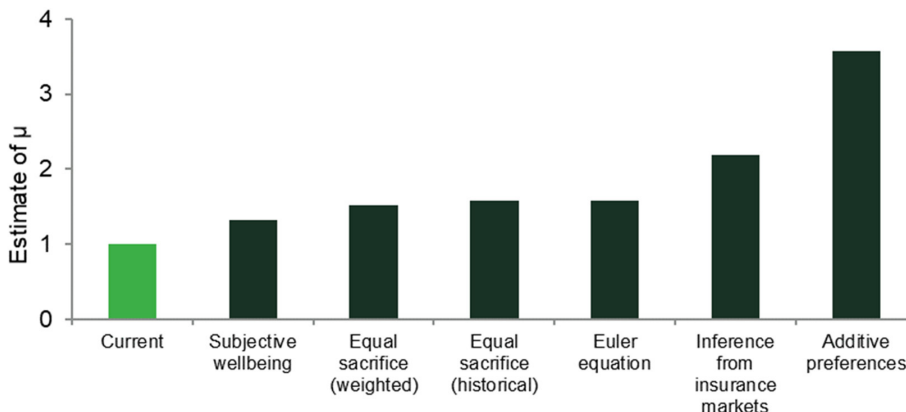


Figure 2 Comparison of current value of μ with empirical estimates by Groom and Maddison (2019)

Source: Groom, B. and Maddison, D. (2019), ‘New estimates of the elasticity of marginal utility for the UK’, *Environmental and Resource Economics*, 72, pp. 1155–82, table 6; HM Treasury (2018), ‘The Green Book’.

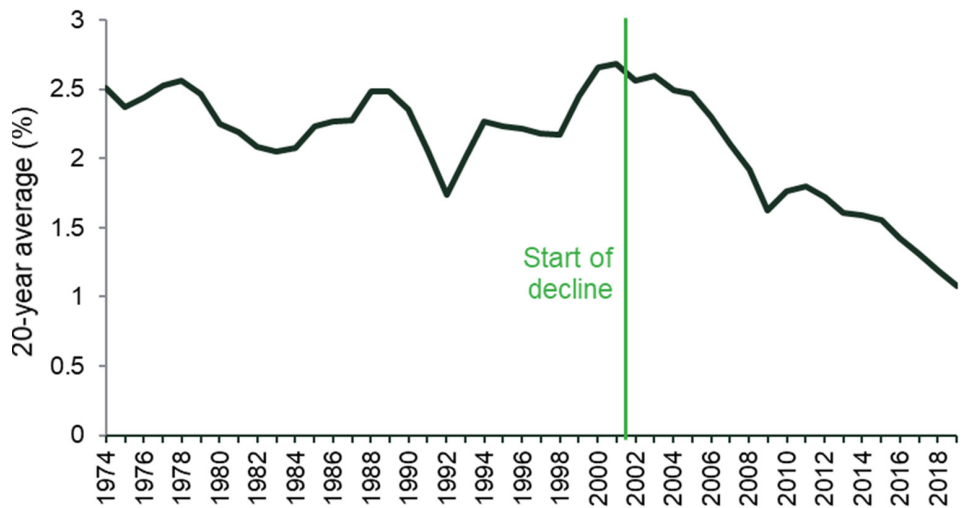


Figure 3 20-year average growth rates

Note: Values on the x-axis reflect the final year of the 20-year growth rate calculation—for example, when the x-axis reading is ‘1974’, the y-axis value is the average growth rate for the 20 years up to (and including) 1974.

Source: Office for National Statistics.

currently 2% in the Green Book. Higher rates of expected consumption growth increase the SDR because they imply that future generations will be wealthier, meaning that the benefit of providing them with additional consumption will be lower.

The 2% estimate is based on a 1.9% long-term real GDP growth forecast from the UK Office for Budget Responsibility (OBR), and historical estimates of consumption growth, which are typically also close to 2%. However, there are two reasons why this estimate may need to be revised downwards.

First, the precise estimate for g is highly dependent on the time period that is analysed. This is shown in Figure 3, which plots the average growth rate for the 20 years prior to the year on the x-axis, and demonstrates that these have declined since around the 1981–2001 period and are now closer to 1.5%. This is consistent with the latest Fiscal Sustainability report from the OBR, which expects consumer spending to increase by 1.4–1.5% per year once the

economy recovers from the COVID-19 pandemic.²⁷

Second, an optimal estimate of g should cover the consumption of both *economic* goods and services (i.e. those that are measured by the GDP and consumption statistics mentioned above) and *non-economic* goods and services (such as biodiversity, which is not measured in official statistics). If consumption of economic goods tends to reduce that of non-economic goods—for example, because the former causes pollution, which reduces biodiversity—then estimates of g based on GDP or consumption data alone would be overestimates, because they would ignore the fact that consumption of non-economic goods was decreasing. This would result in an overstated SDR.

Two possible approaches

Recent research suggests that many of the parameters in the SDR could be changed. In all cases except for μ and L , the indication is that they should be reduced. Table 2 overleaf shows that the net effect of defensible changes would be to reduce the SDR from 3.5% to 2.45%.

The reduction from 3.5% to 2.45% is not necessarily a change that **should** be made, but the material difference between the current approach and this alternative suggests that consideration ought to be given to whether the current approach is still appropriate in light of the latest evidence.

With the exception of L , the changes in Table 2 all represent updated estimates of parameters whose conceptual role has stayed the same. In the case of L , however, the parameterisation reflects the normative approach to systematic risk, where it is considered to be relatively immaterial

Parameter	Current value	Alternative value
δ	0.5%	0%
L	1%	0.2%
μ	1	1.5
g	2%	1.5%
Total	3.5%	2.45%

Table 2 Comparison of current and alternative SDR parameterisations

Note: All alternative parameter values can be found in the text above.

Source: Oxera

(in the public policy context). If the descriptivist approach were taken, the SDR would be considerably higher. The adoption of the normative approach should not be read as an endorsement, but merely as one defensible approach that is selected for illustrative purposes.

The second article in this series will cover further potential changes to the SDR. In combination with the above discussion, we will then present a consolidated picture.

Contact

mateusz.slomka@oxera.com

Mateusz Slomka

matthew.shepherd@oxera.com

Matthew Shepherd

leon.fields@oxera.com

Leon Fields

¹ The formula is derived from a model that assumes no externalities, identical individuals (i.e. no inequality), and no uncertainty.

² Oxera (2002), 'A social time preference rate for use in long-term discounting', 17 December, <https://bit.ly/367ljti>.

³ This is not the same as the relative weight of current and future generations' consumption, which is determined by the SDR as a whole. This means that the rate of pure time preference is conceptually distinct from the risk-free rate used in private project appraisal, often implemented through mechanisms such as the capital asset pricing model (CAPM), which weights intertemporal consumption.

⁴ For long-term interventions, however, analysts performing cost-benefit analyses of policy have to present their results with both $\delta = 0$ and $\delta = 0.5\%$.

⁵ The Green Book refers to a 2018 paper by Freeman, Groom and Spackman, which in turn refers to a 2015 paper by Drupp, Freeman, Groom and Nesje. See HM Treasury (2018), 'The Green Book', <https://bit.ly/368Czhl>, p. 102; Freeman, M., Groom, B. and Spackman, M. (2018), 'Social Discount Rates for Cost-Benefit Analysis: A Report for HM Treasury', <https://bit.ly/2S0QvIw>, p. 8; Drupp, M. A., Freeman, M. C., Groom, B., and Nesje, F. (2015), 'Discounting disentangled', *American Economic Journal: Economic Policy*, 10:4, <https://bit.ly/363vnU5>, pp. 109–34.

⁶ Arrow, K. J. (1996), 'Discounting, morality, and gaming', presented at the EMF-RFF Conference on Discounting, 24 December, <https://stanford.io/3i0jUa5>.

⁷ Stern, N. (2007), 'The Economics of Climate Change: The Stern Review', 30 October.

⁸ Intergovernmental Panel on Climate Change (2014), 'Social, Economic, and Ethical Concepts and Methods', <https://bit.ly/334k7oG>, p. 229.

⁹ Future Generations Commissioner for Wales, 'Well-being of Future Generations (Wales) Act 2015', <https://bit.ly/3i4zhOX>.

¹⁰ More precisely, it measures only those risks that are not explicitly modelled as streams of costs or benefits.

¹¹ HM Treasury (2018), 'The Green Book', <https://bit.ly/32YQ0yO>, p. 102.

¹² Freeman, M., Groom, B., and Spackman, M. (2018), 'Social Discount Rates for Cost-Benefit Analysis: A Report for HM Treasury', <https://bit.ly/33WtVjC>.

¹³ See, for example, the discussion of how to appropriately adjust for different types of risk on p. 4 of Freeman, M., Groom, B. and Spackman, M. (2018), 'Social Discount Rates for Cost-Benefit Analysis: A Report for HM Treasury', <https://bit.ly/2Eypm6x>.

¹⁴ A certainty equivalent is defined as the guaranteed amount of money that an individual would view as equally valuable to an uncertain amount. Projects that provide different levels of net benefit in high- and low-GDP states of the world create uncertainty, and policymakers generally prefer a project that provides lower net benefits on average, but which removes the uncertainty. The net benefit of a project with no uncertainty, that a policymaker is indifferent to relative to a project with uncertainty, is the certainty equivalent of the risky project.

¹⁵ Moore, M. A., Boardman, A. A., and Vining, A. R. (2017), 'Risk in Public Sector Project Appraisal: It Mostly Does Not Matter!', *Public Works Management & Policy*, 22:4, pp. 301–21.

¹⁶ Kochertakota (1996) estimates that this term would equal 0.076%, while Freeman, Groom and Spackman (2018) estimate that it would equal 0.075%. Kochertakota, N. R. (1996), 'The Equity Premium: It's Still a Puzzle', *Journal of Economic Literature*, 34:1, <https://bit.ly/3kLpPBK>, pp. 42–71; Freeman, M., Groom, B. and Spackman, M. (2018), 'Social Discount Rates for Cost-Benefit Analysis: A Report for HM Treasury', <https://bit.ly/2S5az6d>.

¹⁷ This is the estimate used in France. For example, see Freeman, M., Groom, B., and Spackman, M. (2018), 'Social Discount Rates for Cost-Benefit Analysis: A Report for HM Treasury', <https://bit.ly/3i0nZep>, p. 12.

¹⁸ Moore, M. A., Boardman, A. A. and Vining, A. R. (2017), 'Risk in Public Sector Project Appraisal: It Mostly Does Not Matter!', *Public Works Management & Policy*, 22:4, <https://bit.ly/3691fXw>, pp. 301–21.

¹⁹ Greaves, H. (2017), 'Discounting for public policy: A survey', *Economics & Philosophy*, 33:3, <https://bit.ly/2RZMU7i>, pp. 391–439.

²⁰ Stern, N. (2007), 'The Economics of Climate Change: The Stern Review', 30 October, <https://bit.ly/368lGCG>, p. 47.

²¹ This estimate covers only the probability of complete extinction, whereas for an accurate value of L , the probability of other societal collapses should also be included.

²² For precision, we reference consumption rather than income inequality because social welfare functions typically evaluate the welfare of a society based on consumption rather than income levels.

²³ Greaves, H. (2017), 'Discounting for public policy: A survey', *Economics & Philosophy*, 33:3, <https://bit.ly/3i3DtOS>, pp. 391–439.

²⁴ Dasgupta, P. (2008), 'Discounting Climate Change', *Journal of Risk and Uncertainty*, 37, <https://bit.ly/3mQmXoY>, p. 23.

²⁵ Oxera (2002), 'A social time preference rate for use in long-term discounting', 17 December, <https://bit.ly/332Cw51>.

²⁶ Groom, B. and Maddison, D. (2019), 'New estimates of the elasticity of marginal utility for the UK', *Environmental and Resource Economics*, 72, <https://bit.ly/33YS6xN>, pp. 1155–82.