THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Discount rates for long-term projects

Citation for published version:

Armitage, S 2017, 'Discount rates for long-term projects: The cost of capital and social discount rate compared', *The European Journal of Finance*, vol. 23, no. 1-3, pp. 60-79.
<https://doi.org/10.1080/1351847X.2015.1029591>

Digital Object Identifier (DOI):

[10.1080/1351847X.2015.1029591](https://doi.org/10.1080/1351847X.2015.1029591)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

The European Journal of Finance

Publisher Rights Statement:

This is an Accepted Manuscript of an article published by Taylor & Francis in The European Journal of Finance on 22/4/2015, available online: <http://www.tandfonline.com/10.1080/1351847X.2015.1029591>

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

**Discount rates for long-term projects:
the cost of capital and social discount rate compared**

Seth Armitage

University of Edinburgh*

February 2015

Forthcoming, *European Journal of Finance*

Research on the cost of capital and on the social discount rate (SDR) has developed largely along separate paths. This paper offers an overview and comparison of both concepts. The consumption-based theory of discount rates is common to both, but there are striking differences in how the cost of capital and SDR are estimated. A project's cost of capital is inferred in practice from market data, by a well-established package of techniques, and project risk makes a large difference. In contrast, the SDR is estimated by applying judgement about the welfare of future generations, in the setting of consumption-based theory. Project risk has tended to be ignored under the SDR approach.

Keywords: cost of capital; social discount rate; consumption-based theory; risk premium; declining discount rate

JEL codes: G31; H43; G12

*University of Edinburgh Business School, 29 Buccleuch Place, Edinburgh EH8 9JS; email seth.armitage@ed.ac.uk. I am grateful to Mark Freeman, Ian Hirst, Angelica Gonzalez, seminar participants at Heriot-Watt University and the University of Glasgow, and, especially, two anonymous referees and the Editor, Chris Adcock, for helpful comments.

1. Introduction

The discount rate has a massive impact on the present value (PV) of long-lived projects, especially on projects where most of the benefits arise in the fairly distant future, say after 100 years. Many projects have lifetimes of decades or centuries, including investment in infrastructure, housing, energy, research and development, and investment to protect the environment and reduce global warming. Such projects are found extensively in the private as well as the public sector.

The discount rate for projects in the private sector is known as the cost of capital, and is discussed in the finance literature. The discount rate for projects in the public sector is known as the social discount rate (SDR), and is discussed in the literature on public sector economics. Research on the two rates has developed quite separately, especially on the finance side where the SDR is almost never mentioned. Are the two really the same concept, or do they differ in some ways? This paper aims to provide an informative comparison of the two concepts. It should be of interest especially for readers with a finance background who are unfamiliar with the SDR. No systematic comparison has previously been made, though the applicability or otherwise of the cost of capital to public projects is a theme in research on the SDR. The debate about investment to alleviate climate change is a high-profile application of the SDR which is referred to in the paper.

The comparison enables several points to be highlighted. The first point concerns the role of theory. Consumption-based theory provides the analytical framework for thinking about discount rates in both approaches, so in this core respect the two are the same. But they differ very much in their process for the estimation of discount rates, in their use of the theory, and in the tradition that the SDR is a risk-free rate.

The finance approach uses theory to try to understand what determines discount rates, and to predict what discount rates should be observed, given plausible assumptions about risk aversion and the other variables in a consumption-based model. If the predictions from such a model differ from expected rates of return estimated from market data – and they tend to differ – the estimates from market data are the ones which are used for the purpose of estimating the cost of capital in practice. There is a well-developed package of techniques which is used for estimation in practice, outlined in Section 3.

In contrast, the SDR approach uses consumption-based theory to help guide the application of judgement by a public sector executive about the welfare of future people. Judgement involves appeal to evidence, including market data, but market data do not take priority. The amounts of utility from future cash flows can be judged to differ from the

amounts of utility that are implicit in a cost of capital estimated from market data. As a result, the SDR applied to a given project can differ from its estimated cost of capital, with a smaller SDR reflecting a greater weight for future utility. In practice SDRs tend to be lower, sometimes much lower, than estimates of the cost of capital for private-sector projects.

Second, risk and the premium for bearing risk are of central importance in the practical application of the cost of capital. Differences in estimated risk result in large differences in the cost of capital. Risk is much less prominent in the SDR approach. Much of the SDR literature uses the Ramsey formula for the discount rate, which assumes that the project is risk-free. In principle this is a mistake, since public projects do not in general provide risk-free benefits to society. The risk or consumption beta of a public project could be estimated from the covariance of the project's benefits with consumption per head. But even if risk is allowed for in public projects, the risk premiums that arise in standard consumption-based models are small, less than 0.5% per year, compared with the risk premiums estimated in practice for private sector projects. This is a corollary of the famous equity-premium puzzle: consumption-based theory predicts a much smaller premium than is estimated from market data.

Third, discounting using the cost of capital in practice always involves a constant risk premium over multiple future periods. The equity risk premium is usually estimated to be at least three per cent per year. This is the premium that is used in valuing project equity with average risk, ie a beta of around one. The impact of a premium of a few percentage points on the PV of a long-term project is enormous. Compounding a constant premium implies that the valuer's uncertainty about the future cash flows increases steadily with the distance of the cash flows into the future. If the valuer's uncertainty does not, in fact, increase in this way with the time horizon, it is not correct to compound a fixed premium. So the relationship between the valuer's uncertainty and the time horizon is critical. This is not a new point, but we suggest that it is one that deserves more prominence.

Fourth, there are arguments for using a declining discount rate for long-term projects. A declining discount rate means that the rates applied to future periods decline with distance into the future. The possibility of a declining rate for long-term projects is absent from the finance-driven package of techniques used to estimate the cost of capital. However, if the PV of a project is estimated as an average of the PVs using different possible discount rates, the PV using the lowest rate will increasingly dominate the average as project life increases. This procedure has the same impact on the best-guess PV as does discounting by a single rate that is lower for longer-term projects.

2. Consumption-based theory

2.1 Outline

The cost of capital and the SDR share a common theoretical background. An outline is needed in order to follow the subsequent discussion.¹ People are assumed to make decisions in order to maximise their lifetime utility, that is, the utility from their current consumption together with the utility from their expected consumption during the rest of their life. Collective behaviour is modelled by the behaviour of a representative individual, or by what we will call society. It is common, and convenient, to assume that society has a utility function with constant relative risk aversion (CRRA), which is of the form

$$U(C) = C^{1-\eta}/(1-\eta) \quad (1)$$

so
$$MU(C) = C^{-\eta} \quad (2)$$

where $MU(C) = dU(C)/dC$, and η is the elasticity of marginal utility, or the coefficient of relative risk aversion. If $\eta > 0$, the utility function is concave and society is risk-averse. Reasonable values for η are usually considered to be between one and four. They are based on empirical evidence about how risk-averse people are, and on judgement about what seems plausible. To gain a feel for the impact of η , suppose that consumption at the present date will turn out to be either plus or minus 10% (50%) of its expected value, with equal probability. An individual with $\eta = 2.0$ in equation (1) would pay 1% (25%) of the expected value of consumption to avoid this uncertainty. CRRA implies that these proportions are the same whatever the absolute amount of expected consumption.

Society maximises lifetime utility by saving to the point at which the marginal utility from saving, ie from the resulting higher consumption in the future, is the same as the marginal utility from consumption at date 0. Assume first that the future is known for certain. Consider an asset that provides a risk-free payoff t periods ahead. Investing one unit of consumption in the asset at date 0 provides e^{tr_F} units at date t , where r_F is the continuously compounded risk-free interest rate. The condition for maximising lifetime utility is

$$MU_0 = e^{tr_F} MU_t e^{-t\delta} \quad (3)$$

where MU_t is marginal utility at date t , and δ is the discount rate for utility, or the rate of pure time preference. A positive δ means that the utility to be experienced from a given amount of consumption at date t contributes less to lifetime utility than utility from the same amount of consumption today. The standard justification for a positive δ is that people are impatient; they prefer to consume now rather than in the future.

The interest rate is determined as the rate which solves equation (3):

$$r_F = (1/t)\ln(e^{t\delta}MU_0/MU_t) \quad (4)$$

Let be g be the continuously compounded growth rate of consumption, so $C_t = C_0e^{tg}$. With the CRRA utility function, equation (5) can be written as

$$\begin{aligned} r_F &= (1/t)\ln[e^{t\delta}C_0^{-\eta}/(C_0e^{tg})^{-\eta}] \\ &= \delta + (1/t)\{\ln(C_0^{-\eta}) - [\ln(C_0^{-\eta}) - t\eta g]\} \\ &= \delta + \eta g \end{aligned} \quad (5)$$

The interest rate is positive if $C_t > C_0$, and so $MU_t < MU_0$; if society will be better off in the future, the interest rate needs to be positive for society to be indifferent between saving and consumption at date 0. If future utility is discounted ($\delta > 0$), this is a second reason for a positive interest rate. Equation (555) is the formula presented in Ramsey (1928).²

Now assume that future consumption is uncertain. Specifically, assume that the growth rate per period $g_\tau = \ln(C_\tau/C_{\tau-1})$ is independent and identically distributed over time, i.e. it follows a random walk (arithmetic Brownian motion). In this case $C_\tau/C_{\tau-1}$ is lognormally distributed, and g_τ is normally distributed. Let the expected value of g_τ be μ and the variance be σ^2 . μ and σ^2 are the same for each future date. The expected value of consumption at date t is

$$\begin{aligned} E(C_t) &= C_0 E(\prod_{\tau=1}^t e^{g_\tau}) \\ &= C_0 e^{t(\mu + 0.5\sigma^2)} \end{aligned} \quad (6)$$

This uses the fact that, for any variable X (here, e^{g_τ}) which is lognormally distributed,

$$E(X) = \exp[E(\ln X) + 0.5\text{var}(\ln X)] \quad (7)$$

where $\text{var}(\ln X)$ = variance of $\ln X$. Because the covariance between g_τ and $g_{\tau-1}$ is zero under a random walk, the variance of consumption at date t is $e^{t\sigma^2}$.

The utility from investing one unit of consumption in the risk-free asset is now $e^{tr_F}E(MU_t)e^{-t\delta}$. So the equivalent of equation (5) with uncertain growth is

$$\begin{aligned} r_F &= (1/t)\ln[e^{t\delta}MU_0/E(MU_t)] \\ &= \delta + (1/t)\{\ln(C_0^{-\eta}) - [\ln(C_0^{-\eta}) - t\eta\mu + 0.5\eta^2t\sigma^2]\} \\ &= \delta + \eta\mu - 0.5\eta^2\sigma^2 \end{aligned} \quad (8)$$

where the second step uses the last line of (6). Equation (8) has an additional term which implies that the risk-free rate is negatively related to the variance of the growth rate σ^2 . Greater uncertainty about future consumption implies higher expected marginal utility from

risk-free saving, and a lower risk-free rate. However, this additional term is small for conventional values of η and σ^2 .

We now turn to risky assets. Assume that the uncertain payoff Y_{jt} from investing one consumption unit in a risky asset j is given by $Y_{jt} = \Pi_{\tau=1}^t e^{y_{j\tau}}$, where $y_{j\tau}$ is uncertain at each date between $\tau = 1$ and t . If we assume further that $y_{j\tau}$ is independent and identically distributed over time, it follows that

$$\begin{aligned} E(Y_{jt}) &= E(\Pi_{\tau=1}^t e^{y_{j\tau}}) \\ &= e^{t[\bar{y}_j + 0.5\text{var}(y_j)]} \end{aligned} \quad (9)$$

where $\bar{y}_j = E(y_{j\tau})$, and its variance $\text{var}(y_{j\tau})$, are the same for all dates. At this point we introduce the concept of the certainty-equivalent payoff of Y_{jt} , $CE(Y_{jt})$. This is the certain payoff that provides the same utility at date t as the uncertain payoff Y_{jt} . The utility at date t of Y_{jt} is $E[(Y_{jt})(C_t)^{-\eta}]$. The utility of its certainty-equivalent is $CE(Y_{jt})E(C_t)^{-\eta}$. Therefore,

$$CE(Y_{jt}) = E[(Y_{jt})(C_t)^{-\eta}] / E(C_t)^{-\eta} \quad (10)$$

Since $y_{j\tau}$ and g_τ follow a random walk, e^{y_j} , $(e^g)^{-\eta}$ and $e^{y_j}(e^g)^{-\eta}$ are lognormally distributed and the same for each period. It is the case that

$$\begin{aligned} \ln E(XY) &= E(\ln X) + E(\ln Y) \\ &\quad + 0.5[\text{var}(\ln X) + \text{var}(\ln Y) + 2\text{cov}(\ln X, \ln Y)] \end{aligned} \quad (11)$$

if X , Y , and XY are lognormally distributed. Using equations (6) and (11), equation (10) can be written as

$$CE(Y_{jt}) = \exp\{t[\bar{y}_j + 0.5\text{var}(y_j) - \eta\text{cov}(y_j, g)]\} \quad (12)$$

Using equations (9) and (12), the risk premium on the asset expressed as $E(Y_{jt})/CE(Y_{jt})$ is

$$\exp[t\eta\text{cov}(y_j, g)] = \exp(t\beta_j\pi) \quad (13)$$

where $\beta_j = \text{cov}(y_j, g)/\sigma^2$ is the consumption beta of asset j , and $\pi = \eta\sigma^2$ is the premium for systematic risk. Now we can write

$$CE(Y_{jt}) = E(Y_{jt})e^{-t\beta_j\pi} \quad (14)$$

The discount rate of the risky asset, ρ_j , can be found from

$$\begin{aligned} E(Y_{jt})e^{-t\rho_j} &= CE(Y_{jt})e^{-tr_F} \\ &= E(Y_{jt})e^{-t(r_F + \beta_j\pi)} \end{aligned}$$

$$\therefore \rho_j = r_F + \beta_j\pi \quad (15)$$

This is the consumption capital asset pricing model (CCAPM). The term $\beta_j\pi$ quantifies the expected risk premium on the asset. Equation (15) says that an asset with returns which covary positively with changes in consumption will be priced to give a positive risk premium.

Risk does not mean uncertainty about payoff *per se*, but the belief that the asset will pay least in recession.³

2.2 Risk and the time horizon

A further point, possibly less well known, concerns the interplay between risk and the time horizon. This is critical to the valuation of long-lived projects.⁴ Let the discount rate for risky asset j be θ_j , and the risk premium $\theta_j - r_F$ be positive and constant per period, as is normal practice. It does not matter how θ_j is arrived at; it need not be ρ_j in the consumption-based model. It is the case that

$$E(Y_{jt})/CE(Y_{jt}) = e^{t(\theta_j - r_F)} \quad (16)$$

Equation (16) makes a crucial point. A constant risk premium per period involves the implicit assumption that the total premium for risk increases exponentially, the further into the future the cash flow will arise.

For example, suppose $\theta_j = 10\%$ per year, $r_F = 5\%$, and we are considering two payoffs, one with an expected value of \$100 after one year and the other with an expected value of \$100 after five years. Then $CE(\$100_1) = \95.1 and $CE(\$100_5) = \77.9 , so the PV of the compensation for risk increases from $\$4.9/e^{0.1} = \4.4 for the year-one payment to $\$22.1/e^{0.5} = \13.4 for the year-five payment.

In the CCAPM above, compounding a constant premium involves assuming that covariance risk β_j , and the premium π per period for bearing risk, are constant per period, as is clear in equation (14). However, these outcomes are correct in the model only if our uncertainty is such that it is reasonable to assume that $y_{j\tau}$ and g_τ follow a random walk. Putting it approximately, compounding a constant premium is justified only if uncertainty about assets' payoffs and economic growth do in fact increase with the time horizon. Suppose, on the other hand, that we can forecast the payoff Y_{jt} with as much as confidence whether it arises after five years or after ten years. In this case it is clearly not correct to assume that the relation between time horizon and uncertain payoff can be expressed as $Y_{jt} = \prod_{\tau=1}^t e^{y_{j\tau}}$, where $y_{j\tau}$ follows a random walk. Compounding a constant risk premium per period is therefore not justified in this case, and would result in undervaluation of the project. The way in which the forecast risk of the cash flows for a given real-life project increases, as the time horizon lengthens, deserves thought for each project.

2.3 Predictions from consumption-based theory

The predictions depend critically on the values assumed for the four parameters. Weitzman (2007a) suggests that the following ‘quartet of twos’ are representative of the numbers that economists accept as reasonable: $\delta = 2.0\%$; $\eta = 2.0$; $\mu = 2.0\%$; $\sigma = 2.0\%$. The mean real growth rate and its standard deviation, μ and σ , are based on historical evidence from mature economies and on predictions for future growth. The rate of pure time preference and the coefficient of risk aversion, δ and η , are more debatable. Assume for simplicity that returns on the stock market are perfectly correlated with growth of consumption, and that both have the same volatility. In this case $\eta \text{cov}(r_j, g) = \eta \sigma^2$ in equation (13).⁵ Inserting the quartet of twos into equations (8) and (15) then gives the following predictions:

$$\begin{aligned} r_F &= 5.92\% \\ E(r_E) &= 6.00\% \end{aligned}$$

where $E(r_E)$ is the expected return on equity. So the predicted equity premium is 0.08% per year. These predictions compare with the following empirical numbers for mature economies, based on historical arithmetic averages over the last 100 years or so:

$$\begin{aligned} r_F &\approx 1\% \\ r_E &\approx 7\% \end{aligned}$$

and so the historic premium is approximately 6% per year. Many would regard a figure of 3% or 4% as more realistic for the equity premium looking forwards.

The two puzzles that arise from comparing the predicted with the observed numbers are (i) the observed equity premium is much higher than the model can explain, and (ii) the risk-free rate is much lower. There is an important difference between these puzzles. If $\delta = 0.1\%$ instead of 2.0%, and $\eta = 1.0$ instead of 2.0, as in the *Stern Review* (2006) on climate change, the predicted risk-free rate is only 2.1%, much closer to its historical average,⁶ whereas the equity-premium puzzle is still just as great. It is possible to predict a low risk-free rate, similar to the observed number, with low but still-plausible values of δ and η , whereas it is not possible to predict an equity premium of more than about 0.5% even with very high values of η . The fundamental empirical problem for the model is that the historic variance of consumption per head is not enough to justify much of an equity premium.

3. Finance in practice: estimation of the cost of capital

Although the theory just outlined is fundamental in finance, it is not used in practice to estimate the cost of capital. Practice is guided by a well established *practical finance package*

of ideas and methods of estimation, found in finance textbooks below advanced level, and used by companies, regulators of private-sector utility companies, and consultants.⁷ A project's cost of capital is the expected rate of return estimated on an asset with the same risk as the project. The CAPM is the leading model in textbooks and in practice used for estimation. The key point for our purposes is that the ingredients of the CAPM – the risk-free rate, the beta of the asset, and the equity risk premium – are estimated from market data, that is, data about traded financial assets, such as prices and dividends. There are ongoing debates about what the best methods are to estimate these ingredients, and also about the possible use of multifactor models of expected returns. The residual-income method of inferring the expected rate of return on individual shares has become popular in academic papers.

If the project's optimal capital structure includes debt, the project's cost of capital is the tax-adjusted weighted average cost of capital (WACC). The cost of a loan is approximated by the easily-observed promised interest rate (gross of income tax) on the face value of the loan, though this overstates the expected rate of return on risky debt. The WACC is used to discount cash flows from the project that are gross of interest on debt, net of corporation tax ignoring interest on debt, and gross of personal taxes.

Once a discount rate has been estimated, it is applied for each future period of the expected life of the project. Uncertainty about the assumptions underlying the cash flow forecasts, or about the discount rate, is usually captured by means of estimating a range of PVs using a range of assumptions and discount rates. There is little in the finance literature on how to estimate a project's risk, as opposed to the risk of a traded financial asset, nor on how to 'translate' that risk into a discount rate. The main method recommended is to find a listed firm in a similar business to the business of the project, and use estimates of the cost equity for the listed firm as the basis for the project's discount rate.

4.0 The social discount rate

4.1 The SDR approach

The term social discount rate, or rate of social time preference, normally means the rate applicable to long-term projects in the public sector. A public project could in principle be treated like a private sector project. For example, its cost of capital could be inferred using the (degeared) cost of equity estimated for a listed company thought to have similar business risk to that of the project. However, the modern SDR approach does not take the path of the practical finance package of seeking to infer the discount rate from market data. The approach is for a public sector executive to use consumption-based theory to help determine what the

SDR should be, whereas in finance the theory is used to try to explain observed expected rates of return on financial assets. The literature on the SDR takes as its foundation the Ramsey formula, $r_F = \delta + \eta g$. This involves assuming that both the growth rate of consumption and the project's cash flows are certain. These assumptions appear to be maintained for simplicity, and because of a tradition that public projects should be discounted by the risk-free rate (see Section 4.2).

The expected growth rate of consumption, μ , and its standard deviation, σ (if used), are based on empirical evidence in the SDR approach, as in finance. But the rate of pure time preference δ , and the coefficient of risk aversion η , are based explicitly on ethical judgement, as well as, in the case of η , empirical evidence.⁸ Dasgupta (2008, p. 150), for example, writes that the SDR 'has to be derived from an overall conception of intergenerational well-being and the consumption forecast'. Evidence from attempts to infer δ and η 'from the choices people make as they go about their lives' (p. 147) is an input to the estimation of those parameters, but it is not the only consideration. If the SDR for a project as calculated by an executive differs from an estimate of its cost of capital based on market data, the SDR takes priority and the discrepancy is not necessarily a major cause for concern.

One reason for rejecting market data is that market interest rates do not exist beyond a horizon of 30 years, which is the longest maturity of most government bonds. Undated government bonds do exist, but the markets for them are illiquid. Historic data could be used to estimate an expected rate of return over a long horizon from owning a succession of short-term Treasury bills, but the result would be approximate and dependent on assuming that the future will resemble the past. The focus on government bonds is perhaps a corollary of the view that funding for public projects is risk-free, and so the expected returns on risky assets are not relevant. If this view is rejected, the expected rate of return on equity potentially becomes relevant market-based evidence for risky projects, as in Weitzman (2007a).

The most important reason for explicit appeal to ethical judgement in the SDR is probably the view that market data, whether from bonds or equities, might not reflect enough concern for the welfare of future generations. The assumed aim of the government under the SDR approach is to maximise the utility of society, including the utility of people yet to be born (intergenerational utility). The assumed aim of an individual is to maximise her own lifetime utility. Some lifetime utility could come from anticipation of the utility of the individual's heirs, which would explain why people make bequests. Despite this, SDR proponents hold, first, that people alive today might, through their individual behaviour, act to maximise their own lifetime utility rather than intergenerational utility, and second, that they

might, nevertheless, elect a government which would act to benefit future generations at the expense of the current generation. This view is an example of what Sen and Williams (1982, p. 16) call Government House utilitarianism (see also Lind, 1982, pp. 55-9, and Dietz, Hepburn, and Stern, 2008, for a recent defence). Since market data are supposed to reflect the revealed preferences of individuals acting to maximise their lifetime utility, rather than intergenerational utility, market data are not to be trusted for the purpose of making decisions intended to maximise intergenerational utility. Put differently, people tend to ignore or undervalue externalities in their individual behaviour, and the effects of behaviour on future generations are externalities.

An important example is carbon emissions. ‘Business as usual’, with no mitigation of carbon emissions, involves an externality via probable global warming and the resulting imposition of possibly huge costs on future generations. So observed expected rates of return on assets such as equity, and therefore estimates of the cost of capital, exceed the social rates of return on such assets (Dasgupta, 2008; Stern, 2008). The social rate of return on an asset is the expected internal rate of return after the externalities attributable to the asset have been valued and ‘translated’ into cash flows to include in the cash flow forecast. Climate change is an example of a negative externality from current economic activity, one that in principle could be costed and included in cash flow forecasts.⁹ But even ignoring such specific and potentially measurable externalities, someone who believes that the behaviour of people alive today does not place enough weight on the welfare of future generations, will believe that the utility from long-term projects is greater than the utility that is implied by the observed expected rates of return on long-term assets.

Broome (1994), for example, maintains that ‘it is only the disenfranchisement of future generations that gives us the share of the world’s resources that we have’ (p. 152). In support of this, he presents a thought experiment in which a trust fund is set up which would act in the interests of future generations. He regards it as self-evident that ‘from the trust’s point of view... future commodities would be much more valuable than they seem to us who are participating in the market now’, and that the trust would transfer resources from the present to the future. But the trust’s purchases of future commodities would not reduce market interest rates permanently, as he assumes that interest rates are determined by the productivity of the economy’s technology.¹⁰ So if we took proper account of the welfare of future generations, we would use a lower SDR than market interest rates.

There are other possible grounds for incorporating ethical judgement in the SDR. A justification that differs from concern about future utility comes from the concepts of rights

and sustainability (for example, Anand and Sen, 2000). This perspective can lead to doubt about the adequacy of the expected utility approach. It can be argued that a fundamental ethical principal is that people have an equal right to well-being. From this follows the view that current decisions and consumption should be sustainable, meaning that they should not damage the prospects for the well-being of future generations. One version of the sustainability argument is that the Earth's 'capital', its potential to sustain life, should not be depleted. Current consumption should therefore be from 'income', ie from output that does not deplete the planet's capital. However, it is not easy to spell out in practical terms what sustainable development consists of.

There are also various special cases to consider, which invite a low discount rate, or other special treatment of investment decisions, or regulation. Some 'commodities', such as fresh air, or more generally a reasonably healthy environment, could be considered especially important to maintain prospects for well-being. This could justify, for example, a very low discount rate for public projects designed to maintain a reasonably healthy environment. Some activities, such as lifesaving, discussed by Broome (1994), provide utility which does not diminish as society becomes richer. Some commodities might be seen as essential for future well-being, and so as not substitutable at all for other commodities, in which case they will be regarded as necessary to have at almost any cost. Some features of the world might be given a special status because once lost they cannot be replaced, such as a species of animal or an archaeological site.

Finally, Stern (2006, 2008) emphasises that the opportunity cost of capital is a marginal concept. That is, it assumes that the project in question is small in relation to the market, meaning that the relevant market prices are not affected by whether or not the project is undertaken. He emphasises that it is a basic mistake to use this marginal concept in the context of climate change. Global warming is likely to affect market prices in future, and efforts to reduce it could affect market prices today. This argument is perhaps superfluous as a justification of the SDR approach as outlined, since the SDR approach does not rely on appeal to market data even for small projects.

We now consider separately the two parameters in equation (5), δ and η . Many authors support a value for δ of very close to zero, though 1.0% is common in the finance literature and some authors suggest higher values (very little is said about δ in finance, even in advanced texts). The key argument for $\delta \approx 0$ is that it is unethical to weight utility according to when the person is alive. The time a person is alive is, in itself, not a relevant consideration when it comes to weighting utility. The *Stern Review*, for example, concludes that δ ought to

exceed zero only to the extent that it reflects the possibility that humanity might not exist after some future date. The *Review* sets δ at 0.1% per year.

The parameter η provokes further questions. With the CRRA utility function, η measures both the rate at which the utility from marginal consumption declines as consumption grows over time, and aversion to uncertainty about consumption at a given date, ie aversion to risk. In the first of these roles, a higher η implies a desire for less inequality in levels of consumption over time (as does a higher δ). Marginal consumption today provides high utility compared with marginal consumption in the future. A person with a high η chooses higher consumption now and in the near future, and slower growth of consumption, compared with a person with a low η . A lower η implies greater concern about future welfare, and a lower SDR. A consequence of lower η is less consumption today, ie higher saving. Dasgupta (2008) and Nordhaus (2007) argue that a value of η of, say, 1.0, as in the *Stern Review*, implies that the proportion of income which would be saved is very high; 40% or more. This is uncomfortable because it is unrealistic that such high savings ratios will arise, and because it suggests that the current generation should consume much less than it actually does, for the benefit of the future, even though people in the future are already forecast to be substantially richer than are people today.

In its second role of reflecting aversion to risk, we have seen that lower η implies a lower risk premium, which exacerbates the equity-premium puzzle. However, for plausible values of η of up to about 4.0, differences in η make very little difference to the predicted premium in equation (15).

Perhaps more troubling is a comparison with a third role for η , that of reflecting concern about cross-sectional income inequality at a given date. Lower η in this role implies less concern about the welfare of the poor, because the marginal utility from increasing income for the poor increases with η . We have the awkward conclusion that a public sector executive concerned about current income inequality should apply a high η in her project appraisals, whereas she should apply a low η if she is concerned about the welfare of future generations. One answer is to use a utility function in which aversion to risk, and to income inequality at a given date, are separate from aversion to inequality of consumption over time. A number of papers explore such a utility function, including Gollier (2002a).

A general question about the SDR approach is, how do we agree on values of δ and η , and hence set the SDR?¹¹ Or if we are using a declining discount rate (below), how should the decline be determined? The values of δ and η are based partly on individual reflection under the SDR approach, so the values proposed will differ across SDR users. The published

responses to the *Stern Review* show how much disagreement there is about the SDR, as does the range of discount rates applied by different governments to public projects. Appeal to the evidence from clinical studies is rather inconclusive. Salient features about this evidence are as follows (Frederick, Lowenstein, and O'Donoghue, 2002). Measurement of the discount rates used by individuals is fraught with difficulties. The rates reported are highly heterogeneous across studies, and they are generally much larger than the rates of a few per cent used by governments for their SDRs. People apply high discount rates over short horizons, of up to about one year, and lower discount rates for longer future periods. With the important exception of this step reduction after one year, there is little clinical evidence that people apply declining discount rates.

The above discussion risks presenting use of the SDR-by-judgement approach as more firmly rooted than it actually is. The SDR was a rate based on market data in Lind (1982). Portney and Weyant (1999) summarise the deliberations of 20 leading public economists on the SDR in the context of climate change. There was agreement among them that a cost of capital based on market data should be used for projects with a life of up to 40 years, but that the SDR approach should be used for longer-term projects, because of 'discomfort' with the cost-of-capital approach for long-term projects. Weitzman (2007a) believes that a distinct SDR approach is not that of mainstream economics.

4.2 The SDR and risk

It has been 'commonly thought that the risk-free rate of return is appropriate for the appraisal of public projects due to the risk pooling available to governments' (Groom et al, 2005, p. 452; Arrow and Lind, 1970). This view has long been disputed, and the consumption-based theory outlined in Section 2.1 shows that, in general, it is not correct to ignore risk. A public project with uncertain future payoffs should only be discounted at the risk-free rate if the payoffs are uncorrelated with consumption per head, as is in fact assumed by Arrow and Lind (1970), and as is re-iterated by Lind (1982, p. 69). Neither the *Stern Review* (2006) nor any reviews of *Stern* take the view that government investment to alleviate climate change should be treated as risk-free by virtue of being funded by the government.¹²

The finance perspective on the risk of public projects can be summarised as follows. Both portfolio theory and consumption-based theory show that risk-averse investors demand a risk premium for exposure to systematic risk; that is, risk which is not eliminated through diversification by means of holding a portfolio of assets or projects. Each public project has its own discount rate, that depends on its systematic risk, as for private projects. Taxpayers

bear the risk of public projects. ‘Investing’ in public projects via paying taxes does not appear to offer greater elimination of risk than does investing in financial markets. Most unsystematic risk is eliminated by means of holding about 30 randomly selected shares, so a mature stock market provides more than enough opportunities for diversification. It is true that a government with a secure AAA credit rating can borrow at a lower cost of capital than the WACC of even the safest private-sector company. However, a sufficient reason for this is that a stable government can raise funds via taxation, which is a coercive method not available to a company. The coercive nature of taxation does not mean that the government has reduced project risk for the taxpayer.

Despite the importance of risk as a determinant of private sector discount rates, governments apply a single discount rate to all public projects (Spackman, 2008). A few governments now apply declining discounts rates, in which case the rate depends on the lifetime of the project. The fact that a single rate is used by a given government for all projects constitutes an important difference from the cost-of-capital approach. The practical finance package directs attention to differences in risk across companies and projects, which can result in large differences in the cost of capital, of several percentage points.

Risk looms much less large for the SDR, for a cluster of reasons. One is the mistaken tradition, just mentioned, that government funding in itself implies that the discount rate for all public projects is low. Second, the weight given to the welfare of future people, rather than risk, is seen as the primary determinant of the SDR. Guided by the Ramsey formula, future payoffs are discounted because society will be richer, and, if δ is non-negligible, because utility in the future counts for less than utility today. The payoffs are not discounted because they are risky.

Third, the measurement of risk will often be more problematic for public than for private projects. Public projects typically involve non-commercial objectives, and provide ‘payoffs’ that do not arise via cash flows. For example, what is the risk of investment to alleviate climate change? The cost of global warming is usually modelled as a fixed proportion of future output or consumption. In this case the size of the payoff, which is a reduction in that cost, is proportional to output, and the consumption beta is approximately one. A beta of one or more also arises in a model in which output before the impact of temperature increase has a linear positive effect on temperature, investment to reduce carbon emissions has a linear negative effect on temperature, and the marginal negative impact of temperature on output is increasing with temperature (Gollier, 2012). In this case the size of

the benefit of investment to reduce carbon emissions is positively related to underlying output, because the (non-marginal) relation between output and temperature is positive.

On the other hand, global warming might be sufficiently destructive as to have a major impact. Sandsmark and Vennemo (2007) assume simply that the non-marginal relation between temperature and output is negative. Let us assume also that reducing emissions has most impact on temperature if temperature growth will be high. Then this impact is negatively related to output, and the consumption beta is negative. Weitzman (2007a) argues that expenditure to reduce emissions is akin to the purchase of protection, an idea also in Sandsmark and Vennemo (2007). There might be a global catastrophe if emissions are not reduced, ie a large *fall* in consumption per head, with a probability that is non-negligible but impossible to quantify. On this view, the main benefit of reducing emissions is to reduce the probability of climate-induced catastrophe. If it were to turn out that there would be economic collapse under business as usual, and if our investment to reduce emissions were to succeed in averting the collapse, the payoff would be very large, and society's marginal utility without the payoff might have been higher than today's. This perspective implies a low or negative consumption beta for investment to alleviate climate change. This brief discussion suggests that our view about beta hinges on whether emissions, if not curtailed, could cause severe, non-marginal catastrophe in the future with a non-negligible probability.

Finally, the relevant conception of systematic risk is not settled. If the SDR is a rate determined by judgement in the setting of consumption-based theory, the systematic risk for a public project is measured by the correlation of the payoffs with consumption per head. But, then, the premium for risk is predicted to be small in consumption-based theory, as we have seen. The choice of consumption beta makes very little difference to the SDR in the standard model of Section 2.1, assuming that η and σ^2 are chosen from the range of conventional values. This potentially provides a justification for ignoring risk, though it is not a justification which has actually been used much in the SDR literature.

In contrast, if the SDR is estimated from market data, presumably risk is measured by beta in the CAPM, or by the betas in a multifactor model, by techniques that form part of the practical finance package. Use of market data implies an estimate of a premium for systematic risk of at least three per cent per year for a beta of one, and as a result the cost of capital is highly sensitive to the estimate of beta.

Our discussion here also raises the question of the relationship in practice between consumption and CAPM betas. The normal method of estimating the CAPM beta of a project, which is ultimately from the correlation of a share price with a stock market index, is

probably not an adequate method of estimating the consumption beta. There is no evidence about how estimates of the CAPM beta for private sector projects, using the practical finance package, compare with estimates of the consumption beta for the same projects. This is because consumption betas normally are not estimated for private projects. A second point is that the practical finance package is suited to estimating the risk and cost of capital of a project with cash flows. If a public project is undertaken by the private sector, there have to be forecast cash flows to provide payoffs to private finance. But the characteristics of the cash flows to private investors are influenced by the public sector sponsor. The risks faced by private investors, affecting the cost of capital set by investors, may well differ from the underlying risks of the project itself.

None of the above reasons mean that it is correct to assume that all projects are risk-free, or that they all have the same risk. At the same time, the measurement of the risk of real investment projects is a serious challenge for both public and private sector agents. A good deal of judgement is involved, and, for the cost of capital, the judgement made regarding choice of beta has a major impact on the discount rate. Although SDRs used in practice are not explicitly adjusted for risk, evidence from market data is likely to affect one's view of the SDR that is chosen. Plausible numbers, used by governments for the SDR, range from the *Stern Report's* 1.4% per year (in real terms) up to 6.0% or more. If a public sector executive chooses a low SDR, of around 1% to 2%, someone used to the private sector approach, in which risk matters, would see this SDR as implying a belief that public projects are close to risk-free. If the executive chooses a high SDR, of around 5% to 6%, this would imply a belief that public projects have a risk similar to investment in the stock market.

4.3 Declining discount rates

A further difference between the cost of capital and the SDR is that a declining SDR is sometimes proposed. This is another aspect of the use of consumption-based theory to determine what the SDR should be. The cost of capital in practice and in finance texts is always a flat rate, although there is an observable term structure of government bond yields. Conventional discounting, with a flat rate, is sometimes referred to as exponential discounting, and discounting with a declining rate as hyperbolic discounting.

There are a number of arguments for, or models that result in, a declining discount rate. First, if the expected growth rate of consumption per period is assumed to diminish with time into the future, ie the expected g_t for periods in the far future is lower than g_t in the near future, the discount rate per period predicted by equations (5) or (8) will diminish. If a g_t is

negative in a given period, marginal utility increases, and the discount rate for the period could be negative.

Second, the analysis under uncertainty in Section 2.1 assumes that g_τ shows random variation over time. A plausible type of modification is that the changes in g_τ show some persistence. This can result in a declining risk-free rate, since persistence implies that the dispersion of possible future consumption increases with the time horizon by more than is the case under random variation. A greater dispersion of possible future consumption implies a lower discount rate, since it increases expected marginal utility (see equation (8)). Gollier (2013) reviews several types of modification which feature persistence in changes to future growth.

Third, if g_τ follows a random walk, but the utility function exhibits diminishing relative risk aversion (DRRA) instead of CRRA, there will be a declining discount rate for a risk-free project (Gollier, 2002a). DRRA means that as an individual's wealth increases, she would pay a smaller proportion of her wealth to avoid a risk given in absolute terms. Both DRRA and CRRA imply aversion to risk (diminishing marginal utility of consumption), but DRRA implies less aversion than CRRA. Intuitively, there are two opposing forces at work. The expected variance of C_t increases with the time horizon t , which reduces expected utility. This means that the expected utility from a risk-free asset – that is, expected marginal utility – increases with t , which implies a lower discount rate. On the other hand, if g_τ per period is zero or positive, consumption increases over time, which reduces marginal utility, and implies a higher discount rate. DRRA is necessary for the first effect to dominate the second effect. If negative values of g_τ are possible, then DRRA does not guarantee a declining discount rate, and more is involved in specifying the set of utility functions that result in a declining discount rate (Gollier, 2002b).

Fourth, if the discount rate for each future period is fixed but uncertain ex ante, there is a declining discount rate in the sense that, for a given collection of possible discount rates, the single, fixed discount rate that represents the range of possible rates declines as the number of future periods increases (Weitzman, 1998). This argument is easiest to understand by means of an example (Guo et al, 2006). Let the possible discount rates be 1%, 3% and 5% per year, each with equal probability. What Weitzman calls the certainty-equivalent discount rate for t years, $CEDR_t$, is the discount rate which results in the same PV as the average of the PVs which arise from using each of the possible discount rates. $CEDR_t$ is calculated from the

certainty-equivalent discount factor, $CEDF_t$, which is the weighted average of the discount factors for the possible discount rates:

$$CEDF_t = 1/(1 + CEDR_t) = (1/3)[(1/1.01^t) + (1/1.03^t) + (1/1.05^t)] \quad (17)$$

For $t = 10$ years, $CEDF_t = 0.754$ and $CEDR_t = 2.86\%$; for $t = 100$ years, $CEDF_t = 0.143$ and $CEDR_t = 1.96\%$. The mechanism at work here is that, as the future horizon recedes, the lowest discount rate explains an increasing proportion of the PV. The argument only makes a difference if at least one of the possible discount rates is sufficiently low that PV is non-negligible. If the time horizon is 100 years or more, and the lowest of the possible discount rates is around 4%, PV is approximately zero even using the lowest possible rate.

Gollier (2004) criticises this argument on the grounds that there is an implicit assumption of risk neutrality, and that the current generation is exposed to the risk (see also Guo et al, 2005). The future payoff is viewed as certain, while the PV is uncertain until the uncertainty about the discount rate is resolved. The decision about whether to invest in the project is made before this uncertainty is resolved, by comparing the expected NPV with the cost. Risk neutrality comes in because of the expected NPV calculation: each possible PV is weighted by its probability, with no adjustment for risk (or different marginal utilities in different states of the world at the time when uncertainty about the discount rate is resolved). The current generation bears the risk because the PV once uncertainty is resolved could turn out to be higher or lower than its expected value. Gollier (2004) notes that the conclusion is reversed if, in contrast, the aim is to calculate the expected future value, given the cost of the project. For long horizons the expected future value becomes dominated by the outcome in which the rate of return is the highest of the possible rates, and so the rate to assume in calculating the expected future value tends towards the highest rate, not the lowest, as the future lengthens. From this perspective the generation at the terminal date is exposed to the risk, as the future value could turn out to be lower than the expected future value.

The debate can be resolved by considering the question in the consumption-based model. We assume that the uncertain discount rate is the rate of return on saving (Gollier and Weitzman, 2010, with a linear production function) or the marginal rate of return on saving (Gollier, 2013, pp. 103-5). In this case it can be shown that the certainty-equivalent discount rate declines as t increases. The reason is the same as the reason why persistence over time of the uncertain growth rate can result in a declining discount rate. The assumption that the uncertain return on saving is fixed (very persistent) once the uncertainty is resolved means that, from the perspective of date 0, when the discount rate is unknown, the dispersion of possible future consumption outcomes increases with the time horizon by more than is the

case under random variation of g_τ , for a given variance of g_τ equal to the variance of the uncertain discount rate. Since the term structure is flat if g_τ follows a random walk, the greater dispersion of consumption, compared with under a random walk, implies that the discount rate diminishes as t increases (assuming that $\eta > 1$).

This analysis re-establishes the case for a declining discount rate under the Weitzman assumption that the discount rate is uncertain now but will in future become known and fixed in perpetuity thereafter. However, the analysis is quite different from the original Weitzman (1998) argument. The latter is a point about the calculation of present value, and it is implicit in standard valuation procedures. The standard method of allowing for uncertainty about the discount rate is to calculate a range of possible PVs using several different discount rates. It is natural for the average of the possible PVs to be taken as the best-guess PV, so the Weitzman (1998) argument is merely highlighting an aspect of standard procedure. The Gollier-Weitzman analysis is about how the discount rate is determined in the consumption-based model of Section 2.1, given the Weitzman (1998) assumption.

We have outlined several of the conditions under which a consumption-based theory results in a declining discount rate for a risk-free asset.¹³ These ideas are having an impact on practice; the UK, French, Danish and Norwegian governments now apply declining discount rates, and other governments have been prompted to do so (Cropper et al, 2014).

A further question is the term structure of the risk premium. If changes in g_τ show persistence, equation (15) no longer holds; though the risk-free rate per period diminishes with the time horizon, the premium $\pi = \eta\sigma^2$ per period increases (Gollier, 2012). This makes the value of long-term projects more sensitive to systematic risk than under a flat premium. However, recent evidence from dividend strips (Van Binsbergen, Brandt and Koijen, 2013) and from property markets (Giglio, Maggiori and Stroebel, 2014) points to a declining risk premium inherent in market data. Perhaps in these markets uncertainty does not increase with the time horizon in the manner assumed when discounting a constant risk premium (Section 2.2).

4.4 The nature of the SDR

Consider a risky project, which has a single payoff with an expected value of \$100 in real terms in 100 years' time. The project is estimated to have a conventional CAPM beta of one. The discount rate derived from market data – the expected real rate of return on the stock market – is 5.0% per year, and with this discount rate $PV = \$0.8$. A public sector executive

estimates the SDR to be 1.4% per year, and with this rate $PV = \$24.9$. The cost of the project is \$10, so it is acceptable using the SDR but not using the cost of capital. This example reflects the essence of the debate highlighted by the *Stern Review*. Given the estimated future costs of global warming, society should invest much more to reduce carbon emissions than is currently being invested, if a discount rate of 1.4% is used to value the benefits of reduced emissions.

Estimation of the (systematic) risk of carbon-reducing investment is problematic, as we have noted, and the market-based cost of capital is sensitive to estimated risk. If private investors judge such investment to be low- or zero-beta, this would approximately eliminate the discrepancy between the cost of capital and the SDR in the example, assuming that the real risk-free rate is between 1% and 2% per year. However, the reason the executive in the example chooses a lower SDR than the cost of capital is unlikely to be because she estimates the risk of carbon-reducing investment to be lower than the market's estimate of the risk. The executive can agree that the CAPM beta for the project is one, and that the expected rate of return on equity is about 5% per year – and yet still choose a much lower SDR.

The reason the SDR for the project is 1.4% rather than 5.0% is that the executive thinks that society's utility from the uncertain payoff to be received in 100 years time exceeds its utility from \$0.8 today. The market does not 'think' in the same way, otherwise the market rate would be lower than it is. The SDR is a device to reflect judgements about the utility derived from payoffs at different future dates. An SDR which is lower than the market-based cost of capital is a means of indicating that, for reasons that depend on the judgement of the relevant executive, more resources should be transferred from the present to the future than would be the case were discount rates based on market data employed.

The reason for a below-market SDR is not because the cash flow from the project is underestimated. The *Stern* prescription, of much greater investment today to reduce carbon emissions, results both from recognition of the expected costs of unabated global warming, and from application of a low SDR. The low SDR is not to allow for the expected costs of global warming. This externality is recognised via a comprehensive cash flow forecast for emission-reducing investment. The low SDR is because the utility from the comprehensively measured cash flows is judged to be greater than the utility estimate that is implicit in the cost of capital. As a result, δ or η or both are set so that the discount rates produced by equations (5), (8) or (15) are below the discount rates estimated by the practical finance package.

The impact of long-lived projects on future utility is a type of externality, if we make the ethical judgement that a market-based cost of capital involves placing too little weight on

future utility. It is a special type of externality, one that requires adjustment to the discount rate rather than the cash flow forecast, because it is an omission of some of the utility derived from the forecast cash flows. The higher weight for future utility implied by the SDR could be captured by increasing the cash flows to be discounted, and then using the higher market-based cost of capital to discount these augmented cash flows. But then, we would not be using the SDR as the discount rate. Use of the SDR, rather than the cost of capital, *means* that we capture the different weighting of future utility in the discount rate.

The key counter-argument to support use of the cost of capital is that, in our example, the executive would make society better off in 100 years' time by not undertaking the project, and instead investing the \$10 cost for 100 years in the stock market (unless investors turn out to have overestimated long-term cash flows from the private sector). So it seems that the correct discount rate is the market rate. When funds 'can be invested at the market rate of interest [or in the stock market], whether by the private sector or the public sector, it is clearly undesirable to divert them to an investment that will return the lower social discount rate' (Brealey, Cooper and Habib, 1997, p. 20).

From the SDR perspective, the response to this is that society would be colluding in a mistake by investing in the stock market rather than emission reduction. The stock market undervalues projects to reduce carbon emissions, and in general, the market undervalues long-term projects in relation to short-term projects, on the SDR view. It is true that a given investor would be better off in terms of expected wealth by investing in the stock market. But society as a whole would be better off in terms of expected utility with more emission-reducing investment. The misvaluations in the stock market mean that the market fails to provide the incentives for the set of projects to be undertaken which would maximise intergenerational utility.

5.0 Conclusion

This paper has compared two rather separate approaches to the discount rate for long-term projects. The table below draws together the main features of the two approaches, and their points of difference.

Cost of capital

Role of theory	To understand the determinants of historic and expected rates of return on financial assets.
----------------	--

Estimation process	An established package of techniques to infer expected rates of return on assets and projects from market data. Consumption-based theory is not used for estimation.
--------------------	--

Key features	The minimum real WACC for low-risk companies, eg utilities, is currently at least 4% per year. There is uncertainty about the equity risk premium; estimates used in practice range from 3% to 8% per year. An asset's estimated (systematic) risk makes a large difference to WACC, because of the size of risk premium. The WACC applied is the same for each future year (no declining discount rate).
--------------	---

Comments	Mismatch between theory and evidence. Standard theory predicts a much smaller equity risk premium than 3%, and a larger risk-free rate than the historic 1%. A constant WACC per period requires an assumption that uncertainty increases with the time horizon.
----------	--

Social discount rate

Role of theory	To help determine what the discount rate should be.
----------------	---

Estimation process	Primarily judgement regarding the ingredients of the Ramsey formula. But no established procedure. Debate regarding the influence of expected rates of return inferred from market data. The difference in the estimation process between cost of capital and SDR means that the two approaches can produce substantially different estimates of the discount rate.
--------------------	---

Key features	Normal practice is to estimate a common SDR to apply to all public projects. There is no tradition of estimating project risk or adjusting SDR for risk. The concept of declining discount rates
--------------	--

is important in recent SDR theory, and is now incorporated into SDRs applied by some governments.

Comments

Estimates of the SDR tend to be lower than the discount rate that would be estimated using the practical finance package. The low end is around 1% per year, which is massively lower than a low-risk WACC from the practical finance package. This implies that markets undervalue long-term projects.

The message from the finance literature is that there is nothing special about long-term projects, or about public projects. All projects are in principle valued in the same manner, which we have summarised as the practical finance package. The resulting discount rates are high enough that a payoff arising beyond a few decades into the future has a negligible PV, compared with its undiscounted value. Many economists are uneasy about this, and the paper identifies some possible reasons why. Consumption-based theory, which has been developed to explain how the expected rates of return on financial assets are determined, has given rise to a huge academic literature, but has so far made little impact on the practical finance package.

The message from the current SDR literature is that it is probably not appropriate to try to apply the practical finance package to public projects, especially those which affect the welfare of future generations. A different approach is called for, one which is informed by the consumption-based model, and which involves explicit ethical judgement. At the time of writing there is still much doubt about how to arrive at a specific number for the SDR, and about how relevant market-based evidence is to the SDR.

Footnotes

¹ See Armitage (2005), Danthine and Donaldson (2005), or Gollier (2013), for fuller expositions, all relatively accessible in terms of mathematics.

² See Ramsey's equations (3) and (9); also Dasgupta (1982). Ramsey assumes non-increasing marginal utility, but does not specify a utility function.

³ In the derivation of the standard CAPM, the individual's existing consumption at date t comes entirely from the payoff from her holding of the market portfolio, so risk means positive covariance with the market portfolio.

⁴ Lind's (1982) review of the SDR draws attention to this question, but the question has since been somewhat ignored in both the finance and SDR literatures.

⁵ In fact the correlation between returns on the stock market and growth of consumption per head is 0.2 at most, and the standard deviation of annual returns on the market is about 17% (US data; Cochrane, 1997). This gives $\text{cov}(r_j, g) = (0.2)(0.17)(0.02) = 0.07\%$ instead of $\sigma^2 = (0.02)^2 = 0.04\%$.

⁶ In addition, Weitzman (2007a, 2007b) shows that, if g_τ is assumed to have a distribution with fatter tails than the normal distribution, the possibility of 'disaster states of the world' is increased, which reduces the implied risk-free rate.

⁷ See, for example, Graham and Harvey (2001), for survey evidence of company practice. For much more detail about implementation than appears here or in a textbook, see the documentation published by utility regulators and their advisers (for example, NERA, 2009).

⁸ The judgement involves ethics in that it directly involves taking an explicit view about how much the welfare of people in the future matters compared with the welfare of people today. The finance approach is to accept the 'view' about intergenerational welfare that is implicit in market data.

⁹ Dasgupta (2008) and others argue that, in the absence of 'market imperfections', society will maximise its lifetime utility and the social discount rate will be equal to the market rate of return on investment. His conception of market imperfections includes the existence of externalities which are not reflected in the rate of return on investment. There are more conventional types of market imperfection which might also make it difficult to infer from market data the revealed preferences of the current population regarding social decisions, including taxes and poor information on the part of the population.

¹⁰ Assuming a constant return on investment (ROI). The trust would reduce the interest rate if there were diminishing returns on investment. With a constant ROI , $r_F = ROI$, otherwise the

rate of saving would not be optimal. The growth rate in the Ramsey model is then $g = (ROI - \delta)/\eta$. The trust's activities imply a lower δ or η for society than would prevail without the trust, and a higher saving rate and growth rate.

¹¹ There is no established method of estimation from private sector data, ie no equivalent to the practical finance package. Earlier literature does consider which private sector rate of return, or average of such rates, to use in estimating the SDR. The main issue was not risk, but the fact that the (pre-taxes) rate of return on private investment is higher than the (after-taxes) rate on individual saving. Recent reviews, for example Gollier (2013), Dasgupta (2008), and Weitzman (2007a), almost ignore tax.

¹² Stern (2008, pp. 13-14) does argue that a close-to-risk-free rate is relevant for projects to reduce carbon emissions, but on the questionable grounds that they are 'likely to be financed via the diversion of resources from consumption (via pricing) rather than from investment'.

¹³ Aspects of the implementation of the Gollier (2002a) and Weitzman (1998) models are reviewed in Groom et al (2005).

References

- Anand, S. and A. Sen (2000), 'Human development and economic sustainability', *World Development* 28, pp. 2029-49.
- Armitage, S. (2005), *The Cost of Capital: Intermediate Theory*, Cambridge University Press.
- Arrow, K.J. and R.C. Lind (1970), 'Uncertainty and the evaluation of public investment decisions', *American Economic Review* 60, pp. 364-78.
- Brealey, R.A., I.A. Cooper and M.A. Habib (1997), 'Investment appraisal in the public sector', *Oxford Review of Economic Policy* 13, pp. 12-28.
- Broome, J., 'Discounting the future', *Philosophy and Public Affairs* 23, pp. 128-56.
- Cochrane, J.H. (1997), 'Where is the market going? Uncertain facts and novel theories', *Economic Perspectives, Federal Reserve Bank of Chicago* 21, pp. 3-37.
- Cropper, M. L., M.C. Freeman, B. Groom and W.A. Pizer (2014), 'Declining discount rates', *American Economic Review* 104, pp. 538-43.
- Danthine, J-P. and J.B. Donaldson (2002), *Intermediate Financial Theory*, Prentice Hall, Upper Saddle River, NJ.
- Dasgupta, P. (1982), 'Resource depletion, research and development, and the social rate of return', in R.C. Lind (ed.), *Discounting for Time and Risk in Energy Policy*, Resources for the Future, Washington DC, pp. 273-305.
- Dasgupta, P. (2008), 'Discounting climate change', *Journal of Risk and Uncertainty* 37, pp. 141-69.
- Dietz, S., C. Hepburn and N. Stern (2008), 'Economics, ethics and climate change', working paper, London School of Economics: www.personal.lse.ac.uk/dietzs/
- Frederick, S., G. Loewenstein and T. O'Donoghue (2002), 'Time discounting and time preference: a critical review', *Journal of Economic Literature* 40, pp. 351-401.
- Giglio, S., M. Maggiori and J. Stroebe (2014), 'Very long-run discount rates', working paper, ssrn.com/sol3/papers.cfm?abstract_id=2346049
- Gollier, C. (2002a), 'Discounting an uncertain future', *Journal of Public Economics* 85, pp. 149-66.
- Gollier, C. (2000b), 'Time horizon and the discount rate', *Journal of Economic Theory* 107, pp. 463-73.
- Gollier, C. (2004), 'Maximising the expected net future value as an alternative strategy to gamma discounting', *Finance Research Letters* 1, pp. 85-9.

Gollier, C. (2012), 'Evaluation of long-dated investments under uncertain growth trend, volatility and catastrophes', CESifo working paper, ssrn.com/sol3/papers.cfm?abstract_id=2202094

Gollier, C. (2013) *Pricing the Planet's Future: The Economics of Discounting in an Uncertain World*, Princeton University Press.

Gollier, C. and M.L. Weitzman (2010), 'How should the distant future be discounted when discount rates are uncertain?', *Economics Letters* 107, pp. 350-3.

Graham, J.R. and C.R. Harvey (2001), 'The theory and practice of corporate finance: evidence from the field', *Journal of Financial Economics* 60, pp. 187-243.

Groom, B., C. Hepburn, P. Koundouri, and D. Pearce (2005), 'Declining discount rates: the long and the short of it', *Environmental and Resource Economics* 32, pp. 445-93.

Guo, J., C.J. Hepburn, R.S.J. Tol and D. Anthoff (2006), 'Discounting and the social cost of carbon: a closer look at uncertainty', *Environmental Science and Policy* 9, pp. 205-16.

Lind, R.C. (1982), 'A primer on the major issues relating to the discount rate for evaluating national energy options', in R.C. Lind (ed.), *Discounting for Time and Risk in Energy Policy*, Resources for the Future, Washington DC, pp. 21-94.

NERA Economic Consulting (2009), *Cost of Capital for PR09: A Final Report for Water UK*, January, available at www.ofwat.gov.uk

Nordhaus, W.D. (2007), 'A review of the *Stern Review on the Economics of Climate Change*', *Journal of Economic Literature* 45, pp. 686-702.

Portney, P.R. and J.P. Weyant (1999), *Discounting and Intergenerational Equity*, Resources for the Future, Washington, DC.

Ramsey, F.P. (1928), 'A mathematical theory of saving', *Economic Journal* 38, pp. 543-59.

Sandsmark, M. and H. Vennemo (2007), 'A portfolio approach to climate investments: CAPM and endogenous risk', *Environmental and Resource Economics* 37, pp. 681-95.

Sen, A. and B. Williams (1982), *Utilitarianism and Beyond*, Cambridge University Press.

Spackman, M. (2008), 'Time preference, the cost of capital and PPPs', Conference on Discount Rates for the Evaluation of Public Private Partnerships', Queen's University, Ontario:

jdi-legacy.econ.queensu.ca/Files/Conferences/PPPpapers/Spackman-081002-final/pdf

Stern, N. and co-authors (2006), *The Stern Review: The Economics of Climate Change*, HM Treasury.

Stern, N. (2008), 'The economics of climate change', *American Economic Review* 91, 260-71.

Van Binsbergen, J. M. Brandt and R. Koijen (2013), 'On the timing and pricing of dividends', working paper, ssrn.com/sol3/papers.cfm?abstract_id=1710195

Weitzman, M.L. (1998), 'Why the far-distant future should be discounted at its lowest possible rate', *Journal of Environmental Economics and Management* 36, pp. 201-8.

Weitzman, M.L. (2007a), 'A review of *The Stern Review on the Economics of Climate Change*', *Journal of Economic Literature* 45, pp. 703-24.

Weitzman, M.L. (2007b), 'Subjective expectations and asset-return puzzles', *American Economic Review* 97, pp. 1102-30.