A Primer on Discounting Climate Risks

Martin L. Weitzman*

February 28, 2013

Preliminary and Confidential
Comments Appreciated

Abstract

The choice of an overall discount rate for climate change investments depends critically on how different components of investment payoffs are discounted at differing rates reflecting their underlying risk characteristics. Such underlying rates can vary enormously, from ≈1% for idiosyncratic diversifiable risk to ≈7% for systemic non-diversifiable risk. Which risk-adjusted rate is chosen can have a huge impact on cost-benefit analysis. In this paper I attempt to set forth in accessible language with a simple model what I think are some of the basic issues involved in discounting climate risks. The prototype application is calculating the social cost of carbon.

1 Introduction

Consider a long-term public-investment thought experiment used to calculate the social cost of carbon. Start with some baseline path of carbon dioxide emissions, representing some given climate-change policy, along with the corresponding uncertain future trajectory of the economy. Consider a simple variation of climate policy whereby one less ton of carbon dioxide is emitted now. This variation will result in a displaced trajectory of uncertain future outcomes. The displaced trajectory translates, ultimately, into uncertain incremental payoffs in each future period. Suppose that, for each period, some known fraction of payoffs has the systemic non-diversifiable risk profile of the macroeconomy as a whole, while the remaining fraction constitutes independent diversifiable risk. The decision maker wants

*Department of Economics, Harvard University (mweitzman@harvard.edu). Without necessarily tying them to the contents of this paper, I am grateful for useful critical comments on an earlier version to Christian Gollier, Michael Greenstone, Robert Pindyck, and James Poterba.

a simple capital budgeting rule expressed in terms of expected payoffs. At what project-specific risk-adjusted rates should expected future payoffs be discounted? This is a central question for calculating the social cost of carbon, and it is taken as the central question of this paper.

The effects of climate change will be spread out over what might be called the "distant future" – up to centuries and even millennia from now. The logic of compound interest forces us to say that what one might conceptualize as monumental events do not much matter when they occur in the distant future. Perhaps even more disconcerting, when exponential discounting is extended over very long time periods there is a notoriously hypersensitive dependence of cost-benefit analysis (CBA) on the choice of a discount rate. Seemingly insignificant differences in discount rates can make an enormous difference in the present discounted value of distant-future payoffs. In many long-run situations, including climate change, it may not be too much of an exaggeration to say that almost any answer to a CBA question can be defended by one particular choice or another of a discount rate.

A major difficulty with discounting climate change investments concerns the appropriate adjustments for risk. Climate change is characterized by deep structural uncertainties and the possible existence of really bad states we would like to insure against. This effect seems likely to be important in some situations and needs somehow to be incorporated into project-specific discounting.

What is the appropriate risk-adjusted discount rate schedule for a given public investment? Realistically, probably the most we can hope for is a theory that will frame a conceptual answer to this important question in understandable terms of more fundamental constructs, such as an appropriately defined real-project 'beta' along with the economy's underlying risky and riskfree interest rates. Suppose, for the purpose of this paper, that a "descriptive" or "positive" approach to discounting is adopted, which provisionally accepts previous historical values of the marginal product of capital in the real world as a proper guidance for future discount rates. (This assumption has been vigorously challenged in some parts of the literature,[1] but that is the subject of another paper.) The question then becomes: *which* real-world interest rate to use? Here two prototype real-world interest rates stand out. One is the economy-wide average return on all investments. The other is the so-called riskfree rate on safe investments. Unfortunately, the numerical difference between these two focal rates of return is enormous, leading to much debate and confusion about what risk-adjusted discount rates should be used for a particular project. The consequences can be spectacularly important for very long term CBA applications, like investments in mitigating climate change.

[1] See, e.g., Stern (2007).

The average return on all investments in a country is often proxied by the mean real historical return on a comprehensive index of equities traded on that country's stock exchanges. For the U.S., whose stock markets are relatively large in representing the private economy and which have a long uninterrupted historical record, this number is approximately seven percent per year.[2] The U.S. Office of Management and Budget uses 7% as "an estimate of the average pretax real rate of return on private capital in the U.S. economy."[3] Without further ado, for the purposes of this paper I identify the economy-wide average return on all investments as being $r^e = 7\%$.

The riskfree rate on a safe investment is typically proxied by the average real return on very short term U.S. treasury bills. This number is about one percent per year.[4] Once again proceeding without further ado, for the purposes of this paper I identify the relevant riskfree rate on safe investments (real and financial) as being $r^f = 1\%$.

Needless to say, it can make a stunning difference for long-term CBA outcomes whether distant-future payoffs are discounted at $r^e = 7\%$ or at $r^f = 1\%$. If a payoff a century and a half from now is discounted at $r^f = 1\%$ per year, its present discounted value is over *eight thousand* times greater than if the same payoff were discounted at $r^e = 7\%$ per year!

To see the striking effects of different discount rate assumptions, consider the social cost of carbon (SCC). The SCC was estimated recently by the U.S. Government Interagency Working Group on the Social Cost of Carbon, hereafter the USGI WG.[5] The USGI WG employed three "integrated assessment models" (IAMs).[6] An IAM is a computational model with dozens of equations that combine a very basic model of economic growth with a very basic model of climate change. The IAM is first run on the computer for some baseline socioeconomic scenario that specifies some actual path of CO_2 emissions. This will produce a series of outcomes, including a baseline time series of future consumption levels. If the IAM has key uncertain elements built into it, the baseline consumption levels will be uncertain. Tweak the IAM baseline emissions policy by forcing it to emit one less ton of CO_2 now, but otherwise leave climate change policy the same as the base case. This will produce a series of altered outcomes, including an altered time series of uncertain future consumption levels. The benefit payoff in any period is the change in consumption between the tweaked and baseline scenarios for that period. Compute by simulations the average

[2] See Campbell (2003) or Mehra and Prescott (2003).

[3] OMB (2003)

[4] See Campbell (2003) or Mehra and Prescott (2003). This is also very roughly the recent return on U.S. Treasury inflation protected thirty-year bonds.

[5] See US Working Group (2010). Also relevant is the discussion in Greenstone, Kopits, and Wolverton (2011). Nordhaus (2011) provides an interpretation, some criticisms, and some alternative estimates. See also Johnson and Hope (2012).

[6] The acronyms of the three IAMs are DICE, FUND, and PAGE.

benefit payoff (equals average consumption change) in each period. Pick some discount rate schedule and calculate the present discounted value of average benefit payoffs. This is the SCC. The USGI WG averaged five socioeconomic scenarios over three IAMs. The preferred discount rate was r=3%, which generated SCC=$21 per ton of CO_2, but sensitivity analysis was also performed for r=2.5% and r=5%. Table 1 shows the tremendous dependence of SCC on assumed constant values of r.

$r =$	7%	5%	3%	2.5%	2%	1.5%	1%
SCC=	≈$1	$5	$21	$35	$62	$122	$266

Table 1: SCC as function of *constant* discount rate[7]

Among those who look to real-world interest rates for guidance, there is some consensus that when the future payoffs on an uncertain public investment will be essentially proportional to the future level of the macroeconomy, representing non-diversifiable risk, then the appropriate discount rate for the project should be more or less the average rate of return on all investments in the economy, here $r^e = 7\%$. And there is also some consensus that when the future payoffs on an uncertain public investment will be essentially independent of the future level of the macroeconomy, representing diversifiable risk, then the appropriate discount rate for the project should be more or less the riskfree rate on a safe investment, here $r^f = 1\%$. Furthermore, there is a widespread if somewhat more vague sense that "in between" cases should involve "in between" discount rates, where the "in between" relative importance of each of the two prototype components reflects, in some way or another, the degree of covariance of their payoffs with the payoffs of the particular public investment being considered. The central issue for this paper concerns the appropriate value of these "in between" discount rates, in particular their time profile or term structure.

2 A CAPM-Inspired Decomposition of Payoff Risks

In this expository paper I use the standard familiar CAPM (capital asset pricing model) as my workhorse to explore risk adjustment for climate-change investments. The simple-minded CAPM is being used here as an expository device to get a handle on the basic issue of risk-adjusted discount rates, more than as a literal description. The CAPM model is well known, and my hope is that an average reader can relate to it. More sophisticated and more complicated formulations are possible, and they can give rise to different results depending on the specification that is chosen.[8]

[7] Source: For r=2.5%, r=3%, and r=5%, USGI WG. For r=1%, r=1.5% and r=2%, Johnson and Hope (2012). For r=7%, my own rough ballpark calculations.

[8] For example, it is possible to do everything in terms of marginal utility, instead of output, but it requires extra assumptions and complicates the exposition, not least because it is tricky to avoid formulations that

Let the random variable Y_t stand for the average payoff on all investments in the economy at time t. Essentially, output Y_t is intended to represent the systemic non-diversifiable risk of the macroeconomy itself. In the spirit of partial equilibrium analysis, the probability distribution of macroeconomic output $\{Y_t\}$ is treated as given while small variational investment perturbations around $\{Y_t\}$ are considered.

Consider a marginal investment project proposed at the present time zero. The project promises small payoffs of uncertain net benefits during future periods t, which are represented by the random variable B_t. (In the case of climate-change investments, B_t would typically represent the extra effective consumption or output from an extra unit of GHG mitigation.) The question we wish to address is the following. At what project-specific risk-adjusted discount rate should expected benefits $\mathbb{E}B_t$ during time t be discounted back to the present time zero?

Following the CAPM idea, suppose that project benefits can be represented in the two-factor linear form

$$B_t = a_t + b_t Y_t + \varepsilon_t, \tag{1}$$

where the random variable ε_t satisfies $\mathbb{E}\varepsilon_t = \mathbb{E}\varepsilon_t Y_t = 0$.

The parameter

$$b_t = \frac{cov(B_t, Y_t)}{var(Y_t)} \tag{2}$$

can then be interpreted mechanically as the coefficient from a hypothetical OLS regression of B_t on Y_t (while the coefficient a_t is equal to $\mathbb{E}B_t - b_t\mathbb{E}Y_t$). The regression here is largely hypothetical because in most cost-benefit situations we don't typically observe more than one realization of a payoff at any given time. This scarcity of data makes it very difficult to estimate b_t or a_t for public projects having no good analogue in the private sector. Such kind of generic problem seems intrinsic to unique or nearly-unique projects and bedevils empirical attempts to actually quantify values for b_t or a_t for particular projects. The real-project model of this paper therefore constitutes more of a CAPM-like organizing principle or conceptual framework than an actual usable structure ready to be taken to the data, as is much more characteristic of the financial version of CAPM theory. For better or for worse, I am relying on a reader's familiarity with CAPM-style thinking to conceptualize the treatment of climate risks here.

Formula (1) emphasizes that benefit B_t can be conceptualized as if it is a portfolio consisting of two components. The amount b_t of the portfolio replicates the system-wide non-

give rise to the equity-premium and riskfree-rate puzzles. A fully rigorous treatment of the contents of this paper in terms of marginal utility is given in Weitzman (2012). An alternative approach with some different conclusions is given in Gollier (2012).

diversifiable risk characteristics of the aggregate economy, as represented by a comprehensive index of all investment payoffs Y_t, whose return is r^e. In effect, Y_t represents macroeconomic output at time t. The amount $a_t + \varepsilon_t$ of the portfolio has diversifiable risk characteristics that are idiosyncratically independent of the rest of the economy, and whose return is r^f.

The two-factor linear decomposition of B_t represented by equation (1) might appear as innocuous, but it is an assumption nevertheless with consequences. Such a linear decomposition delivers simultaneously a clear portfolio-like conceptualization of benefits, a clean definition of a real-project beta, and a simple general equation for a declining risk-adjusted discount rate schedule expressed neatly in terms of r^e, r^f, and the real-project beta. Other specifications can yield different results, which are not typically solvable in closed form. In any event, the analytically tractable linear specification (1) is a natural point of departure for conceptualizing the risk properties of real-project payoffs and deriving neat results. In essence, I am forcing the problem into the partial equilibrium CAPM-style mold of (1) and then interpreting results in the light of that familiar framework. Assuming, then, a linear decomposition of risk factors, this paper inquires what are its consequences.

3 Risk-Adjusted Discount Rate Schedules

Suppose the following. If $a_t = 0$ in (1), then "everyone agrees" that the appropriate real interest rate for discounting benefits B_t is $r^e = 7\%$. And if $b_t = 0$ in (1), then "everyone agrees" that the appropriate real interest rate for discounting benefits B_t is $r^f = 1\%$. What should be the risk-adjusted discount rate when a_t and b_t are both non-zero? From the assumed linear decomposition of risk factors in (1), it must be the value r_t^* satisfying

$$\mathbb{E}B_t \exp(-r_t^* t) = a_t \exp(-r^f t) + b_t \, \mathbb{E}Y_t \exp(-r^e t). \tag{3}$$

By rough analogy with financial CAPM theory, define the real-project beta at time t to be the *fraction* of expected payoff that on average is due to the non-diversifiable systemic risk of the uncertain macro-economy:

$$\beta_t \equiv \frac{b_t \, \mathbb{E}Y_t}{a_t + b_t \, \mathbb{E}Y_t}. \tag{4}$$

The coefficient β_t is called the *real-project beta* because it plays a role in cost-benefit analysis akin to a financial investment beta in CAPM. As will be shown later, financial investment betas and real-project betas are identical for two-period short-run situations, but otherwise they may differ.

Use (4) and (1) to rewrite (3) as

$$\exp(-r_t^* t) = (1-\beta_t)\exp(-r^f t) + \beta_t \exp(-r^e t). \tag{5}$$

In the story being told by equation (5), β_t is the fractional share of benefit payouts subject to the risky discount rate r^e, while $1-\beta_t$ is the fractional share of benefit payouts subject to the riskfree discount rate r^f. It then stands to reason that the appropriate discount factor $\exp(-r_t^* t)$ for discounting expected benefits is the weighted average of the two discount factors $\exp(-r^f t)$ and $\exp(-r^e t)$, where the weights are the fractional benefit-payoff shares $1-\beta_t$ and β_t.

Equation (5) is the fundamental result of the CAPM-style model of this paper. There is no reason of principle why β_t should be constant over time. We could proceed with a general analysis of $\{r_t^*\}$ in terms of $\{\beta_t\}$ using formula (5), but for expository purposes I think it is more instructive to highlight primarily the benchmark case of a constant real-project beta.

4 The Benchmark of a Constant Real-Project Beta

Henceforth I assume as a benchmark default simplification, which will allow the exposition to focus more sharply on essentials, that $\beta_t = \beta$ is constant for all future periods $t \geq 1$. This is definitely an assumption, but one that seems like a natural point of departure for further analysis because β could always be conceptualized as an "average" future value of β_t serving as a prototype example for learning about the basic properties of risk-adjusted discount rates more generally. Or, β might be conceptualized as the limiting value of β_t with the analysis of this section holding rigorously for large t, which is the case of most interest here anyway. It is difficult enough to assign a real-world value to β even if it were constant, let alone to impose some time-dependent structure β_t on top of β. To be clear, there is no problem with allowing the real-project beta to be time dependent except for the more cumbersome notation and less facile interpretation that goes along with substituting β_t for β in derived formulas.

Without further ado, I therefore make an assumption of *constant proportions of risk*, meaning that the risk-variation characteristics of the payoffs B_t are decomposed into a constant proportion $1-\beta$ of independent idiosyncratic project-specific risk and a constant proportion β of systemic non-diversifiable risk representing the uncertain macro-economy itself. For all future periods $t \geq 1$, then,

$$\beta_t = \beta. \tag{6}$$

To emphasize its dependence upon the assumed value of β, henceforth r_t^* is denoted r_t^β. When the simplification (6) is imposed, equation (5) turns into the relatively neat formula

$$r_t^\beta = -\frac{1}{t}\ln\left((1-\beta)\exp(-r^f t) + \beta\exp(-r^e t)\right). \tag{7}$$

Equation (7) has a sufficiently simple form that the properties of r_t^β are easily analyzed. In what follows, $0 < \beta < 1$ and $r^f < r^e$.

Differentiating (7) and inspecting carefully the resulting expression indicates that the risk-adjusted discount rate becomes ever lower over time

$$\frac{dr_t^\beta}{dt} < 0\,. \tag{8}$$

Using l'Hôpital's rule to evaluate the indeterminate form (7) in the limit as $t \to 0$ gives

$$r_0^\beta = (1-\beta)\,r^f + \beta\,r^e, \tag{9}$$

which is a version of the famous CAPM formula. Thus, financial-investment CAPM betas and real-project betas coincide for a two-period model representing short run situations, but otherwise they may differ.

Using l'Hôpital's rule to evaluate the indeterminate form (7) in the limit as $t \to \infty$ gives

$$r_\infty^\beta = \min\{r^f, r^e\} = r^f. \tag{10}$$

What is the economic story behind the basic result that the risk-adjusted discount rate schedule declines over time (from an initial weighted average given by the CAPM formula down to approaching asymptotically the riskfree rate)? This property comes from a linear specification that results in a beta-weighted average of discount *factors*, rather than discount *rates*. The underlying idea is that having an insurance policy in the form of an investment in a riskfree asset that hedges against really bad tail outcomes is relatively more valuable over time (per unit of expected payoff) than having an investment in an asset replicating the risk characteristics of the economy as a whole.

5 Real-Project Betas and Discount-Rate Schedules

In practical terms, what is perhaps the most difficult stumbling block for applications of the discount rate schedule (7) to public investments in the real world is the estimation of actual project-specific values of the real-project beta coefficient β. This is a very tricky

subject worthy of further research.[9] My own feeling is that in many cases it may be difficult to go much beyond general considerations. However, even if β is in practice knowable only as a rough approximation to an average value, it is still useful to understand how its risk-adjustment role might be properly conceptualized. The attitude of this paper is that it is better to use some theoretically correct risk-adjustment formula for β, augmented by sensitivity analysis of β, than to do nothing about risk adjustment.

The following comments are not much more than speculative ruminations. If the public investment is essentially "privatizable" then I would say that there is a fair default assumption that the project-β is close to one. Perhaps a nationalized transportation sector or a nationalized energy sector belong to this category. If the public investment is in things that are very different from, and not readily substitutable for, material wealth, then, in the absence of any private-sector analogue, I might say that there is a fair default assumption that the project-β is close to zero. Perhaps investments in repairing the ozone layer or improving universal public health belong to this category. For unique one-off projects, like investments in mitigating climate change, it is going to be extremely difficult to estimate real-project betas because there is no close real-world substitute and no historical record from which data could be assembled. For singular public investments that seem strongly "non-privatizable" a constructive rule of thumb might be to *begin* with the default position of the project-β being set at about .5, which is midway between zero and one. This would at least get a conversation going and could always be changed after more serious discussions.

In any event, there is no evading the need to specify a value of β for any given investment and there is no question that this can be more of an art than a science for one-of-a-kind projects. With unique one-off public investments, like climate change, I personally find it somewhat easier to use a kind of "revealed beta" approach to work backwards from some postulated near-term CAPM discount rate r_0^β (which people have used in practice and for which I have some feel) to the underlying CAPM-implied "revealed beta" value of β. Inverting the CAPM equation (9) in this way gives a "revealed beta" value of

$$\beta = \frac{r_0^\beta - r^f}{r^e - r^f}\,, \tag{11}$$

where r^e=7% and r^f=1%. In Table 1 are displayed risk-adjusted discount rate schedules for seven representative near-term CAPM values r_0^β=1%, r_0^β=2%, r_0^β=3%, r_0^β=4%, r_0^β=5%, r_0^β=6%, and r_0^β=7%.

[9] Some relevant thoughts on this subject are expressed in Ewijk and Tang (2003).

t (yrs):	**t=0**	t=25	t=50	**t=100**	t=150	t=200	t=300
β=0	**1%**	1%	1%	**1%**	1%	1%	1%
β=1/6	**2%**	1.6%	1.3%	**1.2%**	1.1%	1.1%	1.1%
β=1/3	**3%**	2.2%	1.8%	**1.4%**	1.3%	1.2%	1.1%
β=1/2	**4%**	**3.0%**	**2.3%**	**1.7%**	**1.5%**	**1.3%**	**1.2%**
β=2/3	**5%**	3.9%	3.0%	**2.1%**	1.7%	1.5%	1.4%
β=5/6	**6%**	5.2%	4.1%	**2.8%**	2.2%	1.9%	1.6%
β=1	**7%**	7%	7%	**7%**	7%	7%	7%

Table 2: Risk-adjusted discount rates r_t^β (% per year, rounded off)

Note that, for "mid-range" values of $1/3 \leq \beta \leq 2/3$, which corresponds to near-term CAPM discount rates $3\% \leq r_0^\beta \leq 5\%$, the benchmark century discount rates are all fairly low, much closer to r^f=1% than to r^e=7%. Even for a real-project beta as high as β=5/6, which corresponds to a near-term CAPM discount rate r_0^β=6%, the century discount rate of 2.8% is appreciably closer to r^f=1% than to r^e=7%. All of this is a consequence of the enormous discrepancy between $r^f = 1\%$ and $r^e = 7\%$, which makes the term structure of risk-adjusted discount rates decline steeply over time in approaching the asymptotic limit of r^f=1%.

While there is no getting around the fact that the time schedule of risk-adjusted discount rates depends upon the assumed value of the real-project beta, the results of Table 1 suggest a strong downward pull over time. This is a basic message of the paper. The large equity premium of $r^e - r^f = 6\%$ tends to cause the time profile of risk-adjusted discount rates to tilt steeply downwards. The standard practice is to use the constant beta-averaged CAPM discount *rates* $r_0^\beta = (1-\beta)\, r^f + \beta\, r^e$ (given by the column t=0 in Table 1) instead of the beta-averaged discount *factors* $\exp(-r_t^\beta t) = (1-\beta)\exp(-r^f t) + \beta \exp(-r^e t)$ (which give rise to the declining risk-adjusted discount rate schedules displayed in Table 1). The message conveyed by Table 1 is that this standard practice (of conceptualizing risk adjustments by modifying via CAPM the discount rate while otherwise allowing it to be constant) could possibly have the potential for significantly biasing CBA against long-term investments whose real-project beta is less than one.

What is the appropriate real-project beta for an investment that reduces by one ton the present emissions of CO_2? This is a key question, the answer to which I don't think anyone knows. About the best we can do here, I fear, may be to tell stories. One insurance-like story would argue for a lower beta on the grounds that climate change itself is part of "effective" consumption, especially for very bad climate outcomes. Unknown uncertainties in climate-change feedbacks could lead to unforeseen catastrophic outcomes. In this bad-fat-

tail scenario, states of lower future effective consumption and higher marginal utility (from bad climate change) will be correlated with higher future benefits from current curtailment of CO_2.[10] This story treats mitigation as a low-beta hedge asset that helps to insure against climate disasters.

In the more standard story, which is told by the conventional IAMs more by default than premeditation, higher growth of conventional future consumption is associated with larger absolute damages from climate change. This occurs directly and mechanically because damages are assumed to be multiplicative in stochastic consumption realizations. It also occurs indirectly because higher growth is assumed to be associated with higher emissions and higher buildup of GHGs. In this standard story, good states of higher future consumption and lower marginal utility will be associated with higher absolute future benefits from current curtailment of CO_2, thereby implying a higher beta.[11] Thus, if standard IAMs are used to calculate a project beta then, with only a thin-tailed or no weighting of catastrophic climate outcomes, they will almost invariably come up with a relatively high value of β.

To summarize, without fat-tailed heavy weights on catastrophic climate damages, the middle-of-the-distribution IAMs will tend implicitly to choose higher values of a real-project beta for GHG mitigation investments. But a model with sufficiently heavy weight on fat-tailed catastrophic climate damages will tend to favor lower values of a real-project beta for GHG mitigation investments, by viewing such investments as insurance against potentially disastrous outcomes with high marginal utility. The key issue is whether to visualize uncertain climate-change damages in the middle-of-the-distribution range where they are likely dwarfed by growth uncertainty, or in the low-probability worst-case tail-risk scenarios of catastrophic climate outcomes sufficiently extreme to dominate the uncertainty about economic growth.

As an exercise, I recalculated the SCC using the risk-adjusted time-varying discount rate schedule from Table 1 that corresponds to the near-term CAPM discount rate of $r_0^\beta \approx 3\%$, which was used in practice as a central estimate by the USGI WG. (This corresponds to a "revealed beta" value of β=1/3). The calculated SCC for this case using USGI WG methodology is \$130 per ton of carbon dioxide, which is quite a significant difference from SCC=\$21.[12] From Table 1, this SCC = \$130 corresponds to an *as if constant* discount rate of $r \approx 1.5\%$.

The determination of the appropriate real project beta for calculating the SCC is a very difficult issue. As yet there is no easy answer to the question of what is the appropriate

[10] This is the approach taken in Sandsmark and Vennemo (2007).

[11] See, e.g., Nordhaus (2008) or Nordhaus (2011).

[12] I am indebted to Laurie Johnson for doing this calculation for PAGE, David Anthoff for FUND, and Antony Millner for DICE.

value of beta.[13] At the end of the day, I think the most we can hope for is to be aware of the basic issues and to try out various values of β in practice.

6 Concluding Comments

This paper suggests several themes.

The paper reinforces the idea that adjusting the discount rate to incorporate project risk represents a significant unsettled issue for CBA of long-term public investments. An economic analysis of climate-change investment policies, for example, depends enormously on what discount rate is chosen. This in turn requires resolution of the issues raised about how best to incorporate project uncertainty into a risk-adjusted discount rate. The resultant indeterminacy is undesirable, but seems unavoidable at this stage.

The default conceptualization of risk-adjusted discounting has mostly envisioned using a constant discount rate given by the CAPM formula $r_0^\beta = (1-\beta)\, r^f + \beta\, r^e$. (In principle it might be acknowledged that β_t should be allowed to depend on time, but in practice the discussion rarely gets this far because it is difficult enough to determine an average β for long-term public projects, much less to specify its time dependence.) The CAPM-style model of this paper is primitive and has a lot of simplistic assumptions built into it, many of which might be challenged. The assumption of a linear decomposition of risk variation suggests that what might be better combined in a beta-weighted average at time t is not the two focal discount rates r^e and r^f, but rather their two corresponding discount factors $\exp(-r^e t)$ and $\exp(-r^f t)$, via the formula $\exp(-r_t^\beta t) = (1-\beta)\exp(-r^f t) + \beta \exp(-r^e t)$. This implies a time-dependent discount rate r_t^β that declines over time from the initial CAPM value $r_0^\beta = (1-\beta)\, r^f + \beta\, r^e$ down to the asymptotic value $r_\infty^\beta = r^f$. Even if β is not known exactly, it is still useful to understand how its risk-adjustment role might be conceptualized.

Because there is such a significant equity-premium difference between discount rates of r^e=7% and r^f=1% per year, there can be an enormous discrepancy between the corresponding discount factors for time spans of a century or more. Other things being equal, this implies a relatively rapid decline of r_t^β and leads to the main empirical implication of the paper. The standard practice of incorporating risk adjustments by modifying a constant discount rate may have the potential for significantly biasing CBA against long-term investments whose real-project beta is less than one.

[13] Gollier (2012) attempts some rough calculations based on IAM specifications, but I am not sure I agree with the results because the procedure is based on the implicit premise that damages are multiplicative in actual realizations of stochastic consumption, which practically guarantees a high value of β. For an interpretation that is closer to my own, see also the very relevant analysis of Litterman (2012).

Finally, this paper is suggesting the importance of a research agenda that might put more effort into determining – if only very roughly and "on average" – the real project-specific betas for long-term public investments. Climate change in general, and the SCC in particular, leap to mind as obvious applications wanting further attention.

References

[1] Campbell, John Y. "Consumption-based asset pricing," in George M. Constantinides, Milton Harris and Rene M. Stulz, eds., *Handbook of the Economics of Finance*. Amsterdam: Elsevier, 2003, pp. 808-887.

[2] Ewijk, Casper van and Paul J.G. Tang (2003). "How to Price the Risk of Public Investment?" *De Economist*, 151 (3), 317-328.

[3] Gollier, Christian and Martin L. Weitzman (2010). "How should the distant future be discounted when discount rates are uncertain?." *Economics Letters*, 107 (3), 350-353.

[4] Gollier, Christian (2012). "Evolution of long-dated investments under uncertain growth trend, volatility, and catastrophes." November 2012 draft.

[5] Greenstone, Michael, Elizabeth Kopits, and Ann Wolverton (2013). "Developing a Social Cost of Carbon for U.S. Regulatory Analysis: A Methodology and Interpretation." *Review of Environmental Economics and Policy*, 7(1): 23-46.

[6] Johnson, Laurie T. and Chris Hope (2012). "The social cost of carbon in U.S. regulatory impact analyses: an introduction and critique." *J EnvironStud Sci*, DOI 10.1007/s134412-012-0087-7.

[7] Litterman, Robert (2012. "Tail risk and the price of carbon emissions." Mimeo dated 12/5/2012, forthcoming in *Regulation* (Cato Institute).

[8] Mehra, Rajnish. "The Equity Premium Puzzle: A Review." *Foundations and Trends in Finance*: 2008 Vol. 2: No 1, pp 1-81.

[9] Mehra, Rajnish and Edward C. Prescott. "The Equity Premium: a Puzzle." *Journal of Monetary Economics*, 1985, 15, pp. 145-161.

[10] Mehra, Rajnish and Edward C. Prescott. "The equity premium in retrospect," in George M. Constantinides, Milton Harris and Rene M. Stulz, eds., *Handbook of the Economics of Finance*. Amsterdam: Elsevier, 2003, pp. 889-938.

[11] OMB (Office of Management and Budget). 2003. *Circular A-4: Regulatory Analysis.* Ececutive Office of the Presicent, Washington, DC. http://www.whitehouse.gov/omb/circulars/

[12] Nordhaus, William D. (2008). *A Question of Balance: Weighing the Options on Global Warming Policies.* New Haven: Yale University Press.

[13] Nordhaus, William D. (2011). "The Social Cost of Carbon." NBER Working Paper 17540.

[14] Sandsmark, M., and H. Vennemo (2007). "A portfolio approach to climate investments: CAPM and endogenous risk." *Environmental and Resource Economics*: 37, 681-695.

[15] Stern, Nicolas (2007). *The Stern Review: the economics of climate change.* UK: Cambridge University Press.

[16] U.S. Interagency Working Group on Social Cost of Carbon (2010). *Technical Support Document: Social Cost of Carbon for Regulatory Impact Analysis Under Executive Order 12866.* Available at http://www.epa.gov/oms/climate/regulations/scc-tsd/pdf.

[17] Weitzman, Martin L. (2012). "Rare Disasters, Tail-Hedged Investments, and Risk-Adjusted Discount Rates." NBER Working Paper # 18496.