

Tail-Hedge Discounting and the Social Cost of Carbon

Martin L. Weitzman*

JEL

July 10, 2013

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 $\approx 1\%$ for idiosyncratic diversifiable risk to $\approx 7\%$ for systematic non-diversifiable risk. Which risk-adjusted rate is chosen can have a huge impact on cost-benefit analysis. In this expository paper I attempt to set forth in accessible language with a simple linear model what I think are some of the basic issues involved in discounting climate risks. The paper introduces a new concept that may be relevant for climate-change discounting: the degree to which an investment hedges against the bad tail of catastrophic damages by insuring positive expected payoffs even under the worst circumstances. The prototype application is calculating the social cost of carbon.

1 Introduction

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

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

value of distant-future payoffs. In a very-long-run situation like 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 unresolved 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 that we would like to insure against. It seems plausible that this insurance effect might be important in some climate-change investments and it needs somehow to be incorporated into project-specific discounting. The paper attempts to get at this aspect via a parameter quantifying the degree to which an investment hedges against the bad tail of catastrophic damages by insuring positive expected payoffs even in worst-case scenarios.

Suppose, for the purpose of this paper, that a “descriptive” or “positive” approach to discounting climate change 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 public project. As this paper will show, the consequences of how these two prototype rates are combined can be spectacularly important for very long term CBA applications – like investments in mitigating climate change.

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

¹See, e.g., Stern (2008).

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

³OMB (2003)

⁴See Campbell (2003) or Mehra and Prescott (2003). This is also very roughly the recent return on U.S.

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 in 2010 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 CO₂ 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 CO₂ 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 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 CO₂ in 2007 dollars, but sensitivity analysis was also performed for $r=2.5\%$ and $r=5\%$. Table 1 shows the tremendous dependence of SCC on the assumed constant value of r .

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

Table 1: SCC as a function of a *constant* discount rate⁷

Note that this SCC computation is very much an exercise in partial equilibrium analysis,

Treasury inflation protected thirty-year bonds.

⁵See US Working Group (2010). Also relevant is the discussion in Greenstone, Kopits, and Wolverton (2013). Nordhaus (2011) provides an interpretation, some criticisms, and some alternative estimates. See also Johnson and Hope (2012). In May 2013, using the same methodology, the USGI WG updated its previous estimates for the SCC. For consistency, I work here with the 2010 numbers.

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

⁷Source: 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 $r=7\%$, see Table 3.

which was required to make comparisons across models and scenarios. Benefits of expected consumption changes are calculated from one source-type (the three IAMs averaged over five socioeconomic scenarios). Then discount rates are exogenously applied (from another source-type, as it were) to calculate present discounted values. For better or for worse, in this paper I follow the same partial equilibrium approach, but I extend the source-type of discount rates to cover a particular form of risk adjustment. The “model” of this paper is extremely crude, its main justification being the extreme importance for climate-change analysis of the underlying set of discounting issues it is attempting to illustrate.

2 An Example: Linear Decomposition of Payoff Risks

Let the random variable C_t stand for “usable” or “effective” net consumption at time t , after subtracting off the damages from climate change. In this setup C_t represents the average payoff on all investments in the economy and can be interpreted at a high level of abstraction as embodying the systematic non-diversifiable risk of the macroeconomy itself. In the spirit of partial equilibrium analysis, the probability distribution of macroeconomic output $\{C_t\}$ is treated as given while small variational investment perturbations around $\{C_t\}$ 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 from an extra unit of CO₂ mitigation.) The critical 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?

Let the random variable A_t represent an idiosyncratic component of project benefit payoffs at time t that is independently distributed from C_t . One possible interpretation is that A_t represents merit goods that are not easily related to C_t like, perhaps, preservation of life, which is postulated to have its own value more or less independent of income or wealth.

The key assumption I now make is that benefits B_t can be linearly decomposed into a component that is proportional to A_t plus a component that is proportional to C_t :

$$B_t = B_t^A + B_t^C, \tag{1}$$

where

$$B_t^A \propto A_t \tag{2}$$

and

$$B_t^C \propto C_t \tag{3}$$

Formula (1) states that benefit payoffs B_t can be conceptualized as if coming from a portfolio consisting of two parts. The amount B_t^C of the portfolio payoff replicates the system-wide non-diversifiable risk characteristics of the aggregate economy, as represented by a comprehensive index of all investment payoffs C_t , whose return is assumed to be r^e . In effect, C_t represents macroeconomic output at time t . The amount B_t^A of the portfolio payoff has diversifiable risk characteristics that are idiosyncratically independent of the rest of the economy, and whose return is r^f .

Crudely inspired by the role of beta in the financial CAPM model, define the “real-project gamma” at time t to be the *fraction* of expected payoff that on average is due to the non-diversifiable systematic risk of the uncertain macro-economy

$$\gamma_t \equiv \frac{\mathbb{E}B_t^C}{\mathbb{E}B_t}. \tag{4}$$

It immediately follows from (1) and (4) that

$$1 - \gamma_t = \frac{\mathbb{E}B_t^A}{\mathbb{E}B_t}. \tag{5}$$

The coefficient γ_t is called the *real-project gamma* (at time t). It plays a role in cost-benefit analysis very roughly analogous to a financial-investment beta in CAPM. As will be shown later, the discount rate formula for financial-investment betas and real-project gammas are of the same form for a two-period short-run situation, but otherwise they may differ.

Combining (1) with (2) and (5) gives

$$B_t^A = \left[(1 - \gamma_t) \frac{\mathbb{E}B_t}{\mathbb{E}A_t} \right] \times A_t, \tag{6}$$

while combining (1) with (3) and (4) gives

$$B_t^C = \left[(\gamma_t) \frac{\mathbb{E}B_t}{\mathbb{E}C_t} \right] \times C_t. \tag{7}$$

The two-factor linear decomposition of B_t represented by equation (1) might appear to be innocuous, but it is an assumption nevertheless with consequences. Such a linear decomposition delivers simultaneously a clear portfolio-like conceptualization of benefit payoffs, a

clean definition of the “real-project gamma,” and a simple closed-form equation for a time-varying risk-adjusted discount rate schedule expressed neatly in terms of r^e , r^f , and γ_t . Other (nonlinear) specifications can yield different results, which typically require additional assumptions and 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.

Equations (1)-(7) will allow the interpretation that $(1 - \gamma_t)$ represents the strength of what might be called the “tail-hedge effect” because it quantifies the fraction of benefits that are expected to be paid even in bad-tail high-damage states of the world. Such an insurance concept seems like it might be especially relevant for discounting climate-change investments like GHG mitigation.

3 Risk-Adjusted Discount Rate Schedules

Suppose the following. If $\gamma_t = 1$ then “everyone agrees” that the appropriate real interest rate for discounting expected benefits $\mathbb{E}B_t$ is $r^e = 7\%$. And if $\gamma_t = 0$ then “everyone agrees” that the appropriate real interest rate for discounting expected benefits $\mathbb{E}B_t$ is $r^f = 1\%$. The question now is: what should be the risk-adjusted discount rate when $0 < \gamma_t < 1$? From the assumed linear decomposition of risk factors in (1), it must be the value r_t^* satisfying

$$\mathbb{E}B_t \exp(-r_t^* t) = \mathbb{E}B_t^A \exp(-r^f t) + \mathbb{E}B_t^C \exp(-r^e t). \quad (8)$$

Use (1), (6), (7), and cancel $\mathbb{E}B_t$ from both sides of the equation to rewrite (8) as

$$\exp(-r_t^* t) = (1 - \gamma_t) \exp(-r^f t) + \gamma_t \exp(-r^e t). \quad (9)$$

Invert equation (9) to obtain

$$r_t^* = -\frac{1}{t} \ln \left((1 - \gamma_t) \exp(-r^f t) + \gamma_t \exp(-r^e t) \right). \quad (10)$$

Equation (10) is the fundamental result of the model of this paper, in which it is assumed that payoff risks can be linearly decomposed into two primary risk factors. There is no reason of principle why γ_t should be constant over time. We could proceed with a general analysis of $\{r_t^*\}$ in terms of $\{\gamma_t\}$ using formula (10), but for expository purposes in an expository article I think it is more instructive to highlight primarily the benchmark case of a constant real-project gamma.

4 The Benchmark of a Constant Real-Project Gamma

Without further ado, for expository purposes I make a benchmark default assumption of *constant proportions of risk*, meaning that the risk-variation characteristics of the payoffs B_t are decomposed into a constant proportion $1 - \gamma$ of independent idiosyncratic project-specific risk $\propto A_t$ and a constant proportion γ of systematic non-diversifiable risk $\propto C_t$ representing the uncertain macro-economy itself. For all future periods $t \geq 1$, then,

$$\gamma_t = \gamma. \quad (11)$$

To emphasize its dependence upon the assumed value of γ , henceforth r_t^* is denoted r_t^γ . When the simplification (11) is imposed, equation (10) turns into the relatively neat formula

$$r_t^\gamma = -\frac{1}{t} \ln \left((1 - \gamma) \exp(-r^f t) + \gamma \exp(-r^e t) \right). \quad (12)$$

Equation (12) has a sufficiently simple form that the properties of r_t^γ are easily analyzed. In what follows, $0 < \gamma < 1$ and $r^f < r^e$.

Differentiating (12) and inspecting carefully the resulting expression indicates that the risk-adjusted discount rate becomes ever lower over time

$$\frac{dr_t^\gamma}{dt} < 0. \quad (13)$$

Using l'Hôpital's rule to evaluate the indeterminate form (12) in the limit as $t \rightarrow 0$ gives

$$r_0^\gamma = (1 - \gamma) r^f + \gamma r^e, \quad (14)$$

which is a version of the famous CAPM formula with γ playing the role of β . In this sense financial-investment CAPM betas and real-project gammas 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 (12) in the limit as $t \rightarrow \infty$ gives

$$r_\infty^\gamma = \min\{r^f, r^e\} = r^f. \quad (15)$$

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-type formula (14) down to approaching asymptotically the riskfree rate)? This property comes from a linear specification that results in a gamma-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 an asset with independent payoffs that hedges against really bad tail outcomes is relatively more valuable over time than having an investment in an asset replicating the risk characteristics of the economy as a whole.⁸

5 Real-Project Gammas and Discount-Rate Schedules

In practical terms, what is perhaps the most difficult stumbling block for applications of the discount rate schedule (12) to public investments in the real world is the estimation of actual project-specific values of the real-project gamma coefficient γ . 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 γ 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 γ , augmented by sensitivity analysis of γ , than to do nothing about risk adjustment.

For unique one-off projects, like investments in mitigating climate change, it is going to be extremely difficult to estimate real-project gammas 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- γ 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 γ 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 gamma” approach to work backwards from some postulated near-term discount rate r_0^γ (which people have used in practice and for which I have some feel) to the underlying “revealed gamma” value of γ . Inverting the near-term CAPM-formula-like equation (14) in this way gives a “revealed gamma” value of

$$\gamma = \frac{r_0^\gamma - r^f}{r^e - r^f}, \quad (16)$$

where $r^e=7\%$ and $r^f=1\%$. In Table 2 are displayed risk-adjusted discount rate schedules

⁸Some other models implying a declining discount rate are given in the comprehensive book of Gollier (2012).

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

for seven representative near-term values $r_0^\gamma=1\%$, $r_0^\gamma=2\%$, $r_0^\gamma=3\%$, $r_0^\gamma=4\%$, $r_0^\gamma=5\%$, $r_0^\gamma=6\%$, and $r_0^\gamma=7\%$.

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

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

Note that, for “mid-range” values of $1/3 \leq \gamma \leq 2/3$, which corresponds to near-term discount rates $3\% \leq r_0^\gamma \leq 5\%$, 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 gamma as high as $\gamma=5/6$, which corresponds to a near-term discount rate $r_0^\gamma=6\%$, 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 gamma, the results of Table 2 suggest a strong downward pull over time. This is a basic message of the paper. The large equity premium of $r^e - r^f = 6\%$ operating over long time periods tends to cause the time profile of risk-adjusted discount rates to tilt steeply downwards. The standard practice is to use the constant gamma-averaged short-term CAPM-formula-like discount rates $r_0^\gamma = (1 - \gamma)r^f + \gamma r^e$ (given by the column $t=0$ in Table 2) instead of the gamma-averaged discount factors $\exp(-r_t^\gamma t) = (1 - \gamma)\exp(-r^f t) + \gamma\exp(-r^e t)$ (which give rise to the declining risk-adjusted discount rate schedules displayed in Table 2). The message conveyed by Table 2 is that this standard practice (of conceptualizing risk adjustments by modifying the discount rate while otherwise allowing it to be constant) could possibly have the potential for significantly biasing CBA against long-term tail-hedge investments whose real-project gamma is less than one.

6 What is the Real-Project Gamma for Climate Change?

What is the appropriate real-project gamma for an investment that reduces by one ton the present emissions of CO₂? This is a key question, the answer to which I don't think anyone knows. About the best we can do here to gain intuition, I fear, may be to tell partial-equilibrium stories.

One insurance-like story would argue for a lower gamma on the grounds that climate change itself is part of “effective” consumption, especially for the very bad climate outcomes that might accompany business-as-usual high CO₂ scenarios. Unknown uncertainties in climate-change feedbacks, for example, could lead to unforeseen catastrophic outcomes with very low values of effective consumption C_t . In such a bad-tail scenario, there is a buffer or hedge built into the CO₂ mitigation investment that is expected to pay benefits B_t^A independent of how catastrophically low is C_t . The strength of this insurance-buffering effect is measured by $1 - \gamma = \mathbb{E}B_t^A / \mathbb{E}B_t$. Like the CAPM story about low-beta investments but with a time-dependent twist, this story views CO₂ mitigation as a low-gamma hedge asset that helps to insure against low-growth climate disasters.¹⁰

In the more standard story about a real-project gamma, economic growth is exogenous and higher growth of conventional future consumption is more strongly associated with larger absolute damages from climate change. This occurs directly and mechanically because damages are assumed to be strictly proportional to stochastic realizations of effective consumption.¹¹ It also occurs indirectly because higher growth is assumed to be associated with higher emissions and higher buildup of CO₂. In this standard story, good states of higher future consumption will be associated with higher absolute future benefits from current curtailment of CO₂, thereby implying a higher gamma.¹² Thus, if standard IAMs are used to calculate a project gamma then, with only a little or no weighting of catastrophic climate outcomes, they will typically come up with a relatively high value of γ .

To summarize, without relatively heavy weights on catastrophic climate damages, the middle-of-the-distribution IAMs will tend implicitly to choose higher values of a real-project gamma for CO₂ mitigation investments. But a model with sufficiently heavy weight on outlier catastrophic climate damages will tend to favor lower values of a real-project gamma for CO₂ mitigation investments, by viewing such investments more as hedge insurance against potentially disastrous outcomes. The key issue is whether to emphasize uncertain climate-

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

¹¹This multiplicative damages assumption is far from being innocuous and accounts for much of the policy-ramp gradualism that emerges from many IAMs because it allows higher output to substitute readily for higher temperatures. See Weitzman (2010) for more details.

¹²See, e.g., Nordhaus (2011).

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.¹³

7 The Social Cost of Carbon, Again

As an exercise, I had the SCC recalculated with the exact same methodology used by USGI WG but, instead of a constant discount rate, I employed the risk-adjusted time-varying discount rate schedule from formula (12) that produces the numbers shown in Table 2.¹⁴ The outcomes are displayed in Table 3 and generally result in quite high values for the SCC. For example, the near-term discount rate of $r_0^\gamma \approx 3\%$, which was used in practice as a central estimate by the USGI WG (and which corresponds to a “revealed gamma” value of $\gamma=1/3$) yields for this case $\text{SCC}=\$183$ per ton of CO_2 , which is quite a significant increase from the USGI WG central estimate of $\text{SCC}=\$21$. Even a near-term discount rate of $r_0^\gamma \approx 6\%$, which corresponds to a relatively high “revealed gamma” value of $\gamma=5/6$, yields for this case $\text{SCC}=\$45$ per ton of CO_2 or about double the central estimate of $\text{SCC}=\$21$.

$r_0^\gamma =$	1%	2%	3%	4%	5%	6%	7%
$\gamma =$	0	1/6	1/3	1/2	2/3	5/6	1
SCC=	\$266	\$228	\$183	\$140	\$92	\$45	\$1

Table 3: SCC (2007 dollars per ton of CO_2) as function of r_0^γ or γ

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

8 Concluding Comments

This paper touches upon 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

¹³This issue is further elaborated in the clear exposition of Litterman (2013).

¹⁴I am indebted to Laurie Johnson for doing these calculations for PAGE, David Anthoff for FUND, and Antony Millner for DICE. Without their help I could not have produced Table 3.

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 short-term CAPM-like equation of the form $r_0^\gamma = (1 - \gamma)r^f + \gamma r^e$. (In principle it might be acknowledged that γ 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 γ for long-term public projects, much less to specify its time dependence.) The model of this paper is an extremely primitive partial-equilibrium exposition with a lot of simplistic assumptions built into it, many of which might legitimately be challenged. The assumption of a linear decomposition of risk variation suggests that what might be better combined in a gamma-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 an equation of the form $\exp(-r_t^\gamma t) = (1 - \gamma)\exp(-r^f t) + \gamma\exp(-r^e t)$. This implies a time-dependent discount rate r_t^γ that declines over time from the initial value $r_0^\gamma = (1 - \gamma)r^f + \gamma r^e$ down to the asymptotic value $r_\infty^\gamma = r^f$. Even if the real-project gamma is not known exactly, it is still useful to understand how its risk-adjustment role might be conceptualized when a climate-change investment includes a component that hedges against the bad tail of catastrophic damages by insuring positive expected benefits even under the worst circumstances.

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^γ 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 gamma is less than one.

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 gammas 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*. Ams-

- terdam: 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 (2012). *Pricing the Planet’s Future: The Economics of Discounting in an Uncertain World*. Princeton: Princeton University Press.
- [4] 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.
- [5] 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.
- [6] Litterman, Robert (2013). “What is the Right Price for Carbon Emissions?” *Regulation*, Summer 2013, 24-29.
- [7] 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.
- [8] OMB (Office of Management and Budget). 2003. *Circular A-4: Regulatory Analysis*. Executive Office of the President, Washington, DC. <http://www.whitehouse.gov/omb/circulars/>
- [9] Nordhaus, William D. (2011). “The Social Cost of Carbon.” *American Economic Review: Papers and Proceedings 2008*, 98(2), 1-37.
- [10] Sandsmark, M., and H. Vennemo (2007). “A portfolio approach to climate investments: CAPM and endogenous risk.” *Environmental and Resource Economics*: 37, 681-695.
- [11] Stern, Nicolas (2008). “The Economics of Climate Change.” *American Economic Review: Papers and Proceedings 2008*, 98(2), 1-37.
- [12] 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-tds/pdf>.
- [13] Weitzman, Martin L. (2010). “What is the Damages Function for Global Warming and What Difference Might it Make?” *Climate Change Economics*, Vol. 1, No. 1, 57-69.