

NBER WORKING PAPER SERIES

RARE DISASTERS, TAIL-HEDGED INVESTMENTS, AND RISK-ADJUSTED DISCOUNT
RATES

Martin L. Weitzman

Working Paper 18496
<http://www.nber.org/papers/w18496>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2012

Without necessarily tying them to the contents of this paper, I am grateful to Bard Halsted, Ian Martin, and Nicholas Stern for useful critical comments. I am especially obliged to interactions with Christian Gollier, who, while having his own views on the subject, has emphasized early on (and well before me) the importance of understanding better the term structure of risk-adjusted discount rates for real investment projects. I also acknowledge stimulating discussions of this subject with Robert Litterman. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2012 by Martin L. Weitzman. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Rare Disasters, Tail-Hedged Investments, and Risk-Adjusted Discount Rates
Martin L. Weitzman
NBER Working Paper No. 18496
October 2012
JEL No. E43,G11,G12,Q54

ABSTRACT

What is the best way to incorporate a risk premium into the discount rate schedule for a real investment project with uncertain payoffs? The standard CAPM formula suggests a beta-weighted average of the return on a safe investment and the mean return on an economy-wide representative risky investment. Suppose, though, that the project constitutes a tail-hedged investment, meaning that it is expected to yield positive payoffs in catastrophic states of nature. Then the model of this paper suggests that what should be combined in a weighted average are not the two discount rates, but rather the corresponding two discount factors. This implies an effective discount rate schedule that declines over time from the standard CAPM formula down to the riskfree rate alone. Some simple numerical examples are given. Implications are noted for discounting long-term public investments and calculating the social cost of carbon in climate change.

Martin L. Weitzman
Department of Economics
Harvard University
Littauer 313
Cambridge, MA 02138
and NBER
mweitzman@harvard.edu

Rare Disasters, Tail-Hedged Investments, and Risk-Adjusted Discount Rates

Martin L. Weitzman*

October 19, 2012

Abstract

What is the best way to incorporate a risk premium into the discount rate schedule for a real investment project with uncertain payoffs? The standard CAPM formula suggests a beta-weighted average of the return on a safe investment and the mean return on an economy-wide representative risky investment. Suppose, though, that the project constitutes a tail-hedged investment, meaning that it is expected to yield positive payoffs in catastrophic states of nature. Then the model of this paper suggests that what should be combined in a weighted average are not the two discount *rates*, but rather the corresponding two discount *factors*. This implies an effective discount rate schedule that declines over time from the standard CAPM formula down to the riskfree rate alone. Some simple numerical examples are given. Implications are noted for discounting long-term public investments and calculating the social cost of carbon in climate change.

1 Introduction via Climate Change

Consider a long-term public-investment thought experiment used to calculate the social cost of carbon. Start with some actual 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

*Department of Economics, Harvard University (mweitzman@harvard.edu). Without necessarily tying them to the contents of this paper, I am grateful to Bard Halsted, Ian Martin, and Nicholas Stern for useful critical comments. I am especially obliged to interactions with Christian Gollier, who, while having his own views on the subject, has emphasized early on (and well before me) the importance of understanding better the term structure of risk-adjusted discount rates for real investment projects. I also acknowledge stimulating discussions of this subject with Robert Litterman.

dioxide is emitted now. This variation will result in a displaced trajectory of uncertain future outcomes. Waving away index number problems, suppose that “comprehensive consumption” is the universal numeraire of benefits or payoffs. The displaced trajectory translates, ultimately, into uncertain payoffs as measured for each period in terms of uncertain comprehensive-consumption increments. Suppose that, for each period, some known fraction of payoffs has the risk profile of the macroeconomy as a whole, while the remaining fraction constitutes independent diversifiable risk. The decision maker wants 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 problem of choosing a discount rate has long bedeviled cost-benefit analysis (CBA), but its relevance has been greatly magnified lately because economists are increasingly being asked to analyze projects or activities whose effects will be felt very far out into the future. Many examples concern environmental activities, including nuclear waste management, loss of biodiversity, groundwater alterations, minerals depletion, and so forth. The most prominent example is the economics of climate change, which will be used throughout this paper as a motivating example.

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 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.

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 critically important question clearly in terms of more fundamental constructs. Here two prototype discount 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 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.¹ 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."² 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.³ 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!

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, 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, 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 correlation 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.

In practical capital budgeting, a single constant discount rate is typically applied to all future cash flows. A standard default interpretation is that "the" appropriate project-specific discount rate is of the famous CAPM (capital asset pricing model) form $r = (1 - \beta)r^f + \beta r^e$. But suppose that the project constitutes a tail-hedged investment, meaning that it is *expected*

¹See Campbell (2003) or Mehra and Prescott (2003).

²OMB (2003)

³See Campbell (2003) or Mehra and Prescott (2003).

to yield positive payoffs in catastrophic states of the world. (Climate change mitigation is a prime example of a tail-hedged investment because payoffs in terms of comprehensive consumption net of climate damages are larger for high-damage outcomes.) Then the model of this paper argues that what should more plausibly be combined in a beta-weighted average at time t are not the two discount *rates* r^e and r^f , but rather the corresponding two discount *factors* $\exp(-r^e t)$ and $\exp(-r^f t)$, in the form $\exp(-r_t t) = (1 - \beta) \exp(-r^f t) + \beta \exp(-r^e t)$. Such a combination implies an effective discount rate schedule that declines over time beginning from the CAPM formula $r_0 = (1 - \beta) r^f + \beta r^e$ (corresponding to $t \rightarrow 0$) and monotonically decreasing towards the asymptotic value $r_\infty = r^f$ (as $t \rightarrow \infty$). The tail-hedged model of this paper provides a new and different story justifying a declining term structure of discount rates. The theory applies to capital budgeting more generally, but its strongest empirical implications are probably for public investments in tail-hedged projects having very long term payoffs, like climate change mitigation.

The next section presents succinctly the key background formulas and applications that are required for the paper, as derived from the simplest standard specification of consumption-based macroeconomic asset pricing. Using this background, a model of risk-adjusted discount rates is then developed and analyzed. Following presentation of the basic theory, the paper gives a few simple numerical examples of risk-adjusted discount rate schedules. I also make a few brief remarks about possible implications, in particular for calculating the social cost of carbon.

2 Consumption-Based Asset Pricing With Fat Tails

The formulation used in this paper is the simplest fruit-tree model of i.i.d. (independent identically distributed) consumption growth combined with an isoelastic utility function. In this pure exchange model of a dynamic competitive general equilibrium, asset markets are like phantom entities because no one actually ends up taking a net position in any of them. Gross returns on an asset are payoffs divided by asset price, with consumption as numeraire. This model is absolutely standard in the macroeconomic finance literature, so I just sketch the main results here very briefly.⁴

The major difference here with the standard presentation is that at the final stage in

⁴The famous fruit-tree model of asset prices in a growing economy traces back to two seminal articles: Lucas (1978) and Mehra and Prescott (1985). For applications, see the survey articles of Campbell (2003) or Mehra and Prescott (2003), both of which also give due historical credit to the other pioneering originators of the important set of ideas and the stylized empirical facts used throughout this paper. Citations for the many sources of these (and related) seminal asset-pricing ideas are omitted here only to save space and because they are readily available in the above two review articles. The presentation here follows most closely my own formulation in Weitzman (2007).

this specification I do not end the story by plugging into the formulas a normal distribution, which, with the usual calibrations, would yield the equity-premium and riskfree rate puzzles. Instead, I use a generic probability distribution that might resemble the calibrated normal probability density function arbitrarily closely but differs in having fatter tails. This implies that large-sigma small-consumption outcomes have relatively higher (but still vanishingly small) probabilities than the super-thin-tailed normal distribution.⁵ In putting a higher probability weight on very bad outcomes, tail fattening reduces the magnitude of both puzzles by simultaneously raising the equity premium and lowering the riskfree rate.

It should be emphasized at the outset that the notion of “consumption” used throughout this paper is intended to be much more inclusive than the ordinary conventional definition. Here consumption is comprehensive in the sense that it is ideally adjusted for externalities, environmental amenities, and so forth. For example, in the presence of climate change net consumption has already been adjusted here by the subtraction of climate-change damages from gross consumption. Such an all-inclusive comprehensive net measure of consumption will sometimes be called “effective” consumption for special emphasis. Effective consumption here is an all-encompassing measure of overall well being that characterizes the state of the economy. Essentially, effective consumption, denoted C , is that scalar index number which gives rise to utility U , and the utility function $U(C)$ is a function only of effective consumption C . Effective consumption is an abstraction, for which conventional consumption may, or may not, be a decent proxy, depending on the application. I assume heroically that conventional consumption and “effective” consumption have been similar in the past, but I allow them to possibly diverge in the presence of future climate change.

Let t denote the time period. The present period is $t = 0$. Effective consumption C_t in future period $t \geq 1$ is uncertain. The evolution of effective consumption is described by the stochastic process

$$C_t = C_{t-1}G_t, \tag{1}$$

where G_t in (1) represents the gross growth of effective consumption in period t . It is assumed that the random variables $\{G_t\}$ are i.i.d. positive.

The representative agent’s utility function is $U(C)$. Let ρ be the non-negative instantaneous rate of pure time preference or utility discount rate. The multiplicative factor for discounting utility of period t into current utility is $\exp(-\rho t)$. The representative agent’s

⁵I call the normal probability distribution “super thin tailed” because it is exponential *quadratic*, as opposed to exponential *linear* (the usual characterization of a thin-tailed distribution). The normal probability of a three-sigma event is .001, which drops to .00003 for a four-sigma event. These probabilities seem to me to be much too small for analyzing seriously many situations of tail-sensitive disastrous impacts.

welfare is

$$W = \mathbb{E} \left[\sum_{t=0}^{\infty} e^{-\rho t} U(C_t) \right], \quad (2)$$

where $\mathbb{E}[\cdot]$ represents the expectation operator.

A critical assumption that will allow enormous analytical simplification is that utility is given by the isoelastic power function

$$U(C) = \frac{C^{1-\eta}}{1-\eta} \quad (3)$$

with corresponding marginal utility

$$U'(C) = C^{-\eta}. \quad (4)$$

The utility function (1) displays constant relative risk aversion (CRRA), where the coefficient of relative risk aversion is the positive constant η .

The riskfree asset promises a fixed unit of effective consumption next period, independent of what will then be the state of the economy. The equilibrium price of the riskfree asset P^f (expressed in units of this period's effective consumption) must satisfy the equilibrium condition

$$P^f U'(C_0) = e^{-\rho} \mathbb{E}U'(C_1). \quad (5)$$

The gross return on the riskfree asset is

$$R^f = \frac{1}{P^f}. \quad (6)$$

Substituting from (1) and (4) into (5) and (6) gives the desired formula

$$R^f = \frac{1}{e^{-\rho} \mathbb{E}G^{-\eta}}, \quad (7)$$

which from symmetry applies to all time periods. The instantaneous riskfree rate corresponding to (7) is

$$r^f = \ln R^f. \quad (8)$$

Broad-based representative equity in this model is a claim on some fraction k of all future consumption dividends. The equilibrium price of representative equity P_t^e at time $t = 0$

(expressed in units of consumption in period 0) must satisfy the equilibrium condition

$$P_0^e U'(C_0) = \mathbb{E} \left[\sum_{t=1}^{\infty} e^{-\rho t} U'(C_t) k C_t \right]. \quad (9)$$

Substituting from (1) and (4) into (9), taking expectations, and using the formula for the sum of a geometric series gives the result

$$P_0^e = k C_0 \Omega, \quad (10)$$

where⁶

$$\Omega \equiv \frac{e^{-\rho} \mathbb{E} G^{1-\eta}}{1 - e^{-\rho} \mathbb{E} G^{1-\eta}}. \quad (11)$$

Symmetrically, the equilibrium price of representative equity P_t^e at time $t = 1$ (expressed in units of consumption in period 1, after C_1 is known) must satisfy the same equilibrium condition as (10), namely

$$P_1^e = k C_1 \Omega. \quad (12)$$

The realized gross return on equity between periods $t = 0$ and $t = 1$ is

$$R_1^e = \frac{k C_1 + P_1^e}{P_0^e}. \quad (13)$$

Substituting (1), (11), (12) into (13) and taking expectations gives

$$\mathbb{E} R_1^e = \frac{k \mathbb{E} C_1 + k \Omega \mathbb{E} C_1}{k H C_0}. \quad (14)$$

Plug (10), (11) into (14) and simplify, giving the desired formula

$$\mathbb{E} R^e = \frac{\mathbb{E} G}{e^{-\rho} \mathbb{E} G^{1-\eta}}. \quad (15)$$

Strictly speaking, formula (15) has been derived for $t = 1$, but by symmetry it applies to *any* time period t . The instantaneous rate of return on equity corresponding to (15) is

$$r^e = \ln \mathbb{E} R^e, \quad (16)$$

⁶It is easy to ensure that $\alpha \mathbb{E}[G^{1-\eta}] < 1$. For example, assume that $\mathbb{E}[G] > 1$ and $\eta > 1$.

while the equity premium is

$$\pi = r^e - r^f. \quad (17)$$

It is critical to realize that the above formulas hold for *any* assumed probability distribution of the random variable G . In what follows next I take a temporary detour to consider calibration of $\ln G$ to a normal probability distribution, but the paper is ultimately interested in fatter-tailed distributions than the super-thin-tailed normal.

Without loss of generality consider a transformation of random variable from G to Z , where $Z = \ln G$, or

$$G = \exp(Z). \quad (18)$$

Following the standard literature, assume, here temporarily and as an aside, that Z is normally distributed

$$Z \sim \mathcal{N}(\mu, \sigma^2), \quad (19)$$

which means that G is lognormally distributed.

Plugging (19), (18) into (7), (8) and taking the expectation of a lognormal distribution obtains the well-known augmented Ramsey formula for the riskfree rate

$$r^f = \rho + \eta \bar{g} - \frac{1}{2} \eta (\eta + 1) \sigma^2, \quad (20)$$

where $\bar{g} \equiv \ln \mathbb{E}G = \mu + \sigma^2/2$.

Next plug (19), (18) into (15), (16) and take the expectation of the resulting lognormal distribution. Combine the result with (20) and (17) to obtain the well-known formula for the equity premium

$$\pi = \eta \sigma^2. \quad (21)$$

As an example of calibration, without further ado I pick the following representative worldwide values: $\eta = 2.5$, $\bar{g} = 2\%$ annually, $\sigma = 4\%$ annually, $\rho = .5\%$ annually. Using these chosen values the riskfree rate calculated from formula (20) is $r^f = 4.8\%$ (instead of $r^f = 1\%$), which discrepancy is the so-called riskfree rate puzzle. With the chosen calibration values the equity premium calculated from formula (21) is $\pi = 0.4\%$ (instead of $\pi = 6\%$), which discrepancy is the so-called equity premium puzzle.

It is critical to realize fully that both of these empirical puzzles arise here under the assumption of a super-thin-tailed normal distribution of Z . A mean-preserving spread of G will simultaneously reduce the magnitude of both puzzles. (Note from (20) that the riskfree rate r^f decreases in σ^2 , while from (21) the equity premium π increases in σ^2 , which is a specific instance of this more general result applied to the particular lognormal family of probability distributions.) The generic argument goes as follows.

The function $G^{-\eta}$ is convex. It follows immediately from equations (7), (8) that a mean-preserving spread of G will *decrease* the riskfree rate r^f .

Using (16), (15), and (8), (7), formula (17) can be rewritten as

$$\exp(\pi) = \frac{\mathbb{E}(G/\mathbb{E}G)^{-\eta}}{\mathbb{E}(G/\mathbb{E}G)^{1-\eta}}. \quad (22)$$

Here and throughout the rest of the paper, assume $\eta > 1$. The convex function $(G/\mathbb{E}G)^{-\eta}$ has everywhere a higher relative curvature than the convex function $(G/\mathbb{E}G)^{1-\eta}$, while both functions are equal to one at $G/\mathbb{E}G = 1$. It follows that a mean preserving spread of G increases $\mathbb{E}(G/\mathbb{E}G)^{-\eta}$ proportionately more than it increases $\mathbb{E}(G/\mathbb{E}G)^{1-\eta}$. The implication from (22) is that a mean-preserving spread of G will *increase* the equity premium π .

Consider what would happen if a Student- t probability distribution were assumed for Z in place of the normal. The Student- t distribution is a sufficient statistic for a normal distribution with unknown parameters when drawn from a finite sample. The size of the sample corresponds to the number of degrees of freedom of the Student- t distribution. With enough degrees of freedom (i.e., a large enough sample size), the Student- t probability distribution appears indistinguishable from the normal. But the Student- t has a fat polynomial tail, instead of the normal's super-thin exponential-quadratic tail, and a CRRA utility function is very sensitive to the possibility of disastrously low levels of effective consumption. It is readily shown by going to the limit of (7), (8) that for a Student- t distribution $r^f \rightarrow -\infty$. When the indefinite integrals containing the Student- t distribution are simultaneously taken to the limit in the numerator and denominator of formula (22), it can be shown that $\pi \rightarrow +\infty$ (essentially because $\mathbb{E}(G/\mathbb{E}G)^{-\eta}$ in the numerator of (22) approaches infinity faster than $\mathbb{E}(G/\mathbb{E}G)^{1-\eta}$ in the denominator of (22)).⁷

At this point I rely on an intuitive continuity argument that there is a way to reverse-engineer a probability distribution whose tail properties are somewhere between a normal and a Student- t , which, with the above calibration, would yield $r^f = 1\%$ and $r^e - r^f = 6\%$. This hybrid between a normal and a Student- t would look almost indistinguishable from the calibrated normal, but would not give rise to the two asset pricing puzzles. A different but related route to tail fattening was taken by Robert Barro (2006), who appended an empirically plausible point distribution of a consumption disaster to the usual normal distribution and was thereby able to diminish significantly the magnitude of the two puzzles. I am thus adopting an approach here that “explains”⁸ the equity-premium and riskfree-rate

⁷These extreme results reflect the fact that the moment generating function of a Student- t distribution is infinite. More details about this entire set of issues are contained in Weitzman (2007), including attribution of the basic ideas.

⁸Many hundreds of articles with scores of different approaches have been used to “explain” the equity-

puzzles by a fatter-than-normal probability distribution of $Z = \ln G$.

However it is accomplished, I henceforth assume that the probability distribution of G has been implicitly chosen (along with implicitly chosen reasonable values of the first two moments of the probability distribution of G and implicitly chosen reasonable values of the key parameters η and ρ) so that formulas (7), (8) imply $r^f = 1\%$ and formulas (15), (16) imply $r^e = 7\%$. Technically G can have *any* probability distribution that is consistent with the desired results $r^f = 1\%$ and $r^e = 7\%$, but is otherwise arbitrary. It is not necessary to explicitly specify the details, but I am implicitly fitting the tails of the probability distribution of $Z = \ln G$ to be fatter than the normal in such a way that the underlying model of this paper is puzzle free. This model is such an over-simplified version of reality that it strains credulity, but at least the fatter-tailed version does not technically carry the stigma of questionable legitimacy that comes from generating the equity-premium and riskfree-rate puzzles.

3 A Model of Cost-Benefit Analysis

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 (expressed in units of future effective consumption C_t). Suppose that the project costs the equivalent of a small sure reduction of effective consumption δ in the present period (expressed in units of present effective consumption C_0). Should the proposed project be undertaken?

We seek a project-specific discount rate schedule, denoted $\{r_t\}$, such that undertaking the proposed investment will improve welfare if and only if the project's present cost δ is less than its expected present discounted benefits

$$\sum_{t=1}^{\infty} \exp(-r_t t) \mathbb{E}B_t. \tag{23}$$

We cannot make further progress on deriving and understanding the appropriate risk-adjusted discount rate schedule without specifying in more detail the risk properties of the uncertain future benefits. We seek to decompose B_t into familiar and analytically-tractable components. We are looking for some simple guideline conceptualization that would allow

premium and riskfree-rate puzzles. See Mehra (2008) for a critical discussion of the various proposals that have been put forth. My favored explanation, which the model of this paper follows, belongs to a subset of explanations that are based on fat-tailed probabilities of outlier catastrophic events. More details are described in Weitzman (2007), including attribution of the basic ideas.

a plausible extension of the “beta” decomposition of financial returns, which is familiar from the CAPM theory of finance, into the domain of real-project payoffs. It turns out that this is not a trivial task because financial returns are different entities from project payoffs, having different units of measurement and requiring some kind of normalization to make them comparable. A financial investment beta is expressed in terms of the risk properties of *financial returns* (which is natural for financial asset pricing), whereas real-project benefits come expressed in terms of the risk properties of *real payoffs* (which is natural for cost-benefit analysis). Financial returns come in a form that is already endogenously normalized via dividing payoffs by their asset prices, whereas real-project payoffs must be exogenously normalized to get them expressed in a form that highlights the decomposition of their essential risk characteristics.

For specificity, the investigation of this paper is focused sharply on CBA of small investments that only incrementally impact distant-future events. However, many climate change investments may themselves be big and involve large impacts. Yet even after some of these big investments have been made, the issue of choosing the appropriate risk-adjusted discount rate for additional marginal investments remains important and is emblematic of a whole series of thorny issues concerning how to conceptualize and how to evaluate risky actions spanning across many generations. The discount rate represents the current price of future benefits or costs, and as such its risk adjustment is a convenient shorthand vehicle for studying the general nature of how investment tradeoffs depend upon underlying risk relationships between the investment and the rest of the economy.

The derivation of the values of r_t to be used in (23) can begin as follows. Start by examining a hypothetical point-input point-output investment project whose single payoff at time t is given by B_t . Let h be the hypothetical loss of effective consumption at time zero which would leave the representative agent exactly indifferent between accepting or rejecting this hypothetical project. Then h must satisfy the condition

$$U'(C_0) h = e^{-\rho t} \mathbb{E}[U'(C_t) B_t]. \quad (24)$$

By definition, the corresponding discount rate r_t appropriate to (23) obeys the equation

$$h = \exp(-r_t t) \mathbb{E}B_t. \quad (25)$$

Combining (26) with (25) and eliminating h gives the basic equation

$$\exp(-r_t t) = \frac{e^{-\rho t} \mathbb{E}[U'(C_t) B_t]}{U'(C_0) \mathbb{E}B_t}. \quad (26)$$

I now make a fundamental distinction between two different types or classes of investments. A “tail-hedged” investment is *expected* to pay positive benefits in catastrophic states of nature. A “tail-exposed” investment is *not expected* to pay positive benefits in catastrophic states of nature. It turns out that there is a fundamental difference in the behavior of risk-adjusted discount rates for tail-hedged as opposed to tail-exposed investments. This paper is mainly interested in tail-hedged investments, with the prototype being projects that mitigate future climate damages via lower current emissions of GHGs. However, to appreciate the distinction between risk-adjusted discount rates for the two types of investments, I first analyze briefly the case of a tail-exposed investment.

4 Discounting a Tail-Exposed Investment

Let $B_t | C_t$ be a random variable representing benefit payoffs B_t conditional on effective consumption being C_t .

I will say that a project constitutes a *tail-exposed* investment if, for all future t ,

$$\lim_{C_t \rightarrow 0} \mathbb{E}[B_t | C_t] = 0. \quad (27)$$

The terminology “tail-exposed” investment comes from the fact that the project is not expected to pay any insurance against catastrophic states of nature. Although it is possible to analyze tail-exposed investments in greater generality, I restrict myself to payoff benefits being of the analytically-tractable “prototype” isoelastic-multiplicative form

$$B_t = \theta_t C_t^{\beta'}, \quad (28)$$

where the exponent β' in (28) is a given parameter satisfying

$$\beta' = \frac{\text{cov}(\ln B_t, \ln C_t)}{\text{var}(\ln C_t)}, \quad (29)$$

while the random variable θ_t is distributed independently of C_t . I assume that the relevant case of interest is $\beta' \geq 0$. It is then readily confirmed that a project whose benefits are of the form (28) constitutes a tail-exposed investment because condition (27) is obviously met.

Substituting from (28) into (26) gives

$$\exp(-r_t t) = \frac{e^{-\rho t} \mathbb{E}[U'(C_t) \theta_t C_t^{\beta'}]}{U'(C_0) \mathbb{E}[\theta_t C_t^{\beta'}]}. \quad (30)$$

Making use of the random variables $\{G_t\}$ being i.i.d. implies

$$\frac{\mathbb{E}[U'(C_t) C_t^{\beta'}]}{U'(C_0) \mathbb{E}C_t^{\beta'}} = \frac{\mathbb{E}C_t^{\beta'-\eta}}{C_0^{-\eta} \mathbb{E}C_t^{\beta'}} = \frac{\mathbb{E}\left(C_0 \prod_{j=1}^t G_j\right)^{\beta'-\eta}}{(C_0)^{-\eta} \mathbb{E}\left(C_0 \prod_{j=1}^t G_j\right)^{\beta'}} = \frac{\mathbb{E}\prod_{j=1}^t G_j^{\beta'-\eta}}{\mathbb{E}\prod_{j=1}^t G_j^{\beta'}} = \left(\frac{\mathbb{E}G^{\beta'-\eta}}{\mathbb{E}G^{\beta'}}\right)^t. \quad (31)$$

From independence of the random variables $\{\theta_t\}$ and $\{C_t\}$, we can cancel $\mathbb{E}\theta_t$ from the numerator and denominator of (30). Combining what remains of (30) with (31) then delivers

$$\exp(-r_t t) = \left(\frac{e^{-\rho} \mathbb{E}G^{\beta'-\eta}}{\mathbb{E}G^{\beta'}}\right)^t. \quad (32)$$

Cancel the exponent t from both sides of (32) and take the natural logarithm, which turns equation (32) into the equivalent expression

$$r_t(\beta') = \rho + \ln\left(\mathbb{E}G^{\beta'-\eta}\right) - \ln\left(\mathbb{E}G^{\beta'}\right). \quad (33)$$

Equation (33) signifies that, for any given β' , the discount rate is *constant* over time. Thus, with the prototype tail-exposed specification (28), the term structure of discount rates is flat. It is readily shown that $r_t(\beta')$ given by expression (33) is monotone increasing in β' , ranging from $r_t(0) = r^f$ to $r_t(1) = r^e$.

In the special case where G is lognormally distributed, (33) becomes⁹

$$r_t(\beta') = (1 - \beta') r^f + \beta' r^e. \quad (34)$$

To summarize this section, when the functional form (28) is imposed, then the risk-adjusted discount rate is constant for all periods. Thus, the main thrust of this paper that the term structure of the risk-adjusted discount rate schedule can decline over time towards the riskfree rate is very much dependent on the assumption that the project is *not* tail-exposed. The rest of the paper concerns how and why risk-adjusted discount rates for tail-hedged investments exhibit a declining term structure.

⁹This is a well-known result in the literature. See, e.g., Mehra (2008). I don't fully trust conclusions based on the lognormality assumption because they give rise to the riskfree rate and equity premium puzzles.

5 Discounting a Tail-Hedged Investment

I will say that a project constitutes a *tail-hedged* investment if

$$\lim_{C_t \rightarrow 0} \mathbb{E}[B_t | C_t] > 0 \tag{35}$$

for all future t . Notice that a project described by (35) is not actually required to *deliver* positive benefits to satisfy (35); it is merely *expected* to deliver positive benefits. In the remainder of the paper I assume that the project under consideration constitutes a tail-hedged investment.

The terminology “tail-hedged investment” should be obvious. I would argue that climate-change projects that lower GHG emissions should plausibly be treated as tail-hedged investments. When (35) holds, then at least some protection is afforded against terrible states of the world by investing in mitigation. Even for dismal states of very low net effective consumption and very high marginal utility caused by catastrophic climate damages, the project is at least expected to deliver a positive benefit payoff.

The story that accompanies (35) is that if disastrous future states of extremely low net effective consumption and extremely high marginal utility occur, then they will likely be caused by huge damages from catastrophic climate change being subtracted from gross consumption. In this scenario, future states of very high marginal utility and very low net effective consumption are identified with strictly positive expected benefits from current curtailment of GHGs, which will reduce future climate-change damages.¹⁰

From the above logic, investments in climate mitigation will pay off in states of nature for which the economy will be severely damaged and utility will be severely diminished by a catastrophic climate-caused disaster. I will model this kind of situation by having a strictly positive floor a_t , below which expected benefits from mitigation may not fall no matter how disastrously low is effective consumption. Although it is possible to analyze tail-hedged investments in greater generality, I restrict myself to payoff benefits being of a highly-tractable “prototype” additive form.

In what follows, I try to parallel the standard CAPM model, only with real-benefit payoffs in place of financial returns. By rough analogy with CAPM theory, define

$$b_t = \frac{\text{cov}(B_t, C_t)}{\text{var}(C_t)}. \tag{36}$$

It is assumed that $b_t > 0$. In this paper I am trying to argue that distant-future

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

discount rates should be low. If $b_t \leq 0$, then this argument is very easy to make, almost too easy, because the relevant project beta is then less than or equal to zero, implying that the relevant discount rate is not higher than r^f . I want to show that the discount rate r_t declines towards r^f over time even when the covariance between B_t and C_t is strictly positive.

With b_t defined by (36), we can then express the relationship between B_t (representing real benefit payoffs at time t) and C_t (representing the systematic non-diversifiable risk of the uncertain macro-economy itself at time t) in the linear regression format

$$B_t = a_t + b_t C_t + \varepsilon_t, \quad (37)$$

with $\mathbb{E}\varepsilon_t = 0$ and $cov(\varepsilon_t, C_t) = 0$. Thus, b_t defined by (36) can be interpreted as the coefficient from a hypothetical OLS regression of B_t on C_t of the form (37). 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 a value for real-project betas. 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.

Condition (35) translates into $a_t > 0$. The random variable

$$A_t \equiv a_t + \varepsilon_t \quad (38)$$

is independent of C_t and can be interpreted as representing idiosyncratic project-specific risk. Using (38), equation (37) can be re-expressed as

$$B_t = A_t + b_t C_t. \quad (39)$$

Define the real-project beta at time t to be the *fraction* of expected payoffs that on average is due to the uncertain macro-economy

$$\beta_t \equiv \frac{b_t \mathbb{E}C_t}{\mathbb{E}A_t + b_t \mathbb{E}C_t}, \quad (40)$$

where $a_t > 0$ and $b_t > 0$ imply $0 < \beta_t < 1$.

The “coefficient of variation” of a random variable X is

$$\nu[X] \equiv \frac{\sigma[X]}{\mathbb{E}X}, \quad (41)$$

where $\sigma[X]$ is the standard deviation of X .¹¹ In the light of (41), it is natural to conceptualize the *variation* of X as the mean-normalized random variable $X/\mathbb{E}X$. Using definition (40), equation (37) can be rewritten in the equivalent but more revealing form of a weighted-average *decomposition of variation* equation

$$\frac{B_t}{\mathbb{E}B_t} = (1 - \beta_t) \frac{A_t}{\mathbb{E}A_t} + \beta_t \frac{C_t}{\mathbb{E}C_t}, \quad (42)$$

The coefficient β_t is called the *real-project beta* here because it plays a role in cost-benefit analysis akin to a financial investment beta in the standard linear version of 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.

Formula (42) emphasizes that the variation of B_t can be conceptualized as if it is a kind of imaginary portfolio with the fraction $1 - \beta_t$ being contributed by a component having the variation properties of A_t while the fraction β_t is contributed by a component having the variation properties of C_t . On average, the fraction β_t of the project-payoff portfolio at time t replicates the non-diversifiable risk characteristics of the system-wide aggregate economy at that time, while the fraction $(1 - \beta_t)$ of the project-payoff portfolio replicates risk characteristics that are idiosyncratically independent of the performance of the aggregate economy.

6 Risk-Adjusted Discount Rate Schedules

Substituting from (42) into (26) gives

$$\exp(-r_t t) = \frac{(1 - \beta_t) e^{-\rho t} \mathbb{E}[U'(C_t) A_t]}{U'(C_0) \mathbb{E}A_t} + \frac{\beta_t e^{-\rho t} \mathbb{E}[U'(C_t) C_t]}{U'(C_0) \mathbb{E}C_t}. \quad (43)$$

From independence of the random variables $\{A_t\}$ and $\{C_t\}$, we have $\mathbb{E}[U'(C_t) A_t]/[U'(C_0) \mathbb{E}A_t] = \mathbb{E}[U'(C_t)]/U'(C_0)$. Analyzing the first term in the right hand side of equation (43) by also

¹¹The coefficient of variation is the reciprocal of a widespread measure of the information ratio, aka signal to noise ratio.

making use of the random variables $\{G_t\}$ being i.i.d. then delivers

$$\frac{\mathbb{E}[U'(C_t) A_t]}{U'(C_0) \mathbb{E}A_t} = \frac{\mathbb{E}[U'(C_t)]}{U'(C_0)} = \frac{\mathbb{E} \left[\left(C_0 \prod_{j=1}^t G_j \right)^{-\eta} \right]}{(C_0)^{-\eta}} = \mathbb{E} \left[\prod_{j=1}^t G_j^{-\eta} \right] = (\mathbb{E}G^{-\eta})^t. \quad (44)$$

Analyzing the second term in the right hand side of equation (43) and making use of the random variables $\{G_t\}$ being i.i.d. delivers

$$\frac{\mathbb{E}[U'(C_t) C_t]}{U'(C_0) \mathbb{E}[C_t]} = \frac{\mathbb{E} \left[\left(C_0 \prod_{j=1}^t G_j \right)^{1-\eta} \right]}{(C_0)^{-\eta} \mathbb{E} \left[C_0 \prod_{j=1}^t G_j \right]} = \frac{\mathbb{E} \left[\prod_{j=1}^t G_j^{1-\eta} \right]}{\mathbb{E} \left[\prod_{j=1}^t G_j \right]} = \left(\frac{\mathbb{E}G^{1-\eta}}{\mathbb{E}G} \right)^t. \quad (45)$$

Use (7) to express (44) in terms of R^f and use (15) to express (45) in terms of $\mathbb{E}R^e$. Substitute the results into (43). Collecting terms, we then have

$$\exp(-r_t t) = \frac{1 - \beta_t}{(R^f)^t} + \frac{\beta_t}{(\mathbb{E}R^e)^t}. \quad (46)$$

Compare (46) with (25). Make use of (8) and (16) to express (46), (25) in the form

$$\exp(-r_t t) = (1 - \beta_t) \exp(-r^f t) + \beta_t \exp(-r^e t). \quad (47)$$

Equation (47) is the fundamental result of this paper. Given the linear decomposition of risk factors (42), equation (47) provides a rigorous welfare basis for using expression (23) as the theoretically correct formula for calculating expected present discounted payoffs in cost-benefit analysis. When r_t is defined by (47), a proposed marginal investment project will improve welfare if and only if the project's present discounted expected benefits (23) are greater than its present costs.

We could proceed with a general analysis of $\{r_t\}$ in terms of $\{\beta_t\}$ using formula (47), but I think it is more instructive to highlight primarily the benchmark case of a constant real-project beta.

7 The Benchmark of a Constant Real-Project Beta

Henceforth I assume as a benchmark default simplification 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 understanding 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.

The case of a constant real-project beta should therefore be viewed as a benchmark specification that will allow us to focus more sharply on essentials. Without further ado, I therefore make an assumption of *constant proportions of payoff 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 systematic non-diversifiable risk representing the uncertain macro-economy itself. For all future periods $t \geq 1$, then,

$$\beta_t = \beta. \quad (48)$$

To emphasize its dependence upon the assumed value of β , henceforth r_t is denoted r_t^β . When the simplification (48) is imposed, equation (47) turns into the relatively neat formula

$$r_t^\beta = -\frac{1}{t} \ln \left((1 - \beta) \exp(-r^f t) + \beta \exp(-r^e t) \right). \quad (49)$$

Equation (49) 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 (49) and inspecting carefully the resulting expression indicates that the risk-adjusted discount rate becomes ever lower over time

$$\frac{dr_t^\beta}{dt} < 0. \quad (50)$$

Using l'Hôpital's rule to evaluate the indeterminate form (49) in the limit as $t \rightarrow 0$ gives

$$r_0^\beta = (1 - \beta) r^f + \beta r^e, \quad (51)$$

which is a continuous-time 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 (49) in the limit as $t \rightarrow \infty$ gives

$$r_\infty^\beta = \min\{r^e, r^f\} = r^f. \quad (52)$$

Because $r^e=7\%$ is so much bigger than $r^f=1\%$, the force of compound interest over a century implies a huge difference in corresponding discount factors, namely $\exp(-100r^e)/\exp(-100r^f) \approx 0.25\%$. This allows an accurate approximation of (49) for even moderately large t by the simple formula

$$r_t^\beta \approx r^f + (-\ln(1 - \beta)) \times \frac{1}{t}. \quad (53)$$

What is the economic story behind the basic result of this paper that the risk-adjusted discount rate schedule declines over time (from an initial weighted average given by the standard CAPM formula (51) down to approaching asymptotically the riskfree rate (52))? The story goes along the following lines. Investing in the independent payoff component $A_t/\mathbb{E}A_t$ (relative to investing in the macro-economy component $C_t/\mathbb{E}C_t$) is like buying an insurance policy against the low-consumption bad states of the world that are inherent in taking a gamble on an investment in $C_t/\mathbb{E}C_t$. The price of having an extra unit of $A_t/\mathbb{E}A_t$ relative to the price of having an extra unit of $C_t/\mathbb{E}C_t$ is $P_t^f/P_t^e = \exp(-r^f t)/\exp(-r^e t) = \exp(\pi t)$, where π is the equity premium. This relative price represents a measure of the value of insurance that hedges against low values of $C_t/\mathbb{E}C_t$.

Having an insurance policy in the form of having an investment in $A_t/\mathbb{E}A_t$ as part of a portfolio of payoffs is relatively more valuable over time than having an investment in $C_t/\mathbb{E}C_t$ (and this insurance value grows at the strong exponential rate π). Why? Because the uncertain future consumption $C_t/\mathbb{E}C_t$ is *more variable* over time and the agent is therefore increasingly fearful of *catastrophic risks*. Thus, it turns out, a declining risk-adjusted discount rate schedule is intimately tied to the increasing value of tail-hedged insurance against the increasingly likely possibility of bad tail outcomes. This insight about the increasingly critical role of disastrous tail risks in discounting is sufficiently important to merit elaborating its economic content somewhat further.

The variance of the product of independent random variables equals the product of their expected squared values minus the square of the product of their expected values. Applied

to $C_t = C_0 \prod_{j=1}^t G_j$ with i.i.d. $\{G_j\}$, this variance formula translates into the equation

$$\nu^2[C_t] = (\nu^2[G] + 1)^t - 1 \quad (54)$$

when expressed in terms of the coefficient of variation (41).

From (54), the coefficient of variation $\nu[C_t]$ increases monotonically in t at an exponential rate, which signifies that uncertain future consumption $C_t/\mathbb{E}C_t$ is becoming ever more variable over time. It turns out mathematically that the random variable $C_t/\mathbb{E}C_t \rightarrow 0$ with probability approaching one as $t \rightarrow \infty$.¹² Because $\mathbb{E}[C_t/\mathbb{E}C_t] = 1$, the tendency of $C_t/\mathbb{E}C_t \rightarrow 0$ along sample paths must be counterbalanced in the other direction by ever-higher explosions of $C_t/\mathbb{E}C_t$ occurring with ever-lower probabilities as $t \rightarrow \infty$. Thus, although $C_t/\mathbb{E}C_t$ has an expected value of one for all horizons t , it tends to zero almost surely as $t \rightarrow \infty$. In this sense, the aversion to $C_t/\mathbb{E}C_t$ in the payoff portfolio (42) is increasingly driven over time by the increasing likelihood of an aggregate catastrophe.¹³ In a nutshell, the declining term structure of risk-adjusted discount rates (50), (51), (52) reflects the ever-growing importance of having some tail-hedged catastrophe insurance against relying on ever-more-variable $C_t/\mathbb{E}C_t$ in the portfolio of long term payoffs.

Most of the literature on time-declining discount rates is concerned with safe investments whose payoffs are *independent* of consumption and whose term structure describes the behavior over time of the *riskfree rate alone*. In this literature the declining term structure of the riskfree discount rate depends, ultimately, on the *persistence* over time of underlying growth rates (or, more directly if more mechanically, on the persistence of interest rates themselves).¹⁴ The explanation in this paper is fundamentally different. There is no persistence in the model of this paper because the random variables $\{G_t\}$ are assumed to be i.i.d.. Declining discount rates emerge here from the fundamental interaction over time of an investment component whose payoffs are essentially *independent* of consumption with an investment component whose payoffs *are* essentially consumption.

The above comments notwithstanding, there is a useful way of conceptualizing (49) *as if*

¹²This result is a special case of what is called in the literature “Kakutani’s theorem on product martingales.” See Martin (2012) for more details.

¹³This is a particular example of the general theory developed recently in an important paper by Ian Martin to explain why the valuation of long-dated assets is dominated by the possibility of catastrophic outcomes. See Martin [2012] for details and numerous other examples.

¹⁴This literature is elegantly synthesized in the forthcoming book by Christian Gollier (2012a). An exception to the persistence explanation is an argument that over time the terms of trade turn in favor of environmental amenities when they enter the utility function separately and have less than unit elasticity of substitution with conventional consumption. This important insight was presented most strikingly in Sterner and Persson (2008). Their model bears a familial reduced-form resemblance to the model of this paper, where over time the terms of trade turn in favor of the independent component of a project’s payoff.

one of two infinitely persistent discount-rate possibilities materializes overnight. In this as-if interpretation, β can be conceptualized heuristically as if representing the probability that some given higher discount rate r^e will persistently apply over all future time, while $1 - \beta$ represents the as-if probability that some given lower discount rate r^f will persistently apply over all future time. Suppose that an investment decision must be made today, exactly one day before it is revealed whether the future discount rate will, for tomorrow and forever after, be r^e (with probability β) or whether it will be r^f (with probability $1 - \beta$). Then, heuristically, the effective rate that should be used today to discount payoffs at time t is the value r_t^β satisfying (49) and having the properties (50), (51), (52). The key intuitive insight here is that what should be probability-averaged over states of the world is not discount *rates* (r^e with probability β and r^f with probability $1 - \beta$) at various times, but discount *factors* ($\exp(-r^e t)$ with probability β and $\exp(-r^f t)$ with probability $1 - \beta$).¹⁵ In the limit as $t \rightarrow \infty$, the properly-averaged discount factor corresponds to the minimum possible discount rate, here r^f . Intuitively or heuristically, the impact of the higher discount rate r^e in formula (49) diminishes over time because the higher rate r^e effectively discounts itself out of existence, leaving the field to the lower discount rate r^f . There is no pretending that this analogy is rigorous, but it does present a vivid image of what is occurring.

8 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 (49) 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.¹⁶ 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 the theoretically correct risk-adjustment formula for constant β , augmented by sensitivity analysis of β , than to do nothing about risk adjustment. In any event, this section explores the consequences for the term structure of discount rates of the linearity assumption (42) coupled with the assumption (48) of a constant real-project beta.

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

¹⁵A rigorous argument is presented in Gollier and Weitzman (2010), which also explains the background of this idea and the relevant literature pertaining to it.

¹⁶Some relevant thoughts on this subject are expressed in Ewijk and Tang (2003).

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 work backwards from some postulated near-term CAPM discount rate r_0^β , which people have been observed to use in practice and for which I have some feel, to the CAPM-implied value of β . The basic motivating idea is that when people are using a short term risk-adjusted CAPM discount rate of r_0^β in practice, then they are implicitly revealing their preferred value of β . Presuming constant proportions of payoff risks, this implicitly revealed value of β determines the corresponding term structure of discount rates by formula (49).

To follow the above logic, invert the CAPM equation (51). This gives

$$\beta = \frac{r_0^\beta - r^f}{r^e - r^f}, \quad (55)$$

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

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

Table 1: Risk-adjusted discount rates r_t^β (% per year, rounded off)

Note that, for short-term CAPM discount rates $3\% \leq r_0^\beta \leq 5\%$, which corresponds to mid-range" values of real-project betas $1/3 \leq \beta \leq 2/3$, the benchmark century discount rates are all fairly low, much closer to $r^f=1\%$ than to $r^e=7\%$. Even for a short-term CAPM discount rate as high as $r_0^\beta=6\%$, which corresponds to a real-project beta as high as $\beta=5/6$, the century discount rate of 2.8% is 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 the basic message of the paper. In the linear case (42) corresponding to the standard CAPM specification, 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 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.

9 Some Brief Remarks on the Social Cost of Carbon

In this section I attempt to apply the preceding framework and results to the discounting issues that lie behind calculating the social cost of carbon. This application is not very

“pure” because what has actually been done so far in practice does not fit snugly within the framework of this paper. Nevertheless, to the extent that long-term discounting is such a central part of computing the social cost of carbon (and the estimates depend so critically on the discount rate that is used), the results of this paper may shed some new light on the key role of risk adjustment. The issues associated with determining a real-project beta for investments in climate change mitigation are quite subjective and very tricky. I can offer little more than some crude reflections here.

The “social cost of carbon” (hereafter SCC) is a useful economic concept that has come to the fore in recent analysis of climate change. For the purposes of this paper, the SCC is the marginal cost inflicted on the future economy by emitting an extra ton of carbon dioxide (CO₂) now. It should be fairly obvious that such a number might be useful in a wide variety of contexts for analyzing various policies that directly or indirectly affect emissions of carbon dioxide, of which there is potentially a very large number. My description of how the SCC is estimated is a highly stylized interpretation of the procedure that was employed by the U.S. Government Interagency Working Group on the Social Cost of Carbon, hereafter the USGI working group.¹⁷ My remarks are focused mainly on the assumed discount rates and the assumed specification of damages.

In the very loosest of terms, what the USGI working group did can be described roughly as follows. Start with a so-called “integrated assessment model” (IAM).¹⁸ An IAM is a computational model with dozens of equations that combines 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 base case of climate change policy that specifies some actual path of CO₂ emissions. This will produce a series of outcomes, including a baseline time series of future effective consumption levels. If the IAM has key uncertain elements built into it, which it should for the model to be credible, the baseline effective 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 effective 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 effective consumption change) in each period. Pick some discount rate schedule and calculate the present discounted value of average benefit payoffs. This is the SCC.

¹⁷See 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.

¹⁸The USGI working group actually used an average from three IAMs, whose acronyms are DICE, FUND, and PAGE. This detail does not seem important for the general purposes of this paper.

At the heart of an IAM is a series of equations describing the relationship between changes in CO₂ emissions, subsequent changes in global mean temperatures, and corresponding changes in welfare-equivalent as-if damages to effective consumption caused by increased global mean temperatures. A majority of IAMs make damages apply to production of output, as opposed to damages applying directly to consumption, but I am not sure this detail is critical and here slide over the distinction.¹⁹

IAMs tend to express damages as if there has been a welfare-equivalent loss of some *fraction* of output (here interpreted as a fraction of ordinary consumption), where the size of the fractional loss increases with the mean global temperature change. But *why* is it so standard to express damages as a fraction of output? The answer is simple: because it is so very difficult to conceptualize damages in any other terms. (If we could imagine that future damages a century from now were somehow given to us in the form of absolute values, then the very first thing we would want to do for conceptualization is to convert this absolute value into the welfare-equivalent fractional loss of output or consumption at that time.) What tends to get obscured is the *as if* aspect of this artificial construct. Climate damages act *as if* some fraction of welfare-equivalent consumption has been lost – they do not typically cause a literal evaporation of consumption.

Now comes an important question. Should the fraction of as-if lost future consumption from temperature damages be conceptualized as a fraction of *expected* future consumption or the same fraction of *actual* future realizations of stochastic consumption, or some weighted average of the two, and if so what should the weights be? In a deterministic setting it makes no difference. But in an uncertain setting this distinction amounts to choosing the real-project beta. Without much thought or discussion, the IAMs presume that temperature damages are proportional to *actual* stochastic consumption. Thus, more by default than by conscious premeditation, the IAMs have inadvertently preempted the discounting issue by choosing in practice a project beta of one. But then consistency might require discounting the expected consumption changes at a high rate of $r^e=7\%$, while what is chosen in practice as a rate to discount reduced consumption damages in SCC calculations is typically in a lower range of $\approx 3-4\%$.²⁰ So the approach that is taken in practice is not directly cognizant of real-project betas. This may be no small matter because the choice of a discount rate can make an enormous difference in calculating the SCC.

This paper seems to be suggesting the following loose thoughts about computing the SCC.

¹⁹If the distinction matters, a cynic might say so much the worse for the IAM, since with climate change it is difficult to provide a good story why the one as-if damages specification rather than the other as-if damages specification should be used.

²⁰The USGI working group (2010) used as a central estimate a discount rate of 3%. Nordhaus (2011) effectively uses 4% in his calculations.

No one is saying that it is easy to estimate the relevant real-project beta for abating one ton of carbon dioxide, but attempts to measure the SCC might try various values of β within a framework of time-declining discount rates. The stochastic nature of future consumption might be modeled more carefully and explicitly. In the specification of damages it could be made more explicit what are the *proportions* in which the postulated fractional loss of consumption is intended to apply to expected consumption as opposed to actual stochastic consumption. The rate used to discount expected future changes in consumption might be made consistent with the real-project beta implicit in the form of the assumed damages function, so that there is overall consistency.

What is the appropriate real-project beta for an investment that reduces by one ton the present emissions of CO₂? This is the 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. In this 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 CO₂.²¹ This story treats mitigation as a low-beta hedge asset helping to insure against climate disasters. In the standard story, which is told by the IAMs more by default than premeditation, higher conventional future consumption is associated with larger absolute damages from climate change (either because damages are assumed to be multiplicative in stochastic consumption realizations or because higher growth is assumed to be associated with higher emissions). In this case, states of higher future consumption and lower marginal utility will be associated with higher absolute future benefits from current curtailment of CO₂, thereby implying a higher beta.²²

In any event, postulating $\beta=1$ seems to me like a somewhat extreme assumption on the intuitive presumption that climate change investment payoffs seem so very different from the average payoffs of private-sector investments. The effects of climate change will impact extremely unevenly among different regions of the world and different sectors of human or natural activity. "Outdoor" activities (like agriculture, recreation, natural landscapes, the existence value of ecosystems, and so forth) will be much more affected than the "indoor" activities that constitute the bulk of a conventional economy. Besides, some component of mitigation payoffs plausibly represents insurance against fat-tailed catastrophic future outcomes due to extreme climate change itself.

If $\beta=1/2$, then the highlighted row of Table 1 suggests a corresponding risk-adjusted

²¹This is the approach taken in Sandsmark and Vennemo (2007).

²²See, e.g., Nordhaus (2008) or Nordhaus (2011).

discount rate that declines from 4% in the immediate future down to 1.7% a century from now and down further to 1.2% three centuries from now. If the CAPM discount rate is $r_0^\beta=3\%$, which might be interpreted as corresponding to the constant value of 3% chosen by the USGI working group, the implied value of beta is 1/3, in which case Table 1 suggests a corresponding risk-adjusted discount rate that declines down to 1.4% a century from now and down further to 1.1% three centuries from now. The most serious climate change damages will be incurred over the course of centuries. Lowering the (constant) discount rate from 3% to 2.5% increased the SCC value in the USGI working group estimates from \$21 to \$35 per ton, so there is quite a bit of sensitivity. It might be interesting to calculate the SCC using the risk-adjusted discount rate schedules from Table 1 that correspond to the short-term CAPM discount rates of $r_0^\beta \approx 3-4\%$ that have been used in practice.

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 value of beta in a climate-change context.²³ 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.

10 Concluding Comments on Tail-Hedged Discount Rates

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 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 model of this paper is primitive and has a lot of simplistic assumptions built into it, many of which might legitimately be challenged. For what is worth, the model suggests that, when a project is a tail-hedged investment, what might better be combined in a beta-weighted average at time t are not the two focal discount rates r^e and r^f , but rather their

²³Gollier (2012b) attempts some rough calculations 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 .

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.

The paper explains how the declining schedule of risk-adjusted discount rates reflects the increasing relative value of having insurance against the increasingly likely possibility over time of disastrous outcomes.

Because there is such a significant equity-premium difference between discount rates of $r^e=7\%$ and $r^f=1\%$ per year, there is 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 tail-hedged investments whose real-project beta 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 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] Barro, Robert J. “Rare Disasters and Asset Markets in the Twentieth Century.” *Quarterly Journal of Economics*, 2006, CXXI, pp. 823-866.
- [2] 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.
- [3] Ewijk, Casper van and Paul J.G. Tang (2003). “How to Price the Risk of Public Investment?” *De Economist*, 151 (3), 317-328.
- [4] Gollier, Christian (2012a). *Pricing the Planet’s Future: The Economics of Discounting in an Uncertain World*. Princeton University Press (forthcoming October 2012).
- [5] Gollier, Christian (2012b). “Term structure of discount rates for risky investments.” October 2012 Preliminary Draft.

- [6] 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.
- [7] Greenstone, Michael, Elizabeth Kopits, and Ann Wolverton (2011). “Estimating the social cost of carbon for use in U.S. federal rulemakings: A summary and interpretation.” NBER Working Paper 16193.
- [8] Lucas, Robert E. “Asset Prices in an Exchange Economy.” *Econometrica*, 1978, 46, pp. 1429-1445.
- [9] Martin, Ian (2012). “On the Valuation of Long-Dated Assets.” *Journal of Political Economy*, 120 (2), 346-358.
- [10] Mehra, Rajnish. “The Equity Premium Puzzle: A Review.” *Foundations and Trends in Finance*: 2008 Vol. 2: No 1, pp 1-81.
- [11] Mehra, Rajnish and Edward C. Prescott. “The Equity Premium: a Puzzle.” *Journal of Monetary Economics*, 1985, 15, pp. 145-161.
- [12] 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.
- [13] 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/>
- [14] Nordhaus, William D. (2008). *A Question of Balance: Weighing the Options on Global Warming Policies*. New Haven: Yale University Press.
- [15] Nordhaus, William D. (2011). “The Social Cost of Carbon.” NBER Working Paper 17540.
- [16] Sandmark, M., and H. Vennemo (2007). “A portfolio approach to climate investments: CAPM and endogenous risk.” *Environmental and Resource Economics*: 37, 681-695.
- [17] Sterner, Thomas, and U. Martin Persson (2008). “An Even Sterner Review: Introducing Relative Prices into the Discounting Debate.” *Review of Environmental Economics and Policy*, 2008 (Winter), 2(1), pp. 61-76.

- [18] 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>.
- [19] Weitzman, Martin L. (2007). "Subjective Expectations and Asset-Return Puzzles." *American Economic Review*, 2007 (March), 97 (4): 1102-1130.