Environmental Protection, Rare Disasters, and Discount Rates*

Robert J. Barro, Harvard University

June 2014

Abstract

The *Stern Review*'s evaluation of environmental protection relies on extremely low discount rates, an assumption criticized by many economists. The *Review* also stresses that great uncertainty is a critical element for optimal environmental policies. An appropriate model for this policy analysis requires sufficient risk aversion and fat-tailed uncertainty to get into the ballpark of explaining the observed equity premium. A satisfactory framework, based on Epstein-Zin/Weil preferences, also separates the coefficient of relative risk aversion (important for results on environmental investment) from the intertemporal elasticity of substitution for consumption (which matters little). Calibrations based on existing models of rare macroeconomic disasters suggest that optimal environmental investment can be a significant share of GDP even with reasonable values for the rate of time preference and the expected rate of return on private capital. Optimal environmental investment increases with the coefficient of relative risk aversion and the probability and typical size of environmental disasters but decreases with the degree of uncertainty about policy effectiveness. The key parameters that need to be pinned down are the proportionate effect of environmental investment on the probability of environmental disaster and the baseline probability of environmental disaster.

*This research has been supported by a grant from the National Science Foundation. I have benefited from comments by Marios Angeletos, John Campbell, Christian Gollier, Matthew Kahn, Robert Litterman, Emi Nakamura, Bob Pindyck, Nick Stern, Jenny Tang, and Marty Weitzman.

Discount rates play a central role in the literature on environmental protection that revolves around the *Stern Review* (Stern [2007]). Spending money now to reduce environmental pollution (including CO2 emissions that enhance global warming) is modeled as generating benefits that arise in the distant future. Therefore, the policy tradeoff depends on whether these future benefits are discounted at a substantial rate, such as the 5-6% per year reflective of average real rates of return on private capital, or at the near-zero social rate advocated by the *Review*.

The *Review's* adoption of extremely low discount rates has been criticized by many economists, including Gollier (2006), Nordhaus (2007), Weitzman (2007), and Mendelsohn (2008), as inconsistent with empirical evidence and theoretical reasoning. Nordhaus's criticism focuses on reasonable calibrations of the standard neoclassical growth model. However, because this model is deterministic, it features a single real rate of return. Therefore, as noted by Weitzman (2007), this framework cannot illuminate the large gap between average real rates of return on equity (or private capital) and risk-free assets. Empirically, the average equity premium is around 5% per year if measured by the difference between long-term averages of real rates of return on unlevered equity and on Treasury Bills (Barro and Ursúa [2008]).

A consideration of the variety of real rates of return and an analysis of which rates are pertinent for environmental issues requires a stochastic model that gets into the ballpark of accounting for the equity premium. My approach relies on a rare-disasters framework of the sort surveyed in Barro and Ursúa (2012). This literature (Rietz [1988], Barro [2006], Barro and Jin [2011]), Barro and Ursúa [2012]) argues that an unlevered equity premium around 5% per year is consistent with the long-term international history of macroeconomic disasters with a reasonable coefficient of relative risk aversion in the neighborhood of 3.0-3.5.

The *Stern Review* argues that uncertainty is central to its analysis of climate change and that this uncertainty strengthens its case for aggressive mitigation. Thus, according to Stern (2007, p. xiv), "Uncertainty about impacts strengthens the argument for mitigation; this Review is about the economics of the management of very large risks." Unfortunately, however, the *Review* lacks a satisfactory framework for analyzing the relevant uncertainty, including probabilities and sizes of potential disasters, the degree of risk aversion, and so on. Moreover, the *Review* does not distinguish uncertainty related to the probability and size distribution of potential environmental disasters from uncertainty about how interventions influence the likely outcomes. In my analysis, the former kind of uncertainty tends to strengthen the case for environmental investment, whereas the latter—concerning policy effectiveness—tends to weaken the case.

One important missing element in the *Stern Review* is a framework—such as that developed by Epstein and Zin (1989) and Weil (1990)—that allows for disentangling of the coefficient of relative risk aversion from the intertemporal elasticity of substitution (*IES*) for consumption. In my analysis, the welfare calculations depend heavily on the coefficient of relative risk aversion but not much on the *IES*.

Weitzman (2007, 2009) emphasizes that a serious treatment of uncertainty is crucial for environmental issues because of the fat-tailed nature of potential environmental crises. Moreover, the inclusion of fat-tailed uncertainty is important for evaluating environmental investment not just because it helps to determine the magnitudes of discount rates relevant for capitalizing likely future paths of social costs and benefits. Weitzman (2007) observes that a central feature of these kinds of social investments is that they influence the probability of the associated rare disasters. Hence, there turn out to be two key relationships: how much is it worth to reduce the probability of an environmental disaster and how much does environmental investment lower this probability? Weitzman (2007, pp. 704-705) puts the matter this way:

> … spending money now to slow global warming should not be conceptualized primarily as being about optimal consumption smoothing so much as an issue about how much insurance to buy to offset the small chance of a ruinous catastrophe …

My analysis follows Weitzman's lead by conceptualizing the optimal choice of environmental policy as a decision about how much society should spend to reduce the probability (or potential size) of environmental disasters. Some of this policy choice looks like spending now to gain something later, because an expenditure of resources to lower today's disaster probability improves likely outcomes not only today but also for the indefinite future. On the other hand, the essential element of the policy tradeoff does not involve a dynamic where the optimal ratio of environmental investment to GDP and the associated disaster probability look a lot different today from tomorrow. In my main model, the optimal environmental investment ratio and the associated disaster probability end up being constant, although the levels of these variables depend on a present-versus-future tradeoff. Extensions of the model may generate a path in which the environmental-investment ratio tends to rise gradually toward a steady-state value.

With respect to discount rates, the connection with environmental investment and disaster probability depends on the underlying source of changes in these rates. If the pure rate of time preference (applying in my analysis to the representative household and the social planner) is lower, so that all expected real rates of return are correspondingly lower, the optimal ratio of environmental investment to GDP is higher, and the equilibrium disaster probability is correspondingly smaller. Thus, in this case, lower discount rates and expected real rates of return associate with higher environmental investment, as in the *Stern Review* literature.

The results are different if lower expected real rates of return reflect some other factors, such as lesser risk aversion or an inward shift in the size distribution of disasters. In these cases, the benefit from lowering disaster probability is decreased, and this force motivates a lower ratio of environmental investment to GDP. Moreover, this effect tends to outweigh the impact from a lower required real rate of return. Thus, contrary to the *Stern Review* literature, in these cases, lower expected real rates of return associate with decreased environmental investment.

I investigate choices of environmental investment within the rare-disasters setting developed in Barro (2009). Although simple, this framework seems adequate to conceptualize the main tradeoff that determines optimal environmental investment and the associated disaster probability. The analysis also brings out parameters needed to deliver quantitative answers about optimal policy. For a given size distribution of disasters and a given coefficient of relative risk aversion, the key parameters are the proportionate impact on disaster probability from a higher ratio of environmental investment to GDP and the baseline environmental disaster probability.

I. Model of Rare Macroeconomic Disasters

The underlying model is an extension of Barro (2009). Details and derivations of the main formulas are in that earlier paper.

The log of real per capita GDP evolves as a random walk with drift:[1]

$$log(Y_{t+1}) = log(Y_t) + g + u_{t+1} + v_{t+1}. \tag{1}$$

The random term u_{t+1} is i.i.d. normal with mean 0 and variance σ^2. This term reflects "normal" economic fluctuations. Given reasonable calibrations, this term is quantitatively unimportant for the results. The parameter $g \geq 0$ is a constant that reflects exogenous productivity growth.

[1]A straightforward extension in Barro (2009) allows real GDP to depend on work effort, which is determined by the representative household's labor-leisure choice.

The random term v_{t+1} picks up low-probability disasters, as in Rietz (1988) and Barro (2006). In these rare events, output jumps down sharply. The probability of a disaster is the constant $p \geq 0$ per unit of time. In a disaster, output contracts instantaneously and permanently by the fraction b, where $0<b<1$.[2] The distribution of v_{t+1} is given by

probability $1\text{-}p$: $v_{t+1} = 0$,

probability p: $v_{t+1} = log(1\text{-}b)$,

where the disaster size, b, follows a time-invariant probability distribution, gauged subsequently by the empirical distribution of these sizes.[3] The idea is that the disaster probability and size distribution capture the fat-tailed nature of the negative range for changes in Y_t. The expected growth rate of Y_t is given, as the (arbitrary) length of the period approaches zero, by

$$g^* = g + (1/2)\sigma^2 - p \cdot Eb, \tag{2}$$

where Eb is the mean of b.

In previous applications of this framework, such as Barro and Ursúa (2008), the disaster realizations correspond empirically to declines in real per capita GDP by 10% or more over short periods for a sample of countries that have annual GDP data starting at least by 1914. In an expanded sample, there are 185 of these events for 40 countries over periods going back as far as 1870.[4] The histogram for these macroeconomic disaster sizes is in Figure 1. The events reflect wartime destruction (notably World Wars I and II), Great Depression-type contractions typically

[2] Nakamura, Steinsson, Barro, and Ursúa (2013) allow for a finite duration of disasters and for a systematic tendency for recoveries, in the sense of sustained above- normal growth rates following disasters. In the subsequent application to environmental disasters, the assumption that the output contractions are permanent may be reasonable. However, technological and other responses to environmental changes may lead to “recoveries” in these cases.

[3] In Barro (2006) and Barro and Ursúa (2008), this distribution corresponds to histograms constructed from the observed data. In Barro and Jin (2011), the distribution is given by a power-law density, with the parameters fit to the data.

[4] The sample has 5349 annual observations on GDP, of which 2977 are for long-term OECD members. (The data are available at rbarro.com.) The main extensions from Barro and Ursúa (2008) are the addition of four countries with newly assembled data—China, Egypt, Russia, and Turkey—and a shift in the end date from 2006 to 2011. The last extension brings in two macroeconomic disasters associated with the Great Recession—for Greece and Iceland. The results discussed later are similar if I use instead the smaller sample for which data are available on a measure of consumption (real per capita personal consumer expenditure).

associated with financial crises, and downturns possibly driven by the Great Influenza Epidemic of 1918-1920. The sizes of the declines in per capita GDP averaged 21% and ranged as high as 60-70% for several countries during the world wars. The sample does not contain any contractions driven by natural disasters, such as tsunamis and earthquakes, or environmental catastrophes.

I now interpret the disaster probability, p, and the distribution of disaster sizes, b, as encompassing potential environmental disasters. Specifically,

(3) $$p = \pi + q,$$

where π is the probability of disasters of the sort isolated in previous empirical work, and q is the probability of an environmental disaster. I assume that an environmental disaster works like the other types of disasters in the sense of contracting real per capita GDP instantaneously and permanently by the fraction b. Crucially, given the lack of direct evidence on the sizes of environmental disasters, I assume that the distribution of disaster sizes, b, is the same as for the observed macroeconomic events. This assumption—implying a mean disaster size of 21% (conditional on being at least 10%)—encompasses the range of environmental costs envisioned by the *Stern Review*:

> "… the Review estimates that if we don't act, the overall costs and risks of climate change will be equivalent to losing at least 5% of global GDP each year, now and forever. If a wider range of risks and impacts is taken into account, the estimates of damage could rise to 20% of GDP or more. … Our actions now and over the coming decades could create risks of major disruption to economic and social activity, on a scale similar to those associated with the great wars and the economic depression of the first half of the 20th century." (Stern [2007, p. xv]).

The model assumes that the economy is closed, with no investment in ordinary capital stock.[5] Output, determined by equation (1), goes to consumption or environmental investment.

II. Household Utility and Optimal Environmental Investment

As in Barro (2009), the representative household's utility depends on the time path of real per capita consumption, C_t. Because it is important to disentangle the coefficient of relative risk aversion from the intertemporal elasticity of substitution (*IES*) for consumption, I generalize from a standard power-utility function to an Epstein-Zin/Weil (Epstein and Zin [1989] and Weil [1990]) recursive specification:[6]

$$(4) \qquad U_t = \frac{\left\{C_t^{1-\theta} + \left(\frac{1}{1+\rho}\right)[(1-\gamma)E_t U_{t+1}]^{(1-\theta)/(1-\gamma)}\right\}^{(1-\gamma)/(1-\theta)}}{(1-\gamma)}.$$

In this formulation, $\rho \geq 0$ is the (constant) rate of time preference, $1/\theta > 0$ is the (constant) *IES*, and $\gamma > 0$ is the (constant) coefficient of relative risk aversion. The power-utility formulation restricts to $\theta = \gamma$. However, this restriction delivers odd properties for equity prices. For example, suppose that $\gamma > 1$ (needed to have any chance to replicate the empirically observed equity premium), so that the $IES < 1$ when $\theta = \gamma$. It then turns out, counter-intuitively, that the price-dividend ratio for equity shares (claims on consumption trees) is lower when the growth rate, g^*, is higher and higher when parameters such as p that describe uncertainty are higher. In the main analysis, I assume $\gamma > 1$ and $\theta < 1$, so that $IES > 1$. In this case, the price-dividend ratio for equity shares relates to g^* and uncertainty parameters in "reasonable" ways.

---

[5]Barro (2009) uses an "AK model" to work out an extension that encompasses endogenous saving and investment for a closed economy. In this model, disasters show up as unusually high depreciation of existing capital.

[6]Environmental studies that use Epstein-Zin/Weil preferences include Ackerman, Stanton, and Bueno (2012); Crost and Traeger (2012); and Cai, Judd, and Lontzek (2013). However, these studies do not calibrate fat-tailed environmental uncertainty based on the empirical evidence on rare macroeconomic disasters.

Let τ be the fraction of GDP that goes to environmental investment—that is, to mitigation of threats to the environment—so that $0 \le \tau \le 1$. The fraction τ may be time varying but is optimally chosen as constant in the present model. Consumption relates to GDP in accordance with

$$C_t = (1-\tau) \cdot Y_t. \tag{5}$$

Environmental investment, if positive, is undertaken by the government and financed in a non-distorting way. (The present model has no scope for tax distortions because all output not used for environmental investment goes to consumption, and there is no labor-leisure choice.) The benevolent government may choose $\tau > 0$ because a higher τ is assumed to reduce the probability, q, of environmental disasters. This relationship is described by the function[7]

$$q = q(\tau), \tag{6}$$

where $q'(\tau) \le 0$. In the main analysis, I assume the functional form:

$$q(\tau) = q(0)e^{-\lambda\tau}, \tag{7}$$

where $\lambda > 0$. In this form, zero environmental investment, $\tau=0$, corresponds to a baseline environmental disaster probability $q(0)>0$.[8] As τ tends to 1, the probability of environmental disaster goes to $q(0) \cdot e^{-\lambda} > 0$. That is, no human action (limited here to $\tau \le 1$) is sufficient to drive this probability down to zero. The parameter λ can be viewed as a measure of policy effectiveness.

In the initial analysis, I assume that the impact of the environmental-investment ratio, τ, on the environmental-disaster probability, q, is known; that is, the policymaker has full

[7]We could also allow for a negative effect of τ on the typical size of a disaster. In the analysis of Barro and Jin (2011), this effect could be represented by a negative impact of τ on the parameters of the power-law densities that govern the thickness of the tail.

[8]The baseline environmental disaster probability, $q(0)$, corresponds to the historical situation, over which I assume τ=0. In this baseline setting, the overall disaster probability is $p(0)=\pi+q(0)$ in equation (3) (taken later to be 0.040 per year). The assumption that τ starts at zero matters only because of the constraint $\tau \ge 0$; that is, the government cannot cut back on the initial level of investment when it starts at zero.

knowledge of λ and $q(0)$. More realistically, there would be a great deal of uncertainty about policy effectiveness. A later section allows for uncertainty about the true value of λ.

Equation (7) implies that the semi-elasticity of q with respect to τ equals the constant $-\lambda$. In the subsequent analysis, the important factor is the derivative of q with respect to τ. Given equation (7), this derivative equals $-\lambda q = -\lambda q(0) \cdot e^{-\lambda\tau}$. Therefore, the key parameters for quantitative analysis will be λ and $q(0)$; $q(0)$ sets the overall level of the environmental disaster probability, whereas λ determines how this probability responds to τ. In the present model, τ will be optimally chosen as a constant, which will be zero if $\lambda \cdot q(0)$ is below a critical value.

It is convenient to frame the results in terms of prices of assets that provide ownership rights over streams of per capita consumption, C_t. These assets correspond to Lucas trees, introduced by Lucas (1978), on which C_t is the fruit that drops each period as a dividend from a tree. A difference from the standard model is that per capita GDP, Y_t, falls from trees each period, with the fraction τ taxed away, leaving $C_t = (1-\tau) \cdot Y_t$ as the net dividend for owners.

Let V be the price-dividend ratio for equity claims on these modified Lucas trees. As in Barro (2009), in this i.i.d. model, the reciprocal of V (the dividend-price ratio) will be determined as the period length tends to zero as the constant

(8) $$1/V = \rho - (1-\theta)g^* + (1/2)\gamma(1-\theta)\sigma^2 + p\left(\frac{1-\theta}{\gamma-1}\right)[E(1-b)^{1-\gamma} - 1 - (\gamma-1)Eb]$$

if $\gamma \neq 1$. For any $\gamma > 0$, the condition $\theta < 1$ implies that, with g^* held fixed, V is lower if uncertainty is greater (higher σ or p or an outward shift of the b-distribution). That is, a once-and-for-all increase in an uncertainty parameter reduces equity prices (as seems plausible) if and only if $\theta < 1$, so that the $IES > 1$. Similarly, $\theta < 1$ implies that a rise in g^* increases V.

The dividend-price ratio, $1/V$, relates to the rate of return in a familiar way:

(9) $$1/V = r^e - g^*,$$

where r^e is the expected rate of return on equity. Equivalently, r^e equals the dividend yield, $1/V$, plus the expected rate of capital gain, g^*, which equals the expected growth rate of dividends (per capita consumption). It also follows that r^e equals the right-hand side of equation (8) after elimination of the term $-g^*$. The condition $r^e > g^*$ is the transversality condition for this model; that is, the right-hand side of equation (8) has to be positive. This condition guarantees that the market value of a tree is positive and finite.

For explaining the equity premium, the dominant effect empirically on the right-hand side of equation (8) is the disaster term, which is proportional to p.[9] The important effect from disaster sizes involves the expression $E(1-b)^{-\gamma}-1$, which can be thought of heuristically (applying directly in the power-utility case) as the difference between the average marginal utility of consumption in a disaster state and in a normal state. Substitution for g^* from equation (2) into equation (8) yields an alternative formula for $1/V$ in terms of g:

$$(10) \qquad 1/V = \rho - (1-\theta)g + (1/2)(1-\theta)(\gamma-1)\sigma^2 + p(\frac{1-\theta}{\gamma-1})[E(1-b)^{1-\gamma} - 1].$$

Using results from Obstfeld (1994), Barro (2009) showed that attained utility in this Epstein-Zin/Weil representative-household model with i.i.d. shocks can be written in a simple way in terms of the price-dividend ratio, V. These results apply to the present setting, with the modification that C_t equals the fraction $1-\tau$ of GDP, Y_t, in accordance with equation (5). The formula for attained utility evaluated at date t is, up to an inconsequential additive constant:[10]

$$(11) \qquad U_t = (\frac{1}{1-\gamma})V^{(1-\gamma)/(1-\theta)}(1-\tau)^{1-\gamma}Y_t^{1-\gamma}.$$

Note that U_t is increasing in V if $\theta<1$, is decreasing in τ (for given V), and is increasing in Y_t.

---

[9]The formula for the equity premium is $r^e - r^f = \gamma\sigma^2 + p\cdot[E(1-b)^{-\gamma} - E(1-b)^{1-\gamma} - Eb]$, where r^f is the risk-free rate. In practice, the disaster term, which involves p, does almost all the work in explaining the equity premium.

[10]Results when γ=1 or θ=1 can be obtained from standard limit arguments.

I think of the government's optimization problem as choosing τ (or, more generally, the path of τ) to maximize U_t, as given by equation (11). That is, the government at each date t is assumed to advance the interests of the representative household alive at date t. Therefore, the government as social planner is assumed to respect the representative household's vision of utility, including the parameters for risk aversion, intertemporal elasticity of substitution, and rate of time preference.

Note that the household's formulation is forward looking and, with a small generalization, takes account of the number, as well as the consumption levels, of descendants. The particular specification may have a positive rate of time preference, $\rho>0$, which may derive from less than one-to-one concern of parents for the welfare of their children. A lot of literature, beginning with Ramsey (1928), considers the ethical basis for this time preference across generations. But from a political perspective, it is hard to see how the choices today by a democratic government would deviate from the wishes of the representative agent currently alive. That is, the prospective utilities of future generations matter today but only because parents care about their children, who care about their children, and so on.[11]

Equation (11) can be used in conjunction with equation (10) to determine the optimal τ (which will be chosen as a constant). Higher τ reduces C_t for given Y_t and, thereby, has a negative effect on U_t in equation (11). However, higher τ lowers the environmental disaster probability, q, in equation (6) or (7) and, hence, reduces the overall disaster probability, p. This change raises the price-dividend ratio, V, in accordance with equation (10) and, thereby, raises U_t

[11]Conceivably, a social planner could internalize the concern of currently living persons about other people's children.

in equation (11).[12] Thus, the tradeoff that determines optimal environmental investment is the direct consumption loss today weighed against the benefits for the entire path of future consumption from a decrease in today's disaster probability.

Although I focus on effects of environmental investment on the environmental-disaster probability, q, one can think analogously of investments that affect the distribution of disaster sizes, b. For example, suppose that each $1-b$ enters into equation (10) as a multiple of a factor $\eta>0$. Then, if an increase in τ raises η—thereby reducing the sizes of disasters—this change works like a reduction in p.

I consider the case from equation (7) in which the semi-elasticity of q with respect to τ is the constant $-\lambda$. When the optimal solution for τ is interior, the optimal value of τ is determined from the first-order condition, which can be written as:[13]

$$\text{(12)} \qquad \frac{1}{V} = r^e - g^* = \left(\frac{1-\tau}{\gamma-1}\right) \cdot [E(1-b)^{1-\gamma} - 1] \cdot \lambda q(0) e^{-\lambda\tau} \, .$$

Recall from equation (9) that the dividend-price ratio, $1/V$, on the far left-hand side equals r^e-g^*, where r^e is the expected real rate of return on unlevered consumption equity and g^* is the expected growth rate of real per capita GDP (and consumption). This dividend-price ratio is the correct capitalization rate in this model for gauging the "present value" of environmental outlays, which are the fraction τ of GDP in each period.

---

[12]This discussion applies when $\gamma>1$ and $\theta<1$. However, the analysis goes through even if these conditions do not hold. For example, if $\theta>1$, a reduction in p lowers V according to equation (10). However, in this case, a decrease in V raises U_t in equation (11). Therefore, a fall in p is still credited with a positive effect on utility.

[13]Setting the derivative of U_t with respect to τ to zero in equation (11) leads to the condition $1 = \frac{(1-\tau)}{(1-\theta)} \cdot \frac{1}{V}\frac{dV}{d\tau}$. Equation (10) implies $\frac{1}{V}\frac{dV}{d\tau} = V \cdot \frac{(1-\theta)}{(\gamma-1)} \cdot [E(1-b)^{1-\gamma} - 1] \cdot \lambda q(0) e^{-\lambda\tau}$. The combination of these two conditions leads to equation (12). If the term (1-b) enters multiplicatively with a factor η, then the right-hand side of equation (12) includes another term, $(1-\tau) \cdot q \cdot E(1-b)^{1-\gamma} \cdot \eta^{-\gamma} \cdot \frac{d\eta}{d\tau}$. Therefore, if an increase in τ lowers disaster sizes, so that $d\eta/d\tau>0$, there is a further force that favors the choice of a higher τ.

The far right-hand side of equation (12) reflects the benefit at the margin from the negative effect of the environmental-investment ratio, τ, on the environmental-disaster probability, q, and, hence, on the overall disaster probability, $p=\pi+q$. Recall that $\lambda q=\lambda q(0)e^{-\lambda\tau}$ is the magnitude of the derivative of q with respect to τ, given the form of equation (7). The marginal benefit on the far right is larger when the coefficient of relative risk aversion, γ, is higher (because the $1-b$ term dominates), or when the distribution of disaster sizes, b, is shifted outward, or when the baseline probability of an environmental disaster, $q(0)$, is higher.

The optimal choice of τ (when the solution is interior) occurs where the marginal benefit on the far right-hand side of equation (12) equals the required rate of return on the left-hand side. In the cases considered later, the various changes shift the marginal benefit on the right-hand side or the required rate of return on the left-hand side or both.

The dividend-price ratio, $1/V$, depends on underlying parameters in accordance with the right-hand side of equation (10). Substitution of this expression into the far left-hand side of equation (12) leads to a condition for determining τ (when the solution is interior) in terms of exogenous parameters:

$$\rho-(1-\theta)g+\left(\frac{1}{2}\right)(1-\theta)(\gamma-1)\sigma^2+\pi\left(\frac{1-\theta}{\gamma-1}\right)[E(1-b)^{1-\gamma}-1] \tag{13}$$
$$=\left(\frac{1}{\gamma-1}\right)[E(1-b)^{1-\gamma}-1][(1-\tau)\lambda-(1-\theta)]q(0)e^{-\lambda\tau}\,.$$

In this analysis, the probability of non-environmental disaster, π, and the baseline probability of environmental disaster, $q(0)$, are taken as given.[14]

The model can be used to compute the consumer surplus from the government's opportunity to carry out environmental investment at the optimal ratio, τ^*, rather than being

[14]A formally parallel analysis could be used to assess the effects of social investments on the probability, π, of non-environmental disasters. For example, one could examine policies that affect the probabilities of wars, depressions (likely related to financial and housing crises), epidemics of disease, natural disasters, and so on. Barro (2009) argues that a prime force behind the creation of the euro as a common currency was the desire to lower the probability of European wars.

constrained to have τ=0. Let Y_t^0 and U_t^0 be the values of Y_t and U_t, respectively, corresponding to an initial position where τ=0. Let Y_t^*be the value of Y_t that yields the same utility, U_t^0, when $\tau=\tau^*$, so that $Y_t^* \leq Y_t^0$. That is, society would be willing to give up some GDP today to have the opportunity to carry out forever environmental investment at the optimal ratio. Equation (11) implies

$$(14) \qquad \frac{Y_t^*}{Y_t^0} = (\frac{V^0}{V^*})^{\frac{1}{1-\theta}} \cdot (\frac{1}{1-\tau^*}),$$

where V^0 and V^* are given from equation (10) using, respectively, $p=\pi+q(0)$ and $p^*=\pi+q^*$, where q^* is the value of q from equation (7) that corresponds to $\tau=\tau^*$. The consumer-surplus ratio, which is non-negative, is defined as $(1-\frac{Y_t^*}{Y_t^0})$. This ratio gives the proportionate fall in today's GDP that society would willingly accept to gain the opportunity to choose τ optimally forever, rather than having $\tau=0$.

III. Calibration

We can assess equation (13) quantitatively using parameter values generated from an updated version of the macro-disaster analysis carried out in Barro and Ursúa (2008). The baseline parameter values associated with the expanded data sample are in Table 1. Note that these parameters come from fitting empirically observed variables, including the frequency and size distribution of macroeconomic disasters, the mean and standard deviation of the growth rate of per capita GDP, the average real rate of return on unlevered equity, and the average real risk-free rate (proxied by returns on Treasury Bills). One finding is that the estimated disaster probability, p (corresponding to events with declines in per capita GDP by 10% or more), is

0.040 per year. This observed value, 0.040, is assumed to equal $p(0)=\pi+q(0)$—that is, the historical situation corresponds to $\tau=0$.

One parameter in Table 1 is the coefficient of relative risk aversion, γ, which is set to 3.3 in the baseline case. This substantial, though not astronomical, degree of risk aversion is important for valuing environmental investments that mitigate disaster risk. Moreover, with the Epstein-Zin/Weil form of preferences, the value γ=3.3 is compatible with an *IES* for consumption, $1/\theta$, that generates reasonable properties for asset prices, as noted above. However, the results on optimal environmental investment turn out to depend relatively little on θ. In other words, it is the coefficient of relative risk aversion, not the intertemporal elasticity of substitution for consumption, that mainly matters for environmental investment.

In contrast, the *Stern Review* uses a standard power-utility formulation that requires $\gamma=\theta=1/IES$ and focuses on the case where γ and θ (represented in the *Review* by the symbol η) are both equal to one (corresponding to log utility). Ironically, given the *Review's* overall tendency to exaggerate the benefits from environmental investments, this low coefficient of relative risk aversion strongly diminishes the benefits from environmental investment in a stochastic setting where these investments reduce disaster risk.

In the sample, with 5349 annual GDP observations for 40 countries, none of the 185 disaster events corresponded to environmental catastrophes. If the country observations were independent, the absence of any realizations would be inconsistent with a substantial annual probability, $q(0)$, of these events. For example, even if $q(0)$were only 0.001 per year, the probability of zero hits in 5349 independent annual observations is only 0.005. However, since some types of environmental disasters, such as negative consequences from global warming, have high positive correlation across countries, we might want to think of the sample from 1870 to 2011 as essentially a single time series of 142 observations.[15] In this case, the probability of zero hits is 0.24 if *q(0)=0.010* per year and 0.49 if *q(0)=0.005* per year. Thus, this perspective can reconcile the sample observations (zero hits) with values of *q(0)* around 0.010, as assumed in the baseline setting. Of course, it is also possible that the probability of environmental disaster was near zero historically but is significant for the future. The probability of non-environmental disasters is given by π=0.040–*q(0)*.

Given the baseline parameter values in Table 1, including *q(0)*=0.010, equation (13) determines the optimal environmental-investment ratio, τ, and the associated disaster probability, *q*, as a function of the parameter λ. Recall that this parameter specifies the proportionate effect of the environmental-investment ratio, τ, on the environmental-disaster probability, *q*, in equation (7). For a given *q(0)*, equation (13) is inconsistent with $\tau>0$ if λ is below a threshold value. For the baseline parameters, this threshold turns out to be 8.63. Therefore, Table 2, Section I, shows that $\tau=0$ is chosen for $\lambda\leq 8.63$.

For values of λ above the threshold, the chosen τ is positive. The selected τ initially rises with λ, then subsequently falls—because a higher λ means that *q* in equation (7) is smaller for given τ (thereby generating a force that diminishes the incentive to choose a high τ). For example, in Table 2, Section I, the chosen τ reaches 0.014 at $\lambda=10$, 0.036 at $\lambda=15$, and 0.042 at $\lambda=20$, but then falls to 0.035 at $\lambda=50$ and 0.025 at $\lambda=100$. However, the environmental-disaster probability, *q*, corresponding to the optimal choice of τ, is monotonically declining with λ.[16]

---

[15]Nakamura, Steinsson, Barro, and Ursúa (2013) allow for spatial dependence for realizations of macroeconomic disasters. A high degree of cross-country correlation applies to global crises such as the world wars, the Great Depression, and the Great Influenza Epidemic.

[16]This analysis, based on equation (13), allows for the effect of τ on $1/V = r^e - g^*$ on the left-hand side of equation (12), reflecting the positive effect on r^e from the disaster probability, $p=\pi+q$. In practice, this channel has only a minor effect on the results. If we base the analysis instead on equation (12), with the left-hand side held fixed at its baseline value of 0.042, the critical value for λ turns out to be 8.71, compared to 8.63 in Table 2, Section I. The values of τ that correspond to higher values of λ are 0.013 for $\lambda=10$, 0.034 for $\lambda=15$, 0.040 for $\lambda=20$, 0.034 for $\lambda=50$, and 0.024 for $\lambda=100$.

The consumer-surplus ratio, computed from equation (14), is in the far right column of Table 2, Section I. This ratio equals zero until λ reaches the threshold of 8.63 and then rises monotonically with λ. At high values of λ, this ratio is substantial—for example, 2.4% of GDP when λ=20 and 6.0% of GDP when λ=50.

IV. Shifts in Exogenous Parameters

Sections II-V of Table 2 show the consequences from (once-and-for-all) variations in the main parameters away from their baseline values. Section II gives the effects from an increase in the coefficient of relative risk aversion, γ, to 5.0, compared to 3.3 in the baseline specification. This change sharply lowers the threshold value of λ for positive environmental investment, $\tau>0$, to 4.81 from 8.63 in the baseline case. Moreover, for λ above the threshold, τ is higher for a given λ. For example, for λ=20, when γ=5, τ=0.072 (q=0.0024), compared to τ=0.042 (q=0.0043) when $\gamma=3.3$. These results reflect the greater incentive to lower the environmental-disaster probability when risk aversion is higher. Note also that a higher γ raises environmental investment while simultaneously increasing the required expected rate of return, r^e-g^*, that applies to this investment in the model. The key mechanism is that a higher γ shifts outward the benefit from environmental investment on the far right-hand side of equation (12), and this effect dominates the impact from the upward shift in r^e-g^* in equation (10).

An outward shift in the distribution of disaster sizes, b, similarly raises the incentive to choose a higher environmental-investment ratio, τ. Table 2, Section III, shows the consequences from a multiplication of each observed disaster size, b, by 1.1. (This analysis holds fixed the baseline disaster probability, $p=\pi+q(0)$, at 0.040 per year.) The outward shift in disaster sizes lowers the threshold value of λ for positive environmental investment, $\tau>0$, to 6.76 from 8.63 in the baseline case. In addition, for λ above the threshold, τ is higher for a given λ. For example, for λ=20, τ=0.054 (q=0.0034) when the disaster sizes are larger by 10%, compared to the baseline value of τ=0.042 (q=0.0043). Thus, this case again features higher environmental investment along with a higher required expected rate of return, r^e-g^* (determined in equation [10]).

Section IV of Table 2 assumes that the baseline environmental disaster probability, $q(0)$, is 0.005, rather than 0.010. The analysis assumes that the overall baseline disaster probability, $p(0)=\pi+q(0)$, is still 0.040—therefore, the probability of a non-environmental disaster, π, is now 0.035, rather than 0.030. The lower value of $q(0)$ reduces the incentive for environmental investment. Therefore, the threshold λ that generates positive investment is sharply higher, 17.3 in Section IV, compared to 8.6 in the baseline case. The reasoning is that the motivation to choose $\tau>0$ depends on the magnitude of the derivative of q with respect to τ at $\tau=0$, and equation (7) implies that this derivative equals $-\lambda\cdot q(0)$. Therefore, when $q(0)$ falls by one-half (from 0.010 to 0.005), λ has to double (from 8.6 to 17.3) to motivate positive environmental investment. The reduction in $q(0)$ also implies, for λ above the threshold, that the chosen τ is much smaller at a given λ. For example, when λ=20, τ=0.007, compared to 0.042 in the baseline case. These results show that a decrease in $q(0)$ from 0.010 to 0.005 produces a large change in the conclusions.

Section V of Table 2 assumes that the rate of time preference, ρ, is 0.030, rather than the baseline value of 0.044.[17] This change generates the pure discounting effect emphasized in the *Stern Review* literature. In particular, equation (10) implies that the dividend-price ratio, $1/V = r^e$-g^*, shifts downward. This effect shifts downward the left-hand side of equation (12),

[17]As discussed in the notes to Table 1, the results in this Epstein-Zin/Weil model depend on an effective rate of time preference, ρ^*, that deviates from ρ. For the parameter values considered in the baseline calibration, ρ^* is 0.029, well below ρ=0.044. If ρ=0.030, then ρ^*=0.015. Intuition about what is a "reasonable" rate of time preference likely applies more to ρ^* than to ρ. In any event, my choice of ρ is dictated by fitting the data on real rates of return, not from an ethical perspective.

implying that the marginal return from environmental investment on the right-hand side has to be lower at the optimum (when the solution is interior). Therefore, the chosen environmental investment ratio, τ, tends to be higher. This effect explains why the threshold λ needed to warrant positive investment, $\tau>0$, declines sharply—to 5.65 in Section V, compared to 8.63 in the baseline case. Moreover, for values of λ above the threshold, the chosen τ is substantially higher than before. For example, for λ=20, τ=0.063, compared to 0.042 in the baseline case.

A change in the *IES*, $1/\theta$, has an ambiguous effect on the dividend-price ratio, $1/V = r^e - g^*$, in equation (10). Section VIa of Table 2 shows that an increase in θ from its baseline value of 0.5 to 1.0 raises the threshold value of λ from 8.63 to 9.20. If λ=20, the chosen τ when θ=1 is 0.037, compared to the baseline value of 0.042. Hence, a change in θ from 0.5 to 1.0 has a minor impact relative to the effects from changes in the coefficient of relative risk, or the size distribution of disasters, or the baseline probability of environmental disaster, or the rate of time preference (Table 2, Section II-V).

Section VIb considers an even larger change in θ to 3.3—which equals γ and, therefore, corresponds to the usual power-utility formulation. In this case, the threshold value for λ rises to 11.79. If λ=20, the chosen τ is 0.022, compared to 0.042 when θ=0.5 and 0.037 when θ=1. Therefore, a very large change in the *IES* matters significantly for the results. However, a θ of 3.3 seems unrealistically high because the implied *IES* of only 0.3 means that the price-dividend ratio, *V*, responds positively to increases in parameters related to uncertainty and negatively to the growth-rate parameter, *g*. Overall, the results support Weitzman's (2007, pp. 704-705) conjecture, quoted in the introduction, that optimal environmental investment is not "primarily … about optimal consumption smoothing" (in particular, about the intertemporal elasticity of substitution for consumption) "so much as an issue about how much insurance to buy to offset the small chance of a ruinous catastrophe" (which brings in the key roles of the coefficient of relative risk aversion and the frequency and size distribution of disasters).

A number of other parametric changes (not shown in Table 2) have effects equivalent to those from changes in ρ. For example, the analysis treated the baseline disaster probability, $p(0)=\pi+q(0)$, as equaling the observed disaster probability of 0.040 per year, so that the probability of non-environmental disasters was fixed at π=0.030 when $q(0)$=0.010. Another approach would fix π at 0.040, because none of the observed disasters were environmental. In this case, $q(0)=0.010$ might apply to the future even if not to the history. This change in specification amounts to an upward shift in $p(0)$, while holding fixed $q(0)$. This change raises the dividend-price ratio, $1/V = r^e - g^*$, in equation (10) and, thereby, shifts outward the left-hand side of equation (12). Since the right-hand side of equation (12) does not shift, the increase in π works in the same way as an increase in the rate of time preference, ρ. The consequence is that the threshold value of λ rises (from 8.63 to 9.15). Moreover, the chosen τ is lower at a given λ—for example, when λ=20, the chosen τ is 0.039, compared to 0.042 in the baseline case. Thus, these effects are comparatively minor.

A higher σ^2 works in a similar way by raising the dividend-price ratio, $1/V = r^e - g^*$, in equation (10) (assuming $\theta<1$ and $\gamma>1$). Therefore, a higher σ^2 tends to reduce the environmental investment ratio, τ, but this effect is quantitatively minor for a reasonable range of σ^2. If $\theta<1$, a higher growth-rate parameter, g, reduces the dividend-price ratio, $1/V = r^e - g^*$, in equation (10) and, therefore, has effects on τ equivalent to those from a lower ρ.

V. Environmental Amenities Enter Utility along with Ordinary Consumption

In the previous analysis, the benefits from environmental investment stem from reductions in the environmental-disaster probability, q. When these disasters occur, they are like other rare disasters that sharply lower real GDP and consumption. Therefore, in this model, environmental disasters covary positively with consumption. This perspective explains why the required expected rate of return on environmental investment on the left-hand side of equation (12) is the dividend-price ratio, r^e-g^*, where r^e is the expected rate of return on unlevered consumption equity (and private capital). This result depends on an underlying specification for environmental disasters that has been reasonably questioned by Weitzman (2007, p. 713):

> … there was never any deep economic rationale in the first place for damages from greenhouse gas warming being modeled as entering utility functions through the particular reduced form route of being a pure production externality …

An alternative specification, consistent with Sterner and Persson (2008), has the flow of amenities from the environment entering into the representative household's utility function along with ordinary consumption. For example, the household may care about an effective consumption flow, C_t^*, that is a *CES* aggregate of ordinary consumption, C_t, and environmental amenities, e_t:

$$C_t^* = [\alpha \cdot C_t^{\frac{\sigma-1}{\sigma}} + (1-\alpha) \cdot e_t^{\frac{\sigma-1}{\sigma}}]^{\sigma/(\sigma-1)}, \tag{15}$$

where $0<\alpha<1$ and the elasticity of substitution between C_t and e_t is $\sigma>0$.

Ordinary consumption, C_t, is again the fraction $1-\tau$ of GDP, Y_t. The stochastic process for Y_t is the same as before, except that only non-environmental disasters are considered. The process for e_t is specified below.

In the previous setting, environmental disasters amounted to sharp declines in GDP and ordinary consumption. A contrasting approach, based on equation (15), assumes that Y_t and e_t are independently distributed (so that environmental disasters do not occur particularly at good or bad economic times) and C_t and e_t are perfect substitutes in the effective consumption flow (σ tends to infinity). In this case, effective consumption is

$$C_t^* = \alpha C_t + (1-\alpha)e_t, \tag{16}$$

and the relative shadow price of e_t and C_t is fixed at $(1\text{-}\alpha)/\alpha$.[18]

The dividend-price ratio, $1/V$, for consumption equity is still given by equation (10), except that the disaster probability, π, applies now only to non-environmental disasters:

$$1/V = \rho - (1-\theta)g + (1/2)(1-\theta)(\gamma-1)\sigma^2 + \pi(\frac{1-\theta}{\gamma-1})[E(1-b)^{1-\gamma} - 1]. \tag{17}$$

I still assume $\pi=0.040$, based on the history of rare macroeconomic disasters (none of which were environmental).

Define $\hat{V}$ to be the dividend-price ratio corresponding to a hypothetical claim on the flow of environmental amenities, e_t. The assumption is that the parameters analogous to g and σ^2 are zero for these amenities. That is, the process excludes trend growth and minor fluctuations, leaving an emphasis on rare environmental disasters. The probability of an environmental disaster is still q, given as a function of the environmental-investment ratio, τ, in equation (7). I assume the same baseline disaster probability, $q(0)$=0.010, as in the main previous analysis. If τ is constant (now no longer exact), q will also be constant. The distribution of environmental disaster sizes, b, is again assumed to be the same as the historically observed distribution for non-environmental disasters. Given these assumptions, the dividend-price ratio for the hypothetical environmental claim is the constant

[18]As discussed in Barro and Sala-i-Martin (2004, p. 68), the *CES* specification in equation (15) can be broadened to include an additional parameter that allows for the arbitrary units of measurement for C_t and e_t. This parameter would also appear in equation (16).

(18) $$1/\hat{V} = \rho + q\left(\frac{1-\theta}{\gamma-1}\right)[E(1-b)^{1-\gamma} - 1].$$

Define V_t^* to be the price-dividend ratio corresponding to a hypothetical claim on the effective consumption flow, C_t^*, given in equation (16). The formula for V_t^* is

(19) $$V_t^* = \left[\frac{(1-\tau)Y_t}{(1-\tau)Y_t+\left(\frac{1-\alpha}{\alpha}\right)e_t}\right] \cdot V + \left[\frac{\left(\frac{1-\alpha}{\alpha}\right)e_t}{(1-\tau)Y_t+\left(\frac{1-\alpha}{\alpha}\right)e_t}\right] \cdot \hat{V},$$

where I used the condition $C_t = (1\text{-}\tau)Y_t$ from equation (5). The terms in brackets are the shares in effective consumption of, respectively, ordinary consumption, C_t, and shadow expenditure on environmental amenities, $\left(\frac{1-\alpha}{\alpha}\right)e_t$. Since V and $\hat{V}$ are constants (if τ is constant), V_t^* varies over time only because of changes in these shares.

Attained utility, denoted U_t^*, relates in a simple way to V_t^*, analogously to equation (11):

(20) $$U_t^* = \left(\frac{1}{1-\gamma}\right)(V_t^*)^{\left(\frac{1-\gamma}{1-\theta}\right)}(C_t^*)^{1-\gamma},$$

where V_t^* is given in equation (19) and C_t^* is given in equation (16). It is again straightforward to work out the first-order condition for an interior solution for the optimal (constant) τ. The result can be approximated as[19]

(21) $$1/\hat{V} \approx \left(\frac{\hat{V}}{V}\right)\left[\frac{\left(\frac{1-\alpha}{\alpha}\right)e_t}{(1-\tau)Y_t}\right]\left(\frac{1-\tau}{\gamma-1}\right)[E(1-b)^{1-\gamma} - 1]\lambda q(0)e^{-\lambda\tau}.$$

This condition is formally similar to equation (12) from the original specification. The difference on the left-hand side is that the dividend-price ratio for a hypothetical claim on environmental amenities appears instead of the ratio for a claim on ordinary consumption. The

[19]The approximation is satisfactory if the share of shadow environmental outlay, $\left(\frac{1-\alpha}{\alpha}\right)e_t$, in effective consumption, C_t^*, is small or if $\hat{V}$ is close to V.

difference on the right-hand side is the multiplication by $\hat{V}/V$ and by the ratio of the shadow outlay on environmental amenities to ordinary consumption.

The baseline calibration corresponding to Table 1 implies that $\hat{V}$ and V do not differ greatly—they equal 21.4 and 24.1, respectively, and the ratio $\hat{V}/V$ is 0.89. If the baseline environmental disaster probability, *q(0)*, were much smaller than 0.010, $\hat{V}$ would be larger, and the ratio $\hat{V}/V$ could exceed 1. The reciprocal, $1/\hat{V}$, would then be close to the risk-free rate, which would be very small, as in the *Stern Review*, if the pure rate of time preference, ρ, were close to zero (rather than the baseline value of 0.044 from Table 1). However, when q is very small, the marginal return from environmental investment on the right-hand side of equation (21) is also very small. For example, as *q(0)* tends to zero, there is obviously no case for investment in environmental protection, and the optimal τ will be zero even though the required rate of return on the left-hand side of equation (21) is small.

Since $\hat{V}/V$ is not far from one in the baseline calibration, the important new element on the right-hand side of equation (21) is the multiplication by the ratio of shadow environmental outlay, $\left(\frac{1-\alpha}{\alpha}\right) e_t$, to ordinary consumption, $(1-\tau)Y_t$. This ratio is hard to pin down but is likely to be small. If I take it to be 0.1,[20] then the incentive for environmental investment is much less than in the initial model. For example, in Section I of Table 2, a positive environmental-investment ratio, τ, was warranted if the parameter λ in equation (7) exceeded 8.6. Now the required λ is 109. Moreover, the maximum value for the optimal τ, occurring for a λ around 300, is only 0.0036, well below the 1% value emphasized by the *Stern Review*. Thus, the bottom line is that this alternative model provides a much weaker case than the original specification for substantial environmental investment.

[20]This value was used by Sterner and Persson (2008) in their main analysis.

Equation (15) can also be used with elasticities of substitution, σ, less than infinity. The case σ=1 would correspond to the original model in which the required rate of return on environmental investment equaled the expected rate of return on consumption equity.[21] The previous analysis showed that substantial environmental investment can be warranted here because the higher marginal benefit from reducing environmental disaster risk more than offsets the higher required rate of return (all compared to values in the model just worked out).

Sterner and Persson (2008) argue for values of σ below 1—they used σ=0.5 in their main analysis. This change would make the required rate of return on environmental investment even higher than the expected rate of return on consumption equity. However, because the benefit from reducing the probability of environmental disaster would also be elevated, this case would likely rationalize higher levels of environmental investment, compared to those in the original model.

VI. Uncertainty about the Impact of Environmental Investment[22]

The results in Table 2 bring out the importance of the policy-effectiveness parameter, λ, which, along with the baseline environmental-disaster probability, $q(0)$, determines the impact of the environmental-investment ratio, τ, on the environmental-disaster probability, q:

$$q(\tau) = q(0)e^{-\lambda\tau}. \tag{7}$$

Table 2 considered a broad range for λ because there seems to be little empirical basis for pinning down this parameter. However, each calculation pretended that λ (and the other parameters) were known precisely by the policymaker. This section allows for uncertainty about

[21]The DICE-2007 model of Nordhaus (2008, appendix, equations [A4] and [A5]) accords with σ=1 in the sense that environmental damages, denoted Ω, affect output proportionately. However, these damages are a non-linear function of mean surface temperature, which depends on the history of emissions, abatements, and other factors.
[22]The analysis in this section follows a suggestion from Jenny Tang.

λ—that is, for policy effectiveness—and shows how this uncertainty impacts the optimal environmental-investment ratio, τ.[23] The main result is that greater uncertainty about λ tends to lower the optimal τ. Thus, this kind of uncertainty differs from the one considered before concerning the likelihood and size of potential environmental disasters. As already shown, an increase in the disaster probability or an outward shift in the size distribution of disasters tends to raise environmental investment.

This section returns to the setting in which environmental disasters are modeled as reductions in GDP, as in equation (2). This specification includes the overall disaster probability, $p=\pi+q$, where π is the probability of a non-environmental disaster (still taken as given) and q is the probability of an environmental disaster.

Suppose that, instead of taking on a known value, the parameter λ can take on a finite array of possible values, each with an associated (subjective) probability. Given $q(0)$ and a choice of τ, each λ maps into $q(\tau)$ in accordance with equation (7). The overall environmental-disaster probability is the mean of these values, computed using the probability density for λ.

It is straightforward to show that, with uncertainty about λ, the interior first-order condition for the optimal τ generalizes from equation (12) to:

$$(22) \qquad \frac{1}{V} = r^e - g^* = \left(\frac{1-\tau}{\gamma-1}\right)[E(1-b)^{-\gamma} - 1]q(0) \cdot E(\lambda e^{-\lambda\tau}) \; .$$

That is, the term $\lambda e^{-\lambda\tau}$ is replaced by its expectation. Substituting out for $1/V$ on the left-hand side from a form of equation (10) that allows for uncertainty about λ leads to a generalization of equation (13):

[23]The setting is formally similar to the analysis of macroeconomic policy uncertainty in Brainard (1967). In Brainard's model, which assumes a quadratic objective function for a target variable, greater uncertainty about the impact of a policy instrument on the target variable tends to diminish the optimal extent of policy intervention.

(23) $$\rho-(1-\theta)g+\left(\frac{1}{2}\right)(1-\theta)(\gamma-1)\sigma^2+\pi\left(\frac{1-\theta}{\gamma-1}\right)[E(1-b)^{1-\gamma}-1]$$

$$=\left(\frac{1}{\gamma-1}\right)[E(1-b)^{1-\gamma}-1]q(0)\cdot[(1-\tau)E(\lambda e^{-\lambda\tau})-(1-\theta)E(e^{-\lambda\tau})]\ .$$

Hence, the terms $\lambda e^{-\lambda\tau}$ and $e^{-\lambda\tau}$ are replaced by their expectations.

There are two effects from uncertainty about λ on the optimal τ. First, for a given mean of λ and a given τ, a greater spread in possible values of λ raises uncertainty in the sense of the term $E(e^{-\lambda\tau})$ (because this term is convex in λ). If $\theta<1$, as I assume, this effect raises the dividend-price ratio, $1/V$, on the left-hand side of equation (22) and tends, thereby, to decrease the optimal τ. Equivalently, in equation (23), the increase in $E(e^{-\lambda\tau})$ lowers the right-hand side.

The second, more important, effect is that greater uncertainty about λ makes it harder to match the selected τ with the true value of λ. This effect shows up on the right-hand sides of equations (22) and (23) in the term $E(\lambda e^{-\lambda\tau})$, which determines the expected magnitude of the derivative of q with respect to τ. A greater spread in possible values of λ lowers this term (because it is concave in λ, assuming $\lambda\tau<2$). This effect also tends to reduce the optimal τ.

Table 3 provides some quantitative guidance about the effect of uncertainty in λ on the optimal τ. The analysis assumes that λ can take on two possible values, λ_1 and λ_2, each with probability one-half. For a given mean of the λ's, a larger magnitude of the spread between λ_1 and λ_2 represents greater uncertainty about the true value. Aside from the treatment of λ, the specification in Table 3 corresponds to the baseline case considered in part I of Table 2. This specification includes a baseline environmental-disaster probability of $q(0)$=0.010.

In the first part of Table 3, the mean of λ is 10. When there is no spread between the two possible values of λ (as in Table 2), the optimal τ is 0.0140, corresponding to an environmental-disaster probability of q=0.00869 and a consumer-surplus ratio (indicating the benefit of being able to choose positive environmental investment) of 0.00108. The optimal τ falls to 0.0133 when the spread in λ is (7.5, 12.5) and 0.0114 when the spread is (5, 15). Correspondingly, q rises to 0.00876 and 0.00894, and the consumer-surplus ratio falls to 0.00102 and 0.00088. The main result is that if the extent of uncertainty about policy effectiveness is represented by a λ of 10±5, the optimal τ declines to 1.1%, compared to 1.4% when λ is known to equal 10.

In the second part of Table 3, the mean of λ is 20. With no spread in the λ values, the optimal τ is 0.0415, corresponding to an environmental-disaster probability of q=0.00436 and a consumer-surplus ratio of 0.0237. The optimal τ falls to 0.0400 when the spread in λ is (15, 25) and 0.0353 when the spread is (10, 30). Correspondingly, q rises to 0.00458 and 0.00525, and the consumer-surplus ratio falls to 0.0226 and 0.0197. Thus, if the extent of uncertainty about policy effectiveness is represented by a λ of 20±10, the optimal τ declines to 3.5%, compared to 4.2% when λ is known to equal 20.

VII. Conclusions and Extensions

One conclusion is that the reasoning in the *Stern Review* that lower discount rates warrant greater environmental investment is correct when the source of the lower discount rate is a decline in the pure rate of time preference, ρ (or to other changes described in the text that have effects equivalent to decreases in ρ). However, this reasoning is incorrect when the source of reduced expected returns on equity is lower risk aversion (Table 2, Section II) or smaller disaster sizes (Section III). These changes lower the required expected rate of return as gauged by the price-dividend ratio, $1/V=r^e-g^*$, which is determined in equation (10) and appears on the left-hand side of equation (12). However, these changes also shift inward the return on environmental investment, as implied by the right-hand side of equation (12). In the quantitative analysis, the latter effects dominate.

The numerical results in Table 2 depend on several disaster-related parameters, for which baseline estimates appear in Table 1. Although the estimates of these parameters are imprecise, the analysis in Table 2 is suggestive about how the findings depend quantitatively on reasonable variations in these parameters. Moreover, recent research on rare macroeconomic disasters helps to pin down plausible ranges for some of the key parameters.

However, the results are sensitive to variations in two parameters, λ and $q(0)$, for which specifications of reasonable ranges of variation are much more challenging. The parameter λ represents policy effectiveness; it equals the semi-elasticity of the environmental-disaster probability, q, with respect to the environmental-investment ratio, τ (in accordance with equation [7]). For the cases considered in Table 2, the critical value of λ that generates an interior solution where $\tau>0$ ranges between 5 and 17. In thinking about which values of λ seem reasonable, we can take as an example the intermediate value $\lambda=10$. This value means that, starting from the baseline value $q(0)=0.010$, an increase in τ from 0 to 0.01 would lower q by roughly 10%; that is, from 0.010 to 0.009. Unfortunately, I cannot judge at this point whether this response in environmental-disaster probability is roughly correct or way too big or way too small.

An extension of the analysis dealt with uncertainty about the true value of λ. If λ is precisely 10, the optimal environmental-investment ratio, τ, in the baseline case is 1.4%. In contrast, if λ is 10±5, the optimal τ falls to 1.1%.

The results depend also on the baseline environmental-disaster probability, $q(0)$, taken to be 0.010 per year in the main analysis. However, the results are substantially different if this baseline probability is much smaller; for example, 0.005 in Table 2, Section IV. Given the absence of environmental catastrophes in the sample of rare macroeconomic disasters, it is impossible to use the history to pin down reasonable magnitudes for $q(0)$ that apply looking forward.

My analysis implies that the environmental-investment ratio, τ, is optimally chosen as constant over time. Thus, the results do not feature the ramp-up property emphasized by Nordhaus (2007, p. 687):

> One of the major findings in the economics of climate change has been that efficient or "optimal" economic policies to slow climate change involve modest rates of emissions reductions in the near term, followed by sharp reductions in the medium and long term. We might call this the *climate-policy ramp*, in which policies to slow global warming increasingly tighten or ramp up over time.

The model fails to yield this kind of time variation in the optimal environmental-investment ratio, τ, because the environmental-disaster probability, q, depends only on the contemporaneous value of τ (in equations [6] and [7]). If this linkage involved learning or other "adjustment costs," then the equilibrium might feature a rising path of τ, with the optimal τ tending to approach a long-run or steady-state value. A dynamic path for the optimal τ may also emerge from a model in which the environmental disaster probability, q, depends non-linearly on the cumulation of past levels of environmental pollution, which relate to the history of GDP. These issues could be addressed within the dynamic model of environmental damages contained in Nordhaus's (2008, Appendix) DICE model (dynamic integrated model of climate and the economy). However, the stochastic structure of that model would have to be laid out and would have to include a fat-tailed distribution for potential environmental disasters.

I should stress that the present analysis, even with a substantial rate of time preference and a substantial required expected rate of return on environmental investment, may be consistent with a "large" optimal ratio of environmental investment to GDP, τ. Depending particularly on the values of the key parameters λ and $q(0)$, the results in Table 2 may support the environmental investment of 1% or more of GDP that the *Stern Review* offers as a benchmark (Stern [2007, p. xv]).

One reason that the optimal environmental-investment ratio may be high in my analysis is that the coefficient of relative risk aversion, γ, is set at the reasonably high value of 3.3 in the baseline case. This value is much higher than the 1.0 used in the *Stern Review* (although, with the power-utility formulation, the *Review's* model cannot disentangle risk aversion from intertemporal substitution for consumption). With my calibration for risk aversion, along with a specification for disaster probability and size distribution based on the history of non-environmental disasters, the model can support substantial environmental investment without having to invoke an unrealistically low rate of time preference and a correspondingly low expected real rate of return on private capital.

Table 1 **Baseline Parameter Values**	
Parameter	**Value**
γ (coefficient of relative risk aversion)	3.3
θ (inverse of *IES* for consumption)	0.5
σ (standard deviation of normal shock per year)	0.020
g (growth rate parameter per year)	0.025
$g*$ (expected growth rate per year of per capita GDP)	0.017
Eb (expected disaster size in disaster state)	0.21
$E(1-b)^{-\gamma}$ (expected "marginal utility" in disaster state)	2.11
$p(0)=\pi+q(0)$ (historical probability per year of disaster)	0.040
$q(0)$ (baseline probability of environmental disaster)	0.010
r^f (risk-free rate per year)	0.010
r^e (expected rate of return per year on unlevered equity)	0.059
ρ (pure rate of time preference per year)	0.044
$\rho*$ (effective rate of time preference per year)	0.029

Note: The sample of macroeconomic disasters, updated from Barro and Ursúa (2008), is used to determine the distribution of disaster sizes, b, and the probability per year, $p(0)$, of entering into these disasters. The standard deviation, σ, of normal shocks is set to 0.020 (and is quantitatively unimportant for the results). The growth rate parameter, g, is 0.025, and the corresponding expected growth rate, $g*$, of per capita GDP and consumption is 0.017 (from equation [2]). The coefficient of relative risk aversion, γ, is determined, given the frequency and size distribution of macroeconomic disasters and the other parameters, so that the model fits the observed average equity premium. This premium is estimated to be 0.049, given by 0.059 (average real rate of return on unlevered equity) less 0.010 (average real rate of return on Treasury Bills). The pure rate of time preference, ρ, is set so that the model fits the risk-free rate of 0.010 (that is, so that the model gets right the overall level of rates of return). However, with the Epstein-Zin/Weil form of preferences, the solution for rates of return depends not on ρ but on an effective rate of time preference, $\rho*$, shown in Barro (2009) to be
$\rho* = \rho - (\gamma-\theta)\left\{ g - (1/2)(\gamma-1)\sigma^2 - (\frac{p}{\gamma-1})[E(1-b)^{1-\gamma} - 1] \right\}$. The *IES* for consumption is set, as in Barro (2009), at $1/\theta$=2, so that θ=0.5.

Table 2 **Optimal Environmental-Investment Ratios**			
***λ*: semi-elasticity of environmental disaster probability with respect to environmental-investment ratio**	***τ*: environmental-investment ratio**	***q*: environmental disaster probability**	**consumer-surplus ratio**
I (baseline): γ=3.3, empirical size distribution of disasters, q(0)=0.010, ρ=0.044			
≤ 8.63	0	0.010	0
10	0.014	0.0087	0.001
15	0.036	0.0058	0.012
20	0.042	0.0044	0.024
50	0.035	0.0017	0.060
100	0.025	0.0008	0.080
II: γ (coefficient of relative risk aversion) = 5.0			
≤ 4.81	0	0.010	0
7	0.051	0.0070	0.011
10	0.071	0.0049	0.034
15	0.076	0.0032	0.065
20	0.072	0.0024	0.087
50	0.048	0.0009	0.139
100	0.031	0.0004	0.163
III: disaster sizes multiplied by 1.1			
≤ 6.76	0	0.010	0
7	0.005	0.0097	0.000
10	0.038	0.0068	0.009
15	0.052	0.0046	0.028
20	0.054	0.0034	0.044
50	0.041	0.0013	0.088
100	0.027	0.0007	0.109
IV: q(0) (baseline environmental disaster probability) = 0.005			
≤ 17.3	0	0.005	0
20	0.007	0.0043	0.001
50	0.021	0.0018	0.017
100	0.018	0.0008	0.030
V: ρ (rate of time preference) = 0.030			
≤ 5.65	0	0.010	0
7	0.029	0.0082	0.003
10	0.055	0.0058	0.019
15	0.064	0.0038	0.045
20	0.063	0.0028	0.064
50	0.044	0.0011	0.112
100	0.029	0.0006	0.135

λ	τ	q	cons. surp.
VIa: θ (1/IES) = 1.0			
≤ 9.20	0	0.010	0
10	0.008	0.0092	0.0003
15	0.031	0.0063	0.009
20	0.037	0.0048	0.019
50	0.033	0.0019	0.053
100	0.024	0.0009	0.072
VIb: θ (1/IES) = γ = 3.3			
≤ 11.79	0	0.010	0
15	0.013	0.0082	0.002
20	0.022	0.0064	0.007
50	0.026	0.0027	0.031
100	0.020	0.0014	0.046

Note: The parameter λ determines the proportionate effect of the environmental investment ratio, τ, on the environmental disaster probability, q, in equation (7). The coefficient of relative risk aversion, γ, is 3.3 in Sections I and III-VI and 5.0 in Section II. (This change raises $[E(1-b)^{1-\gamma}-1]$ to 5.57, compared with the baseline value of 2.11.) The size distribution of disasters, b, is given in Sections I-II and IV-VI by the historical pattern, corresponding to the histogram in Figure 1, and the sizes are multiplied by 1.1 in Section III. (This change raises Eb to 0.23 and $[E(1-b)^{1-\gamma}-1]$ to 2.55, compared to the baseline values of 0.21 and 2.11.) In Section III, the baseline disaster probability per year is maintained at $p(0)=\pi+q(0)=0.040$. The baseline environmental disaster probability, $q(0)$, is 0.010 per year in Sections I-III and V-VI and 0.005 in Section IV. The rate of time preference, ρ, is 0.044 in Sections I-IV and VI and 0.030 in Section V. (Note from n.15 that ρ^*, the effective rate of time preference in this model with Epstein-Zin/Weil preferences, is well below ρ.) The parameter θ (the reciprocal of the *IES*) is 0.5 in Sections I-V, 1.0 in Section VIa, and 3.3 (γ) in Section VIb. The threshold values of λ that generate positive environmental investment, corresponding to $\tau>0$, are 8.63 in Section I, 4.81 in Section II, 6.76 in Section III, 17.3 in Section IV, 5.65 in Section V, 9.20 in Section VIa, and 11.79 in Section VIb. For values of λ above the thresholds, the solutions for the optimal τ are interior and satisfy equation (13). The consumer-surplus ratios in the far right column are computed from equation (14). These ratios indicate the proportionate decline in today's GDP that society would willingly accept to gain the opportunity to choose forever the optimal environmental-investment ratio, τ, rather than having τ=0.

Table 3 Effects of Uncertainty about Policy Effectiveness				
Semi-elasticities of environmental-disaster probability with respect to τ		***τ:*** **environmental-investment ratio**	***q:*** **environmental-disaster probability**	**consumer-surplus ratio**
$\boldsymbol{\lambda_1}$	$\boldsymbol{\lambda_2}$			
10	10	0.0140	0.00869	0.00108
7.5	12.5	0.0133	0.00876	0.00102
5	15	0.0114	0.00894	0.00088
20	20	0.0415	0.00436	0.0237
15	25	0.0400	0.00458	0.0226
10	30	0.0353	0.00525	0.0197

The environmental-disaster probability is given by $q=0.01 \cdot e^{-\lambda\tau}$, where τ is the ratio of environmental investment to GDP. The parameter λ takes on the two possible values λ_1 and λ_2 with probability one-half each. The other parameter values correspond to the baseline case from Table 2. For each specification for λ, Table 3 shows the optimal choice of τ and the corresponding values of q (the mean corresponding to the two possible values of λ) and the consumer-surplus ratio.

Figure 1

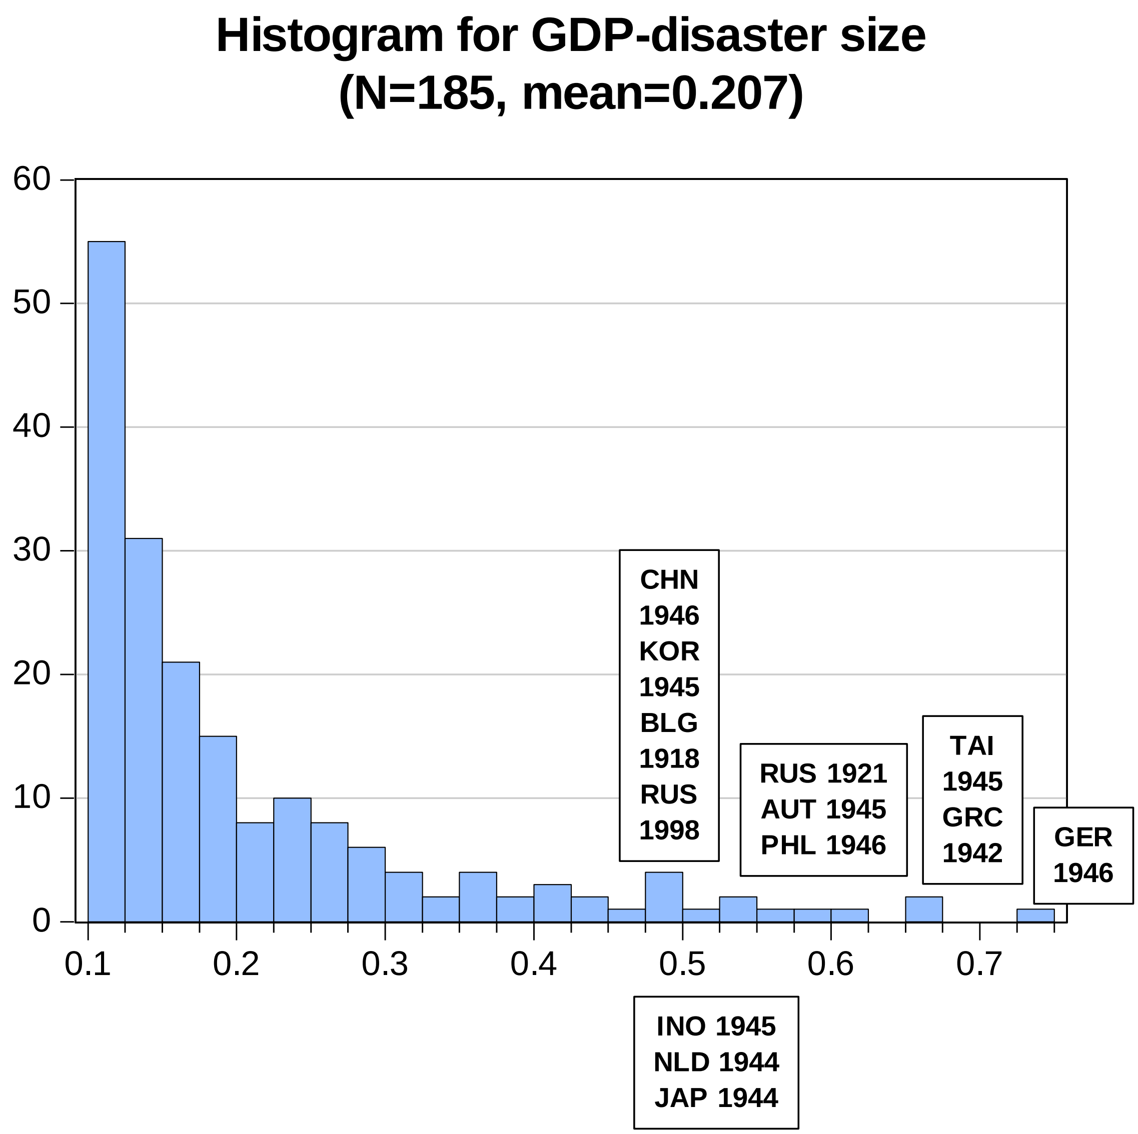

Note: The horizontal axis has the proportionate decline in real per capita GDP for the 185 macroeconomic disasters. The vertical axis has the number of observations. The countries and trough years for the largest disasters are indicated. GER is Germany, TAI is Taiwan, GRC is Greece, RUS is Russia, AUT is Austria, PHL is the Philippines, INO is Indonesia, NLD is the Netherlands, JAP is Japan, CHN is China, KOR is South Korea, and BLG is Belgium.

References

Ackerman, Frank, Elizabeth A. Stanton, and Ramon Bueno (2012), "Epstein-Zin Utility in DICE: Is Risk Aversion Irrelevant to Climate Policy?," unpublished, Stockholm Environmental Institute, March.

Barro, Robert J. (2006). "Rare Disasters and Asset Markets in the Twentieth Century, *Quarterly Journal of Economics*, 121, August, 823–866.

Barro, Robert J. (2009). "Rare Disasters, Asset Prices, and Welfare Costs," *American Economic Review*, 99, March, 243–264.

Barro, Robert J. and Tao Jin (2011). "On the Size Distribution of Macroeconomic Disasters," *Econometrica*, 57, September, 1567-1589.

Barro, Robert J. and Xavier Sala-i-Martin (2004). *Economic Growth*, 2nd edition, Cambridge MA, MIT Press.

Barro, Robert J. and José F. Ursúa (2008). "Macroeconomic Crises since 1870," *Brookings Papers on Economic Activity*, 1, 255–335.

Barro, Robert J. and José F. Ursúa (2012). "Rare Macroeconomic Disasters," *Annual Review of Economics*, 4, 10.1-10.17.

Brainard, William C. (1967). "Uncertainty and the Effectiveness of Policy," *American Economic Review*, 57, May, 411-425.

Cai, Yongyang, Kenneth L. Judd, and Thomas S. Lontzek (2013). "The Social Cost of Stochastic and Irreversible Climate Change, National Bureau of Economic Research, working paper no. 18704, January.

Crost, Benjamin and Christian Traeger (2012). "Risk and Aversion in Assessing Climate Policy," unpublished, Department of Agricultural & Resource Economics, University of California Berkeley, December.

Epstein, Larry G. and Stanley E. Zin (1989). "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 57, July, 937–69.

Gollier, Christian (2006). "An Evaluation of Stern's Report on the Economics of Climate Change," IDEI working paper no. 464, University of Toulouse.

Lucas, Robert E., Jr. (1978). "Asset Prices in an Exchange Economy," *Econometrica*, 46, November, 1429–1445.

Mendelsohn, Robert (2008). "Is the *Stern Review* an Economic Analysis?" *Review of Environmental Economics and Policy,* Winter, 45-60.

Nakamura, Emi, Jon Steinsson, Robert J. Barro, and José F. Ursúa (2013). "Crises and Recoveries in an Empirical Model of Consumption Disasters," *American Economic Journal: Macroeconomics*, 5, July, 35-74.

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

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

Obstfeld, Maury (1994). "Evaluating Risky Consumption Paths: The Role of Intertemporal Substitutability," *European Economic Review*, 38, August, 1471-1486.

Ramsey, Frank (1928). "A Mathematical Theory of Saving," *Economic Journal*, 38, December, 543-559.

Rietz, Thomas A. **(**1988). "The Equity Risk Premium: A Solution," *Journal of Monetary Economics*, 22, July, 117–131.

Stern, Nicholas (2007). *The Economics of Climate Change, the Stern Review,* Cambridge, Cambridge University Press.

Sterner, Thomas and U. Martin Persson (2008). "An Even Sterner Review: Introducing Relative Prices into the Discounting Debate," *Review of Environmental Economic and Policy,* Winter, 61-76.

Weil, Philippe (1990). "Nonexpected Utility in Macroeconomics," *Quarterly Journal of Economics*, 105, February, 29–42.

Weitzman, Martin L. (2007). "A Review of the 'Stern Review on the Economics of Climate Change'," *Journal of Economic Literature*, 45, September, 703-724.

Weitzman, Martin L. (2009). "On Modeling and Interpreting the Economics of Catastrophic Climate Change," *The Review of Economics and Statistics*, 91, February, 1-19.