

UC Berkeley

CUDARE Working Papers

Title

Discounting and Climate Change Policy

Permalink

<https://escholarship.org/uc/item/5sm6j36x>

Authors

Karp, Larry
Tsur, Yacov

Publication Date

2007-05-17

Discounting and Climate Change Policy

Larry Karp* Yacov Tsur[◇]

May 17, 2007

Abstract

A constant social discount rate cannot reflect both a reasonable opportunity cost of public funds and an ethically defensible concern for generations in the distant future. We use a model of hyperbolic discounting that achieves both goals. We imbed this discounting model in a simple climate change model to calculate “constant equivalent discount rates” and plausible levels of expenditure to control climate change. We compare these results to discounting assumptions and policy recommendations in the Stern Review on Climate Change.

Keywords: discounting, climate change modeling, Stern Review, Markov Perfect Equilibria

JEL Classification numbers: C61, C73, D63, D99, Q54

*Department of Agricultural and Resource Economics, 207 Giannini Hall, University of California, Berkeley CA 94720 email:karp@are.berkeley.edu

[◇]Department of Agricultural Economics and Management, The Hebrew University of Jerusalem, P.O. Box 12 Rehovot 76100, Israel (tsur@agri.huji.ac.il).

1 Introduction

Public projects that provide benefits in the distant future rise or fall on discounting assumptions. The recent debate on The Stern Review on Climate Change (Stern 2006) is a case in point. The Review recommends expenditures on greenhouse gas (GHG) abatement substantially larger than those suggested by other Integrated Assessment Models (IAMs – numerical models that combine economic and climate modules). Several commentators, including Dasgupta (2007a,b), Nordhaus (2006) and Weitzman (2007) claim that these results are due to unreasonable assumptions about the social discount rate. We use a model of hyperbolic discounting to assess the plausibility of both the assumptions and the policy recommendations in the Stern Review.

The controversy turns on the parameters of the social discount rate. Let $\theta(\tau)$ be the present value of a unit of utility τ time periods in the future, and let $u'(c_{t+\tau})$ equal marginal utility of consumption at that time. The discount factor for consumption, $\theta(\tau)u'(c_{t+\tau})$, equals the number of dollars that a social planner would be willing to sacrifice today (time t) in order to obtain one more unit of consumption at time $t + \tau$. The social discount rate, $r(\tau, t + \tau)$, is the rate of decrease of this quantity:

$$r(\tau, t + \tau) = \rho(\tau) + \eta(c_{t+\tau})g(c_{t+\tau}), \quad (1)$$

where $\rho = -\dot{\theta}/\theta$ is the pure rate of time preference, $\eta = -(u''/u')c$ is the elasticity of marginal utility and $g = \dot{c}/c$ is the growth rate in per-capita consumption. The baseline Stern parameters are $\rho = 0.1\%$, $\eta = 1$ and $g = 1.3\%$, implying $r = 1.4\%$, which is far lower than rates used in other IAMs (see, e.g., Nordhaus 1999).

The debate about discounting in the Stern Review has been framed in a manner that is unlikely to lead to the emergence of a consensus. The problem is that the standard representation of the social discount rate is too parsimonious. The social discount rate is a valuable parameter for

summarizing a model, but a single number (resulting from constant values of the parameters on the right side of equation (1)) cannot capture both inter-generational equity over a long time scale and the opportunity cost of public funds.

Our suggestion for assessing discounting parameters in climate change models is to begin with a time-varying social discount rate that reflects both a reasonable degree of inter-generational equity (the long run considerations) and the opportunity cost of public funds. We imbed this discount function in a transparent climate change model. This model captures the risk of climate-related damage and the inertia in the climate system. That inertia leads to a delayed relation between current actions and future reductions in risk. This model extends Karp and Tsur (2007) by including exogenous growth in consumption, g , a necessary feature in order for equation (1) to be useful.

In some cases it is possible to construct a “constant equivalent discount rate”, i.e. a single number that, if used as the social discount rate, would lead to policy prescriptions identical to those obtained under the time-varying social discount rate. This constant-equivalent discount rate depends on the time-varying social discount rate, which should be the same function for all public projects. The constant-equivalent discount rate also depends on the specifics of the problem, in particular the longevity of the public project. For example, decisions about climate policy may affect welfare over many centuries, while a decision about a bridge may affect welfare over a century. The differing time scale of these two types of public projects means that the constant-equivalent discount rates corresponding to them may be very different, even though both are based on the same time-varying social discount rate. The constant-equivalent discount rate for the climate project is more sensitive to social discount rates in the distant future. A comparison of the constant equivalent discount rates with the rates used in IAMs helps

to assess the reasonableness of the latter.¹

There are at least three plausible reasons why the social discount rate is a decreasing function of time. First, uncertainty about future growth rates means that the “certainty equivalent” discount rate is likely to fall over time; therefore, the constant-equivalent discount rate tends to be smaller for projects that produce benefits at a more distant point in the future (Weitzman 1998, 1999). Second, if there is imperfect substitutability between produced goods and the natural environment, and if the supply of produced goods increases while the stock of the natural environment falls, the social discount rate may fall over time (Traeger 2004, Hoel and Sterner 2007).

We emphasize the third source of time-dependence: hyperbolic discounting due to a decrease in the pure rate of time preference. This rate is positive if we prefer an increase in utility sooner rather than later; $1 + \rho(\tau)$ is (approximately) the number of units of utility we would be willing to surrender at time $\tau + d\tau$ to obtain one more unit of utility at time τ . If we prefer the youngest generation currently alive to the next, not-yet-born generation, ρ should be positive in the near term. We may be less able to distinguish between two contiguous generations in the distant future – they appear “more similar” than two contiguous generations in the near future, implying that ρ falls over time. Arrow (1999) notes that “agent-relative ethics suggests that each generation will maximize a weighted sum of its own utility and the sum of utility of all future generations, with less weight on the latter. At the very least, really distant generations are treated all alike.” Cropper et al. (1994) find that questionnaire respondents weight returns to people living

¹Our approach to evaluating the choice of social discount rates in some respects resembles the approach in Weitzman (1999). Both recognize that the social discount rate depends on the time scale of the project. There are important differences, however. Weitzman emphasizes non-constant discounting arising from uncertain growth rates, whereas we begin with hyperbolic discounting, and the uncertainty in our model is due to a climate-related event. In addition, he is interested in the certainty equivalent social discount rate used to evaluate a payoff at a point in time in the future. We are interested in a constant equivalent discount rate used to evaluate a trajectory of payoffs.

100 years from now only slightly more than people living 200 years from, a result consistent with hyperbolic but not with exponential discounting.

Nordhaus (1999) and Mastrandrea and Schneider (2001) imbed hyperbolic discounting in IAMs, in order to show the importance of long run discounting in determining climate change policy. These authors assume that the decision-maker in the current period can choose the entire policy trajectory, thus solving by assumption the time-inconsistency problem (which is a consequence of hyperbolic discounting).

The time-inconsistency of optimal programs is a plausible feature of the policy problem: politicians, like other mortals, tend to procrastinate in solving difficult problems. Because of the long time scale over which policies must be implemented, we think that the restriction to Markov perfection provides a reasonable description of the policy problem. There are multiple equilibria in this setting, but these can be bounded in a simple manner. We compare this set of equilibria to a benchmark (called “restricted commitment”) in which the policymaker’s set of feasible policies is restricted in such a way as to cause the resulting optimal choice to be time consistent. This outcome is not plausible, but it provides a useful comparison to the set of Markov perfect equilibria (MPE), and because of its optimality and time consistency it has a simple welfare interpretation.

The next section describes our extension of Karp and Tsur (2007), providing only enough detail for this paper to be self-contained. The following section calibrates the hazard and discounting parameters; the resulting time-varying social discount rate reflects a reasonable opportunity cost of public funds and a degree of inter-generational equity. We then calculate the constant equivalent social discount rates and suggest plausible levels of expenditure to reduce the risk of climate-related damage.

2 The model

The model incorporates hyperbolic discounting, exogenous growth, and uncertainty about the timing of a catastrophic event (Tsur and Zemel 1996, Alley et al. 2003). We first describe the payoff corresponding to different actions and states, and then discuss the relation between actions and risk. The final subsection describes the two types of equilibria corresponding to the different assumptions about the regulator’s ability to commit to future actions.

2.1 The payoff

The “catastrophe” occurs abruptly at a random date, reducing income by the fraction Δ (the income-at-risk) from the occurrence time onward. At each point in (continuous) time, society has a choice between two actions: “business as usual (BAU)” and “stabilization”. The next section describes the effects of these two actions; here we describe their costs. BAU costs nothing, and stabilization costs the fraction $X\Delta$ of the flow of income. The parameter X is the fraction of the income-at-risk that we need to spend in order to stabilize. The control variable, $w(t) \in \{0, 1\}$ is a switch: $w(t) = 0$ means that at time t society follows BAU, and $w(t) = 1$ means that at time t society stabilizes.

Consumption grows at an exogenous constant rate g , and the utility of consumption is iso-elastic, with the constant elasticity η .² With initial consumption normalized to 1, the flow of consumption prior to the catastrophe is $e^{gt}(1 - \Delta X w_t)$ and utility is

$$u(c(t), w(t)) = \frac{(e^{gt}(1 - \Delta X w_t))^{1-\eta}}{1 - \eta}. \quad (2)$$

²The assumption of exogenous growth in consumption means that there is no distinction in this model between consumption and income. The model is consistent with a constant savings rate if all expenditures for climate change mitigation come from consumption rather than savings.

After the event occurs, income falls to $e^{gt}(1-\Delta)$ and there is no role for stabilization.

Conditional on the event occurring at time T , the payoff under policy $w(t)$ is

$$\int_0^T \theta(t) \frac{(e^{gt}(1-\Delta X w(t)))^{1-\eta}}{1-\eta} dt + \int_T^\infty \theta(t) \frac{(e^{gt}(1-\Delta))^{1-\eta}}{1-\eta} dt =$$

$$\int_0^T \theta(t) \left(\frac{(e^{gt}(1-\Delta X w(t)))^{1-\eta}}{1-\eta} - \frac{(e^{gt}(1-\Delta))^{1-\eta}}{1-\eta} \right) dt + \text{constant}$$

where

$$\text{constant} = \int_0^\infty \theta(t) \frac{(e^{gt}(1-\Delta))^{1-\eta}}{1-\eta} dt. \quad (3)$$

Ignoring the constant term, the conditional payoff can be written as

$$\int_0^T \theta(t) e^{-g(\eta-1)t} \left(\frac{(1-\Delta X w(t))^{1-\eta}}{1-\eta} - \frac{(1-\Delta)^{1-\eta}}{1-\eta} \right) dt$$

$$\equiv \int_0^T \theta(t) e^{-g(\eta-1)t} U(w(t)) dt,$$

where

$$U(w) \equiv \frac{(1-\Delta X w)^{1-\eta}}{1-\eta} - \frac{(1-\Delta)^{1-\eta}}{1-\eta}. \quad (4)$$

Prior to the catastrophe, the current generation's payoff is

$$E_T \int_0^T \theta(t) e^{-g(\eta-1)t} U(w(t)) dt, \quad (5)$$

where E_T denotes expectation with respect to the random occurrence time.

We express the equilibria as a function of x , a parameter proportional to the utility cost of stabilization:

$$x \equiv 1 - \frac{U(1)}{U(0)} = 1 - \frac{(1-\Delta X)^{1-\eta} - (1-\Delta)^{1-\eta}}{1 - (1-\Delta)^{1-\eta}}.$$

We invert this relation to write X as a function of x :




Figure 1: Graph of $X(x)$ for $\Delta = 0.2$ and $\eta = 4$ (solid) and $\eta = 1$ (dotted).

$$X = \frac{1}{\Delta} \left[1 - \{1 - x [1 - (1 - \Delta)^{1-\eta}]\}^{\frac{1}{1-\eta}} \right]. \quad (6)$$

Figure 1 shows the graphs of $X(x)$ for $\eta = 1$ and $\eta = 4$ when $\Delta = 0.2$.³

The utility discount factor is the convex combination of exponential functions

$$\theta(t) = \beta e^{-\gamma t} + (1 - \beta) e^{-\delta t}, \quad (7)$$

with $0 < \beta < 1$. This function permits a simple calibration, it allows for substantial flexibility, and it implies a declining pure rate of time preference. In view of expression (5), we define the “effective discount factor”, a function that incorporates both the pure rate of time preference and the effect of η and g :

$$\tilde{\theta}(t) \equiv \theta(t) e^{g(1-\eta)t} = \beta e^{-\tilde{\gamma}t} + (1 - \beta) e^{-\tilde{\delta}t}, \quad (8)$$

where

$$\tilde{\gamma} = \gamma + g(\eta - 1) \quad (9a)$$

and

$$\tilde{\delta} = \delta + g(\eta - 1). \quad (9b)$$

³For the logarithmic case ($\eta = 1$), equation (6) reduces to $X = (1 - (1 - \Delta)^x)/\Delta$.

The “effective discount rate” is the rate of decrease of $\tilde{\theta}(t)$.

When $\eta \neq 1$ and $g \neq 0$ this model has one degree of freedom: for given β , the “effective discount rate” depends on $\tilde{\gamma}$ and $\tilde{\delta}$, determined by two equations in three unknowns, δ , γ , and g . These parameters, unlike η , do not enter the function U , defined in equation (4). We normalize by setting $\gamma = 0$.⁴ This normalization implies that the long run pure rate of time preference is 0, i.e. it means that we are unwilling to transfer utility between two agents living in the infinitely distant future at a rate other than one-to-one. It also implies that the long run effective discount rate is $g(\eta - 1)$.

2.2 The relation between actions and risk

The endogenous hazard rate, $h(t)$, determines the probability distribution of the occurrence time of the catastrophe. The change in the hazard is

$$\dot{h}(t) = \mu(a - h(t))(1 - w(t)); h(0) \text{ given.} \quad (10)$$

This model is a simplified representation of the following situation. The actions that we take at a point in time (e.g. abatement, levels of consumption) determine greenhouse gas (GHG) emissions levels at that time. These flows, and existing GHG stocks, determine the evolution of the stock of GHG. The risk of a climate-related catastrophe, given by $h(t)$, is a monotonic function of the stock of GHG. We can invert this function to write the time derivative of h as a function of h and w , as in equation (10).⁵

⁴When $\eta = 1$ the equilibrium is always independent of g . For $\eta = 1$ or $g = 0$, γ and δ equal $\tilde{\gamma}$ and $\tilde{\delta}$. In this case, setting $\gamma = 0$ is an assumption, not a normalization.

When $g > 0$, the constant defined in equation (3) is finite if and only if $\eta > 1$. In contrast, the maximand in expression (5) is defined even for some values of $\eta < 1$, because the hazard has an effect similar to discounting. For $\eta \leq 1$ we can adopt the “overtaking criterion” to evaluate welfare.

⁵Let S_t be the stock of GHG and z_t the control variable at time t ; $\dot{S} = q(S, z)$ is the equation of motion. Suppose that the hazard associated with the catastrophe is a monotonic function $h_t = H(S_t)$; define the inverse function $S = K(h) \equiv H^{-1}(h)$, so $\dot{h} = H'(S)q(S, z) = H'(K(h))q(K(h), z)$.

Equation (10) implies that (for $h < a$, as we assume throughout) the hazard grows most quickly when h is small. This feature means that society is willing to spend more (i.e. tolerate a larger value of X) to stabilize when h is small. For hazards close to the steady state, there is little benefit in incurring costs in order to prevent the hazard from growing.⁶

This model implies that the level of the hazard, not simply the occurrence of the catastrophe, is irreversible. Given the level of inertia in the climate system, this assumption seems reasonable. It also simplifies the class of equilibria, because it eliminates the possibility of non-monotonic outcomes (e.g. allowing the hazard to grow and then causing it to fall).

Defining $y(t) = \int_0^t h(\tau) d\tau$, the payoff (expression (5)) is evaluated as

$$E_T \int_0^T \theta(t) e^{-g(\eta-1)t} U(w(t)) dt = \int_0^\infty \theta(t) e^{-g(\eta-1)t - y(t)} U(w(t)) dt. \quad (11)$$

The simplicity of equation (10) is important. There are conjectures on the level of risk for different types of events (such as a reversal of the thermohaline circuit or a rapid increase in sea level) corresponding to different policy trajectories (e.g. BAU or specific abatement trajectories). We can use these kinds of conjectures to suggest reasonable magnitudes for the parameters of the risk model ($h(0)$, a and μ). There is little empirical basis for calibrating a more complicated model.

2.3 Types of equilibria

If the decision-maker at time 0 were able to choose an arbitrary policy trajectory to maximize the payoff in expression (11) subject to the constraint in equation (10), a typical solution involves procrastination: the decision-maker decides to begin stabilization at some time in the future. For well-known reasons, this policy is time-inconsistent.

⁶The results in a model in which \dot{h} is non-monotonic in h would change in fairly obvious ways. For example, if \dot{h} is small when h is close to both 0 and the steady state level, stabilization would not be worthwhile either for very small or for very large levels of h .

We emphasize the situation where the decision-maker at time t can choose the current policy $w \in \{0, 1\}$, but not future policies. The decision-maker understands how this choice affects the evolution of the hazard, and forms beliefs about how future regulators' decisions depend on the future level of the state variable h . The resulting MPE is a Nash equilibrium to the sequential game played by the succession of regulators. Each regulator chooses the current decision and wants to maximize the present discounted value of the stream of future payoffs, given by expression (11). (Each regulator rewinds the time at which he makes the decision to $t = 0$.)

There are in general multiple MPE because the optimal decision for the current regulator depends on his beliefs about the actions of subsequent regulators. The equilibrium beliefs of the current regulator (i.e. those that turn out to be correct) depend on his beliefs about the beliefs (and thus the actions) of successors. There is an infinite sequence of these higher order beliefs, leading to generic multiplicity of equilibria.

However, due to the binary nature of the problem, and the specific form of the hazard equation, the equilibrium set has a simple characterization. There are upper and lower limits of X , which we denote as $X^U(h)$ and $X^L(h)$, such that: the unique MPE is perpetual stabilization for $X \leq X^L(h)$; the unique MPE is perpetual BAU for $X \geq X^U(h)$; and there are MPE with either perpetual stabilization or perpetual BAU for $X^L(h) < X < X^U(h)$. For a subset of $(X^L(h), X^U(h))$ there are additional MPE that involve delayed stabilization, but we do not discuss those here.

We also consider a benchmark equilibrium, called "limited commitment", in which the current policymaker decides between perpetual BAU and perpetual stabilization (i.e. he can commit future generations to one of the two policies without the possibility to switch between them in the future). This equilibrium is time-consistent, but it is not particularly plausible. It involves either an unreasonable amount of commitment, or an arbitrary restriction of the set of feasible actions. In this equilibrium, the policymaker chooses to

stabilize if and only if $X \leq X^C(h)$. The equilibrium is useful as a benchmark for two reasons: (i) it has the same structure (a critical bound) as the MPE, and (ii) the fact that it is a constrained optimum means that it has a simple welfare interpretation.

Our companion paper (Karp and Tsur 2007) gives the formulae for the boundary functions corresponding to the utility cost of stabilization $x^U(h)$, $x^L(h)$, and $x^C(h)$. These functions are independent of Δ ; they depend on the growth rate g and the utility parameter η via the discounting parameters $\gamma + g(\eta - 1)$ and $\delta + g(\eta - 1)$. We use those formulae and equation (6) to obtain the boundaries in terms of consumption, $X^L(h)$, $X^U(h)$ and $X^C(h)$. These functions depend on all of the parameters of the model.

This model shows the potentially offsetting effects of an increase in η . This parameter affects the equilibrium by altering the “effective discount rate” $\rho(t) + g(\eta - 1)$ and it also enters the function $X(x)$ defined in equation (6). For $g > 0$, an increase in η increases the “effective discount rate” $\rho(t) + g(\eta - 1)$, which in turn decreases the critical values $x^U(h)$, $x^L(h)$, and $x^C(h)$; that is, the change makes the decision-maker less willing to sacrifice current utility for future reduction in risk. However, the larger value of η makes the decision-maker more risk averse; it shifts up the graph of $X(x)$, as shown in Figure 1, so the smaller value of the critical x (resulting from the increase in η) might correspond to a larger value of the critical X . Thus, in general the effect of η on X is ambiguous. For our calibration, discussed in the next section, an increase in η reduces X .

3 Policy bounds and constant equivalent rates

We discuss the calibration of the model and then present the critical values of X corresponding to the MPE and the limited commitment equilibrium. We also present the corresponding constant equivalent pure rates of time

preference; these are the rates that would yield the same policy bounds if ρ were constant (i.e., if $\beta = 0$ or $\beta = 1$ or $\delta = \gamma$).

We obtain an exact constant equivalent discount rate because of the simplicity of the model. The exact equivalence occurs if the decision rules under both hyperbolic and constant discounting can be characterized by a single parameter. In our problem, the single parameter is the critical value of X below which stabilization occurs.⁷

3.1 Calibration

We choose the hazard parameters $h(0)$, μ and a in order to satisfy: (i) under stabilization the probability of occurrence within a century is 0.5%; (ii) in the BAU steady state, where $h = a$, the probability of occurrence within a century is 50%; and (iii) under BAU it takes 120 years to travel half way between the initial and the steady state hazard levels. These assumptions imply $a = 0.00693147$, $h_0 = 0.000100503$ and $\mu = 0.00544875$. With these values, the probability of occurrence within a century is 15.3% under BAU, compared to 0.5% under stabilization.

In order to be able to compare the damage estimates under our calibration with those used by IAMs, we define $P^B(t) \equiv Pr\{T \leq t|BAU\}$ as the probability that the catastrophe occurs by time t under BAU, and $P^S(t) \equiv Pr\{T \leq t|Stabilization\}$ as the corresponding probability under stabilization. The future (time t) expected increase in damages from following BAU rather than stabilization, as a percentage of future income, is $D(t) = (P^B(t) - P^S(t)) 100\Delta\%$. For all calibrations where $h(0) > 0$,

⁷Barro (1999) also obtains a constant equivalent discount rate, because the single parameter in his logarithmic model is the slope of the decision rule. When the decision rules cannot be described by a single parameter, it is possible only to obtain an approximate constant equivalent discount rate. For example, in the linear-quadratic model there exists a linear equilibrium control rule under both constant and hyperbolic discounting. Because this control rule involves two parameters – the slope and the intercept – it is in general not possible to find a constant equivalent discount rate for the hyperbolic model (Karp 2005).




Figure 2: Percentage expected increased loss of income under BAU: $\Delta = 0.05 =$ dashed; $\Delta = 0.1 =$ solid; $\Delta = 0.2 =$ dotted.

$\lim_{t \rightarrow \infty} D(t) = 0$, because both probabilities converge to 1.⁸ Figure 2 shows the graphs of $D(t)$ over the next millennium for $\Delta = 0.05$, 0.1 and 0.2. The corresponding damages after 100 and 200 years are $D(100) = \{0.72, 1.43, 2.88\}$ and $D(200) = \{2.03, 4.01, 8.11\}$.

The Stern Review provides a range of damage estimates. Their second-lowest damage scenario (“market impacts + risk of catastrophe”) assumes that climate-related damages equal about 1% in one century, and 5% after two centuries. Our calibration with $\Delta = 0.05$ implies significantly lower damages over the next two centuries. The Stern Review also describes scenarios in which damages might be as high as 15-20% of income, a level considerably above our scenario with $\Delta = 0.2$ (for the next two centuries).

The Stern Review assumes that climate-related damages are zero after 200 years, whereas in our calibration damages continue to rise for 800 years and then decrease asymptotically to 0. The maximum level of $D(t)$ equals $91\Delta\%$, i.e. 4.5%, 9.1% and 18.2% for the three values of Δ . In view of the different profiles of damages in the Stern Review and in our calibration, exact

⁸Using equation (10), $P^B(t) = 1 - e^{-at + (a-h(0))(1-e^{-\mu t})/\mu}$ and $P^S(t) = 1 - e^{-h(0)t}$. For $h_0 = 0$, $D(t) = P^B(t) 100\Delta$, which converges to $100\Delta\%$.

matching is not possible. However, our case $\Delta = 0.2$ approximates one of the high (but not the highest) Stern damage scenarios; the value $\Delta = 0.1$ approximates the Stern “market impacts + risk of catastrophe” scenario, and the value $\Delta = 0.05$ corresponds to a much lower damage scenario.

Using equation (7), the pure rate of time preference is

$$\rho(t) \equiv \frac{-\dot{\theta}(t)}{\theta(t)} = \frac{\beta\gamma e^{-\gamma t} + \delta e^{-\delta t}(1 - \beta)}{\beta e^{-\gamma t} + e^{-\delta t}(1 - \beta)}. \quad (12)$$

We set $\gamma = 0$, so that the long-run pure rate of time preference is 0, and choose β and δ in order to satisfy

$$\rho(0) = 0.03 \quad \text{and} \quad \rho(30) = 0.01.$$

Our value of $\rho(30)$ is ten times greater than the Stern Review’s constant pure rate of time preference.

3.2 Results

For a variety of parametric and equilibrium assumptions, we calculated upper and lower bounds on X – the fraction of income-at-risk that society spends to stabilize risk. These values were insensitive to choices of Δ over the interval $(0.1, 0.2)$, so the tables below report only results for $\Delta = 0.2$. We also report the corresponding constant equivalent discount rates. We discuss results for $g \in [1\%, 2\%]$ and $\eta \in [1, 4]$. As we noted above, when $\eta = 1$ the long run effective discount rate is 0 and the integral in equation (3) (which plays no role in determining the equilibrium) does not converge; moreover, the results are extremely sensitive in the neighborhood of $\eta = 1$. This value is therefore a special case. Consequently, we emphasize the case where $1.1 \leq \eta \leq 4$ and consider the case $\eta = 1$ separately.

Tables 1 – 3 show the (X) policy bounds and constant-equivalent ρ values for the 6 cases corresponding to $\eta \in \{1.1, 2, 4\}$ and $g \in \{0.01, 0.02\}$. In each case the constant equivalent social discount rate (not shown) equals

the constant equivalent value of ρ plus ηg . We emphasize the case where $\eta = 2$ and compare the results for $g = 1\%$ and $g = 2\%$ across the different equilibria.

We begin with the limited commitment equilibrium, which is both time consistent and constrained optimal. For $\eta = 2$, the maximum fraction of the income at risk that society would forgo in order to stabilize ranges between 6% and 17% as g changes from 2% to 1%. For these experiments, where $\Delta = 0.2$, these bounds imply expenditures of between 1.2% and 3.4% of GWP. If $\Delta = 0.1$, the corresponding values of X^C are 5.4% and 15.5%, implying an expenditure of between 0.54% and 1.5% of GWP. These values bracket the Stern recommendation to spend 1% of GWP annually on climate change policy. For $\Delta = 0.2$, the constant equivalent values of ρ range from 0.13% and 0.32%, so the constant equivalent social discount rate ranges between 2.13% and 4.32%.

Table 1: Limited commitment upper bounds X^C and constant-equivalent ρ values for $\eta \times g = \{1.1, 2, 4\} \times \{0.01, 0.02\}$.

η	$g = 1\%$		$g = 2\%$	
	X^C (%)	Cons-equiv ρ (%)	X^C (%)	Cons-equiv ρ (%)
1.1	76.22	0.01	60.71	0.02
2	17.1	0.13	6.07	0.32
4	3.78	0.53	1.06	1.09

For $g = 1\%$ and $\eta = 2$ the upper and lower bounds of X in a MPE are 17.8% and 9.9%, with corresponding constant equivalent values of ρ of 0.1% and 0.7% (Tables 2 and 3). In this case, for $17\% < X < 17.8\%$ of the value at risk, the optimal policy is to follow BAU, but there are MPE that result in excessive stabilization. For $9.9\% < X < 17\%$ the optimal policy is to stabilize, but there are MPE that result in BAU. Thus, a MPE may result in either excessive or insufficient stabilization (although, in a sense, the latter

is more likely). The broad range of values for which there are multiple MPE indicates the importance of establishing commitment devices that enable the current generation to lock in the desired policy trajectory.

Table 2: MPE upper bounds X^U and constant-equivalent ρ values for $\eta \times g = \{1.1, 2, 4\} \times \{0.01, 0.02\}$.

η	$g = 1\%$		$g = 2\%$	
	X^U (%)	Cons-equiv ρ (%)	X^U (%)	Cons-equiv ρ (%)
1.1	94.15	-0.08	81.55	-0.12
2	17.80	0.1	5.44	0.49
4	3.33	0.8	0.98	1.4

For $g = 2\%$ the upper and lower bounds (5.4% and 4%) are much closer (compared to when $g = 1\%$), and both lie below the upper bound under restricted commitment. In this case, for any X such that stabilization is a MPE, stabilization also maximizes welfare. For $5.4\% < X < 6.1\%$ all MPE involve BAU even though stabilization is optimal. With $g = 2\%$ the constant equivalent ρ in a MPE ranges from 0.5% to 1%. As expected, the bounds are lower under the higher ($g = 2\%$) growth rate, since the current generation is less willing to sacrifice for the sake of much wealthier future generations.

Table 3: MPE lower bounds X^L and constant-equivalent ρ values for $\eta \times g = \{1.1, 2, 4\} \times \{0.01, 0.02\}$.

η	$g = 1\%$		$g = 2\%$	
	X^L (%)	Cons-equiv ρ (%)	X^L (%)	Cons-equiv ρ (%)
1.1	38.3	0.37	30.86	0.41
2	9.89	0.69	3.98	1.00
4	2.73	1.25	0.9	1.76

Table 4 in the appendix shows that in the neighborhood of $\eta = 1$ society is willing to spend nearly all of the value at risk simply in order to reduce the future risk. However, the case $\eta = 1$ and neighboring cases are actually not informative. Recall that our model contains one free parameter when $\eta \neq 1$, and we used this degree of freedom to set $\gamma = 0$. Had we chosen a positive value of γ , we could have adjusted the values of g in all the previous tables (where $\eta \neq 1$) without changing the results. However, the positive value of γ would have caused a significant change in the results for $\eta = 1$.

Table 5 reports some bounds when we change the discounting calibration so that $\rho(30) = 2\%$ (rather than 1% as above). This rather modest change in discounting assumptions leads to large percentage changes in the bounds on X . This experiment reinforces the point that policy prescriptions can be quite sensitive to assumptions about discount rates.

4 Conclusion

A constant discount rate cannot reflect both a reasonable opportunity cost of public funds and an ethically defensible concern for distant future generations. One way to reconcile these two goals is to let the social discount rate vary over time. There are several reasons why the social discount rate might be a function of time. We emphasized the role of a declining pure rate of time preference (hyperbolic discounting), which gives rise to time inconsistency and a tendency to defer difficult decisions.

In our binary action model, a reduction in current consumption (“stabilization”) reduces the future hazard rate of a random event that causes permanent loss of utility. There are multiple Markov Perfect equilibria (MPE) for an interval of stabilization costs. The upper bound of this interval is the maximum cost consistent with a MPE involving stabilization; the lower bound is the maximum cost consistent with a MPE involving BAU. For each of these bounds we calculate a constant equivalent discount rate,

i.e. a constant discount rate that leads to the same decision rule as the time-varying discount rate. We compared the set of MPE to a time-consistent (constrained optimal) reference equilibrium. The MPE equilibrium set indicates how much society would be willing to spend to stabilize the risk, if it managed to solve the intragenerational collective action problem; the reference equilibrium indicates how much society should be willing to spend.

Our risk and damage calibration includes the moderate and the high damage estimates in the Stern Review. If the catastrophe reduces income by 10-20%, the calibration implies a range of expected damages (under BAU) of 1.4 - 2.9% after 100 years and 4 - 8% after 200 years. Our discounting calibration assumes that the pure rate of time preference begins at 3%, falls to 1% over the first 30 years, and then asymptotically declines to 0. As η (the elasticity of marginal utility) ranges between 2 and 4 and g (the growth rate) ranges between 1% and 2%, the constant equivalent pure rate of time preference ranges between 0.1% and 1.8% (depending on the equilibrium assumption). For $\eta = 2$ and $g = 2\%$, society is willing to spend between 0.5% – 1% of GWP per year to reduce the risk in a MPE; society is willing to spend between 0.6% – 1.2% under limited commitment. A lower η or lower g implies higher expenditures.

As the Stern Review emphasizes, quantitative models are valuable in providing a reality check for policy prescriptions only when used in conjunction with other qualitative and quantitative information. The limited empirical basis for these models, and the fact that they inevitably incorporate value judgements, means that they are unlikely to resolve policy disagreements. Our model incorporates two important aspects of the policy problem: the risk of abrupt events, and hyperbolic discounting (in order to reflect both the opportunity costs of public funds and a concern for future generations). Moreover, it does so in a simple framework, making it easy to see the relation between the policy conclusions presented above and our choice of parameters, and to explore the effects of other parameter choices.

References

- Alley, R. B., Marotzke, J., Nordhaus, W. D., Overpeck, J. T., Peteet, D. M., Pielke Jr., R. S., Pierrehumbert, R. T., Rhines, P. B., Stocker, T. F., Talley, L. D. and Wallace, J. M.: 2003, Abrupt climate change, *Science* **299**, 2005–2010.
- Arrow, K. J.: 1999, Discounting, morality and gaming, in P. R. Portney and J. P. Wayant (eds), *Discounting and Intergenerational Equity*, Resources for the Future, Washington, DC, pp. 13–22.
- Barro, R.: 1999, Ramsey meets Laibson in the neoclassical growth model, *Quarterly Journal of Economics* **114**, 1125–52.
- Cropper, M. L., Aydede, S. and Portney, P.: 1994, Preferences for life saving programs: How the public discounts time and age, *Journal of Risk and Uncertainty* **8**, 243–265.
- Dasgupta, P.: 2007a, Commentary: The stern review’s economics of climate change, *National Institute Economic Review* **199**(1), 4–7.
- Dasgupta, P.: 2007b, Discounting climate change. University of Cambridge.
- Hoel, M. and Sterner, T.: 2007, Discounting and relative prices, web unpublished.
- Karp, L.: 2005, Global warming and hyperbolic discounting, *Journal of Public Economics* **89**, 261–282.
- Karp, L. and Tsur, Y.: 2007, Climate policy when the distant future matters: Catastrophic events with hyperbolic discounting, *Working paper*, Giannini Foundation.
- Mastrandrea, M. D. and Schneider, S. H.: 2001, Integrated assessment of abrupt climatic changes, *Climate Policy* **1**, 433–449.

- Nordhaus, W.: 2006, The stern review on the economic of climate change.
URL: <http://nordhaus.econ.yale.edu/SternReviewD2.pdf>
- Nordhaus, W. D.: 1999, Discounting and public policies that affect the distant future, *in* P. R. Portney and J. P. Weyant (eds), *Discounting and intergenerational equity*, Resources for the Future, Whashington, DC.
- Stern, N.: 2006, Stern review on the economics of climate change, *Technical report*, HM Treasury, UK.
- Traeger, C.: 2004, Should environmental goods be discounted hyperbolically? a general perspective on individual discount rates, unpublished.
- Tsur, Y. and Zemel, A.: 1996, Accounting for global warming risks: Resource management under event uncertainty, *Journal of Economic Dynamics & Control* **20**, 1289–1305.
- Weitzman, M.: 2007, The stern review of the economics of climate change, forthcoming, *Journal of Economic Literature*.
- Weitzman, M. L.: 1998, Why the far-distant future should be discounted at its lowest possible rate, *Journal of Environmental Economics and Management* **36**, 201–208.
- Weitzman, M. L.: 1999, Just keep discounting, but..., *in* P. R. Portney and J. P. Weyant (eds), *Discounting and intergenerational equity*, Resources for the Future, Washington, DC.

Appendix: Additional sensitivity results

Table 4: Additional runs to illustrate the sensitivity of results in the neighborhood of $\eta = 1$. Policy bounds and constant-equivalent ρ values when $\eta \times g = \{1.01, 1.001, 1\} \times \{0.01, 0.02\}$.

	X-bound (%)	Cons-equiv ρ (%)
<u>$\eta = 1.01, g = 1\%$</u>		
Limited commitment upper bound	94.96	0.00
MPE upper bound	99.79	-0.02
MPE lower bound	47.15	0.34
<u>$\eta = 1.001, g = 1\%$</u>		
Limited commitment upper bound	97.18	0.00
MPE upper bound	99.94	-0.01
MPE lower bound	48.19	0.336609
<u>$\eta = 1.01, g = 2\%$</u>		
Limited commitment upper bound	92.57	0.00235185
MPE upper bound	99.53	-0.0281171
MPE lower bound	45.99	0.343802
<u>$\eta = 1.001, g = 2\%$</u>		
Limited commitment upper bound	96.92	0.00
MPE upper bound	99.92	-0.01
MPE lower bound	48.07	0.34
<u>$\eta = 1$</u>		
Limited commitment upper bound	97.43	0.00
MPE upper bound	99.95	-0.01
MPE lower bound	48.31	0.34

Table 5: Sample of results when $\rho(30) = 2\%$.

	X -bound (%)	Cons-equiv ρ (%)
1. <u>$\eta = 2, g = 2\%$</u>		
Limited commitment upper bound	4.22	0.90
MPE upper bound	3.52	1.22
MPE lower bound	2.68	1.77
2. <u>$\eta = 2, g = 1\%$</u>		
Limited commitment upper bound	12.53	0.43
MPE upper bound	10.60	0.61
MPE lower bound	5.93157	1.35916
7. <u>$\eta = 4, g = 1\%$</u>		
Limited commitment upper bound	2.70	1.28
MPE upper bound	2.38	1.58
MPE lower bound	2.02	2.03