

Time perspective and climate change policy*

Larry Karp[◇] Yacov Tsur[§]

March 10, 2011

Abstract

The tendency to foreshorten time units as we peer further into the future provides an explanation for hyperbolic discounting at an intergenerational time scale. We study implications of hyperbolic discounting for climate change policy, when the probability of a climate-induced catastrophe depends on the stock of greenhouse gasses. We characterize the set of Markov perfect equilibria (MPE) of the intergenerational game amongst a succession of policymakers. Each policymaker reflects her generation's preferences, including its hyperbolic discounting. For a binary action game, we compare the MPE set to a "restricted commitment" benchmark. We compare the associated "constant equivalent discount rates" and the willingness to pay to control climate change with assumptions and recommendations in the Stern Review on Climate Change.

"... My picture of the world is drawn in perspective... I apply my perspective not merely to space but also to time" –Ramsey..

Keywords: Hyperbolic discounting, Markov Perfect Equilibria, catastrophic climate change, uncertainty

*We benefitted from comments of an anonymous referee and Charles Mason and from seminar participants at UC Santa Barbara, UC Berkeley (ARE), University of Southampton, Hebrew University of Jerusalem, and Iowa State University.

[◇]Corresponding author: Department of Agricultural and Resource Economics, 207 Giannini Hall, University of California, Berkeley CA 94720. Phone: (925) 451-6744; Fax: (510) 643-8911; Email: karp@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). Financial support by the Center for Agricultural Economics Research is gratefully acknowledged.

1 Introduction

The long delay between the times when society incurs the cost and reaps the benefit of climate policy may make cost-benefit analysis sensitive to the discount rate [30, 6, 7, 22, 35]. Few people would defend the view that today's generation should ignore the welfare of all generations in the distant future, but standard discounting assumptions imply approximately that attitude. Our alternative to the standard analysis incorporates a plausible view of how people evaluate trade-offs across distant generations, and is also consistent with observable market discount rates. The analysis emphasizes the danger of catastrophic change. The model we develop is sufficiently tractable that we can analytically characterize equilibria. By choosing a few key parameter values, we can numerically assess whether it is reasonable to incur a particular level of expenditure to reduce a particular risk.

There are many political-economy processes (intra-generational games) that could explain how the social planner in a generation aggregates her generation's preferences. The precise intra-generational game is unimportant for our purposes, so we focus instead on inter-generational issues.

In each generation there is a representative agent whose "... picture of the world is drawn in perspective... [applied] not merely to space but also to time" [27, p.291]. This perspective gives rise to hyperbolic discounting at the level of the individual agent (Section 2). These agents care less about future generations' utility than about their own, so their pure rate of time preference (at the generational time scale) is positive and over some interval may be large. However, they make smaller distinctions between successive generations in the distant future, compared to successive generations in the near future, so their pure rate of time preference falls. Arrow [2] describes this attitude as "agent-relative ethics". Cropper et al. [5] and Section 8 of Heal [11] provide empirical evidence that individuals discount utility in this manner. Hyperbolic discounting leads to a model that is flexible enough to produce short and medium run social discount rates equal to market discount

rates, and which also gives non-negligible weight to the well-being of distant generations.

The optimal program for any generation is time-inconsistent under hyperbolic discounting [31, 25]. This time-inconsistency 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 focus on Markov Perfect equilibria (MPE), in which the current generation cannot commit to future actions.

Nordhaus [23] and Mastrandrea and Schneider [20] imbed hyperbolic discounting in integrated assessment models of climate change. 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. Karp [14, 9] studies MPE in climate change models with hyperbolic discounting and deterministic damages.

The current paper is the first to imbed hyperbolic discounting (with a MPE) in a model of catastrophic climate-related damages.¹ Our sequential game captures the risk of abrupt climate change [4, 33, 1, 30, 13] and the inertia in the climate system. That inertia leads to a delay between current actions and future reductions in risk.

There are multiple MPE, because the optimal policy today depends on beliefs about the policies that future regulators will choose. We obtain a closed form characterization for a binary action model in which the feasible actions are either to stabilize atmospheric greenhouse gas concentration or to follow business-as-usual (BAU). The MPE set to this game can be bounded in a simple manner. We compare it to a benchmark, “restricted commitment”, in which the policymaker’s feasible policies are restricted in order to cause the resulting optimal choice to be time consistent. This outcome is not plausible but it has an obvious welfare interpretation and therefore provides a useful comparison to the MPE set. A MPE may result in either too much or too little stabilization, relative to the benchmark.

For the binary action model we calculate a “constant equivalent” discount rate; this is a constant rate that would lead to policy prescriptions identical to a particular MPE in the sequential game. This constant-equivalent discount rate depends on the individual agents’ time-varying pure rate of time preference, 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 affect welfare over centuries, while a decision about a bridge affects welfare over decades. 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 pure rate of time preference.

After showing the relation between time perspective and hyperbolic discounting, we discuss damages associated with abrupt climate change. We then explain the relation between risk and climate policy in our model, and describe the payoff. The analytic results characterize the MPE and its relation to a benchmark equilibrium with restricted commitment. The numerical results provide a new perspective on policy recommendations in the Stern Review. Proofs and technical derivations are relegated to an online supporting material document, available in JEEM’s online supplementary materials archive at <http://aere.org/journals/>.

2 Time perspective and discounting

There are a number of ways to motivate hyperbolic discounting at the generational level. Heal [10, 11] proposes “logarithmic discounting”, based on the Weber-Fechner “Law”, a statement that human response to a change in a stimulus (e.g. vocal or visual) is inversely related to the magnitude of the pre-existing stimulus. Sumaila and Walters [32] use an overlapping generations model to justify hyperbolic discounting. They assume that each

generation discounts its own future consumption stream at a constant “intra-generational” rate. In each period a new generation arrives that discounts its own future consumption at the same rate. The generation born in period t discounts the stream of subsequent generations’ consumption at a different, intergenerational rate. The greater is $t' - t$, the more generations have been born between periods t and t' . This fact causes the discount rate used by generation t to evaluate future consumption to fall over time, leading to hyperbolic discounting.

Our explanation of hyperbolic discounting is based on time perspective - the tendency to foreshorten time periods as we peer further into the future. A function $s(n)$ captures time perspective by assigning a perceived length to a year that begins n years from now. This function satisfies $s(0) = 1$, $s'(\cdot) \leq 0$ and $s(\infty) \geq 0$; undistorted time corresponds to $s(\cdot) \equiv 1$. The relation between real time (t) and perceived (foreshortened) time is

$$S(t) = \int_0^t s(\zeta) d\zeta.$$

From the standpoint of today, the time period from now until t “looks like” a period from now until $S(t)$.

The constant pure rate ρ_0 represents impatience as applied to the perceived time S . From today’s perspective, the present value of a utility stream $U(c(S))$, $S \geq 0$, is

$$\int_0^\infty U(c(S)) e^{-\rho_0 S} dS.$$

Making a change of variables from S to t (i.e. from foreshortened time to real time), the payoff expressed in real time is

$$\int_0^\infty \exp\left(-\rho_0 \int_0^t s(\zeta) d\zeta\right) U(c(t)) s(t) dt.$$

The utility discount factor is therefore

$$\theta(t) = \exp\left(-\rho_0 \int_0^t s(\zeta) d\zeta\right) s(t)$$

and the corresponding pure rate of time preference is

$$\rho(t) \equiv -\frac{\dot{\theta}(t)}{\theta(t)} = \rho_0 s(t) - \frac{\dot{s}(t)}{s(t)}. \quad (1)$$

Equation (1) shows how the pure rate of time preference originates from impatience ρ_0 and from “time perspective” $s(\cdot)$. A constant pure rate of preference occurs in the following special cases: when $s(t) = 1$ identically for all t (undistorted time perspective), in which case $\rho(t) = \rho_0$; or when

$$s(t) = \frac{\alpha}{\rho_0 + (\alpha - \rho_0)e^{\alpha t}}, \quad \alpha > \rho_0,$$

in which case $\rho(t) = \alpha$.

In order to focus on the time-perspective motive of discounting we set $\rho_0 = 0$, so $s(t) = \theta(t)$. We choose a functional form for $s(t)$ to accommodate the situation where the pure rate changes little during the near future (e.g. the next 20 - 30 years) then decreases rapidly for a while and finally tapers off towards a limiting (vanishing or positive) rate. The following specification exhibits this pattern:

$$s(t) = \theta(t) = \beta e^{-\gamma t} + (1 - \beta) e^{-\delta t}, \quad \delta > \gamma. \quad (2)$$

The corresponding pure rate of discount is

$$\rho(t) \equiv \frac{-\dot{\theta}(t)}{\theta(t)} = \frac{\gamma\beta e^{-\gamma t} + \delta(1 - \beta)e^{-\delta t}}{\beta e^{-\gamma t} + (1 - \beta)e^{-\delta t}}, \quad (3)$$

which decreases from $\rho(0) = \gamma\beta + \delta(1 - \beta)$ to $\rho(\infty) = \gamma$ when $\beta \in (0, 1)$. An increase in β lowers the discount rate, i.e., increases the concern for the future. The constant rates $\rho = \delta$ or $\rho = \gamma$ correspond to the special cases where $\beta = 0$ or $\beta = 1$, respectively. This functional form is flexible and tractable; Ekeland and Lazrak [8] use this form to study an overlapping generations model.

Other functional forms for hyperbolic discounting are consistent with the time perspective explanation. For example, logarithmic discounting

is obtained by setting $s(t) = \frac{1}{1+kt}$, $k > 0$, with the resulting pure rate $\rho(t) = \frac{\rho_0+k}{1+kt}$. Barro [3] uses the discount factor $e^{-(\rho(t-\tau)+\phi(t-\tau))}$, with ρ a constant; for (our parameter) $\rho_0 = 0$, Barro's formulation corresponds to $s(t - \tau) = e^{-(\rho(t-\tau)+\phi(t-\tau))}$.

3 Catastrophic climate change

Recent evaluations suggest that global warming could result in catastrophic damages [13, 30]. The current atmospheric GHG concentration is estimated at 380 ppm CO₂ (430 ppm of CO₂e), compared with 280 ppm CO₂ at the onset of the Industrial Revolution. Under BAU, the concentration could double the pre-Industrial level by 2035 and treble this level by the end of the century. The recent IPCC report gives 2 – 4.5°C as a likely range for the increase in equilibrium global mean surface air temperature due to doubling of atmospheric GHG concentration with a non-negligible chance of exceeding this range [13, p. 749]. The Stern Review gives 2–5°C and 3–10°C as likely ranges for equilibrium global mean warming due to doubling and trebling of GHG concentration, respectively [30]. The probability of outcomes that significantly exceed the most likely estimates are not negligible: under doubling of GHG concentration, there is a 50% chance that the global mean warming will exceed 5°C (close to the warming since the last ice age) in the long term [30, Summary and Conclusion, p. vi]. Global warming can therefore give rise to catastrophic events, including the reversal of the thermohaline circulation, a sharp rise in sea level, the spread of lethal diseases and massive species extinction.

Each link in this chain, leading from changing GHG concentration to the ensuing damage, involves uncertain elements [28]. Following Clarke and Reed [4] and Tsur and Zemel [33, 34], we account for this uncertainty by assuming that a catastrophic climate event occurs at random time T with a distribution that depends on the GHG concentration, $Q(t)$. Denote the

distribution and density functions of T by $F(t)$ and $f(t)$, respectively. This distribution induces a hazard rate function $h(Q(t)) = -d[\ln(1 - F(t))]/dt$, the conditional probability that the catastrophe will occur during $[t, t + dt]$ given that it has not occurred by time t when atmospheric GHG concentration is $Q(t)$. When $h(Q)$ is an increasing function, there is one-to-one relation between the hazard and the atmospheric GHG concentration and we can use the hazard as the state variable.

A common modeling practice uses post-event scenarios that are easy to understand, e.g., a reduction in GDP or in the growth rate. These scenarios provide a basis for evaluating a policy that spends a certain amount today to decrease the expected damage. In our model, the event reduces income by a constant known share, Δ , from the occurrence date onward. Most climate change models assume a continuous relation between GHG stocks and damages. In our setting, which includes only abrupt changes, there is a continuous relation between GHG stocks and *expected* damages.

4 Risk and climate policy

The actions that society takes at a point in time (e.g. abatement, levels of consumption) determine greenhouse gas (GHG) emissions at that time. These flows, and existing GHG stocks, Q , determine the evolution of the stock, dQ/dt . The risk of a climate-related catastrophe, or hazard rate, h , is a strictly increasing function of the stock of GHG: $h = H(Q)$. The time derivative of the hazard rate is $H'(Q) dQ/dt$. The monotonicity of $H(\cdot)$ enables us to write the time derivative of the hazard rate as a function of the current hazard rate and society's current action, which we denote $w(t)$. We adopt the following functional form

$$\dot{h}(t + \tau) = \mu(a - h(t + \tau))(1 - w(t + \tau)), \quad h(t) \text{ given.} \quad (4)$$

We restrict $0 \leq w(t) \leq 1$; $w(\cdot) = 1$ corresponds to abatement that stabilizes the hazard (equivalently, the GHG stock) and $w(\cdot) = 0$ corresponds to no

abatement, i.e. BAU. We let X measure the cost of stabilization as a fraction of the income at risk, Δ . An abatement effort w costs $wX\Delta$.

In equation (4), a represents the maximal hazard rate that $h(\cdot)$ approaches under BAU (as τ increases) and μ measures the rate of convergence to a . The hazard grows most quickly when h is small. This feature means that each dollar spent on abatement effort leads to a larger reduction in expected damages when h is small. For hazards close to the steady state a , there is little benefit in incurring the abatement costs in order to prevent the hazard from growing.²

The model implies that the level of the hazard, not simply the occurrence of the catastrophe, is irreversible. This assumption reflects the considerable inertia in the climate system, and it simplifies the characterization of equilibria because it prevents non-monotonic hazard processes.

The simplicity of equation (4) 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 equation (4) (the initial value of h , and the constants a and μ). There is little empirical basis for calibrating a more complicated model.

5 The payoff

The payoff of the generation alive at time t , “generation t ”, is the expectation of the present discounted value of current and future generations’ utility, using the discount factor $\theta(t)$. Consumption grows at an exogenous constant rate g and the utility of consumption is iso-elastic, with the constant elasticity η .³ With initial (time 0) consumption normalized to 1, the flow of consumption from time t onward prior to the event occurrence is $e^{g(t+\tau)}(1 - \Delta Xw(t + \tau))$. After the occurrence date there is no role for

abatement, and consumption equals $e^{g(t+\tau)}(1-\Delta)$. The corresponding pre- and post-event utility flows are, respectively,

$$\frac{(e^{g(t+\tau)}(1-\Delta Xw(t+\tau)))^{1-\eta}-1}{1-\eta} \quad \text{and} \quad \frac{(e^{g(t+\tau)}(1-\Delta))^{1-\eta}-1}{1-\eta}.$$

Conditional on the event occurring T periods from now, i.e., at time $t+T$, the present (time t) value under policy $w(t+\tau)$ is

$$\int_0^T \theta(\tau) \frac{(e^{g(t+\tau)}(1-\Delta Xw(t+\tau)))^{1-\eta}-1}{1-\eta} d\tau + \int_T^\infty \theta(\tau) \frac{(e^{g(t+\tau)}(1-\Delta))^{1-\eta}-1}{1-\eta} d\tau =$$

$$\int_0^T \theta(\tau) e^{-g(\eta-1)(t+\tau)} \frac{(1-\Delta Xw(t+\tau))^{1-\eta} - (1-\Delta)^{1-\eta}}{1-\eta} d\tau + \varphi$$

where the constant φ is

$$\varphi = \int_0^\infty \theta(\tau) \frac{(e^{g(t+\tau)}(1-\Delta))^{1-\eta}-1}{1-\eta} d\tau. \quad (5)$$

Ignoring the constant φ , the present value at time t can be written as

$$e^{-g(\eta-1)t} \int_0^T \theta(\tau) e^{-g(\eta-1)\tau} U(w(t+\tau)) d\tau \quad (6)$$

where

$$U(w) \equiv \frac{(1-\Delta Xw)^{1-\eta} - (1-\Delta)^{1-\eta}}{1-\eta}. \quad (7)$$

We now introduce expectations. Let

$$y(t, \tau) = \int_t^{t+\tau} h(\zeta) d\zeta = \int_0^\tau h(t+\zeta) d\zeta. \quad (8)$$

Taking expectation of (6) conditional on $T > t$, recalling that $Pr\{T > t\} = 1 - Pr\{T \leq t\} = e^{-y(0,t)}$, gives the expected payoff at time t :

$$e^{-g(\eta-1)t} \int_0^\infty \theta(\tau) e^{-g(\eta-1)\tau - y(t,\tau)} U(w(t+\tau)) d\tau.$$

Multiplying by $e^{g(\eta-1)t}$ (to re-scale time- t BAU consumption to unity) gives the payoff to generation t , conditional on $h(t)$ and the sequence of current and future policies:

$$J(h(t), w(\cdot)) = \int_0^\infty \theta(\tau) e^{-g(\eta-1)\tau - y(t,\tau)} U(w(t+\tau)) d\tau. \quad (9)$$

In view of equations (2) and (9), we define the “effective discount factor”, a function that incorporates both the pure rate of time preference and the effect of η and g :

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

where

$$\tilde{\gamma} \equiv \gamma + g(\eta - 1) \quad \text{and} \quad \tilde{\delta} \equiv \delta + g(\eta - 1).$$

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

6 Equilibria

Different assumptions about commitment ability and about the set of feasible policies lead to different equilibrium sets. If the decisionmaker at time 0 can commit to an arbitrary function $w(t)$ (conditional on the event not having occurred before t), the solution is obtained by solving a standard non-stationary optimal control problem. This “full commitment” solution is time-inconsistent (unless it happens to involve the boundary solution $w(t) \equiv 0$ or $w(t) \equiv 1$, i.e. never begin stabilization, or begin full stabilization immediately). Since “full commitment” over a long period of time is implausible, we do not consider it further and focus instead on Markov Perfect Equilibria (MPE) to a sequential game. The agents in this game consist of a sequence of policymakers. We study the limiting game where each agent acts for an arbitrarily short period of time, leading to a continuous time model [16].

In a MPE, the current regulator cannot commit future generations to a specific course of action but she can *influence* successors’ actions by affecting the world they inherit, i.e. by changing the payoff-relevant state variable. The MPE recognizes the difference between influencing future policies and

choosing those policies. In a MPE agents condition their actions on (only) the payoff-relevant state variable, and they understand that their successors do likewise. Therefore, an agent's beliefs about future policies depend on her beliefs about the future trajectory of the state variable. An agent's action has an immediate effect on her current flow payoff and it also affects the continuation value via its influence on the state variable. We provide the necessary condition for a MPE for the general case and then analyze a binary action specialization. In order to provide a benchmark for the set of MPE in this binary case, we then consider an equilibrium involving "restricted commitment".

6.1 MPE in the general model

The state variable is the vector $z \equiv (h, y)$. A policy function maps the state z into the control w . The decision-maker at time t chooses the current policy $w(t)$ but not future policies. She understands how the current choice affects the evolution of the state variable and forms beliefs about how future regulators' decisions depend on the future level of the state variable. Each regulator chooses the current decision and wants to maximize the present discounted value of the stream of future payoffs, given by expression (9). A MPE policy function $\hat{\chi}(z)$ satisfies the Nash property: $w(t) = \hat{\chi}(z(t))$ is the optimal policy for the regulator at time t given the state $z(t)$ and given the belief that regulators at $\tau > t$ will choose their actions according to $w(\tau) = \hat{\chi}(z(\tau))$.

The state variable h is standard: at a future time $t + \tau, \tau > 0$, the value of $h(t + \tau)$ depends on the current hazard $h(t)$ and the intervening decisions $w(t + \xi), 0 \leq \xi \leq \tau$. The probability of survival until time $t + \tau$, conditional on $T > t$, is $Pr\{T > t + \tau | T > t\} = e^{-y(t, \tau)}$, which also depends on $h(t)$ and the intervening decisions. However, if the regulator at time t is in a position to make a decision, the event has not yet occurred: $y(t, 0) = 0$. Therefore, a *stationary* equilibrium depends only on the current

hazard, $h(t)$. Conditional on survival at time t , $h(t)$ is the only payoff-relevant state variable. We restrict attention to stationary pure strategies.

Let $q(h(t + \tau), w(t + \tau))$ denote the right-hand side of equation (4) and let h and w stand for $h(t)$ and $w(t)$, respectively. The following Lemma gives the necessary condition for a MPE (proofs can be found in an online supporting material document, accessed via JEEM's online supplementary materials archive at <http://aere.org/journals/>):

Lemma 1. *Consider the game in which the payoff at time t equals expression (9); the regulator at time t chooses $w(t) \in \Omega \subset R$, taking as given her successors' control rule $\hat{\chi}(z)$; and the state variables h and y obey equations (4) and (8). Let $V(h)$ equal the value of expression (9) in a MPE (the value function). A MPE control rule $\chi(h) \equiv \hat{\chi}(z)$ satisfies the (generalized) dynamic programming equation (DPE):*

$$K(h) + (\tilde{\gamma} + h)V(h) = \max_{w \in \Omega} \{U(w) + q(h, w)V'(h)\}, \quad (10)$$

with the "side condition"

$$K(h) \equiv (\delta - \gamma)(1 - \beta) \int_0^\infty e^{-(\tilde{\delta}\tau + y(t, \tau))} U(\chi(h(t + \tau))) d\tau. \quad (11)$$

Remark 1. *The control rule that maximizes the right-hand side of equation (10) depends on the payoff relevant state h , but not on y . This control rule also depends on the current regulator's beliefs about her successors' policies. Those policies affect the shadow value of the hazard, $V'(h)$.*

Remark 2. *The DPE is "generalized" in the sense that it collapses to the standard model with constant discounting in the two limiting cases $\beta = 1$ and $\beta = 0$. The former case is obvious from equation (11). To demonstrate the latter case, note that for $\beta = 0$,*

$$K(h) = (\delta - \gamma) \int_0^\infty e^{-(\tilde{\delta}\tau + y(t, \tau))} U(\chi(h(t + \tau))) d\tau = (\delta - \gamma)V(h).$$

Substituting this equation into (10) produces the DPE corresponding to the constant discount rate $\tilde{\delta}$.

6.2 A binary action specialization

We focus on the situation where $w(t)$ is limited to either full stabilization ($w = 1$) or BAU ($w = 0$). There are in general multiple MPE because the optimal decision for the current regulator depends on her 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 her 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, the equilibrium set has a simple characterization.

We now develop some notation needed for this characterization. Recall that Δ is the reduction in income due to the climate event, and $X\Delta$ is the fractional reduction in income due to complete stabilization ($w = 1$); X is a measure of the income cost of stabilization. It is convenient to describe the equilibrium set using the “utility cost of stabilization”, denoted x . To derive the relation between x and X , we use equation (7) to define

$$U(1) = \frac{(1 - \Delta X)^{1-\eta} - (1 - \Delta)^{1-\eta}}{1 - \eta} \quad \text{and} \quad U(0) = \frac{1 - (1 - \Delta)^{1-\eta}}{1 - \eta}. \quad (12)$$

Recall that $U(0)$ is (proportional to) the difference in the flow of utility under BAU before and after the climate event, so $U(0)$ is a measure of the utility at risk. The utility cost of stabilization, x , equals the fraction of utility at risk sacrificed to achieve full stabilization:

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

The relation between the income cost of stabilization, X , and the utility cost of stabilization, x , is

$$X = \frac{1}{\Delta} \left[1 - \left\{ 1 - x \left[1 - (1 - \Delta)^{1-\eta} \right] \right\}^{\frac{1}{1-\eta}} \right]. \quad (13)$$

The elasticity of marginal utility, η , affects the equilibrium in offsetting ways. First, an increase in η increases the “effective discount rate” $\rho(t) +$

$g(\eta - 1)$, which tends to reduce the amount that society is willing to spend to stabilize the hazard rate. Second, η is a measure of risk aversion. We hold the utility cost x constant and use equation (13) to find the effect of η on the income cost of stabilization. Substituting plausible values into equation (12) and using numerical methods, we find that $\frac{\partial X}{\partial \eta} < 0$ tends to hold. This inequality means that an increase in η decreases the monetary cost X of a given utility cost x , a fact that makes society willing to spend more to stabilize the hazard.

6.2.1 Markov Perfect Equilibria

The control space is $w(t) \in \{0, 1\}$, the flow payoffs are given in equation (12) and the hazard evolves according to equation (4). Let $\chi(h)$ be a MPE decision rule. Using the equilibrium condition (10) and the convention that in the event of a tie the regulator chooses stabilization, in the binary setting χ satisfies

$$\chi(h) = \begin{cases} 1 & \text{if } U(1) \geq U(0) + \mu(a - h)V'(h) \\ 0 & \text{if } U(1) < U(0) + \mu(a - h)V'(h) \end{cases}. \quad (14)$$

A particular control rule corresponds to a division of the state space $[0, a]$ into a “stabilization region” (where $\chi(h) = 1$) and a “BAU region” (where $\chi(h) = 0$).

The next lemma provides conditions under which perpetual stabilization or perpetual BAU are MPE. It makes use of the following functions:

$$\pi(h) \equiv \frac{1}{1 - \mu(a - h)\xi'(h)} \quad (15)$$

and

$$\sigma(h) \equiv 1 + \mu(a - h)\nu'(h), \quad (16)$$

where

$$\xi(h) \equiv \int_0^\infty e^{-ht} \tilde{\theta}(t) dt = \frac{(1 - \beta)(\gamma + g(\eta - 1)) + h + \beta(\delta + g(\eta - 1))}{(\delta + g(\eta - 1) + h)(h + \gamma + g(\eta - 1))} \quad (17)$$

and

$$\nu(h) \equiv \int_0^\infty \exp\left(\frac{-a\mu t + (a-h)(1-e^{-\mu t})}{\mu}\right) \tilde{\theta}(t) dt. \quad (18)$$

Lemma 2. (i) *Perpetual stabilization is a MPE if and only if*

$$\frac{U(1)}{U(0)} = 1 - x \geq \pi(h). \quad (19)$$

(ii) *Perpetual BAU is a MPE if and only if*

$$\frac{U(1)}{U(0)} = 1 - x < \sigma(h). \quad (20)$$

Proof. The proof contains notation and definitions that will be used below and is therefore presented here. We begin by introducing some notation. Superscripts B and S denote functions under perpetual BAU or stabilization, respectively. Under BAU, using equation (4), the probability of disaster by time t is

$$F^B(t) = 1 - \exp\left(\frac{-at\mu + (a-h_0)(1-e^{-\mu t})}{\mu}\right).$$

Substituting $F^B(t)$ into equation (9) gives the expected payoff under perpetual BAU:

$$V^B(h_0) \equiv U(0) \int_0^\infty (1 - F^B(t)) \tilde{\theta}(t) dt = U(0)\nu(h_0). \quad (21)$$

Under perpetual stabilization, the probability of disaster by time t is $1 - e^{-h_0 t}$ and the expected payoff is

$$V^S(h_0) \equiv U(1) \int_0^\infty e^{-h_0 t} \tilde{\theta}(t) dt = U(1)\xi(h_0). \quad (22)$$

Using equation (14) we see that in order for perpetual stabilization to be a MPE, the current regulator must want to stabilize when she believes that all future regulators will stabilize. Under this belief, $V(h) = V^S(h)$

and $V'(h) = V^{S'}(h) = U(1)\xi'(h)$. Thus, using the equilibrium rule (14), $U(1) \geq U(0) + \mu(a-h)U(1)\xi'(h)$ must hold for the initial value of h in order for stabilization to be a MPE, as stated in equation (19).

Similarly, for perpetual BAU to be a MPE, it must be the case that $U(1) < U(0) + \mu(a-h)V^{B'}(h) = U(0) + \mu(a-h)U(0)\nu'(h)$. Recalling equations (16)-(18), the condition under which perpetual BAU is a MPE can be written as in equation (20). ■

The properties of $\pi(h)$ and $\sigma(h)$ are summarized in

Lemma 3. *The functions $\pi(h)$ and $\sigma(h)$ are increasing over $(0, a)$ with $\pi(a) = \sigma(a) = 1$, and $\sigma(h)$ is concave.*

The following proposition provides a condition for existence of MPE and characterizes the class of MPE in which regulators never switch from one type of policy to another:

Proposition 1. *There exists a pure strategy stationary MPE for all $0 < x < 1$ and all initial conditions $h = h_0 \in (0, a)$ if and only if*

$$\pi(h) < \sigma(h), \quad h \in (0, a). \quad (23)$$

Under inequality (23), there exists a MPE with perpetual stabilization ($w \equiv 1$) if and only if at the initial hazard h the cost of stabilization satisfies

$$x < x^U(h) \equiv 1 - \pi(h); \quad (24)$$

there exists a MPE with perpetual BAU ($w \equiv 0$) if and only if at the initial hazard h the cost of stabilization satisfies

$$x > x^L(h) \equiv 1 - \sigma(h). \quad (25)$$

The fact that we have closed form expressions for the functions $\pi(h)$ and $\sigma(h)$ means that it is straightforward to determine when the conditions of Proposition 1 are satisfied. Figure 1 illustrates the proposition, showing

the graphs of $1 - \sigma(h)$ and $1 - \pi(h)$ with $\pi(h) < \sigma(h)$ for $h \in (0, a)$. The curves divide the rectangle $\{0 \leq h \leq a, 0 \leq x \leq 1\}$ into three regions. For points above the curve $1 - \sigma(h)$ there is a MPE trajectory with perpetual BAU, and for points beneath the curve $1 - \pi(h)$ there is a MPE trajectory with perpetual stabilization. For points between the curves, both perpetual stabilization and perpetual BAU are MPE.

Figure 1

Because the region between these two curves has positive measure (when inequality (23) is satisfied), the existence of multiple equilibria is generic in this model.⁴ The multiplicity of equilibria stems from the fact that the optimal action today depends on the shadow value $V'(h)$, which depends on future actions *that the current regulator does not choose*. If future regulators will stabilize, the shadow cost of the state ($-V'(h)$) is high, relative to the shadow cost when future regulators follow BAU. The current regulator has more incentive to stabilize if she believes that future regulators will also stabilize. Actions are “strategic complements”, a circumstance common to coordination games. Our problem resembles the dynamic coordination game familiar from the “history versus expectations” literature [21, 17]. In those coordination games, the optimal decision for (non-atomic) agents in the current period depends on actions that will be taken by agents in the future. The non-convexity in the payoffs in these problems typically leads to multiple rational expectations equilibria for a set of initial conditions of the state variable. These equilibria are in general not Pareto efficient. We show that inter-generational coordination problems in our game can lead to either too little or too much stabilization, relative to a benchmark under restricted commitment.

Proposition 1 characterizes only equilibrium trajectories in which the action never changes. It is clear that a switch from stabilization to BAU is impossible, since the hazard remains constant under stabilization and the

decisionmaker uses a pure strategy. However, the proposition does not rule out the possibility of a MPE with delayed stabilization, i.e. an equilibrium beginning with BAU and switching to stabilization once the hazard reaches a threshold. The next proposition shows that such equilibria exist.⁵ We use the following definition

$$\Theta(h) \equiv \frac{\mu(a-h) \left(\frac{\beta}{\tilde{\gamma}+h} + \frac{1-\beta}{\tilde{\delta}+h} \right)}{h + \beta\tilde{\gamma} + \tilde{\delta}(1-\beta) + \mu(a-h) \left(\frac{\beta}{\tilde{\gamma}+h} + \frac{1-\beta}{\tilde{\delta}+h} \right)}. \quad (26)$$

Proposition 2. *Suppose that Condition (23) is satisfied. (i) For $x > 1 - \pi(h)$ the unique (pure strategy) MPE is perpetual BAU. (ii) There are no equilibria with “delayed BAU”. (iii) A necessary and sufficient condition for the existence of equilibria with delayed stabilization is*

$$\Theta(h) < x < 1 - \pi(h). \quad (27)$$

(iv) For all parameters satisfying $0 \leq h \leq a$, $0 < \beta < 1$, $\delta \neq \gamma$, and $\mu > 0$, a MPE with delayed stabilization exists for some $x \in (0, 1)$.

Recall that x equals the utility cost of stabilizing the hazard (or the atmospheric GHG concentration) as a fraction of the value-at-risk $U(0)$. Relation (27) defines the lower and upper bounds of x for a delayed stabilization MPE to exist. We verify (see the Online Supplementary Material document) that

$$1 - \pi(h) - \Theta(h) = \frac{(\tilde{\delta} - \tilde{\gamma})^2(2h + \tilde{\gamma} + \tilde{\delta})}{(h + \tilde{\gamma})^2(h + \tilde{\delta})^2} \beta(1 - \beta). \quad (28)$$

Thus, these bounds form a non-empty interval when $0 < \beta < 1$ and $\gamma \neq \delta$, i.e., when the discount rate is non-constant.

6.2.2 Restricted commitment: a benchmark

We saw in the previous section that a class of MPE decision rules leads to either perpetual BAU or perpetual stabilization. Here we consider a

“restricted commitment” benchmark in which the decision-maker at time 0 behaves as if she could commit future generations to either perpetual stabilization or perpetual BAU. In contrast, “full commitment” permits switches between BAU and stabilization – or vice-versa. The restricted commitment outcome requires solving a standard optimization problem, leading to a unique solution.

Restricted commitment is not a plausible equilibrium concept, but it provides a useful benchmark for welfare comparisons.⁶ Suppose, for example, we find that for some initial value of h all MPE involve BAU, but the restricted commitment involves perpetual stabilization. In that case, there is an obvious sense in which there is “too little” stabilization in the MPE. Alternatively, if we find that there exist MPE involving perpetual stabilization, and the restricted commitment outcome involves perpetual BAU, then there is a sense in which there can be “too much” stabilization in a MPE. We show that both of these outcomes are possible.

Under restricted commitment there exists a critical function $x^C(h)$ such that initial decision-maker chooses perpetual stabilization if $x \leq x^C(h)$ and she chooses perpetual BAU if $x > x^C(h)$. To determine the function $x^C(h)$, we note that the regulator chooses to stabilize if and only if $V^S \geq V^B$, where V^B and V^S are, respectively, the values under BAU and stabilization defined in (21) and (22). This inequality is equivalent to $\frac{U(1)}{U(0)} \geq \lambda(h_0)$, where

$$\lambda(h) \equiv \frac{\nu(h)}{\xi(h)}.$$

Noting that $\frac{U(1)}{U(0)} = 1 - x$, the condition $V^S \geq V^B$ holds if and only if $x \leq x^C(h_0)$, where

$$x^C(h) \equiv 1 - \lambda(h). \tag{29}$$

A restricted commitment policy that involves stabilization is obviously time consistent, since under stabilization the hazard does not change. Under a restricted commitment policy of BAU, the hazard h increases. In our

model, a larger value of h decreases the value of stabilization, because the growth rate of the hazard falls as its steady state approaches. Therefore, with restricted commitment, if the regulator wants to follow BAU for a given initial value of h , all of her successors would make the same choice at the larger future values of h . Consequently, a restricted commitment policy that involves BAU is also likely to be time-consistent. We summarize this discussion in⁷

Proposition 3. *Given the initial hazard $h \in [0, a]$, the optimal restricted-commitment policy is to stabilize if and only if $x \leq x^C(h)$. This policy is time consistent for all $h \in [0, a]$ and $x \in [0, 1]$ if and only if $\frac{dx^c}{dh} \leq 0$. A sufficient condition for this inequality is $\mu \geq a + \delta + g(\eta - 1)$.*

The last part of the proposition provides a condition under which the policy is time consistent. When this condition is satisfied, a larger value of h decreases the range of x for which the policy-maker wants to stabilize. Here, stabilization is “more likely” at lower values of h , as noted above. In exploring numerical examples, we found no parameter values that violate the time-consistency condition $\frac{dx^c}{dh} \leq 0$, suggesting that time consistency is “typical” for this model. The optimal plan under full commitment is, in general, time inconsistent. By reducing the set of possible plans that a regulator can announce, we also reduce the temptation for subsequent regulators to deviate from the plan announced by the initial regulator.

6.2.3 Constant discounting

The specialization with constant discounting is useful for interpreting numerical results in the next section, and more generally for understanding the MPE when β is near one of its boundaries. Because our empirical application involves a small value of β , we consider the case where $\beta = 0$. Analysis of the case $\beta = 1$ requires only replacing $\tilde{\delta}$ with $\tilde{\gamma}$. With $\beta = 0$, the constant discount rate is $\tilde{\delta}$, so the distant future is “heavily discounted”. Following

the standard procedure to obtain the DPE, or invoking Remark 2, we have the following DPE:

$$\left(\tilde{\delta} + h\right) V(h) = \max_{w \in \{0,1\}} \{U(w) + \mu(a - h)(1 - w)V'(h)\}. \quad (30)$$

Let $\pi^0(h)$ and $\sigma^0(h)$ denote the functions $\pi(h)$ and $\sigma(h)$, defined in equations (15) and (16), evaluated at $\beta = 0$. The following proposition describes the optimal solution to the control problem with $\beta = 0$.

Proposition 4. *Under constant discounting (with $\beta = 0$), it is optimal to stabilize in perpetuity when $x \leq 1 - \sigma^0(h)$ and it is optimal to follow BAU in perpetuity when $x > 1 - \sigma^0(h)$. The function $\sigma^0(h)$ determines the boundary between the BAU and stabilization regions and $\pi^0(h)$ is irrelevant.*

The proposition has two implications. First, there can be MPE involving “excessive stabilization”. The functions $\pi(h)$ and $\sigma(h)$ are continuous in β , so $\pi^0(h)$ and $\sigma^0(h)$ are the limits of these functions as $\beta \rightarrow 0$. Consider a value of β that is positive but close to 0 and values of h and x that satisfy $1 - \pi(h) > x > 1 - \sigma(h)$. (Such values exist because $\pi(h)$ and $\sigma(h)$ are continuous in β , and there exists h, x that satisfy $1 - \pi^0(h) > x > 1 - \sigma^0(h)$, as shown in the proof of Proposition 4.) For this combination of parameters and state variable, there are two MPE, involving either perpetual stabilization or perpetual BAU (by Proposition 1), but the payoff under perpetual BAU is higher than under stabilization (by continuity and Proposition 4). That is, there are MPE that involve *excessive stabilization* relative to the benchmark under restricted commitment.

The second implication is that $1 - x^C(h) = \sigma(h)$ under constant discounting. This equality means that the optimal solution when the regulator is restricted to making a commitment (in perpetuity) at time 0, is equal to the solution when the regulator has the opportunity to switch between BAU and stabilization. For abrupt events, the regulator is tempted to delay stabilization (i.e. the “restriction” in restricted commitment binds) only under

hyperbolic discounting. The ability to switch between policies is of no value for abrupt events under constant discounting. The economic explanation for this result is simply that BAU is the optimal policy only if the hazard is sufficiently large; under BAU the hazard increases, whereas it remains constant under stabilization.

7 Policy bounds and constant equivalent rates

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 (7). We normalize by setting $\gamma = 0$.⁸ 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 $g(\eta - 1)$.

We discuss the calibration of the model and then present the three critical values of X that characterize the MPE and the restricted commitment equilibrium. We also present, for each critical X value, the constant equivalent (“observationally equivalent”) pure rate of time preference; each of these is the rate that would yield the same policy bound if ρ were constant. For each critical level of X we obtain an exact constant equivalent discount rate because each bound is a single number.⁹

7.1 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 $\mu = 0.00544875$. (The unit of time is one year.) 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 other models, 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 $h(0) > 0$, $\lim_{t \rightarrow \infty} D(t) = 0$, because both probabilities converge to 1.¹⁰ 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% of annual consumption 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).

Figure 2

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

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.

We set $\gamma = 0$, so that the long-run pure rate of time preference is 0, and use equation (3) to choose β and δ in order to satisfy

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

This parameterization implies that the pure rate of time preference begins at 3% and falls to 1% by 30 years, eventually declining to 0. Our value of $\rho(30)$ is ten times greater than the Stern Review’s constant pure rate of time preference. An ethical concern for generations in the distant future requires a small pure rate of time preference *only in the case of a constant pure rate of time preference*. A declining pure rate of time preference is consistent with both ethical considerations and a large pure rate of time preference in the near and medium term. This flexibility means that the model is compatible with both a reasonable ethical view and also with market discount rates.

7.2 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 pure rate of time preference (ρ). We discuss results for $g \in [1\%, 2\%]$ and $\eta \in [1.1, 4]$. Due to the iso-elastic utility functional form, the numerical results of the X bounds are sensitive to parameter changes in the neighborhood of $\eta = 1$ (cf. equation (13)).

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.

Table 1

We begin with the restricted 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 the hazard 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 $\Delta = 0.2$ and $\eta = 2$, 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 2

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 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 3

For $g = 2\%$ and $\eta = 2$, 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 $g = 2\%$ and $\eta = 2$, the constant equivalent ρ in a MPE ranges between 0.5% to 1% (the upper and lower bounds that correspond to the MPE set). As expected, higher growth rates make the current generation less willing to sacrifice for the sake of wealthier future generations, decreasing the X bounds.

The online appendix contains additional numerical analysis, showing the circumstances where stabilization in a MPE is either excessive or insufficient, relative to the limited commitment level. This analysis also shows that the possibility of delayed stabilization in a MPE is non-negligible.

8 Conclusion

Individuals may care less about the utility of future generations than about their own, but make smaller distinctions between the utility of successive distant generations, compared to the utility of the current and next generation. “Time perspective” is consistent with this kind of agent-relative ethics, and it leads to hyperbolic discounting across generations. In a sequential game, each of a succession of policymakers aggregates the preferences of her generation and chooses the policy for that generation. In a MPE to this sequential game, each policymaker takes as given her successors’ stationary decision rule, a function of the current economic fundamental (the GHG concentration).

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 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 minimum cost consistent with a MPE involving BAU. For each of these bounds we calculated a constant equivalent pure rate of time preference, i.e. a constant rate that leads, in the control problem, to the same decision rule as does the time-varying pure rate of time preference in the sequential game. We compared the set of MPE to a time-consistent 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 *intra-generational* but not the *intergenerational* collective action problem; the reference equilibrium indicates how much society should be willing to spend, if it solves both the intra- and the inter-generational problems.

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.

Across most dimensions, our model is vastly simpler than the integrated assessment models typically used for policy recommendations. However, catastrophic risk is central to our model, and we take seriously the fact that future policies are not chosen at the current time, but will instead be conditioned on future fundamentals. In addition, our model of the pure rate of time preference provides a reasonable description of ethics while also being consistent with observed market rates. Ethical concern does not require a

small pure rate of time preference in the near and medium run; it requires that the pure rate of time preference eventually become small. Our numerical results concerning the acceptable level of expenditure to reduce the threat of climate-related catastrophe bracket the recommendations in the Stern Review. The simplicity and parsimony of the model make it easy for other researchers to examine the sensitivity of those results.

Notes

¹Weitzman [36] builds a model from which he concludes that uncertainty about catastrophic changes can (i) make cost-benefit analysis of climate change impractical and (ii) render the effect of discounting a second-order issue. Nordhaus [24], Horowitz and Lange [12] and Pindyck [26] argue that Weitzman's model does not imply the first conclusion, and Karp [15] argues that in addition it does not imply the second conclusion.

²The 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 .

³This model does not contain capital, so it does not distinguish between income and consumption. The model is consistent with a neoclassical growth model in which capital and income grow at a constant rate, and the savings rate is constant. It is also consistent with a model in which all expenditures for climate control are deducted from consumption, so that climate policy does not affect aggregate savings or the trajectory of income.

⁴Laibson [19] shows that there are multiple equilibria to this kind of sequential game under non-Markov policies. Krussel and Smith [18] show the existence of a continuum of MPE when agents use step functions. Elements of this equilibrium set involves an infinite sequence of steps, and the step sizes are endogenous. Our setting contains a single, exogenously determined step size. Karp [14, 16] shows the existence of multiple *candidates* solving the necessary conditions for MPE, due to an indeterminacy in the steady

state conditions. Ekeland and Lazrak [8] show that these candidates are in fact equilibria. In our setting, the multiplicity arises because of a non-convexity in the game. Section 7 elaborates on this observation, showing the resemblance between the problem under constant discounting and the familiar “Skiba problem” in optimal control.

⁵From the proof of the proposition it is evident that for initial conditions such that delayed stabilization equilibria exist, there are a continuum of such equilibria, indexed by the threshold at which the decisionmaker begins to stabilize.

⁶Since we are interested in a situation that unfolds over many decades or centuries, it is not reasonable for the current regulator to act as if she can commit future generations to follow the plan that she announces. The problem with such a policy as an equilibrium concept (in our setting) is *not* that it requires commitments that subsequent generations would want to break. When policies are time consistent, future generations are happy to abide by the choice made by a previous generation, provided that they can make the same choice for their successors. Instead, commitment is an unsatisfactory equilibrium concept because it is based on an assumption that is patently false, namely that the current generation can commit future generations to a specific course of action.

⁷The proof of this proposition (see “Online Supporting Material”) shows that the shadow value of h is negative and decreasing (in absolute value) under either policy, and $1 - x^C(h) \leq 1$, with equality holding only when $h = a$. Since $U(1) < U(0)$, the regulator does not want to stabilize for h sufficiently close to the steady state value a .

⁸When $\eta = 1$ 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 (5) is finite if and only if $\eta > 1$. In contrast, the maximand in expression (9) 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.

⁹The exact equivalence occurs if the decision rules under both hyperbolic and constant discounting can be characterized by a single parameter. Barro

[3] 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 an exact constant equivalent discount rate for the hyperbolic model [14].

¹⁰Using equation (4), $P^B(t) = 1 - e^{-at+(a-h(0))(1-e^{-\mu t})/\mu}$ and $P^S(t) = 1 - e^{-h(0)t}$. For $h_0 = 0$, $D(t) = P^B(t) 100\Delta$, which converges to $100\Delta\%$.

References

- [1] R. B. Alley, J. Marotzke, W. D. Nordhaus, J. T. Overpeck, D. M. Peteet, R. S. Pielke Jr., R. T. Pierrehumbert, P. B. Rhines, T. F. Stocker, L. D. Talley, and J. M. Wallace. Abrupt climate change. *Science*, 299:2005–2010, 2003.
- [2] K. J. Arrow. Discounting, morality and gaming. In Paul R. Portney and John P. Weyant, editors, *Discounting and Intergenerational Equity*, pages 13–22. Resources for the Future, Washington, DC, 1999.
- [3] R.J. Barro. Ramsey meets Laibson in the neoclassical growth model. *Quarterly Journal of Economics*, 114:1125–52, 1999.
- [4] H. R. Clarke and W. J. Reed. Consumption/pollution tradeoffs in an environment vulnerable to pollution-related catastrophic collapse. *Journal of Economic Dynamics and Control*, 18(5):991–1010, 1994.
- [5] M. L. Cropper, S.K. Aydede, and P.R. Portney. Preferences for life saving programs: How the public discounts time and age. *Journal of Risk and Uncertainty*, 8:243–265, 1994.

- [6] P. Dasgupta. Commentary: The Stern review's economics of climate change. *National Institute Economic Review*, 199(1):4–7, January 2007.
- [7] P. Dasgupta. Discounting climate change. University of Cambridge, April 2007.
- [8] I. Ekeland and A. Lazrak. The golden rule when preferences are time inconsistent, 2010.
- [9] T. Fujii and L. Karp. Numerical analysis of non-constant pure rate of time preference: a model of climate policy. *Journal of Environmental Economics and Management*, forthcoming, 56:83–101, 2008.
- [10] G. M. Heal. *Valuing the Future: Economic Theory and Sustainability*. Columbia University Press, 1998.
- [11] G. M. Heal. Intertemporal welfare economics and the environment. In K.-G. Mäler and J. R. Vincent, editors, *Handbook of environmental economics*, volume 3, chapter 21, pages 1105–1145. Elsevier, 2005.
- [12] J. Horowitz and A. Lange. What's wrong with infinity - a note on Weitzman's dismal theorem. University of Maryland Working Paper, 2008.
- [13] Intergovernmental Panel on Climate Change. Climate change 2007. <http://www.ipcc.ch/SPM040507.pdf>, 2007.
- [14] L. Karp. Global warming and hyperbolic discounting. *Journal of Public Economics*, 89:261–282, 2005.
- [15] L. Karp. Sacrifice, discounting and climate policy: five questions, 2009. Giannini Foundation working paper.
- [16] L. Karp. Non-constant discounting in continuous time. *Journal of Economic Theory*, 132(1):557–568, January 2007.

- [17] P. Krugman. History versus expectations. *Quarterly Journal of Economics*, (106):651–67, 1991.
- [18] P. Krusell and Anthony Smith. Consumption-saving decisions with quasi-geometric discounting. *Econometrica*, 71(1):365–75, 2003.
- [19] D. Laibson. Self control and saving. Harvard University Working Paper, 1994.
- [20] M. D. Mastrandrea and S. H. Schneider. Integrated assessment of abrupt climatic changes. *Climate Policy*, 1:433–449, 2001.
- [21] K. Matsuyama. Increasing returns, industrialization, and indeterminacy of equilibrium. *Quarterly Journal of Economics*, 106:587–597, 1991.
- [22] W. D. Nordhaus. A review of the Stern Review on the economics of climate change. *Journal of Economic Literature*, 45:686–702, September 2007.
- [23] W. D. Nordhaus. Discounting and public policies that affect the distant future. In Paul R. Portney and John P. Weyant, editors, *Discounting and intergenerational equity*. Resources for the Future, Washington, DC, 1999.
- [24] W. D. Nordhaus. An analysis of the dismal theorem. Cowles Foundation Discussion Paper No 1686, 2009.
- [25] E. S. Phelps and R. Pollak. On second-best saving and game-equilibrium growth. *Review of Economic Studies*, 35(2):185–199, 1968.
- [26] R. S. Pindyck. Fat tails, thin tails and climate change policy, 2010. NBER WP 16353.
- [27] F. P. Ramsey. *Foundations of Mathematics and Other Logical Essays*. Routledge and Kegan Paul, 1931.

- [28] T. C. Schelling. Climate change: the uncertainties, the certainties, and what they imply about action. *Economists' Voice*, July:1–5, 2007.
- [29] A. Skiba. Optimal growth with a convex-concave production function. *Econometrica*, 46:527–39, 1978.
- [30] N. Stern. *The Economics of Climate Change*. Cambridge University Press, 2007.
- [31] R. H. Strotz. Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies*, 23(3):165–180, 1956.
- [32] U. Sumaila and C. Walters. Intergenerational discounting: a new intuitive approach. *Ecological Economics*, 52:135–142, 2005.
- [33] Y. Tsur and A. Zemel. Accounting for global warming risks: Resource management under event uncertainty. *Journal of Economic Dynamics & Control*, 20:1289–1305, 1996.
- [34] Y. Tsur and A. Zemel. Pollution control in an uncertain environment. *Journal of Economic Dynamics & Control*, 22:967–975, 1998.
- [35] M. L. Weitzman. A review of the Stern Review on the economics of climate change. *Journal of Economic Literature*, 45(3):703–724, 2007.
- [36] M. L. Weitzman. Reactions to the Nordhaus Critique. unpublished manuscript, 2009.




Figure 1: There is a MPE with perpetual stabilization for parameters below the graph of $1 - \pi$. There is a MPE with perpetual BAU for parameters above the graph of $1 - \sigma$. Both types of MPE exist for parameters between the graphs.




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

Table 1: Restricted commitment upper bounds X^C and constant-equivalent ρ values for $\eta \times g = \{1.1, 2, 4\} \times \{0.01, 0.02\}$ and $\Delta = 0.2$.

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

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

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