Climate Policy when the Distant Future Matters: Catastrophic events with hyperbolic discounting

Larry Karp[*] Yacov Tsur[◇]

November 10, 2006

Abstract

We study climate change policy with hyperbolic discounting under the risk of catastrophic events. Using a binary action model, we compare the set of Markov Perfect Equilibria (MPE) to the optimal policy under (time consistent) commitment. For some initial conditions there are multiple MPE that may involve either excessive or insufficient stabilization efforts relative to the policy under commitment. Numerical analysis shows that standard Integrated Assessment Models may have significantly understated the optimal level of near-term effort to reduce GHG emissions. Even if the free-rider problem amongst contemporaneous decision-makers were solved, the coordination problem amongst different generations complicates the design of optimal policy.

Keywords: abrupt climate change, event uncertainty, catastrophic risk, hyperbolic discounting, Markov Perfect Equilibria

---

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

Integrated Assessment Models (IAMs) – numerical models that combine climate and economic modules – typically endorse modest near-term efforts to reduce greenhouse gas (GHG) emissions (Nordhaus and Boyer 2000, IPCC 2001, Manne 2003, Edmonds and Sands 2003). These models have been influential in guiding policy recommendations because of their empirical foundation. (See, for example, the open letters from economists to President Bush in the appendix to Griffin (2003), and the chapter on the Kyoto Protocol in Barrett (2003).) Although few of these models formally incorporate abrupt catastrophic events, discounting causes these events to have little effect on optimal climate policy (Mastrandrea and Schneider 2001). Low probability events (those with low hazard rates) are not likely to occur until the distant future, and discounting at a non-negligible rate makes the distant future almost irrelevant to policy today. The conclusion that low-probability catastrophic events are relatively unimportant for climate change policy probably sounds strange to the representative citizen (or climate scientist).

We study catastrophic threats with hyperbolic discounting to provide another view of the importance of low probability climate events. Our main policy conclusion is that standard IAMs may significantly understate the optimal reduction in GHG emissions. In reaching this conclusion, we show how to solve a new type of optimal control/dynamic game model, and we explain its relation to more familiar dynamic coordination problems.

A low probability event occurs infrequently, but its probability of occurrence over a few centuries is non-negligible. For example, the thermohaline circuit (THC) (the ocean currents that contribute to the moderate climate in Western Europe) has shut down or reversed more than once during the geological record (Broecker 1997). Suppose that such an "event" reduces the annual flow of utility (or income) by 1 util (the value-at-risk). Let the constant hazard rate be h, so that the probability of the event occurring during the next year, conditional on it not having yet occurred, is $1-\exp(-h)$.

When the future is discounted at a constant annual rate r, the expected present value of the value-at-risk is $\frac{1}{h+r}$. Suppose that "stabilization" eliminates the hazard, at a *flow cost* (perhaps caused by a perpetual reduction of economic activity) of $x(100\%)$, leading to a present value of $\frac{1-x}{r}$. The maximal flow cost, x^*, that society would pay for stabilization equates these two expressions, leading to $x^* = \frac{h}{h+r}$. If the annual discount rate is 5% ($r = .05$) and the probability of the event occurring within a century is 5% (equivalent to the annual hazard $h = 5.129 \times 10^{-4}$), then $x^* = 0.01015$. In this case, society would be willing to sacrifice about 1% of the value-at-risk to eliminate the threat.[1]

Under constant discounting, society will not devote substantial resources to reduce a risk unless the hazard rate and discount rate are of a similar order of magnitude. For low probability events, this similarity requires a discount rate near zero. However, a negligible constant discount rate requires excessive saving by current generations, and it is inconsistent with empirical estimates of time preferences (Frederick et al. 2002).

A positive social discount rate requires a positive pure rate of time preference and/or the anticipation that the future growth rate of income will be positive. Although both of these conditions plausibly hold for the near term, they may not hold for the distant future. We may feel closer to our children than to our unborn grandchildren, but it is less likely that we care much more about the 10^{th} future generation than about the 11^{th}. There is also less basis for assuming that income will be higher in the 11^{th} than in the 10^{th} future generation. The longest-term bonds mature within 30 years and there are no financial instruments that reflect (very) long-term discount rates.

There is a large literature dealing with the present valuation of future welfare, motivated by considerations of inter-generational equity and sustain-

---

[1] In the case of climate change, where inertia is important, current actions could alter future but not current risk. By assuming that the policy has an immediate effect on the hazard, this example overstates the amount that society would be willing to spend.

ability (Solow 1974, Hartwick 1977, Chichilnisky 1996, Arrow 1999, Asheim and Buchholz 2004) and uncertainty (Weitzman 2001, Gollier 2002, Dasgupta and Maskin 2005). An important strand of this literature uses hyperbolic discounting, where the discount rate falls over time. This model of discounting is consistent with the theoretical and empirical arguments for a non-negligible near-term discount rate, and the ethical arguments for giving positive weight to the distant future. Cropper and Laibson (1999) discuss hyperbolic discounting in the context of climate change, and Mastrandrea and Schneider (2001) use hyperbolic discounting with a variation of the DICE model (a prominent IAM), assuming that the current regulator is able to make commitments centuries into the future. Karp (2005) uses quasi-hyperbolic discounting under the assumption that regulators use Markov Perfect policies.

We focus on the type of low probability catastrophic climate events that climate scientists have identified as among the most worrisome consequences of climate change. In addition to the THC weakening, these events include a sudden rise in sea level and mass extinction of species (Chichilnisky and Heal 1993, IPCC 2001, Alley et al. 2003, Thomas et al. 2004, United Nations 2005). There is currently little basis for assessing the numerical magnitude of the various risks, so quantitative modeling is still speculative.[2] We know of only two model-based attempts to obtain numerical values of the risk of the collapse or significant weakening of the THC within the next century (Challenor et al. 2006, Schlesinger et al. 2006). Both studies estimate the risk at more than 30%, under projected business-as-usual (BAU) emissions.[3] However, some climate scientists think that this risk has been wildly exagger-

---

[2] In contrast, there has been substantial work in assessing the likely economic costs of gradual climate change (Chakravorty et al. 1997, Mendelsohn 2003, Schlenker et al. 2006). These kinds of studies form the empirical basis for IAMs.

[3] In an informal survey (Challenor 2005), three of five (self-identified) climate experts estimated that the probability of a significant THC weakening before 2100 was less than 5%, and two estimated that the probability was between 20% and 30%. Non-experts thought that the risk was considerably larger.

ated (Wunsch 2006a,b). At this stage, we know little about the magnitude of these risks, limiting the usefulness of complicated numerical models. Our use of a simple model of climate change makes it possible to incorporate more realistic behavioral assumptions.

The present work is also related to the literature on dynamic management under event uncertainty (Cropper 1976, Clarke and Reed 1994, Tsur and Zemel 1996, 1998). We extend these models by replacing constant discounting with hyperbolic discounting, using methods from Karp (2006).[4]

The next section describes our general model and a binary action specialization, in which the regulator chooses either to stabilize or to follow Business as Usual (BAU). Section 3 considers the situation in which the current regulator can make a commitment to all future policies. This assumption is not plausible, because we are interested in events that unfold over many decades and possibly centuries. In our setting, the solution with commitment is time consistent. This model is useful as a benchmark for comparison with the subgame perfect equilibrium without commitment.

Section 4, which contains our theoretical contribution, studies the Markov Perfect Equilibrium (MPE), where regulators cannot commit to future actions. We provide the necessary conditions for a MPE in a general setting, and then obtain a closed form characterization of the equilibrium set in the binary action setting. For some initial conditions of the state variable there are two MPE, involving either perpetual stabilization or perpetual BAU. The model under hyperbolic discounting, like some dynamic coordination games, has multiple subgame perfect equilibria because the optimal policy today depends on beliefs about the policies that will be chosen by future regulators. A MPE might result in either too much or too little stabilization, relative to our benchmark. Section 5 studies the optimization problem under constant discounting, in order to highlight the effect of hyperbolic discounting. A nu-

---

[4] Harris and Laibson (2004) study a model in which a random event (a jump process) leads to the replacement of the current regulator by her successor, thus providing a different motivation for hyperbolic discounting.

merical application in Section 6 shows the importance of risk, commitment, and discounting.

The temptation to free ride on other nations' abatement effort makes it difficult to achieve agreement about climate policy amongst current decision-makers. Our results show that even if this free-riding problem could be solved, an inter-generational (i.e., dynamic) coordination problem sometimes makes it difficult to follow an optimal policy.

2 The model

We use a stationary model, in which the business-as-usual (BAU) flow of per capita consumption before the catastrophic event occurs is a constant $c+\Delta$. After the event occurs, the constant flow of consumption is c, so $\Delta > 0$ is the income-at-risk, expressed as a perpetual loss in the flow of consumption. Society can take an action $w(t)$ to reduce the probability of the event, e.g. by reducing greenhouse gas emissions. This action requires abatement expenditures and therefore reduces instantaneous utility; the control $w = 0$ corresponds to BAU ("no action"). Once the disaster has occurred it is too late to act, so $w = 0$ is optimal during the post-event period. The flow of utility prior to the catastrophe is $u(c+\Delta, w(t))$ and utility after the catastrophe is $u(c, 0)$.

The discount factor is

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

with $\delta > \gamma$, implying the discount rate

$$r(t) \equiv \frac{-\dot{\theta}(t)}{\theta(t)} = \frac{\beta\gamma e^{-\gamma t} + \delta e^{-\delta t}(1-\beta)}{\beta e^{-\gamma t} + e^{-\delta t}(1-\beta)}. \tag{2}$$

Equation (2) implies that $\frac{dr(t)}{d\beta} < 0$: an increase in β lowers the discount rate, i.e. increases the concern for the future. For $\beta = 0$ the constant discount rate is δ and with $\beta = 1$ the constant discount rate is γ.

Let T represent the random time when the event occurs, with the probability distribution and density functions $F(t)$ and $f(t)$, respectively. The hazard rate is defined as $h(t) = f(t)/(1-F(t)) = -d[\ln(1-F(t))]/dt$, yielding

$$F(t) = 1 - e^{-y(t)} \text{ and } f(t) = h(t)e^{-y(t)}, \tag{3}$$

where

$$y(t) = \int_0^t h(\tau)d\tau. \tag{4}$$

Conditional on the disaster not having yet occurred, $y(0) = 0$. Thus,

$$\dot{y}(t) = h(t), \quad y(0) = 0. \tag{5}$$

The present value associated with catastrophe at T and a policy $w(t)$ is

$$\begin{gathered} \textstyle\int_0^T \theta(t)u(c+\Delta, w(t))dt + \int_T^\infty \theta(t)u(c,0))dt = \\ \textstyle\int_0^T \theta(t)\left(u(c+\Delta, w(t)) - u(c,0)\right)dt + \int_0^\infty \theta(t)u(c,0)dt = \\ \textstyle\int_0^T \theta(t)U\left((w(t), c, \Delta\right)dt + \text{ constant,} \end{gathered} \tag{6}$$

where $U(w(t), c, \Delta) \equiv u(c+\Delta, w(t)) - u(c,0)$. Hereafter we refer to $U(0, c, \Delta)$ as the "value-at-risk", the change in the flow of utility if society takes no action to stabilize the risk ($w = 0$) and the event happens. Due to our stationarity assumption, the parameters Δ and c are constants and will be suppressed when convenient. We also ignore the constant term in the payoff, $\int_0^\infty \theta(t)u(c,0)dt$, since this term is independent of the control $w(t)$ and the time of catastrophe T. At time $t = 0$ the expected present value of the future flow of utility is

$$\begin{gathered} E_T\left\{\textstyle\int_0^T \theta(t)U(w(t))dt\right\} = \\ \textstyle\int_0^\infty (1-F(t))\,\theta(t)U(w(t))dt = \int_0^\infty e^{-y(t)}\theta(t)U(w(t))dt. \end{gathered} \tag{7}$$

The hazard rate is an increasing function of the stock of greenhouse gasses (GHGs). Current consumption and abatement decisions (the control, w)

affect the evolution of this stock, thus affecting the evolution of the hazard. In view of the (assumed) one-to-one relation between the hazard rate and the stock of GHGs, there is no loss of generality in treating the hazard rate rather than the pollution stock as the state variable. The control variable (prior to the catastrophe) at time t, $w(t)$ affects the flow of utility at time time t, $U(w(t))$, and also the change in the hazard rate, according to

$$\dot{h} = g(h, w),\ \ h(0) = h_0 \text{ (given).} \tag{8}$$

We suppress the time argument when there is no confusion. The optimal policy $w(t)$ maximizes (7) subject to (5) and (8). We define a low probability event as one for which $h_0 < r(0)$.

2.1 A binary action specialization

Given the current level of uncertainty about risks (and the corresponding difficulty of calibrating *any* model), we think that simplicity trumps complexity. We therefore study a particular *binary action model*: society can either stabilize the stock of greenhouse gasses at the current level, thus stabilizing the hazard at the current level, or follow BAU and allow the hazard to increase (with the stock of GHGs). The stabilization policy (corresponding to $w(t) = 1$) costs society the fraction X of the income-at-risk Δ, so the flow of consumption under stabilization (before the catastrophe occurs) is $c + \Delta(1 - X)$. Under BAU (corresponding to $w(t) = 0$), the flow of consumption is $c + \Delta$. The flow payoffs in this binary-action model are

$$U(1) = u(c + \Delta(1 - X)) - u(c) \ \text{ and } \ U(0) = u(c + \Delta) - u(c). \tag{9}$$

We define $x \equiv 1 - \frac{U(1)}{U(0)}$, the fractional reduction in the value-at-risk under stabilization (prior to the catastrophe); x summarizes all of the pertinent information about the utility functions and the parameters Δ and X. In the case where $u(\cdot)$ is linear, $X = x$.

We assume that under BAU the hazard rate approaches the steady state level a at a constant rate ρ:

$$\dot{h} = \rho(a - h). \tag{10}$$

If the current hazard rate (at time 0) is h_0 and society follows BAU until time t, the hazard rate at time t is

$$h\left(t\right) = a - \left(a - h_0\right) e^{-\rho t}. \tag{11}$$

We also assume that the hazard (not just the event) is irreversible. If society follows BAU until time t and then switches (forever) to stabilization, the hazard rate remains constant at the level $h(t)$. Provided that $h_0 < a$ (as we assume), the hazard never falls under BAU. The three parameters of the hazard function, h_0, a and ρ provide measures of the current risk, the eventual risk under BAU, and the speed of adjustment of the risk.

For all of the equilibria that we study, a larger value of h makes it "less likely" that the decision-maker chooses to stabilize. As the hazard approaches the steady state level a, its growth rate approaches 0.[5] There is little benefit from stabilization when the hazard is close to its steady state, so stabilization does not occur unless the cost is low (x is small). Obviously there is no benefit from incurring stabilization costs if $h = a$, since in this case the hazard rate does not increase under BAU.

3 Restricted commitment: a benchmark

This section analyzes the binary action model under restricted commitment. Here the current ($t = 0$) regulator decides whether to adopt stabilization or BAU in perpetuity. This policy menu is "restricted", because the current regulator commits to one of two policies in perpetuity. "Unrestricted" commitment, in contrast, allows the current regulator to announce a trajectory in

---

[5] Different growth functions for the hazard are discussed in a separate note, available upon request.

which the policy switches at a specified time in the future (conditional on the event not having occurred). For example, under unrestricted commitment the current regulator is able to delay stabilization until a positive but finite future time. Whenever the optimal time to switch from BAU to stabilization is greater than 0 and less than infinity, non-exponential discounting causes the policy announced at time 0 to be time-inconsistent. A regulator in the future would want to deviate from the policy announced by the regulator at time 0, by delaying the switching time (i.e., by procrastinating). Hereafter we consider only restricted commitment.

Under BAU, noting (11), the probability of disaster by time t is

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

Substituting $F^{BAU}(t)$ into equation (7) gives the expected payoff under BAU in perpetuity:

$$V^B(h_0) \equiv U(0)\int_0^\infty \left(1 - F^{BAU}(t)\right)\theta(t)dt = U(0)\nu(h_0), \tag{12}$$

where

$$\nu\left(h\right) \equiv \int_0^\infty e^{\left(\frac{-a\rho t + (a-h)\left(1 - e^{-\rho t}\right)}{\rho}\right)}\left(\beta e^{-\gamma t} + \left(1 - \beta\right)e^{-\delta t}\right)\, dt. \tag{13}$$

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

$$V^S(h_0) \equiv U(1)\int_0^\infty e^{-h_0 t}\theta(t)dt = U\left(1\right)\xi\left(h_0\right) \tag{14}$$

where

$$\xi\left(h\right) \equiv \int_0^\infty e^{-ht}\theta(t)dt = \frac{\left(1 - \beta\right)\gamma + h + \beta\delta}{\left(\delta + h\right)\left(h + \gamma\right)} \tag{15}$$

The regulator chooses to stabilize if and only if $V^S \geq V^B$. (We assume that a tie results in stabilization.) Noting that $\frac{U(1)}{U(0)} = 1 - x$, $V^S \geq V^B$ holds if and only if $x \leq \bar{x}^c(h_0)$, where is the solution to

$$\bar{x}^c(h) \equiv 1 - \lambda\left(h\right) \tag{16}$$

with

$$\lambda(h) \equiv \frac{\nu(h)}{\xi(h)}. \tag{17}$$

The superscript on $\bar{x}^c$ is a mnemonic for "commitment", and the over-bar indicates that this variable is an upper bound. Note that λ is a function of the hazard rate and of the parameters of the growth equation, but is independent of the utility function.

The following Proposition describes the optimal policy under restricted commitment. (All proofs are in the appendix.)

Proposition 1. *(i) The functions $\nu(h)$ and $\xi(h)$ are positive, decreasing and convex for $h \geq 0$. (ii) $\nu(a) = \xi(a)$ and $\nu(h) < \xi(h)$ for $0 \leq h < a$ and $\rho > 0$. (iii) The optimal policy under restricted commitment is to stabilize if and only if $x \leq \bar{x}^c(h)$ (iv) The optimal policy under restricted commitment is time consistent for all initial hazard values $0 \leq h \leq a$ and $0 < x < 1$ if and only if $\lambda'(h) \geq 0$. (v) $\rho \geq a + \delta$ is sufficient for $\lambda'(h) \geq 0$.*

Part (i) implies that the shadow value of h is negative and decreasing (in absolute value) under either policy. Part (ii) implies that $\lambda(h) \leq 1$ and $\lambda(a) = 1$. Since $U(1) < U(0)$, the regulator does not want to stabilize for h sufficiently close to the steady state value, a. Part (iii) is simply a restatement of the earlier derivation, and part (iv) 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 in Section 2.1.

In exploring numerical examples, we found no parameter values that violate the necessary and sufficient condition $\lambda'(h) \geq 0$, suggesting that time consistency is "typical" for this model. As noted above, the optimal plan under unrestricted 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.

Since 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 restricted commitment 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, restricted 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.

4 Markov Perfect Equilibrium

This section studies the Markov Perfect Equilibria (MPE), where a regulator at a point in time is not able to commit to future actions. The current regulator chooses the optimal current action, recognizing that future actions depend on the payoff relevant state variable.

We explained above why restricted commitment, despite its time-consistency, is an unsatisfactory equilibrium concept. In the MPE the current regulator cannot commit future generations to a specific course of action, but the current regulator is able to *influence* successors' actions by affecting the world that successors 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 use results from Karp (2006), who obtains the necessary conditions for a MPE under hyperbolic discounting by taking the limit of a discrete time sequential game amongst a succession of regulators. Regulators are indexed by the time at which they move, $t = i\varepsilon$, $i = 0, 1, 2, 3...$. Regulator i's decision lasts for ε units of time. Regulator i chooses the current policy and understands that future policies depend on the payoff-relevant state variable. She takes the current hazard rate as given and recognizes that her action (possibly) affects the evolution of the hazard, thereby affecting the actions of her successors. Each regulator in the infinite sequence of regulators cares about current and future welfare (discounted using the hyperbolic discount function $\theta(t)$) but not about her predecessors' welfare: bygones are bygones. Regulators use stationary, Markov policies. Beginning with the equilibrium condition to this discrete stage problem and formally letting $\varepsilon \to 0$, produces the equilibrium condition for the continuous time game, in which each regulator is active for an infinitesimal length of time. We begin with the general model of Section 2 and then specialize to the binary action model of Section 2.1.

4.1 The general model

The state variable is the vector $z \equiv (h, y)$. A policy function maps the state z into the control w. An equilibrium 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 $t > 0$, the value of $h(t)$ depends on $h(0)$ and intervening decisions $w(\tau)$, $0 \leq \tau < t$. The probability of survival until time t also depends on $h(0)$ and intervening decisions. However, *conditional on survival at time* t, $y(t) = 0$. If the regulator at time t is in a position to make a decision, the event has not yet occurred. This fact means that there is a *stationary* equilibrium that depends only on the current hazard,

$h(t)$. Conditional on survival at time t, $h(t)$ is the only payoff-relevant state variable.

Throughout this paper we restrict attention to stationary pure strategies. The following Proposition gives the necessary condition for a MPE:

Proposition 2. *Consider the game in which the payoff at time t equals expression (7); 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 y and h obey equations (5) and (8). Let $V(h)$ equal the value of expression (7) in a MPE (the value function). A MPE control rule $\hat{\chi}(z) \equiv \chi(h)$satisfies the (generalized) dynamic programming equation (DPE):*

$$K(h) + (\gamma + h)\, V(h) = \max_{w \in \Omega} \left\{ U(w) + g(h,w)V'(h) \right\}, \tag{18}$$

with the "side condition" (a definition)

$$K(h) \equiv (\delta - \gamma)(1 - \beta) \int_0^\infty e^{-(\delta s + y(s))} U(\chi(s))\, ds. \tag{19}$$

Remark 1. *The control rule that maximizes the right-hand side of equation (18) 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 hazard shadow value $V'(h)$.*

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 (19). To demonstrate the latter case, note that for $\beta = 0$, $K(h) = (\delta - \gamma) \int_0^\infty e^{-(\delta s + y(s))} U(\chi(s)) ds = (\delta - \gamma) V(h)$. Substituting this equation into (18) produces the DPE corresponding to the constant discount rate δ.*

4.2 The MPE for the binary action model

We now specialize to the binary model, where the control space is $\Omega = \{0, 1\}$. The payoff $U(w)$ is given by equation (9), and the equation of motion for the hazard is $\dot{h} = \rho\,(a-h)\,(1-w)$. Under stabilization ($w=1$) the flow of consumption is $c+\Delta(1-X)$ and the hazard remains constant; under BAU ($w=0$) the flow of consumption is $c+\Delta$ and the hazard changes according to equation (10). Let $\chi\,(h)$ be an equilibrium MPE decision rule. Using the equilibrium condition (18) and the convention that in the event of a tie the regulator chooses stabilization, in the binary setting χ satisfies

$$\chi\,(h) = \left\{ \begin{array}{l} 1 \text{ if } U\,(1) \geq U\,(0) + \rho\,(a-h)\,V'(h) \\ 0 \text{ if } U\,(1) < U\,(0) + \rho\,(a-h)\,V'(h) \end{array} \right\}. \tag{20}$$

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 $\chi\,(h)=0$).

For perpetual stabilization to be a MPE it must be in the interest of the current regulator to stabilize when she believes that all future regulators will stabilize, in which case $V(h)=V^S(h)$ and $V'(h)=V^{S\prime}(h)=U(1)\xi'(h)$, where $V^S(h)$ and $\xi(h)$ are defined in equations (14) and (15), respectively, and

$$\xi'(h) = -\int_0^\infty te^{-ht}\theta(t)dt. \tag{21}$$

Thus, observing the equilibrium rule (20), $U(1) \geq U(0)+\rho(a-h)U(1)\xi'(h)$ must hold for stabilization to be a MPE. Defining

$$\pi(h) \equiv (1-\rho(a-h)\xi'(h))^{-1}, \tag{22}$$

this condition can be stated as $\frac{U(1)}{U(0)} \geq \pi(h)$.

Using the same line of argument, for perpetual BAU to be a MPE, it must be that $U(1) < U(0)+\rho(a-h)V^{B\prime}(h) = U(0)+\rho(a-h)U(0)\nu'(h)$. Defining

$$\sigma(h) \equiv 1+\rho(a-h)\nu'(h) \tag{23}$$

with $\nu(h)$ defined in equation (13) and

$$\nu'(h) = -\int_0^\infty \frac{1-e^{-\rho t}}{\rho} \exp\left(-at+(a-h)\frac{1-e^{-\rho t}}{\rho}\right)\theta(t)dt, \tag{24}$$

this condition can be rendered as $\frac{U(1)}{U(0)} < \sigma(h)$.

The following lemma summarizes properties of $\pi(h)$ and $\sigma(h)$.

Lemma 1. *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 characterizes the MPE:

Proposition 3. *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*

$$\pi(h) < \sigma(h), \ h \in (0,a). \tag{25}$$

Under (25), there exists a MPE with perpetual stabilization ($w \equiv 1$) if and only if the initial condition $h_0 = h$ satisfies

$$x < \bar{x}^S(h) \equiv 1 - \pi(h); \tag{26}$$

there exists a MPE with perpetual BAU ($w \equiv 0$) if and only if the initial condition $h_0 = h$ satisfies

$$x > \underline{x}^B(h) \equiv 1 - \sigma. \tag{27}$$

Figure 1 illustrates Proposition 3. The figure shows $1-\sigma(h)$ and $1-\pi(h)$ with $\pi(h) < \sigma(h)$ for $h \in (0,a)$, which is necessary and sufficient for the existence of (pure strategy stationary) MPE for all initial hazard rates 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 equilibria.

Because the region between these two curves has positive measure, the existence of multiple equilibria is generic in this model. This situation provides a simple example of the existence of multiple MPE under hyperbolic

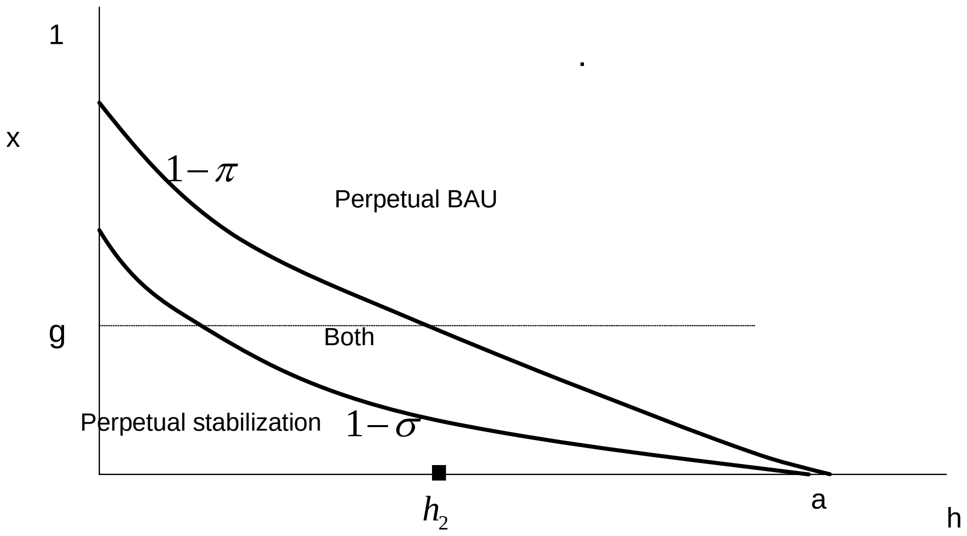

Figure 1: Graph of $1-\pi(h)$ and $1-\sigma(h)$, the thresholds in a MPE

discounting – a possibility previously noted by Krusell and Smith (2003) and Karp (2005, 2006). 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".

When the current regulator cannot commit to future policies, and each regulator in the infinite sequence of regulators follows Markov Perfect policies and has hyperbolic discounting, the equilibrium problem resembles the dynamic coordination game familiar from the "history versus expectations" literature (Matsuyama 1991, Krugman 1991). 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; there are typically multiple rational expectations equilibria for a set of initial conditions of the state variable. In addition, the competitive equilibrium trajectory is not necessarily efficient (relative to the social planner benchmark).

The temptation to free ride on other countries' abatement efforts makes it difficult to induce nations to participate in a climate change treaty such as the Kyoto Protocol. The next two sections show that inter-generational coordination problems can lead to either too little or too much stabilization, relative to the level under restricted commitment.

Proposition 3 considers only equilibrium trajectories in which the action never changes. The following proposition establishes a necessary and sufficient condition for an equilibrium with "delayed stabilization" (i.e., an equilibrium beginning with BAU and switching to stabilization once the hazard reaches a threshold). The proposition uses the following definition

$$h_2(x) = \left\{ \begin{array}{c} \text{solution to } x = 1 - \pi(h) \text{ for } x < 1 - \pi(0) \\ 0 \text{ for } x \geq 1 - \pi(0) \end{array} \right\}. \tag{28}$$

Figure 3 shows the value of h_2 for $x = g$ (where g is a constant less than $1 - \pi(0)$). We state the proposition using a (complicated) function $\Theta(h)$ which is defined in the proof. Here we provide the first order approximation of that function, since we are interested in circumstances where h is small.

Proposition 4. *Suppose that Condition (25) is satisfied. (i) For $h > h_2(x)$ 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 an equilibrium with delayed stabilization is that $h < h_2(x)$ and*

$$x > \Theta(h) \approx 1 - \frac{\delta\gamma}{\rho a \left(\beta\delta + (1-\beta)\gamma\right)} h + o(h). \tag{29}$$

Using the formulae of $h_2(x)$ and $\Theta(h)$, and numerical values for $(\beta, \delta, \gamma, \rho, a)$, it is easy to determine whether there exists $0 < x < 1$ and $0 < h < a$ that satisfy the necessary and sufficient conditions for an equilibrium with delayed stabilization. It is also easy to see that there are no equilibria with delayed stabilization when x is sufficiently large. Equation (22) shows that

$0 < \pi(0) < 1$; using the definition of h_2 in equation (28) $h_2(x) = 0$ for $x > 1 - \pi(0)$. Thus, for $x > 1 - \pi(0)$ the necessary and sufficient condition for delayed stabilization (29) simplifies to $x > 1$, which never holds.

5 Exponential (constant) discounting

Even with constant discounting, the binary action model is not entirely standard. Understanding this model 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.

Since 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 δ with γ.) The constant discount rate is δ and the distant future is heavily discounted. Following the standard procedure to obtain the DPE, or invoking Remark 2, we have the following DPE:

$$(\delta + h) V(h) = \max_{w \in \{0,1\}} \{U(w) + \rho(a - h)(1 - w) V'(h)\}. \tag{30}$$

Let $\pi^0(h)$ and $\sigma^0(h)$ denote the functions $\pi(h)$ and $\sigma(h)$ (defined in equations (22) and (23)) evaluated at $\beta = 0$. The following proposition describes the optimal solution to the control problem with $\beta = 0$.

Proposition 5. *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, note that $\pi(h)$ and $\sigma(h)$ are continuous in β, so $\pi^0(h)$ and $\sigma^0(h)$ are the limits of these functions as $\beta \to 0$. Consider a value of β that is positive but close to 0 and a value of h that satisfies $1 - \pi(h) > x > 1 - \sigma(h)$. (Such a state exists because $\pi(h)$ and $\sigma(h)$ are continuous in β, and there exists h that satisfies

$1-\pi^0(h) > x > 1-\sigma^0(h)$, as shown in the proof of Proposition 5.) For this combination of parameters and state variable, there are two MPE, involving either perpetual stabilization or perpetual BAU (by Proposition 3), but the payoff under perpetual BAU is higher than under stabilization (by continuity and Proposition 5). That is, there are MPE that involve *excessive stabilization* relative to the first best. It is not, however, true in general that when $1-\pi(h) > x > 1-\sigma(h)$, BAU yields a higher payoff than stabilization. The argument used in the proof of Proposition 5 shows that where there are two solutions to the DPE, the solution associated with BAU gives a higher payoff. Inspection of the proof shows that this argument does not carry over to the case where β is bounded away from 0, because in this situation the DPE under hyperbolic discounting is not close to the DPE under constant discounting.

The second implication is that $\lambda(h) = \sigma(h)$ under constant discounting. That is, 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 low probability 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 low probability 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.

6 Numerical illustration

We illustrate the binary action model of Section 2.1 using two climate scenarios that differ with regard to the initial hazard. Under the "pessimistic" initial hazard the probability of the catastrophe occurring within a century is 5% under stabilization (the policy that keeps the hazard constant) and the probability under BAU is 18%. Under the "optimistic" initial hazard the probability of occurrence is 0.5% under stabilization, and 15.3% under BAU. For both scenarios the maximal hazard a implies a 50% occurrence probability within a century. We choose ρ so that under BAU it takes 100 years to travel half way between the "pessimistic" initial hazard h_0 and a. Table 1 presents the resulting hazard parameter values.

Table 1: Hazard parameter values.

a	0.00693147	
ρ	0.00544875	
	Optimistic	Pessimistic
h_0	0.000100503	0.000512933

Table 2 lists values of the hyperbolic discounting parameters, β, δ and γ. We use three long run discount rates, $\gamma = 0.0005$, $\gamma = 0.00005$ and $\gamma = 0$ (corresponding to long run discount rates of 0.05%, 0.005% and 0%, respectively). We choose the parameters β, δ so that the short run discount rate is 5% ($r(0) = 0.05$) and the discount rate a century in the future is 4% ($r(100) = 0.04$) for each of the three long run rates.

Table 2: Discounting parameter values.

γ	5×10^{-4}	5×10^{-5}	0
β	0.00178999	0.00169212	0.00168159
δ	0.0500888	0.0500847	0.0500842

Figure 2 shows the graphs of the discount rates and discount factors corresponding to a constant discount rate of 5% and to the hyperbolic rate in Table 2 with $\gamma = 0.0005$. Under hyperbolic discounting, the discount

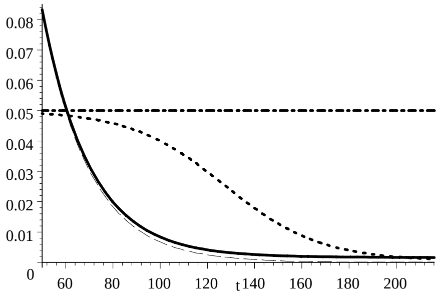

Figure 2: Discount rates and factors: dotted curve = hyperbolic discount rate; solid curve - discount factor under hyperbolic discounting; dashed curve = discount factor under constant 0.05 discount rate

rate is greater than 4% during the first century. Despite the similarity of the discount factors under constant and hyperbolic discounting during this period, the policy implications differ markedly: The future lasts a long time. A slightly smaller discount factor in the distant future has a large effect on current policy.

The parameter values in Tables 1 and 2 do not satisfy the sufficient condition in Proposition 1. Nevertheless, as Figure 3 reveals, $\lambda(h)$ is strictly increasing (and it remains so for all climate and discounting scenarios considered here). By Proposition 1, the policy under restricted commitment is time consistent. Figure 3 also shows that $\sigma(h) \geq \pi(h)$, so the necessary and sufficient condition for existence of the MPE (cf. Proposition 3) is satisfied.

Recall that $\frac{U(1)}{U(0)} \equiv 1-x$, so x is the loss in utility flow under stabilization, as a fraction of the value-at-risk. For example, if utility is linear in income, and if the event reduces the flow of income by \$100 billion annually, the value $x = 0.15$ means that stabilization costs \$15 billion per year. We defined $\bar{x}^c = 1 - \lambda$ as the maximum fractional reduction in the flow of value-at-risk that the decision-maker is willing to sacrifice (in order to stabilize the hazard) under restricted commitment. We use equation (26) to define $\bar{x}^S \equiv 1 - \pi$ as the maximal fractional reduction in the flow of value-at-risk that is consistent

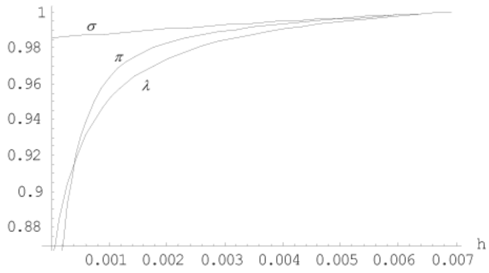

Figure 3: The graphs of $\lambda(h), \sigma(h)$, and $\pi(h)$

with stabilization in a MPE. Using equation (27), let $\underline{x}^B \equiv 1-\sigma$ represent the minimal fractional reduction in the flow of value-at-risk that is consistent with BAU in a MPE. Table 3 shows these thresholds under the different discounting and hazard scenarios.

Table 3: Policy bounds.

Discounting mode	$\underline{x}^B$	$\bar{x}^S$	$\bar{x}^C$
		Optimistic	
Hyperbolic $\gamma = 0.0005$	0.0143888	0.166323	0.122572
Hyperbolic $\gamma = 0.00005$	0.0144876	0.736511	0.355016
Hyperbolic $\gamma = 0$	0.0144996	0.861326	0.451631
Constant 4%	0.0193835		0.0193835
		Pessimistic	
Hyperbolic $\gamma = 0.0005$	0.0132928	0.0694618	0.0737647
Hyperbolic $\gamma = 0.00005$	0.0133795	0.166933	0.123743
Hyperbolic $\gamma = 0$	0.01339	0.191701	0.13406
Constant 4%	0.0179322		0.0179322

Consider for example the pessimistic initial hazard and the long run discount rate $\gamma = 0.00005$. In this scenario, stabilization is the only MPE if stabilization reduces the utility flow by less than 1.34% of the value-at-risk. If the reduction caused by stabilization lies between 1.34% and 16.69%, both stabilization and BAU are MPE. If stabilization reduces utility by more than 16.69% of the value-at-risk, BAU is the only MPE. If the present generation can commit to actions taken by future generations, stabilization is the optimal policy only if the loss does not exceed 12.37%.

Table 3 shows that it is possible to have a MPE with stabilization even though the socially optimal policy (under restricted commitment) requires BAU. When that occurs, the MPE leads to excessive stabilization effort. This possibility requires $\bar{x}^S(h) > x > \bar{x}^C(h)$, which can happen when $\lambda(h) > \pi(h)$ (see Figure 3.) It is also possible that $\bar{x}^S(h) < \bar{x}^C(h)$, which happens when $\lambda(h) < \pi(h)$. In this case if $x^S(h) < x < \bar{x}^C(h)$, stabilization is the optimal (restricted commitment) policy, but all MPE lead to BAU. In all of our experiments, $\underline{x}^B < \bar{x}^C$ (i.e., $\lambda < \sigma$) so there always exist parameter values ($\underline{x}^B < x < \bar{x}^C$) for which there is a MPE leading to BAU, even though the optimal policy requires stabilization.

Several other features of the table are noteworthy. Stabilization is an equilibrium for larger values of x under the optimistic (low) initial hazard h_0 compared to the pessimistic (high) h_0 scenario. BAU is *not* an equilibrium for larger x under optimistic compared to the pessimistic scenario. In other other words, a low initial value of h encourages stabilization. We discussed the reason for that result in Section 2.1. The growth rate of the hazard is a decreasing function of the hazard. Since stabilization keeps the risk from growing, it is more attractive to incur costs to stabilize when the growth rate is large, i.e. when h is small.

The upper bounds, $\bar{x}^S$ and $\bar{x}^C$, are quite sensitive to changes in the long run discount rate γ. The lower bound, $\underline{x}^B$, is relatively insensitive to these changes. The set of utility-related parameter values for which the MPE is indeterminate (i.e., where $\underline{x}^B < x < \bar{x}^S$) varies from about 5.5% to 85% of the feasible range $(0, 1)$. It is clear from these examples that indeterminacy of the MPE is not a knife-edge phenomenon.

The two rows in Table 3 labelled “constant 4%” show the critical threshold of x under a constant discount rate of 4%, in the optimistic and pessimistic scenarios. As we explained below Proposition 5, this bound equals $\underline{x}^B = \bar{x}^c$. (Since $\bar{x}^S$ is irrelevant, it is not reported.) For example, in the optimistic scenario, society is willing to sacrifice 1.9% of the value-at-risk in order to stabilize the risk, when the discount rate is constant at 4%. With a hyperbolic discount rate that begins at 5%, is greater than 4% for a century, and asymptotically approaches .05%, society is willing (under restricted commitment) to sacrifice over 12% of the value-at-risk in order to stabilize the risk. Thus, even though the short run discount rate is higher under hyperbolic discounting in this example, the fact that the long run discount rate is much smaller, increases society's willingness to pay for stabilization by a factor of more than six.

Table **??** reports "constant-equivalent" discount rates. These rates lead to the same decision rules (the same threshold levels of x) as in the Markov Perfect and restricted commitment equilibria with hyperbolic discounting and $\gamma = .0005$ (a long run discount rate of 0.05%). For example, a constant discount rate of 0.0126 (1.26%) leads to the same threshold as under limited commitment with hyperbolic discounting in the optimistic scenario. The constant-equivalent discount rates corresponding to the MPE with the least likelihood of stabilization are about 4.7%, close to the short run discount rate under hyperbolic discounting. In contrast, the constant-equivalent discount rates corresponding to the MPE with the greatest likelihood of stabilization, and to the restricted commitment case, lie between 1.2% and 2.1%.

Table 4: Constant discounting equivalent rates corresponding to hyperbolic discounting with $\gamma = 0.0005$.

	$\underline{x}^B$	$\bar{x}^S$	$\bar{x}^C$
Optimistic	0.0471568	0.009964	0.012628
Pessimistic	0.0472403	0.0176443	0.016947

A standard public finance position is that an externality (or some other market failure) that justifies undertaking a project should be incorporated directly into the cost-benefit analysis, rather than captured by an adjustment in the discount rate. In our view, the discount rate in a 30 year bond (or some other short-lived financial instrument) tells us very little about the "correct" long run social discount rate, and it certainly does not tell us that the discount rate should be constant. We think that hyperbolic discounting models provide a better representation of society's view of far-distant generations, but these models are difficult to work with. Therefore, it is worth having a sense of how constant discount rates should be adjusted to mimic a hyperbolic discounting model. Table **??** provides some evidence on this point.

7 Conclusion

Most integrated assessment models that are used to evaluate climate policy either do not consider catastrophic events or introduce them in an *ad hoc* manner. The damage estimates used in many of these models suggest that climate change is a less serious problem than many climate scientists and environmentalists think. There appears to be a widespread view amongst environmental economists that taking into account (more systematically) the risk of catastrophic climate-related events would not fundamentally alter the recommendations implied by mainstream models. There are two reasons for this view (apart from general skepticism of climate science): (i) the event is unlikely, so the probability of it happening in the near future is too small to worry about; (ii) the inertia in the climate system means that current policy changes would affect the risk only in the distant future. These arguments are persuasive only if the far-distant discount rate remains substantial. In view of our inability to distinguish between generations in the distant future, we think that a model with a declining discount rate provides a better description of how most people regard the distant future, and therefore provides a better normative guide for climate policy.

We studied a model in which changes in the profile of GHG emissions affect the future risk of abrupt climate events. To account for the inertia in the climate system, different policy scenarios lead to gradually diverging risks, with finite steady state differences. In this setting, a normative model with constant discounting (at a "plausible" rate) might conclude that the policy change is not warranted. Such a conclusion reflects the judgement that the current generation should be indifferent to the welfare of generations in the distant future. This judgement should not be mistaken for a scientific conclusion. Market rates for financial instruments that mature in 30 years tell us little about our willingness to transfer consumption between two very distant generations. Hyperbolic discounting forces us to make an explicit judgement about trade-offs in the long run, while still respecting the empirical evidence about short and medium run discount rates.

We obtained the necessary condition for Markov Perfect Equilibria in a general setting with hyperbolic discounting and event uncertainty, and then specialized to a binary action model. At each point in time the regulator follows BAU or stabilizes the risk. In general, there are multiple MPE because the optimal decision for the current regulator depends on the shadow value of the hazard, which in turn depends on the strategies used by succeeding regulators.

We emphasized the situation where the event is "low-risk", i.e. where the hazard rate is smaller than anything that (most) economists would recognize as a plausible short run social discount rate. A model of constant discounting has little that is useful to say about such events, simply because a non-negligible constant discount rate places so little weight on future welfare, and the low-risk event is not likely to occur until the distant future. We used a hyperbolic discount rate that remains above 4% for a century into the future, and eventually falls to a level similar to the steady state hazard or below.

The scientific evidence is currently inadequate to reliably estimate the risk of specific climate-related events. We chose parameters so that the current risk is very low. Under stabilization (in perpetuity) the risk remains constant; under BAU in perpetuity, it increases at a diminishing rate, reaching half way toward its maximal level within a century. Under plausible parameter values, it is optimal to forgo a substantial fraction (e.g., 74%) of the value-at-risk in order to stabilize the hazard. In view of the limited empirical basis for the risk calibration, these numerical results are only suggestive, yet they indicate that a systematic accounting of catastrophic risk might warrant a more aggressive climate policy, compared to the prescriptions of integrated assessment models.

In our experiments, there always exist utility parameters (summarized by x) for which there is a MPE that results in too little stabilization (relative to the restricted commitment benchmark). For some combinations of x and risk-related parameter values, *all* MPE result in too little stabilization efforts. For other combinations of x and risk-related parameter values, there are MPE that result in excessive stabilization effort.

The difficulty in solving the free-rider problem amongst decision-makers in the current generation presents a serious impediment to optimal climate policy. It may also be difficult to achieve coordination amongst decision-makers in succeeding generations.

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, pp. 13–22.

Asheim, G. B. and Buchholz, W.: 2004, A general approach to welfare measurement through national income accounting, *Scandinavian Journal of Economics* **106**(2), 361–384.

Barrett, S.: 2003, *Environment and Statecraft*, Oxford University Press.

Broecker, W. S.: 1997, Thermohaline circulation, the Achilles heel of our climate system: Will man-made CO_2 upset the current balance?, *Science* **278**, 1582–1588.

Chakravorty, U., Roumasset, J. and Tse, K.: 1997, Endogenous substitution among energy resources and global warming, *Journal of Political Economy* **105**, 1201–1234.

Challenor, P.: 2005, Analysis of survey at RAPID annual meeting 2005, unpublished.

Challenor, P. G., Hankin, R. K. S. and Marsh, R.: 2006, Toward the probability of rapid climate change, *in* H. J. Schellnhuber (ed.), *Avoiding Dangerous Climate Change*, Cambridge University Press, pp. 55–64.

Chichilnisky, G.: 1996, An axiomatic approach to sustainable development, *Social Choice and Welfare* **13**, 231 – 257.

Chichilnisky, G. and Heal, G.: 1993, Global environmental risks, *The Journal of Economic Perspectives* **7**, 65–86.

Clarke, H. R. and Reed, W. J.: 1994, Consumption/pollution tradeoffs in an environment vulnerable to pollution-related catastrophic collapse, *Journal of Economic Dynamics and Control* **18**(5), 991–1010.

Cropper, M. L.: 1976, Regulating activities with catastrophic environmental effects, *Journal of Environmental Economics & Management* **3**, 1–15.

Cropper, M. L. and Laibson, D.: 1999, The implications of hyperbolic discounting for project evaluation, *in* P. R. Portney and J. P. Wayant (eds), *Discounting and Intergenerational Equity*, Resources for the Future, p. 163Ű172.

Dasgupta, P. and Maskin, E.: 2005, Uncertainty and hyperbolic discounting, *The American Economic Review* **95**, 1290–1299.

Edmonds, J. A. and Sands, R. D.: 2003, What are the costs of limiting CO_2 concentrations?, *in* J. M. Griffin (ed.), *Global Climate Change*, Edward Elgar, pp. 140–186.

Frederick, S., Loewenstein, G. and O'Donoghue, T.: 2002, Time discounting and time preference: A critical review, *Journal of Economic Literature* **40**, 351–401.

Gollier, C.: 2002, Discounting an uncertain future, *Journal of Public Economics* **85**, 149–166.

Griffin, J. M.: 2003, *Global Climate Change*, Edward Elgar.

Harris, C. and Laibson, D.: 2004, Instant gratification, Cambridge University working paper; forthcoming, Econometrica.

Hartwick, J. M.: 1977, Intergenerational equity and the investing of rents from exhaustible resources, *American Economic Review* **67**(5), 972–74.

IPCC: 2001, *Climate change 2001: Impacts, adaptation and vulnerability*, Cambridge University Press.

Karp, L.: 2005, Global warming and hyperbolic discounting, *Journal of Public Economics* **89**, 261–282.

Karp, L.: 2006, Non-constant discounting in continuous time, *Journal of Economic Theory, forthcoming* .

Krugman, P.: 1991, History versus expectations, *Quarterly Journal of Economics* (106), 651–67.

Krusell, P. and Smith, A.: 2003, Consumption-saving desisions with quasi-geometric discounting, *Econometrica* **71**(1), 365–75.

Manne, A. S.: 2003, Energy,the environment and the economy: hedging our bets, *in* J. M. Griffin (ed.), *Global Climate Change*, Edward Elgar, pp. 187–203.

Mastrandrea, M. D. and Schneider, S. H.: 2001, Integrated assessment of abrupt climatic changes, *Climate Policy* **1**, 433–449.

Matsuyama, K.: 1991, Increasing returns, industrialization, and indeterminancy of equilibrium, *Quarterly Journal of Economics* **106**, 587–597.

Mendelsohn, R.: 2003, Assessing the mrket damages from climate change, *in* J. M. Griffin (ed.), *Global Climate Change*, Edward Elgar, pp. 92–113.

Nordhaus, W. D. and Boyer, J.: 2000, *Warming the World: Economic Models of Global Warming*, The MIT Press.

Schlenker, W., Hanemann, W. M. and Fisher, A. C.: 2006, The impact of global warming on u.s. agriculture: an econometric analysis of optimal growing conditions, *Review of Economics and Statistics* **88**(1), 113–125.

Schlesinger, M. E., Yin, J., Yohe, G., Andronova, N. G., Melyshev, S. and Li, B.: 2006, Assessing the risk of a collapse of the atlantic thermohaline circulation, *in* H. J. Schellnhuber (ed.), *Avoiding Dangerous Climate Change*, Cambridge University Press, pp. 37–48.

Skiba, A.: 1978, Optimal growth with a convex-concave production function, *Econometrica* **46**, 527–39.

Solow, R. M.: 1974, Intergenrational equity and exhaustible resources, *The Review of Economic Studies* **41**, 29–45.

Thomas, C. D. et al.: 2004, Extinction risk from climate change, *Nature* **427**, 145–148.

Tsur, Y. and Zemel, A.: 1996, Accounting for global warming risks: Resource management under event uncertainty, *Journal of Economic Dynamics & Control* **20**, 1289–1305.

Tsur, Y. and Zemel, A.: 1998, Pollution control in an uncertain environment, *Journal of Economic Dynamics & Control* **22**, 967–975.

United Nations: 2005, Milennium ecosystem assessment, *Report*, United Nations, http://maweb.org//en/Products.Global.Condition.aspx.

Weitzman, M. L.: 2001, Gamma discounting, *American Economic Review* **91**, 260–271.

Wunsch, C.: 2006a, Abrupt climate change: an alternative view, *Quarternary Research* **65**, 191–203.

Wunsch, C.: 2006b, Letters to editor, The Economist, September 30 - October 6.

Appendix: Proofs

Proof of Proposition 1 (i) This claim follows from differentiating the functions $\nu(h)$ and $\xi(h)$ and by inspection. (ii) From equations (13), (15) and (33),

$$\nu(h) - \xi(h) = \int_0^\infty \theta(t) \left(e^{-y(t,h)} - e^{-ht}\right) dt. \tag{31}$$

It is easy to verify that $\frac{1-e^{-\rho t}}{\rho}$ is strictly decreasing in ρ for $\rho > 0$ and equals t at $\rho = 0$. Therefore, $y(t,h) > ht$ when $h < a$ and $\rho > 0$, and the right-hand side of equation (31) is negative. (iii) This claim is merely a summary of the derivation in the text above equation (16).

(iii) (Sufficiency) Suppose that $\lambda(h)$ is non-decreasing. Then for any $1 - x \geq \lambda(h)$ it is optimal to stabilize. Since h does not change under stabilization, it is also optimal to stabilize at any point in the future. For any $1 - x < \lambda(h)$ it is optimal to follow BAU. Since h increases along the BAU trajectory, the inequality $1 - x < \lambda(h)$ continues to hold along this trajectory and BAU remains optimal. (Necessity). Suppose that λ is strictly decreasing over some interval $0 \leq h_1 < h < h_2 \leq a$. Choose a value of h in this interval (the initial condition $h(0)$), and choose $1 - x = \lambda(h(0)) - \epsilon$, where ϵ is small and positive. At this initial condition and for this value of $1 - x$, it is optimal to follow BAU, causing h to increase. Because λ is decreasing in this neighborhood, there is a future time $t > 0$ at which $1 - x = \lambda(h(t))$. At this time, it becomes optimal to stabilize, so the initial decision to pursue BAU in perpetuity is not time consistent.

(iv) Using (12) and (14), we express $\lambda(h)$ as

$$\lambda(h) = \frac{\int_0^\infty e^{-y(t,h)}\theta(t)dt}{\int_0^\infty e^{-ht}\theta(t)dt}, \tag{32}$$

where $\theta(t) = \beta e^{-\gamma t} + (1-\beta)e^{-\delta t}$ (equation (1)) and

$$y(t,h) \equiv \int_0^t (a - (a-h)e^{-\rho\tau})d\tau = at - (a-h)\frac{1-e^{-\rho t}}{\rho} \tag{33}$$

satisfies

$$y_h(t,h) \equiv \partial y(t,h)/\partial h = \frac{1-e^{-\rho t}}{\rho}. \tag{34}$$

($y(t,h)$ of (33) is a specialization of $y(t)$, defined in (4), when the hazard process under BAU evolves according to (11). The argument h in $y(t,h)$ signifies the initial hazard.) Differentiating (32) with respect to h, we see that $\lambda'(h) > 0$ if and only if

$$\int_0^\infty e^{-y(t,h)}\theta(t)dt \int_0^\infty e^{-ht}t\theta(t)dt > \int_0^\infty e^{-ht}\theta(t)dt \int_0^\infty e^{-y(t,h)}y_h(t,h)\theta(t)dt. \tag{35}$$

Noting $\int_0^\infty e^{-ht}\theta(t)dt = \frac{\beta}{h+\gamma} + \frac{1-\beta}{h+\delta}$ and $\int_0^\infty e^{-ht}t\theta(t)dt = \frac{\beta}{(h+\gamma)^2} + \frac{1-\beta}{(h+\delta)^2}$ and using (34), we express (35) as

$$\begin{aligned} &\left(\tfrac{\beta}{(h+\gamma)^2} + \tfrac{1-\beta}{(h+\delta)^2}\right)\textstyle\int_0^\infty e^{-y(t,h)}\theta(t)dt > \\ &\left(\tfrac{\beta}{h+\gamma} + \tfrac{1-\beta}{h+\delta}\right)\textstyle\int_0^\infty e^{-y(t,h)}\theta(t)\tfrac{1-e^{-\rho t}}{\rho}dt. \end{aligned} \tag{36}$$

Since $\delta > \gamma$, the right-hand side of inequality (36) is smaller than

$$\left(\frac{\beta}{(h+\gamma)^2} + \frac{1-\beta}{(h+\delta)^2}\right)\int_0^\infty e^{-y(t,h)}\theta(t)\frac{(h+\delta)(1-e^{-\rho t})}{\rho}dt. \tag{37}$$

Thus, it suffices to show that the left-hand side of (36) exceeds (37), i.e., that

$$\int_0^\infty e^{-y(t,h)}\theta(t)\left(1 - \frac{(h+\delta)(1-e^{-\rho t})}{\rho}\right)dt > 0,$$

which is guaranteed to hold if $\rho > h + \delta$. Since $h \leq a$ and h approaches a under BAU, the inequality holds at all $h \in [0,a]$ if $\rho > a + \delta$. $\square$

Proof of Proposition 2 We use Proposition 1 and Remark 2 in Karp (2006). In that paper the state variable is a scalar, but the same results hold (making obvious changes in notation) when the state is a vector, as in the present case. Our state variable is $z \equiv (h, y)$ and the flow of utility (prior to the event) is

$e^{-y(t)}U(w(t))$. Specializing equation (5) of Karp to our setting, and using the hyperbolic discount factor in equation (1), yields the generalized DPE

$$\hat{K}(z) + \gamma W(z) = \max_{w \in \Omega} \left(e^{-y(t)} U(w(t)) + W_h g + W_y h \right), \tag{38}$$

where $W(z)$ is the value function (with subscripts denoting partial differentiation) and

$$\hat{K}(z) = (\delta - \gamma)(1 - \beta) \int_0^\infty e^{-(\delta t + y(t))} U(\hat{\chi}(z))\, dt \tag{39}$$

is implied by equation (4) and Remark 2 of Karp (2006)

Use the "trial solution" $W(z) = e^{-y}V(h)$ and $\hat{K}(z) = e^{-y}K(h)$, so $W_y = -e^{-y}V(h)$ and $W_h = e^{-y}V'(h)$. Substituting these expressions into equation (38), cancelling e^{-y} and rearranging, yields equation (18). Conclude that $\hat{\chi}(z) = \chi(h)$: the equilibrium control depends only on the hazard rate.

Conditional on survival up to time t, the probability of survival until time $s > t$ equals $\exp\left(-\int_t^s h(\tau)d\tau\right) = \exp(-y(s) + y(t))$. Use this fact and the trial solution to rewrite equation (39) as

$$\begin{aligned} K(h(t)) &= (\delta - \gamma)(1 - \beta)\, e^{y(t)} \int_t^\infty e^{-\delta(s-t)} \exp\left(-\int_t^s h(\tau)d\tau\right) e^{-y(t)} U(\chi(h(s)))\, ds \\ &= (\delta - \gamma)(1 - \beta) \int_t^\infty e^{-\delta(s-t)} \exp\left(-\int_t^s h(\tau)d\tau\right) U(\chi(h(s)))\, ds \end{aligned} \tag{40}$$

Setting $t = 0$ in equation (40) produces equation (19). □

Proof of Lemma 1 Define

$$\varpi(h) \equiv \pi(h)^{-1} = 1 - \rho(a - h)\xi'(h). \tag{41}$$

Differentiating, noting (21), we obtain

$$\varpi'(h) = \rho\xi'(h) - \rho(a - h)\xi''(h) < 0. \tag{42}$$

Thus,

$$\pi'(h) = -\varpi'(h)/\varpi(h)^2 > 0. \tag{43}$$

Differentiating (23), noting (24), gives

$$\sigma'(h) = -\rho\nu'(h) + \rho(a-h)\nu''(h) > 0. \tag{44}$$

To establish $\sigma''(h) < 0$, note from (24) that $\nu'''(h) < 0$ and differentiating (44) gives

$$\sigma''(h) = -2\rho\nu''(h) + \rho(a-h)\nu'''(h) < 0.$$

By inspection $\pi(a) = \sigma(a) = 1$. □

Proof of Proposition 3 We first establish sufficiency using a constructive proof, and then necessity using a proof by contradiction.

Sufficiency Suppose that $\sigma > \pi$ for $h \in (0, a)$. We first show that there exists a MPE that satisfies $w \equiv 1$ (perpetual stabilization) if and only if the initial condition $h_0 = h$ satisfies equation (26). In a MPE with perpetual stabilization, it is optimal for the current regulator to stabilize given that she believes that future values of h lie in the stabilization region (so she believes that all subsequent regulators will stabilize). The belief that future values of h lie in the stabilization region (a belief we test below) means that for initial conditions in the interior of the stabilization region the value function is given by $V^S(h)$, defined in equation (14), and

$$V^{S\prime}(h) = U(1)\xi'(h) \tag{45}$$

with $\xi'(h)$ given by equation (21).

Using equation (18) (and the belief that future values of h lie in the stabilization region), it is optimal for the current regulator to stabilize if and only if

$$U(1) \geq U(0) + \rho(a-h)U(1)\xi'(h) \tag{46}$$

or

$$\frac{U(1)}{U(0)} \geq \pi(h). \tag{47}$$

If inequality (47) is satisfied with *strict inequality* (as the Proposition requires) at the current time, then regardless of whether the current regulator uses stabilization of BAU, the inequality is satisfied at neighboring times (the near future). Thus, the current regulator's beliefs that future regulators will stabilize are consistent with equilibrium, regardless of the actions taken by the current regulator. If inequality (47) is not satisfied, then clearly perpetual stabilization is not an equilibrium. We consider below the case where equation (47) holds with equality.

We turn now to the equilibrium with perpetual BAU. In a MPE with perpetual BAU, it is optimal for the current regulator to follow BAU given that she believes all subsequent regulators will follow BAU. This belief implies that the value function is given by $V^B(h)$, defined in equation (12). It is optimal for the current regulator to pursue BAU if and only if $U(0) + \rho(a-h)U(0)\nu'(h) > U(1)$ or, equivalently, if and only if

$$\frac{U(1)}{U(0)} < \sigma(h) \equiv 1 + \rho(a-h)\nu'(h), \tag{48}$$

establishing condition (27).

To complete the demonstration that perpetual stabilization is an equilibrium, it is necessary to confirm that if equation (27) is satisfied at time t when the hazard is h, then it is also satisfied at all subsequent times, so that the regulator's beliefs are confirmed. The hazard is increasing on the BAU equilibrium path (and non-decreasing on any feasible path), so it is sufficient to show that $\sigma'(h) > 0$. This inequality was established in Lemma 1.

Now we return to the case where inequality (47) is satisfied with equality. We want to show that in this case, stabilization is not an equilibrium action. That is, the curve $\pi(h)$ does not lie in the stabilization set. Suppose to the contrary that $\pi(h)$ does lie in the stabilization. From equation (20), the current regulator wants to use BAU if and only if $U(1) < U(0) + \rho(a-h)V'(h)$. In order to evaluate the right side of this inequality, we need to know the value of $V'(h)$; this (shadow) value of course depends on the behavior of future regulators.

Because $\pi'(h) > 0$ from Lemma 1, if the current regulator uses BAU, h

increases and the state is driven out of the stabilization region. Therefore, the current regulator can discard the possibility that (if she were to use BAU) all future regulators would stabilize. Future actions could lead to only one of two possible equilibrium trajectories: (i) All future regulators will follow BAU; or (ii) future regulators will follow BAU until the state h reaches a threshold, say $h_0 < \tilde{h} < a$, after which all regulators stabilize. There are no other possibilities, because once the state enters a stabilization region it does not leave it. This fact is a consequence of our restriction to pure strategy equilibria. However, alternative (ii) cannot occur, because $\tilde{h}$ lies to the right of the curve $\pi(h)$, and therefore is not an element of the stabilization region. Thus, the only equilibrium belief for the current regulator is that the use of BAU (and the subsequent increase in h) will cause all future regulators to use BAU. Consequently, where inequality (47) is satisfied with equality, it must be the case that $V'(h) = V^{B\prime}(h) = U(0)\nu'(h)$. The assumption that $\sigma(h) > \pi(h)$ implies that $\pi(h)$ lies in the region where perpetual stabilization is an equilibrium strategy. Thus, $\pi(h)$ does not lie in the stabilization region, as asserted by the proposition.

Necessity: We use a proof by contradiction to establish necessity. Suppose that for some interval $\sigma(h) < \pi(h)$. Figure 4 helps to simplify the proof. This figure shows a situation where $\sigma(h) < \pi(h)$ for small h, but it is clear from the following argument that the region over which $\sigma(h) < \pi(h)$ is irrelevant. (An obvious variation of the following argument can be used regardless of the region over which $\sigma < \pi$, because both of these curves are monotonic.) Suppose that the value of $\frac{U(1)}{U(0)}$ lies between the vertical intercepts of the curves, as shown in the figure; e.g. $\frac{U(1)}{U(0)} = d$. Define h_1 implicitly by $\sigma(h_1) = d$. we want to establish that for any initial condition $h_0 = h < h_1$ there are no pure stationary MPE. Perpetual stabilization is not an equilibrium because $d < \pi(h_1)$, and perpetual BAU is not an equilibrium because $d > \sigma(h_1)$. The only remaining possibility is to follow BAU until the hazard reaches a level $\bar{h} < h_1$ and then begin perpetual stabilization. (Recall that

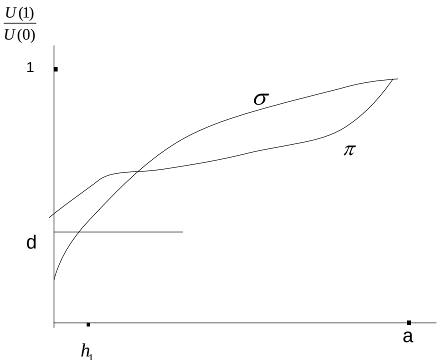

Figure 4: Graphs of $\sigma(h)$ and $\pi(h)$ that do not satisfy inequality (25).

once the state enters the stabilization set it cannot leave that set.) However, this trajectory cannot be an equilibrium because the subgame beginning at $\bar{h}$ cannot lead to perpetual stabilization (because the point (h_1, d) lies below the curve π). $\square$

Proof of Proposition 4 (i) The stabilization set is absorbing, because if a (pure strategy) MPE calls for a regulator to stabilize, the hazard never changes. By Proposition 3, there are no equilibria with perpetual stabilization when $h(0) \geq h_2$, and there is an equilibrium with perpetual BAU. The latter is therefore the unique equilibrium. Claim (ii) follows immediately from the fact that the stabilization set is absorbing

(iii) We now consider the case where $h(0) < h_2$ – the maintained assumption for this part of the proof From Proposition 3 we know that there is an equilibrium with perpetual stabilization for these initial conditions; and we know that there is an equilibrium with perpetual BAU if x lies between the curves $1 - \pi$ and $1 - \sigma$. Since the stabilization set is absorbing, we do not need to consider the possibility of equilibria that begin with stabilization and then switch to BAU. Thus, we need only find a necessary and sufficient condition under which there is a "delayed stabilization" equilibrium, i.e. one that begins with BAU and switches to stabilization when the state reaches a threshold $\tilde{h} > h(0)$. To conserve notation, throughout the remainder of this proof we use h to denote an initial condition, and use $h(\tau)$, with $\tau \geq 0$, to denote a subsequent value of the hazard when regulators use a MPE.

Define two sets, $A = \left\{h \mid h_a \leq h < \tilde{h}\right\}$ and $B = \left\{h \mid \tilde{h} \leq h < h_b\right\}$, where $h_a < \tilde{h} < h_b < h_2$. The MPE for initial conditions in set B is to stabilize, and the MPE for initial conditions in set A is to follow BAU. The existence of B follows from the fact that it is an equilibrium to stabilize for any initial conditions in $[0, h_2)$ (in view of Proposition 3). In addition, h remains constant when the regulator stabilizes. Therefore, any subset of the interval $[0, h_2)$ qualifies as the set B.

The existence of A is not obvious. We cannot rely on the proof of Proposition 3, since that proof applies to the case where the regulator follows BAU in perpetuity. Here we are interested in the case where the regulator switches from BAU to stabilization at a finite time. We obtain the necessary and sufficient condition for the existence of a non-empty set A.

Suppose (provisionally) that the non-empty set A exists. We define the value function for initial conditions in $A \cup B$ as $V(h;\tilde{h})$. We include the second argument in order to emphasize the dependence of the payoff on the switching value $\tilde{h}$. For convenience, we repeat the definition of the value function, given the initial condition $h \in A \cup B$.

$$\begin{aligned} V\left(h;\tilde{h}\right) &= \int_0^\infty e^{-y(\tau)}\theta(\tau)\,U(\chi(h(\tau))\,d\tau \quad \text{with } \chi(h) = \left\{\begin{array}{c} 0 \text{ for } h \in A \\ 1 \text{ for } h \in B \end{array}\right\}, \\ y(\tau) &= \int_0^\tau h(s)ds, \quad h(s) = \left\{\begin{array}{c} \min\left(a - (a-h)\,e^{-\rho s}, \tilde{h}\right) \text{ for } h \in A \\ h \text{ for } h \in B \end{array}\right\}. \end{aligned}$$

Note that for $h(\tau) \in A$, $h(\tau)$ is a function of the initial condition, h.

For $h \in A$ the regulator chooses BAU (under the candidate program). Using equation (20), this action is part of an equilibrium if and only if

$$U(0) - U(1) > -\rho(a-h)\,V_h(h;\tilde{h}). \tag{49}$$

In order to determine when this inequality holds, we need to evaluate $V_h(h;\tilde{h})$. For $h \in A$ the value function can be split into two parts: the payoff that arises from following BAU, and the subsequent payoff under stabilization. We state some intermediate results before discussing this two-part value function.

Define $T(h;\tilde{h})$ as the amount of time it takes to reach the stabilization threshold (the "time-to-go"), given the current state $h \in A$; T is the solution to

$$\tilde{h} = a - (a-h)\, e^{-\rho T} \Rightarrow \tag{50}$$

$$T\left(\tilde{h};\tilde{h}\right) = 0 \text{ and } \frac{dT}{dh} = \frac{-1}{\rho\,(a-h)}. \tag{51}$$

For $h \in A$ and for $\tau \leq T$

$$\frac{dy\,(\tau)}{dh} = \frac{d\int_0^\tau h(s)ds}{dh} = \int_0^\tau \frac{dh(s)}{dh}ds = \int_0^\tau e^{-\rho s}ds = \frac{1-e^{-\rho\tau}}{\rho}. \tag{52}$$

In addition, for $h \in A$ and for $\tau > T$

$$\begin{aligned} &\tfrac{dy(\tau)}{dh} = \tfrac{d\left(\int_0^T h(s)ds + \tilde{h}(\tau - T)\right)}{dh} = \\ &\int_0^T \tfrac{dh(s)}{dh}ds + \left(h(T) - \tilde{h}\right)\tfrac{dT}{dh} = \int_0^T e^{-\rho s}ds \end{aligned} \tag{53}$$

The last equality uses the fact that $h(T) = \tilde{h}$, from the definition of T. Using equation (50) and (51), we can invert the function $T(h;\tilde{h})$ to write the initial condition h as a function of the time-to-go T and the threshold $\tilde{h}$. Using this fact, equation (52) and the definition of $y(\tau)$, we have

$$\begin{aligned} y(T) &= \int_0^T h(s)ds \Rightarrow \\ \frac{dy(T)}{dT} &= h\,(T) + \int_0^T \frac{dh(s)}{dh}\frac{dh}{dT}ds \end{aligned} \tag{54}$$

We now discuss the value function for $h \in A$. Splitting the payoff into the parts before and after the threshold is reached, this function equals

$$V\left(h;\tilde{h}\right) = \int_0^T e^{-y(\tau)}\theta\,(\tau)\,U(0)dt + \int_T^\infty e^{-y(\tau)}\theta\,(\tau)\,U(1)dt,$$

and its derivative with respect to the initial condition (using equation (52) is

$$\begin{gathered}V_h\left(h;\tilde{h}\right)=\left(U\left(0\right)-U\left(1\right)\right)e^{-y(T)}\theta\left(T\right)\tfrac{dT}{dh}+\\ \int_0^T\tfrac{d\left(e^{-y(\tau)}\right)}{dh}\theta\left(\tau\right)U(0)dt+\int_T^\infty\tfrac{d\left(e^{-y(\tau)}\right)}{dh}\theta\left(\tau\right)U(1)dt\\ =\tfrac{-(U(0)-U(1))}{\rho(a-h)}e^{-y(T)}\theta\left(T\right)-\\ \left(\int_0^T\left(\tfrac{1-e^{-\rho\tau}}{\rho}\right)e^{-y(\tau)}\theta\left(\tau\right)U(0)dt+\int_T^\infty\left(\tfrac{1-e^{-\rho T}}{\rho}\right)e^{-y(\tau)}\theta\left(\tau\right)U(0)dt\right).\end{gathered}\tag{55}$$

Using this expression, we can write the optimality condition (49) as

$$\begin{gathered}U\left(0\right)-U\left(1\right)>\left(U\left(0\right)-U\left(1\right)\right)e^{-y(T)}\theta\left(T\right)+\\ \rho\left(a-h\right)\left(\int_0^T\left(\tfrac{1-e^{-\rho\tau}}{\rho}\right)e^{-y(\tau)}\theta\left(\tau\right)U(0)dt+\int_T^\infty\left(\tfrac{1-e^{-\rho T}}{\rho}\right)e^{-y(\tau)}\theta\left(\tau\right)U(0)dt\right).\end{gathered}\tag{56}$$

It is convenient to treat T as the independent variable, recognizing that the initial condition h is a function of T (from equation (50)): $h=h(T)$. The existence of a non-empty set A requires that inequality (56) hold for small positive values of T, i.e. for initial conditions h close to but smaller than $\tilde{h}$.

The first order Taylor expansion of the first term on the right side of inequality (56) is

$$\left(U\left(0\right)-U\left(1\right)\right)-\left(U\left(0\right)-U\left(1\right)\right)\left(\tilde{h}+r(0)\right)T+o\left(T\right).\tag{57}$$

This expansion uses equations (2) and (54) and the fact that $\theta\left(0\right)=1$. Using the fact that $1-e^{-\rho T}=0$ at $T=0$, the first order Taylor expansion of the second term in inequality (56) is

$$\begin{gathered}\rho\left(a-\tilde{h}\right)T\int_0^\infty e^{-y(\tau)}\theta\left(\tau\right)U(1)dt+o\left(T\right)=\\ \rho\left(a-\tilde{h}\right)T\int_0^\infty e^{-\tilde{h}\tau}\theta\left(\tau\right)U(1)dt+o\left(T\right)=\\ \rho\left(a-\tilde{h}\right)T\tfrac{(1-\beta)\gamma+\beta\delta+\tilde{h}}{\left(\tilde{h}+\gamma\right)\left(\tilde{h}+\delta\right)}U(1)+o\left(T\right).\end{gathered}\tag{58}$$

Substituting expressions (57) and (58) into inequality (56), dividing by T and letting $T \to 0$ produces the inequality

$$(U(0) - U(1))\left(\tilde{h} + r(0)\right) > \rho\left(a - \tilde{h}\right)\frac{(1-\beta)\gamma + \beta\delta + \tilde{h}}{\left(\tilde{h} + \gamma\right)\left(\tilde{h} + \delta\right)}U(1). \tag{59}$$

Using the definition $x = 1 - \frac{U(1)}{U(0)}$ and $r(0) = \beta\gamma + \delta(1-\beta)$ (from equation (2)), we write inequality (59) as

$$x > \frac{\rho(a-h)((1-\beta)\gamma+\beta\delta+h)}{h^3+(\delta+\gamma-\rho)h^2+((a-(1-\beta)\gamma-\beta\delta)\rho+\gamma\delta)h+((1-\beta)\gamma+\beta\delta)a\rho} \equiv \Theta(h). \tag{60}$$

Equation (29) provides the first order Taylor expansion of $\Theta(h)$, evaluated at $h = 0$.

Proof of Proposition 5 We first point out that existence of a solution to the optimal control problem requires that $\sigma^0(h) \geq \pi^0(h)$ over $h \in [0, a]$. We then show that there is no solution to the regulator's optimization problem that involves delayed stabilization. We then show that stabilization is optimal if and only if $x \leq 1 - \sigma^0(h)$.

If $\sigma^0(h) \geq \pi^0(h)$ over $h \in [0, a]$ were not satisfied, then (using the argument in the proof of Proposition 3) there would be some initial h and values $0 < \frac{U(1)}{U(0)} < 1$ for which there is no Markov perfect solution. However, the objective function under constant discounting is bounded and solution to the optimal control problem exists. Therefore, $\sigma^0(h) \geq \pi^0(h)$.

Now we show that there can be no solution with delayed stabilization. Define

$$\Theta^0(h) = \rho\frac{a-h}{h^2 - \rho h + h\delta + \rho a},$$

the function Θ evaluated at $\beta = 0$. It is easy to establish that

$$\Theta^0(0) = 1,\ \Theta^0(a) = 0,\ \text{and}\ \frac{d\Theta^0(h)}{dh} < 0 \text{ for } 0 < h < a. \tag{61}$$

The limiting case (as $\beta \to 0$) of the argument in the proof of Proposition 4 implies that there is a solution (not necessarily the optimal solution) to the

DPE when $\beta = 0$ if and only if $h < h_2(x)$ (using the definition in equation (28), setting $\beta = 0$) and $x > \Theta^0(h)$. Here it is convenient to treat h_2 as primitive, and use equation (28) (with $\beta = 0$) to solve for x. This operation gives

$$x(h) = \rho \frac{a - h}{(h + \delta)^2 + \rho (a - h)}.$$

Substituting this expression into the necessary condition $x > \Theta^0(h)$ implies

$$\rho \frac{a - h}{(h + \delta)^2 + \rho (a - h)} > \rho \frac{a - h}{h^2 - \rho h + h\delta + \rho a}.$$

Simplifying, this inequality implies $-h\delta(1 + \delta) > 0$, which is false. Therefore, the necessary condition for delayed stabilization never holds when $\beta = 0$.

We now turn to the main part of the proof. For h close to but smaller than a, $\sigma^0(h) > \pi^0(h)$. This claim uses a Taylor expansion. The Taylor expansion uses the facts that $\sigma^0(a) = \pi^0(a) = 1$ and the derivatives evaluated at $h = a$:

$$\sigma_h^0(a) = \frac{\rho}{(a + \rho + \delta)(\delta + a)} < \frac{\rho}{(\delta + a)^2} = \pi_h^0(a).$$

Thus, for some parameter values and initial conditions, $\pi^0(h) < \frac{U(1)}{U(0)} < \sigma^0(h)$ holds. In view of Proposition 3, the DPE (30) admits two solutions. With constant discounting, however, the solution to the optimization problem is unique. The possibility that there are multiple solutions to the necessary condition (the DPE), even though there is a unique optimal policy, also occurs in other control problems (e.g., Skiba 1978). We use the same line of reasoning as in the "Skiba problem" to identify the optimal policy.

Consider the situation where $\pi^0(h) < \frac{U(1)}{U(0)} < \sigma^0(h)$. Denote $V^S(h)$ and $V^B(h)$ as the value functions that satisfy the DPE (30) under stabilization and BAU, respectively, and let $V(h) = \max\left\{V^S(h), V^B(h)\right\}$ denote payoff under the optimal decision. The arguments used in the proof of Proposition 3 imply that for $\frac{U(1)}{U(0)} < \sigma^0(h)$, $V^B(h)$ satisfies

$$
\begin{aligned}
V^B(h) &= \tfrac{1}{\delta+h}\max\left\{U(1), U(0)+\rho(a-h)V_h^B(h)\right\} \\
&= \tfrac{1}{\delta+h}\left(U(0)+\rho(a-h)V_h^B(h)\right) > \tfrac{1}{\delta+h}U(1).
\end{aligned} \tag{62}
$$

Similarly, for $\frac{U(1)}{U(0)} > \pi^0(h)$, $V^S(h)$ satisfies

$$
\begin{aligned}
V^S(h) &= \tfrac{1}{\delta+h}\max\left\{U(1), U(0)+\rho(a-h)V_h^S(h)\right\} \\
&= \tfrac{1}{\delta+h}U(1) \geq \tfrac{1}{\delta+h}\left(U(0)+\rho(a-h)V_h^S(h)\right)
\end{aligned} \tag{63}
$$

From (62) and (63) we see that $V^B(h) > V^S(h)$ when $\pi^0(h) < \frac{U(1)}{U(0)} < \sigma^0(h)$. Therefore, when $\pi^0(h) < \frac{U(1)}{U(0)} < \sigma^0(h)$ the (unique) optimal policy is BAU.

Again using the arguments in Proposition 3, $V^S(h)$ is the only solution to the DPE when $\frac{U(1)}{U(0)} > \sigma^0(h)$; when this inequality is satisfied, the optimal solution is to stabilize. $V^B(h)$ is the only solution when $\frac{U(1)}{U(0)} < \pi^0(h)$; when this inequality is satisfied, BAU is the optimal solution. By convention, we break the tie, which occurs when $\frac{U(1)}{U(0)} = \sigma^0(h)$, by choosing stabilization. $\square$

Comments on alternate growth functions

(This section is not intended for publication.)

We discuss several generalizations of the growth equation. The fact that the equation of motion, equation (10), is decreasing in h implies that it in all of the equilibria we study, it is optimal to stabilize only if the hazard is sufficiently small. It is easy to see that in a more general setting it might be optimal to stabilize (i) for intermediate values of h (but not for large or small values) or (ii) only for large values. To illustrate case (i), suppose that $\dot{h} \approx 0$ (but positive) for $h < a_1$, and $\dot{h}$ is first increasing and then decreasing for $a_1 < h < a_2$, with $\dot{h} = 0$ for $h \geq a_2$. Since there is very little (or no) value of incurring the cost of stabilizing when $\dot{h} \approx 0$ (or $\dot{h} = 0$), it is obvious that stabilization occurs only for intermediate values of h. For case (ii), suppose that $\dot{h} \approx 0$ for $h < a_1$ a constant, and that $\dot{h}$ increases for $h \geq a_1$. In this case, stabilization would occur only for large h.

It is probably not possible to immediately stabilize the stock of GHGs at the current level. Moreover, due to inertia of various climatic processes, stabilizing GHG concentration is unlikely to immediately stabilize the hazard. Our model assumes that instantaneous stabilization of the hazard is possible. A simple generalization involves letting the hazard under BAU and under stabilization follow different trajectories. For example, the equations that determine these trajectories might both have the form of equation (10), but with different steady states and growth rates.

If the hazard under "stabilization" were non-constant, it would be natural to allow reversible hazard. For example, it would be possible to follow BAU until the hazard rate exceeded the steady state level under stabilization and then switch to stabilization, causing a decrease in the hazard rate. The rate of change of the hazard is probably slower when it is above rather than below its steady state, because the stock of GHG (and thus the hazard) can increase more quickly than it can decrease. Therefore, allowing reversibility would require an asymmetric model of the change in the hazard. Other possibilities depart from the binary action model. When the control variable w is an argument of the growth function, as in equation (8), the model allows for the possibility that (partial) stabilization causes the hazard to increase slowly, rather than to stop immediately. When the simple model (that we study) is better understood, and when there is a stronger scientific basis for calibration, these kinds of extensions will be worthwhile.