AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

*Department of Agricultural &
Resource Economics, UCB*
CUDARE Working Papers
(University of California, Berkeley)

Year 2007

Paper 1037

Climate Policy When the Distant Future
Matters: Catastrophic Events with
Hyperbolic Discounting

Larry Karp *

Yacov Tsur †

*University of California, Berkeley and Giannini Foundation

†The Hebrew University of Jerusalem

This paper is posted at the eScholarship Repository, University of California.

http://repositories.cdlib.org/are_uchb/1037

Copyright ©2007 by the authors.

Climate Policy When the Distant Future Matters: Catastrophic Events with Hyperbolic Discounting

Abstract

Low probability catastrophic climate change can have a significant influence on policy under hyperbolic discounting. We compare the set of Markov Perfect Equilibria (MPE) to the optimal policy under time-consistent commitment. For some initial levels of risk there are multiple MPE; these may involve either excessive or insufficient stabilization effort. These results imply that even if the free-rider problem amongst contemporaneous decision-makers were solved, there may remain a coordination problem amongst successive generations of decision-makers. A numerical example shows that under plausible conditions society should respond vigorously to the threat of climate change.

Climate policy when the distant future matters: Catastrophic events with hyperbolic discounting

Larry Karp* Yacov Tsur[◇]

February 2, 2007

Abstract

Low probability catastrophic climate change can have a significant influence on policy under hyperbolic discounting. We compare the set of Markov Perfect Equilibria (MPE) to the optimal policy under time-consistent commitment. For some initial levels of risk there are multiple MPE; these may involve either excessive or insufficient stabilization effort. These results imply that even if the free-rider problem amongst contemporaneous decision-makers were solved, there may remain a coordination problem amongst successive generations of decision-makers. A numerical example shows that under plausible conditions society should respond vigorously to the threat of climate change.

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

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

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

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

1 Introduction

Low probability events – those with low hazard rates – are unlikely to occur until the distant future. A constant (non-negligible) discount rate makes this kind of event appear insignificant, even if it causes substantial and long-lasting future damage. We study the set of Markov perfect equilibria (MPE) to a model with catastrophic events and hyperbolic discounting. This model provides a new way to analyze policy where low probability irreversible events are important, as with climate change. The analysis increases the scope of application of hyperbolic discounting, and it illuminates the relation between hyperbolic discounting and intertemporal coordination games.

We use the model to study climate change policy. An example helps to illustrate the problem. Suppose that an “event” reduces in perpetuity the annual flow of utility by 1 util, and the constant annual hazard rate is h . A policy eliminates the risk at a *flow cost* (caused by abatement efforts or a perpetual reduction of economic activity) of x and the constant annual pure rate of time preference is r . The largest flow cost that society would accept in order to eliminate the hazard is $x^* = \frac{h}{h+r}$. This example shows that if the hazard rate is small relative to the discount rate – the likely case when the event is low probability and we use a constant discount rate – society is not willing to spend much to eliminate the risk.¹

Climate change modelers have used both constant and hyperbolic discounting. The near-zero pure rate of time preference that Stern (2006) uses is not consistent with the empirical evidence that individuals have non-negligible discount rates (Frederick et al. 2002), and the fact of positive real

¹The expected present value of the value-at-risk is $\frac{1}{h+r}$ and the present value under the policy is $\frac{1-x}{r}$. Equating these values and solving for x gives the largest flow cost that society would accept in order to eliminate the hazard. If the annual discount rate is 5% and the probability of the event occurring within a century is 5% then $x^* = 0.01015$. 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.

interest rates. A near-zero discount rate also implies that current generations should make implausible sacrifices for the future (Nordhaus 2006). Models that use constant rates similar to current medium run real interest rates assume that these accurately reflect our long run discount rate, and essentially ignore the future beyond a century or so. A compromise, using a smaller constant discount rate, is vulnerable to the criticism that externalities should be modeled explicitly rather than by reducing the discount rate.

Nordhaus (1999) and Mastrandrea and Schneider (2001) use a declining pure rate of time preference, thereby respecting the evidence concerning short run rates and the ethical imperative to put non-negligible weight on future welfare. The declining pure rate of time preference is reasonable: We may feel closer to our children than to our unborn grandchildren, but it is less likely that we make a distinction between the 10th and the 11th future generation. However, these papers assume that the current policy-maker can make commitments about future actions. Since the policy horizon extends for centuries, this is a strong assumption. In addition, the full commitment outcome under hyperbolic discounting (typically) exhibits procrastination, which might be confounded with the sensible proposal to defer action until technological improvements make it cheaper.^{2, 3}

Climate scientists have identified several low probability catastrophic consequences of climate change, including a sudden rise in sea level, a mass extinction of species, or a weakening of the Thermohaline Circuit (the THC, which moderates European climate) (Chichilnisky and Heal 1993, IPCC 2001, Alley et al. 2003, Thomas et al. 2004, Millennium Ecosystem Assessment 2005). Two recent studies estimate a 30% chance of the collapse or sig-

²Cropper and Laibson (1999) discuss hyperbolic discounting in the climate change context. Karp (2005) uses quasi-hyperbolic discounting in a deterministic setting under the assumption that regulators use Markov Perfect policies.

³A large literature addresses the valuation of future welfare, motivated by considerations of inter-generational equity and sustainability (Solow 1974, Hartwick 1977, Chichilnisky 1996, Arrow 1999, Asheim and Buchholz 2004) and uncertainty (Weitzman 2001, Gollier 2002, Dasgupta and Maskin 2005).

nificant weakening of the THC within the next century under Business as Usual (BAU) (Challenor et al. 2006, Schlesinger et al. 2006).⁴ Some climate scientists think that the risk is negligible (Wunsch 2006).

Models of dynamic management under event uncertainty (Cropper 1976, Clarke and Reed 1994, Tsur and Zemel 1996, 1998) provide our starting point. We extend these models by replacing constant discounting with hyperbolic discounting, using methods from Karp (2007).⁵ Our stationary model (where the pure rate of time preference equals the social discount rate) incorporates risk and hyperbolic discounting using realistic behavioral assumptions.

The next section describes our general model and a binary action specialization, in which the regulator chooses either to stabilize the hazard or to follow BAU. Section 3 provides a benchmark, in which the current regulator can make a commitment to all future policies. Section 4, which contains our theoretical contribution, studies the set of 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 in the binary action setting. 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 may 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. Section 6 shows numerically the importance of risk, commitment, and discounting.

⁴There has been substantial work in assessing the likely economic costs of gradual climate change (Chakravorty et al. 1997, Mendelsohn 2003, Schlenker et al. 2006).

⁵Harris 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.

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 due to the event occurrence. 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

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

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

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

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

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

where

$$y(t) = \int_0^t h(\tau) d\tau. \quad (4)$$

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

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

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

$$\begin{aligned} & \int_0^T \theta(t) u(c + \Delta, w(t)) dt + \int_T^\infty \theta(t) u(c, 0) dt = \\ & \int_0^T \theta(t) (u(c + \Delta, w(t)) - u(c, 0)) dt + \int_0^\infty \theta(t) u(c, 0) dt = \\ & \int_0^T \theta(t) U(w(t), c, \Delta) dt + \text{constant}, \end{aligned} \quad (6)$$

where $U(w(t), c, \Delta) \equiv u(c + \Delta, w(t)) - u(c, 0)$. We refer to $U(0, c, \Delta)$ as the “value-at-risk”—the change in the flow of utility due to the event occurrence if society takes no action to stabilize the risk ($w = 0$). 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 it 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

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

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 GHG stock as the state variable. The control variable prior to the catastrophe, $w(t)$, affects the flow of utility at time t , $U(w(t))$, as well as the evolution of the hazard rate:

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

(We suppress the time argument when there is no confusion.) The optimal policy $w(t)$ maximizes (7) subject to (5) and (8).

2.1 A binary action specialization

We emphasize a *binary action* specialization of the above general model: society can either stabilize the stock of greenhouse gasses at the current level, thus stabilizing the hazard, 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

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

We let $x \equiv 1 - \frac{U(1)}{U(0)}$ represent the fractional reduction in the value-at-risk under stabilization (prior to the catastrophe). Hereafter we refer to x as simply the “cost of stabilization”. This parameter summarizes all of the pertinent information regarding the utility function and the parameters Δ and X . In the case where $u(\cdot)$ is linear, $X = x$. For example, with linear utility, 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 assume that under BAU the hazard rate approaches the steady state level a at a constant rate ρ :

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

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

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

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 0.⁶ 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 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 positive and finite, non-exponential discounting causes the policy announced at time 0 to be time-inconsistent. A future regulator would want to deviate from the policy announced by the current regulator, 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

$$F^{BAU}(t) = 1 - \exp\left(\frac{-at\rho + (a - h_0)(1 - e^{-\rho t})}{\rho}\right). \quad (12)$$

⁶Different growth functions for the hazard are discussed in a separate note, available upon request.

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

$$V^B(h_0) \equiv U(0) \int_0^\infty (1 - F^{BAU}(t)) \theta(t) dt = U(0) \nu(h_0), \quad (13)$$

where

$$\nu(h) \equiv \int_0^\infty \exp\left(\frac{-a\rho t + (a-h)(1-e^{-\rho t})}{\rho}\right) \theta(t) dt. \quad (14)$$

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

$$V^S(h_0) \equiv U(1) \int_0^\infty e^{-h_0 t} \theta(t) dt = U(1) \xi(h_0) \quad (15)$$

where

$$\xi(h) \equiv \int_0^\infty e^{-ht} \theta(t) dt = \frac{(1-\beta)\gamma + h + \beta\delta}{(\delta+h)(h+\gamma)}. \quad (16)$$

The regulator chooses to stabilize if and only if $V^S \geq V^B$. This inequality is equivalent to $\frac{U(1)}{U(0)} \geq \lambda(h_0)$, where

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

(We assume that a tie results in stabilization.) Noting that $\frac{U(1)}{U(0)} = 1 - x$, the condition $V^S \geq V^B$ holds if and only if $x \leq \bar{x}^C(h_0)$, where

$$\bar{x}^C(h) \equiv 1 - \lambda(h). \quad (18)$$

The superscript on $\bar{x}^C$ is a mnemonic for “commitment”, and the over-bar indicates that this variable is an upper bound.

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. (Another view is that restricted

commitment is an unsatisfactory equilibrium concept because the restriction is ad hoc.)

4 Markov Perfect Equilibrium

This section studies the Markov Perfect Equilibria (MPE), where a regulator at a point in time is unable 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 a MPE trajectory, the current regulator cannot commit future generations to a specific course of action but she can *influence* the 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 use results from Karp (2007), 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, $t = i\varepsilon$, $i = 0, 1, 2, 3, \dots$. 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 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 game and formally letting $\varepsilon \rightarrow 0$, 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 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.

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 $\chi(h) \equiv \hat{\chi}(z)$ satisfies the (generalized)*

dynamic programming equation (DPE):

$$K(h) + (\gamma + h) V(h) = \max_{w \in \Omega} \{U(w) + g(h, w) V'(h)\}, \quad (19)$$

with the “side condition”

$$K(h) \equiv (\delta - \gamma) (1 - \beta) \int_0^\infty e^{-(\delta s + y(s))} U(\chi(h_s)) ds. \quad (20)$$

Remark 1. The control rule that maximizes the right-hand side of equation (19) 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 (20). To demonstrate the latter case, note that for $\beta = 0$, $K(h) = (\delta - \gamma) \int_0^\infty e^{-(\delta s + y(s))} U(\chi(h_s)) ds = (\delta - \gamma) V(h)$. Substituting this equation into (19) 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 a MPE decision rule. Using the equilibrium condition (19) and the convention that in the event of a tie the regulator chooses stabilization, in the binary setting χ satisfies

$$\chi(h) = \begin{cases} 1 & \text{if } U(1) \geq U(0) + \rho(a - h)V'(h) \\ 0 & \text{if } U(1) < U(0) + \rho(a - h)V'(h) \end{cases}. \quad (21)$$

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. With this belief, $V(h) = V^S(h)$ and $V'(h) = V^{S'}(h) = U(1)\xi'(h)$, where $V^S(h)$ and $\xi(h)$ are defined in equations (15) and (16), respectively, and

$$\xi'(h) = - \int_0^\infty t e^{-ht} \theta(t) dt = - \left(\frac{\beta}{(h + \gamma)^2} + \frac{1 - \beta}{(h + \delta)^2} \right). \quad (22)$$

Thus, using the equilibrium rule (21), $U(1) \geq U(0) + \rho(a - h)U(1)\xi'(h)$ must hold for stabilization to be a MPE. Defining

$$\pi(h) \equiv \frac{1}{1 - \rho(a - h)\xi'(h)}, \quad (23)$$

the condition under which perpetual stabilization is a MPE can be stated as $\frac{U(1)}{U(0)} \geq \pi(h)$.

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

$$\sigma(h) \equiv 1 + \rho(a - h)\nu'(h), \quad (24)$$

with $\nu(h)$ given by equation (14) and

$$\nu'(h) = - \int_0^\infty \frac{1 - e^{-\rho t}}{\rho} \exp \left(-at + (a - h)\frac{1 - e^{-\rho t}}{\rho} \right) \theta(t) dt, \quad (25)$$

the condition under which perpetual BAU is a MPE can be written as $\frac{U(1)}{U(0)} < \sigma(h)$.

We summarize properties of $\pi(h)$ and $\sigma(h)$ in

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 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 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*

$$\pi(h) < \sigma(h), \quad h \in (0, a). \quad (26)$$

Under equation (26), there exists a MPE with perpetual stabilization ($w \equiv 1$) if and only if the initial condition $h_0 = h$ satisfies

$$x < \bar{x}^S(h) \equiv 1 - \pi(h); \quad (27)$$

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

$$x > \underline{x}^B(h) \equiv 1 - \sigma(h). \quad (28)$$

Figure 1 illustrates Proposition 3. The figure shows $1 - \sigma(h)$ and $1 - \pi(h)$ with $\pi(h) < \sigma(h)$ for $h \in (0, a)$. 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 discounting – a possibility previously noted by Krusell and Smith (2003) and Karp (2005, 2007). 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

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.

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”. 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.

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, and these equilibria are in general not Pareto efficient.

Proposition 3 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.⁷ We use the following definition

$$\Theta(h) \equiv \frac{\rho(a-h) \left(\frac{\beta}{\gamma+h} + \frac{1-\beta}{\delta+h} \right)}{h + \beta\gamma + \delta(1-\beta) + \rho(a-h) \left(\frac{\beta}{\gamma+h} + \frac{1-\beta}{\delta+h} \right)}. \quad (29)$$

Proposition 4. *Suppose that Condition (26) 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*

$$\Theta(h) < x < 1 - \pi(h). \quad (30)$$

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

Recall that x , the “cost of stabilization”, equals the utility cost of stabilizing the hazard (or the atmospheric GHG concentration) as a fraction of the value-at-risk $U(0) = u(c + \Delta) - u(c)$. Relation (30) defines the lower and upper bounds of x for a delayed stabilization MPE to exist. We verify in the appendix that

$$1 - \pi(h) - \Theta(h) = \frac{(\delta - \gamma)^2(2h + \gamma + \delta)}{(h + \gamma)^2(h + \delta)^2} \beta(1 - \beta). \quad (31)$$

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.

⁷From 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.

5 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 , so the distant future is “heavily discounted”. Following the standard procedure to obtain the DPE, or invoking Remark 2, we have the following DPE:

$$(\delta + h) V(h) = \max_{w \in \{0,1\}} \{U(w) + \rho(a - h)(1 - w) V'(h)\}. \quad (32)$$

Let $\pi^0(h)$ and $\sigma^0(h)$ denote the functions $\pi(h)$ and $\sigma(h)$ (defined in equations (23) and (24)) 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 $\beta \rightarrow 0$. 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 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 benchmark under restricted commitment.

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 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.

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 proba-

bility 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 for these two scenarios.

Table 1: Hazard parameter values.

a	0.00693147	
ρ	0.00544875	
	<u>Optimistic</u>	<u>Pessimistic</u>
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 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.

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

The parameter values in Tables 1 and 2 do not satisfy the sufficient condition in part iv of 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 $\sigma(h) > \pi(h)$, so the necessary and sufficient condition for existence of the MPE (cf. Proposition 3) is satisfied. The fact that the graphs of λ and π intersect implies that in some MPE there is excessive stabilization, relative to the restricted commitment benchmark.

Equations (18), (27), and (28) define the three critical values of x : $\bar{x}^c = 1 - \lambda$ is the maximum cost that the decisionmaker with restricted commitment is willing to incur in order to stabilize; $\bar{x}^S = 1 - \pi$ is the maximum cost that is consistent with stabilization in a MPE; $\underline{x}^B = 1 - \sigma$ is the minimum cost that is consistent with BAU in a MPE. Table 3 shows these thresholds under the different discounting and hazard scenarios.

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 costs less than 1.34% of the value-at-risk. If stabilization costs between 1.34% and 16.69% of the value-at-risk both stabilization and BAU

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

are MPE. If stabilization costs 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 it costs no more than 12.37% of the value-at-risk.

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. In this respect, there can always be insufficient stabilization in a MPE.

Comparison of the optimistic and pessimistic scenarios (low and high h_0 , respectively) illustrates that a low initial value of h encourages stabilization: it expands the range of x for which stabilization is an (or the only) equilibrium. We discussed the reason for this result in Section 2.1. The growth

Table 3: Policy bounds.

Discounting mode	$\underline{x}^B$	$\bar{x}^S$	$\bar{x}^C$
		<u>Optimistic</u>	
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
		<u>Pessimistic</u>	
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

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 commit-

ment) 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 4 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 $\gamma = .0005$. 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 lie between 1% and 1.8%.

Table 4: “Constant equivalent” discount 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-

Figure 4: The graphs of $1 - \pi(h)$, $\Theta(h)$, and $1 - \sigma(h)$ with $\gamma = 5 \times 10^{-4}$.

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 4 provides some evidence on this point.

Figure 4 shows the graphs of $1 - \pi(h)$, $\Theta(h)$ and $1 - \sigma(h)$ under hyperbolic discounting with $\gamma = 5 \times 10^{-4}$. For any given initial hazard h , a delayed stabilization MPE exists when x falls between $\Theta(h)$ and $1 - \pi(h)$ (Proposition 4). The figure shows that the delayed stabilization region $(\Theta(h), 1 - \pi(h))$ is a strict subset of $(1 - \sigma(h), 1 - \pi(h))$, which is the region where both perpetual BAU and perpetual stabilization are MPE (Proposition 3). Since the two sets are nearly the same, in “most” cases where multiple “single action MPE” exist, there also exist “delayed stabilization MPE”. The existence of a MPE with perpetual stabilization is a necessary condition for existence of a MPE with delayed stabilization.

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 due to climate change is typically assumed to be gradual, allowing for adjustment and adaptation. 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 main reasons for this view: (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 long-run 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 policies 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 stabilization is too expensive. 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. We then specialized to a binary action model, in which at each point in time the regu-

lator 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. These equilibria involve either perpetual BAU or perpetual stabilization. We provided a necessary and sufficient condition under which there is also an equilibrium in which policymakers follow BAU until the hazard reaches a threshold, and then switch to stabilization. By considering a limiting problem with constant discounting, we showed that there can be MPE with excessive stabilization (relative to the social optimum).

We emphasized the situation where the event is “low-risk”, i.e. the hazard rate is much smaller than anything that (most) economists would recognize as a plausible short run social discount rate. A model of constant discounting (at a non-negligible rate) has little that is useful to say about such events. Our numerical example used a hyperbolic discount rate that remains above 4% for a century into the future, and eventually falls to a level below the steady state hazard.

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 low. Under perpetual stabilization the risk remains constant; under perpetual BAU it increases at a diminishing rate, reaching half way to its maximal level within a century. Under plausible parameter values, it is optimal to forgo a substantial fraction 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 most integrated assessment models.

The numerical experiments illustrate the possibility that there can be MPE that result in excessive stabilization, relative to the restricted commitment benchmark. However, there always exist utility parameters for

which there is a MPE that results in too little stabilization, relative to this benchmark. For some combinations of the utility and risk-related parameter values, *all* MPE result in too little stabilization effort.

The free-rider problem amongst decision-makers in the current generation presents a serious and well-understood impediment to optimal climate policy. The present analysis illuminates the problem of 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, Washington, DC, 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.
- 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. 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, Cambridge, UK, 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, Washington, DC, p. 163–172.
- Dasgupta, P. and Maskin, E.: 2005, Uncertainty and hyperbolic discounting, *The American Economic Review* **95**, 1290–1299.
- 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.
- 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, Cambridge, UK.
- Karp, L.: 2005, Global warming and hyperbolic discounting, *Journal of Public Economics* **89**, 261–282.

- Karp, L.: 2007, Non-constant discounting in continuous time, *Journal of Economic Theory* (132), 557–568.
- Krugman, P.: 1991, History versus expectations, *Quarterly Journal of Economics* (106), 651–67.
- Krusell, P. and Smith, A.: 2003, Consumption-saving decisions with quasi-geometric discounting, *Econometrica* **71**(1), 365–75.
- 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 indeterminacy of equilibrium, *Quarterly Journal of Economics* **106**, 587–597.
- Mendelsohn, R.: 2003, Assessing the market damages from climate change, in J. M. Griffin (ed.), *Global Climate Change*, Edward Elgar, Cheltenham, UK, pp. 92–113.
- Millennium Ecosystem Assessment: 2005, *Millennium Ecosystem Assessment*, Island Press, Washington, DC.
URL: <http://maweb.org/en/Products.Global.Condition.aspx>
- Nordhaus, W.: 2006, The stern review on the economic of climate change.
URL: <http://nordhaus.econ.yale.edu/SternReviewD2.pdf>
- Nordhaus, W. D.: 1999, Discounting and public policies that affect the distant future, in P. R. Portney and J. P. Weyant (eds), *Discounting and intergenerational equity*, Resources for the Future, Washington, DC.
- 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, Cambridge, UK, pp. 37–48.
- Skiba, A.: 1978, Optimal growth with a convex-concave production function, *Econometrica* **46**, 527–39.
- Solow, R. M.: 1974, Intergenerational equity and exhaustible resources, *The Review of Economic Studies* **41**, 29–45.
- Stern, N.: 2006, Stern review on the economics of climate change, *Technical report*, HM Treasury, UK.
- 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.
- Weitzman, M. L.: 2001, Gamma discounting, *American Economic Review* **91**, 260–271.
- Wunsch, C.: 2006, Abrupt climate change: An alternative view, *Quaternary Research* **65**, 191–203.

Appendix: Proofs

Proof of Proposition 1 (i) This claim follows from differentiating the functions $\nu(h)$ and $\xi(h)$ and by inspection. (ii) We begin with

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

where $y(t, h)$ is a specialization of $y(t)$, defined in (4), when the hazard process under BAU evolves according to equation (11). From equations (14), (16) and (33),

$$\nu(h) - \xi(h) = \int_0^\infty \theta(t) (e^{-y(t, h)} - e^{-ht}) dt. \quad (34)$$

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 (34) is negative. (iii) This claim is merely a summary of the derivation in the text above equation (18).

(iv) (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.

(v) Using (13) and (15), we express $\lambda(h)$ as

$$\lambda(h) = \frac{\int_0^\infty e^{-y(t, h)} \theta(t) dt}{\int_0^\infty e^{-ht} \theta(t) dt}. \quad (35)$$

Using equation (33) we have

$$y_h(t, h) \equiv \partial y(t, h) / \partial h = \frac{1 - e^{-\rho t}}{\rho}. \quad (36)$$

The argument h in $y(t, h)$ is the initial hazard. Differentiating (35) with respect to h , we see that $\lambda'(h) > 0$ if and only if

$$\int_0^\infty e^{-y(t, h)} \theta(t) dt \int_0^\infty e^{-ht} t \theta(t) dt > \int_0^\infty e^{-ht} \theta(t) dt \int_0^\infty e^{-y(t, h)} y_h(t, h) \theta(t) dt. \quad (37)$$

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 (36), we express (37) as

$$\begin{aligned} & \left(\frac{\beta}{(h+\gamma)^2} + \frac{1-\beta}{(h+\delta)^2} \right) \int_0^\infty e^{-y(t, h)} \theta(t) dt > \\ & \left(\frac{\beta}{h+\gamma} + \frac{1-\beta}{h+\delta} \right) \int_0^\infty e^{-y(t, h)} \theta(t) \frac{1-e^{-\rho t}}{\rho} dt. \end{aligned} \quad (38)$$

Since $\delta > \gamma$, the right-hand side of inequality (38) is smaller than

$$\left(\frac{\beta}{(h+\gamma)^2} + \frac{1-\beta}{(h+\delta)^2} \right) \int_0^\infty e^{-y(t, h)} \theta(t) \frac{(h+\delta)(1-e^{-\rho t})}{\rho} dt. \quad (39)$$

Thus, it suffices to show that the left-hand side of (38) exceeds (39), 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 (2007). 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 ? to our setting, and using the hyperbolic discount factor in equation (1), yields the generalized DPE

$$\hat{K}(z) + \gamma W(z) = \max_{w \in \Omega} (e^{-y(t)} U(w(t)) + W_h g + W_y h), \quad (40)$$

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

$$\hat{K}(z) = (\delta - \gamma)(1 - \beta) \int_0^\infty e^{-(\delta t + y(t))} U(\hat{\chi}(z)) dt \quad (41)$$

is implied by equation (4) and Remark 2 of Karp (2007)

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 (40), cancelling e^{-y} and rearranging, yields equation (19). 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 $\exp(-\int_t^s h(\tau)d\tau) = \exp(-y(s) + y(t))$. Use this fact and the trial solution to rewrite equation (41) as

$$\begin{aligned} K(h(t)) &= (\delta - \gamma)(1 - \beta) e^{y(t)} \int_t^\infty e^{-\delta(s-t)} \exp(-\int_t^s h(\tau)d\tau) e^{-y(t)} U(\chi(h(s))) ds \\ &= (\delta - \gamma)(1 - \beta) \int_t^\infty e^{-\delta(s-t)} \exp(-\int_t^s h(\tau)d\tau) U(\chi(h(s))) ds \end{aligned} \quad (42)$$

Setting $t = 0$ in equation (42) produces equation (20). $\square$

Proof of Lemma 1 Define

$$\varpi(h) \equiv \pi(h)^{-1} = 1 - \rho(a - h)\xi'(h). \quad (43)$$

Differentiating, noting (22), we obtain

$$\varpi'(h) = \rho\xi'(h) - \rho(a - h)\xi''(h) < 0. \quad (44)$$

Thus,

$$\pi'(h) = -\varpi'(h)/\varpi(h)^2 > 0. \quad (45)$$

Differentiating (24), noting (25), gives

$$\sigma'(h) = -\rho\nu'(h) + \rho(a - h)\nu''(h) > 0. \quad (46)$$

To establish $\sigma''(h) < 0$, use equation (25) and differentiate twice to obtain $\nu'''(h) < 0$. Differentiating equation (46) gives

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

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

Proof of Proposition 3 We first establish sufficiency of inequality (26) using a constructive proof, which also establishes the claims associated with inequalities (27) and (28). We then show necessity of inequality (26) using a proof by contradiction.

Sufficiency Suppose that $\sigma > \pi$ for $h \in (0, a)$. We 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 (27). 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 (15), and

$$V^{S'}(h) = U(1)\xi'(h) \tag{47}$$

with $\xi'(h)$ given by equation (22).

Using equation (19) (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

$$U(1) \geq U(0) + \rho(a-h)U(1)\xi'(h) \tag{48}$$

or

$$\frac{U(1)}{U(0)} \geq \pi(h). \tag{49}$$

If inequality (49) is satisfied with *strict inequality* (as the Proposition requires) at the current time, then regardless of whether the current regulator

uses stabilization or 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 (49) is not satisfied, then clearly perpetual stabilization is not an equilibrium. We consider below the case where the weak inequality (49) 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 (13). 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

$$\frac{U(1)}{U(0)} < \sigma(h) \equiv 1 + \rho(a - h)\nu'(h), \quad (50)$$

establishing condition (28).

To complete the demonstration that perpetual stabilization is an equilibrium, it is necessary to confirm that if equation (28) 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 (49) is satisfied with equality. We want to show that in this case, stabilization is not an equilibrium action. Suppose to the contrary that it is optimal to stabilize when inequality (49) is satisfied with equality. From equation (21), 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 (49) is satisfied with equality, it must be the case that $V'(h) = V^{B'}(h) = U(0)\nu'(h)$. The assumption that $\sigma(h) > \pi(h)$ implies that $\pi(h)$ lies in the region where perpetual BAU 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, consisting of two parts, to establish necessity. The first part shows that $\sigma(h) < \pi(h)$ cannot hold, and the second part shows that it cannot be the case that $\sigma(h) = \pi(h)$ at any points in $(0, a)$.

For the first part, suppose that for some interval $\sigma(h) < \pi(h)$. Figure 5 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

Figure 5: Graphs of $\sigma(h)$ and $\pi(h)$ that do not satisfy inequality (26).

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 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).

For the second part, suppose that $\sigma(h) \geq \pi(h)$ with equality holding at one or more points in $(0, a)$ (that is, the graphs are tangent at one or more points). Let $\hat{h}$ be such a point. The argument above under “sufficiency” establishes that if $\frac{U(1)}{U(0)} = \pi(\hat{h})$, then at $h = \hat{h}$ (where equation (49) holds with equality) neither perpetual stabilization nor perpetual BAU are MPE. The only remaining possibility would be to follow BAU for a time and then switch to stabilization in perpetuity. However, that cannot be an equilibrium trajectory, because the initial period of BAU drives the h above $\hat{h}$, where $\frac{U(1)}{U(0)} < \pi(h)$, so the subsequent stabilization period cannot be part of a MPE. Therefore, at $h = \hat{h}$ there is no MPE if $\frac{U(1)}{U(0)} = \pi(\hat{h})$.

□

Proof of Proposition 4 We use the following definition

$$h_\pi(x) \equiv \begin{cases} \pi^{-1}(1-x) & \text{for } x \in [0, 1 - \pi(0)) \\ 0 & \text{for } x \in [1 - \pi(0), 1] \end{cases} \quad (51)$$

Hazard rates that satisfy $h > h_\pi(x)$ lie above the curve $1 - \pi$ in Figure 1.

(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 $h(0) \geq h_\pi$, 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 $h(0) < h_\pi$; equivalently, $x < 1 - \pi(h)$. 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, $A = \{h \mid h_a \leq h < \tilde{h}\}$ and $B = \{h \mid \tilde{h} \leq h < h_b\}$, where $h_a < \tilde{h} < h_b < h_\pi$. 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 $[0, h_\pi)$ (in view of Proposition 3). In addition, h remains constant when the regulator stabilizes. Therefore, any subset of the interval $[0, h_\pi)$ 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 set A with positive measure.

Suppose (provisionally) that the 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 $h \in A \cup B$.

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

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 (21), this action is part of an equilibrium if and only if

$$U(0) - U(1) > -\rho(a - h) V_h(h; \tilde{h}). \quad (52)$$

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 until reaching the threshold $\tilde{h}$, 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

$$\tilde{h} = a - (a - h)e^{-\rho T} \Rightarrow \quad (53)$$

$$T(\tilde{h}; \tilde{h}) = 0 \quad \text{and} \quad \frac{dT}{dh} = \frac{-1}{\rho(a - h)}. \quad (54)$$

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

$$\frac{dy(\tau)}{dh} = \frac{d \int_0^\tau h(s) ds}{dh} = \int_0^\tau \frac{dh(s)}{dh} ds = \int_0^\tau e^{-\rho s} ds = \frac{1 - e^{-\rho \tau}}{\rho}. \quad (55)$$

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

$$\begin{aligned} \frac{dy(\tau)}{dh} &= \frac{d(\int_0^T h(s)ds + \tilde{h}(\tau-T))}{dh} = \\ \int_0^T \frac{dh(s)}{dh} ds + \left(h(T) - \tilde{h} \right) \frac{dT}{dh} &= \int_0^T e^{-\rho s} ds \end{aligned} \quad (56)$$

The last equality uses the fact that $h(T) = \tilde{h}$, from the definition of T . Using equation (53) and (54), 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 (55) and the definition of $y(\tau)$, we have

$$\begin{aligned} y(T) &= \int_0^T h(s)ds \Rightarrow \\ \frac{dy(T)}{dT} &= h(T) + \int_0^T \frac{dh(s)}{dh} \frac{dh}{dT} ds \end{aligned} \quad (57)$$

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

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

and its derivative with respect to h (using equation (55)) is

$$\begin{aligned} V_h(h; \tilde{h}) &= (U(0) - U(1)) e^{-y(T)} \theta(T) \frac{dT}{dh} + \\ \int_0^T \frac{d(e^{-y(\tau)})}{dh} \theta(\tau) U(0) dt + \int_T^\infty \frac{d(e^{-y(\tau)})}{dh} \theta(\tau) U(1) dt \\ &= \frac{-(U(0) - U(1))}{\rho(a-h)} e^{-y(T)} \theta(T) - \end{aligned} \quad (58)$$

$$\left(\int_0^T \left(\frac{1-e^{-\rho\tau}}{\rho} \right) e^{-y(\tau)} \theta(\tau) U(0) dt + \int_T^\infty \left(\frac{1-e^{-\rho T}}{\rho} \right) e^{-y(\tau)} \theta(\tau) U(0) dt \right).$$

Using this expression, we can write the optimality condition (52) as

$$\begin{aligned} U(0) - U(1) &> (U(0) - U(1)) e^{-y(T)} \theta(T) + \\ \rho(a-h) &\left(\int_0^T \left(\frac{1-e^{-\rho\tau}}{\rho} \right) e^{-y(\tau)} \theta(\tau) U(0) dt + \int_T^\infty \left(\frac{1-e^{-\rho T}}{\rho} \right) e^{-y(\tau)} \theta(\tau) U(0) dt \right). \end{aligned} \quad (59)$$

It is convenient to treat T as the independent variable, recognizing that the initial condition h is a function of T (from equation (53)): $h = h(T)$. The existence of a set A with positive measure requires that inequality (59) holds 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 (59) is

$$(U(0) - U(1)) - (U(0) - U(1))(\tilde{h} + r(0))T + o(T). \quad (60)$$

This expansion uses equations (2) and (57) and the fact that $\theta(0) = 1$. Using the fact that $1 - e^{-\rho T} = 0$ at $T = 0$, the first order Taylor expansion of the second term on the right side of inequality (59) is

$$\begin{aligned} & \rho(a - \tilde{h})T \int_0^\infty e^{-y(\tau)}\theta(\tau)U(1)d\tau + o(T) = \\ & \rho(a - \tilde{h})T \int_0^\infty e^{-\tilde{h}\tau}\theta(\tau)U(1)d\tau + o(T) = \\ & \rho(a - \tilde{h})T \frac{(1-\beta)\gamma + \beta\delta + \tilde{h}}{(\tilde{h} + \gamma)(\tilde{h} + \delta)}U(1) + o(T). \end{aligned} \quad (61)$$

Substituting expressions (60) and (61) into inequality (59), dividing by T and letting $T \rightarrow 0$ (from above) produces the inequality

$$(U(0) - U(1))(\tilde{h} + r(0)) > \rho(a - \tilde{h}) \frac{(1-\beta)\gamma + \beta\delta + \tilde{h}}{(\tilde{h} + \gamma)(\tilde{h} + \delta)}U(1). \quad (62)$$

Using $x \equiv 1 - \frac{U(1)}{U(0)}$ and $r(0) = \beta\gamma + \delta(1-\beta)$ (from equation (2)), and replacing $\tilde{h}$ with h , inequality (62) can be expressed as

$$\frac{x}{1-x}(h + \beta\gamma + \delta(1-\beta)) > \rho(a - h) \left(\frac{\beta}{h + \gamma} + \frac{1-\beta}{h + \delta} \right) \quad (63)$$

or, equivalently,

$$x > \Theta(h), \quad (64)$$

where $\Theta(h)$ is defined in equation (29), establishing part (iii).

(iv) Using

$$-\xi'(h) = \int_0^\infty t e^{-ht} \theta(t) dt = \frac{\beta}{(h+\gamma)^2} + \frac{1-\beta}{(h+\delta)^2},$$

we express $\pi(h)$, defined in (23), as

$$\pi(h) = \frac{1}{1 + \rho(a-h) \left(\frac{\beta}{(h+\gamma)^2} + \frac{1-\beta}{(h+\delta)^2} \right)}.$$

Expanding $1 - \pi(h) - \Theta(h)$ as a polynomial in β and collecting terms gives (after some algebraic manipulations) equation (31). $\square$

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 a solution to the optimal control problem exists. Therefore, $\sigma^0(h) \geq \pi^0(h)$.

Constant discounting occurs when $\beta = 0$ or $\beta = 1$ or $\gamma = \delta$. It is clear from equation (31) that condition (30) is not satisfied in any of these cases, implying, in view of Proposition 4 Part (iii), that there can be no equilibrium with delayed stabilization.

We now turn to the main part of the proof. For h close to but smaller than a , $\sigma^0(h) > \pi^0(h)$. (We established the weak inequality above; here we need the strict inequality.) 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. For parameters that satisfy this inequality, in view of Proposition 3, the DPE (32) 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 (32) under stabilization and BAU, respectively, and let $V(h) = \max\{V^S(h), V^B(h)\}$ 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

$$\begin{aligned} V^B(h) &= \frac{1}{\delta+h} \max\{U(1), U(0) + \rho(a-h)V_h^B(h)\} \\ &= \frac{1}{\delta+h} (U(0) + \rho(a-h)V_h^B(h)) > \frac{1}{\delta+h} U(1). \end{aligned} \quad (65)$$

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

$$\begin{aligned} V^S(h) &= \frac{1}{\delta+h} \max\{U(1), U(0) + \rho(a-h)V_h^S(h)\} \\ &= \frac{1}{\delta+h} U(1) \geq \frac{1}{\delta+h} (U(0) + \rho(a-h)V_h^S(h)). \end{aligned} \quad (66)$$

From (65) and (66) 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. $\square$