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 2008

Paper 1062

Time perspective and climate change
policy

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/1062

Copyright ©2008 by the authors.

Time perspective and climate change policy

Abstract

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

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

Time perspective and climate change policy^{*}

Larry Karp[◇] Yacov Tsur[§]

July 2008

Abstract

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

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

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

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

^{*}We benefitted from comments from seminar participants at UC Santa Barbara, UC Berkeley (ARE), University of Southampton, Hebrew University of Jerusalem, and Iowa State University.

[◇]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

Global warming is arguably the most pressing environmental problem we currently face. The long time-span of climate policies and the uncertainties of climate change and related damages complicate policy negotiations. Cost-benefit evaluations are sensitive to discounting due to the long delay between paying the cost of a climate policy and reaping the benefits of reduced damages. The Stern Review on Climate Change (Stern 2007) gave rise to a controversy that turns largely on the choice of parameters underlying the social discount rate (Dasgupta 2007a,b, Nordhaus 2007, Weitzman 2007). A constant social discount rate that reflects market rates, and thus accounts for the opportunity costs of public funds, gives little weight to the welfare of generations in the distant future. The Stern Review’s defense of a low pure rate of time preference rests on concern for inter-generational equity.

We construct a positive model based on preferences that are consistent with both observed market behavior and a measure of intergenerational equity. Individual agents in this model view their world in the manner described by Ramsey (1931, p.291) “...[their] picture of the world is drawn in perspective.... [applied] not merely to space but also to time.” This perspective gives rise to hyperbolic discounting at the level of the individual agent. These agents care less about future generations’ utility than about their own, so their pure rate of time preference (at the generational time scale) is positive and over some interval may be large. However, they make smaller distinctions between successive generations in the distant future, compared to successive generations in the near future, so their pure rate of time preference falls. Arrow (1999) describes this attitude as “agent-relative ethics”. Cropper et al. (1994) and Section 8 of Heal (2005) provide empirical evidence that individuals discount utility in this manner. Hyperbolic discounting leads to a model that is flexible enough to produce short and medium run social discount rates equal to market discount rates, and which also gives non-negligible weight to the wellbeing of distant generations.

The optimal program for any generation is time-inconsistent under hyperbolic discounting (Strotz 1956, Phelps and Pollak 1968). This time-inconsistency is a plausible feature of the policy problem: politicians, like other mortals, tend to procrastinate in solving difficult problems. Because of the long time scale over which policies must be implemented, we focus on Markov Perfect equilibria (MPE). In a MPE the current generation cannot commit to future actions.¹

In each generation, a policymaker aggregates the preferences of agents in her generation and chooses an action that is optimal with respect to that generation's preferences. We model the policy problem as a sequential game amongst a succession of these policymakers. Because we study the equilibrium of an intergenerational game, rather than the outcome of a global optimizer, our analysis is positive rather than normative.

There are many political-economy processes (intra-generational games) that could explain how the social planner in a generation aggregates her generation's preferences. For example, the social planner may adopt the preferences of the median voter in her generation, or use a convex combination of individuals' preferences. The precise intra-generational game is unimportant for our purposes and we do not model that game.

We imbed the sequential game in a transparent climate change model that captures the risk of abrupt climate change (Clarke and Reed 1994, Tsur and Zemel 1996, Alley et al. 2003, Stern 2007, Intergovernmental Panel on Climate Change 2007) and the inertia in the climate system. That inertia leads to a delayed relation between current actions and future reductions in risk.

There are multiple MPE, because the optimal policy today depends on

¹Nordhaus (1999) and Mastrandrea and Schneider (2001) imbed hyperbolic discounting in integrated assessment models (numerical models that integrate climate and economic modules) of climate change. These authors assume that the decision-maker in the current period can choose the entire policy trajectory, thus solving by assumption the time-inconsistency problem.

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

For the binary action model we calculate a “constant equivalent” (or “observationally equivalent” Barro (1999)) discount rate, i.e. a constant rate that, if used as the pure rate of time preference, would lead to policy prescriptions (in an optimal control problem) identical to a particular MPE in the sequential game. This constant-equivalent discount rate depends on the individual agents’ time-varying pure rate of time preference, which should be the same function for all public projects. The constant-equivalent discount rate also depends on the specifics of the problem, in particular the longevity of the public project. For example, decisions about climate policy affect welfare over centuries, while a decision about a bridge affects welfare over decades. The differing time scale of these two types of public projects means that the constant-equivalent discount rates corresponding to them may be very different, even though both are based on the same time-varying pure rate of time preference.

The next section derives hyperbolic discounting as an outcome of time perspective. Section 3 discusses damages associated with abrupt climate change, and Section 4 explains how we model the relation between risk and climate policy. Section 5 describes the payoff. We characterize the MPE and a benchmark equilibrium with restricted commitment in Section 6. We

compare the equilibria under constant and hyperbolic discounting in Section 7. Section 8 illustrates numerically the importance of risk, commitment, and discounting, providing a different perspective on the Stern Review.

2 Time perspective and discounting

Heal (1998, 2005) proposes “logarithmic discounting”, based on the Weber-Fechner “Law”, a statement that human response to a change in a stimulus (e.g. vocal or visual) is inversely related to the pre-existing stimulus. Our explanation of hyperbolic discounting is based on time perspective - the tendency to foreshorten time periods as we peer further into the future.

A function $s(n)$ captures time perspective by assigning a perceived length to a year that begins n years from now. This function satisfies $s(0) = 1$, $s'(\cdot) \leq 0$ and $s(\infty) \geq 0$; undistorted time corresponds to $s(\cdot) \equiv 1$. The relation between real time (t) and perceived (foreshortened) time is

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

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

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

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

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

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

The utility discount factor is therefore

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

and the corresponding pure rate of time preference is

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

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

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

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

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

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

The corresponding pure rate of discount is

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

which decreases from $\rho(0) = \gamma \beta + \delta (1 - \beta)$ to $\rho(\infty) = \gamma$ when $\beta \in (0, 1)$. An increase in β lowers the discount rate, i.e., increases the concern for the future. The constant rates $\rho = \delta$ or $\rho = \gamma$ correspond to the special cases where $\beta = 0$ or $\beta = 1$, respectively.

Other functional forms for hyperbolic discounting are consistent with the time perspective explanation. For example, logarithmic discounting is obtained by setting $s(t) = \frac{1}{1+kt}$, $k > 0$, with the resulting pure rate $\rho(t) = \frac{\rho_0 + k}{1+kt}$. Barro (1999) uses the discount factor $\exp -(\rho(t - \tau) + \phi(t - \tau))$, with ρ a

constant; for (our parameter) $\rho_0 = 0$, Barro's formulation corresponds to $s(t - \tau) = \exp - (\rho(t - \tau) + \phi(t - \tau))$.

3 Catastrophic climate change

Recent evaluations of likely outcomes of global warming are alarming (Stern 2007, Intergovernmental Panel on Climate Change 2007). The current atmospheric GHG concentration is estimated at 380 ppm CO₂ (430 ppm of CO₂e), compared with 280 ppm CO₂ at the onset of the Industrial Revolution. Under BAU, the concentration could double the pre-Industrial level by 2035 and treble this level by the end of the century. The recent IPCC report gives 2 – 4.5°C as a likely range for the increase in equilibrium global mean surface air temperature due to doubling of atmospheric GHG concentration with a non-negligible chance of exceeding this range (Intergovernmental Panel on Climate Change 2007, p. 749). The Stern Review gives 2 – 5°C and 3 – 10°C as likely ranges for equilibrium global mean warming due to doubling and trebling of GHG concentration, respectively (Stern 2007). Even more disturbing is the observation that the probability of outcomes that significantly exceed the most likely estimates is far from negligible: under doubling of GHG concentration, there is a 50% chance that the global mean warming will exceed 5°C (close to the warming since the last ice age) in the long term (Stern 2007, Summary and Conclusion, p. vi). Global warming can therefore give rise to truly catastrophic events; the usual list includes the reversal of the thermohaline circulation, a sharp rise in sea level, the spread of lethal diseases and massive species extinction.

Each link in this chain, leading from changing GHG concentration to the ensuing damage, involves uncertain elements (Schelling 2007). Following Clarke and Reed (1994) and Tsur and Zemel (1996), we account for this uncertainty by assuming that a catastrophic climate event occurs at random time T with a distribution that depends on the GHG concentra-

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

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

4 Risk and climate policy

A climate policy that begins at time t consists of an abatement process $w(t + \tau) \in [0, 1]$, $\tau \geq 0$; $w(\cdot) = 0$ corresponds to no abatement (BAU), and $w(\cdot) = 1$ corresponds to abatement that stabilizes the hazard (the GHG concentration). We let X measure the cost of stabilization as a fraction of the income at risk, Δ . An abatement effort w costs $wX\Delta$.

The abatement process $\{w(t + \tau), \tau \geq 0\}$ induces both the cost process $\{w(t + \tau)X\Delta, \tau \geq 0\}$ and the hazard process

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

This specification is a simplified representation of the following situation. The actions that we take at a point in time (e.g. abatement, levels of consumption) determine greenhouse gas (GHG) emissions at that time. These flows, and existing GHG stocks, determine the evolution of the stock of GHG.

The risk of a climate-related catastrophe, given by h , is a monotonic function of the stock of GHG. We can invert this function to write the time derivative of h as a function of h and w , as in equation (4).

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

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

The simplicity of equation (4) is important. There are conjectures on the level of risk for different types of events (such as a reversal of the thermohaline circuit or a rapid increase in sea level) corresponding to different policy trajectories (e.g. BAU or specific abatement trajectories). We can use these kinds of conjectures to suggest reasonable magnitudes for the parameters of equation (4) (the initial value of h , and the constants a and μ). There is little empirical basis for calibrating a more complicated model.

5 The payoff

The payoff of the generation alive at time t , “generation t ” is the expectation of the present discounted value of current and future generations’ utility, using the discount factor $\theta(t)$. Consumption grows at an exogenous constant rate g and the utility of consumption is iso-elastic, with the

²The results in a model in which $\dot{h}$ is non-monotonic in h would change in fairly obvious ways. For example, if $\dot{h}$ is small when h is close to both 0 and the steady state level, stabilization would not be worthwhile either for very small or for very large levels of h .

constant elasticity η .³ With initial (time 0) consumption normalized to 1, the flow of consumption from time t onward prior to the event occurrence is $e^{g(t+\tau)}(1 - \Delta X w(t + \tau))$. After the occurrence date there is no role for abatement, and consumption equals $e^{g(t+\tau)}(1 - \Delta)$. The corresponding pre- and post-event utility flows are, respectively,

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

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

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

$$\int_0^T \theta(\tau) e^{-g(\eta-1)(t+\tau)} \frac{(1 - \Delta X w(t + \tau))^{1-\eta} - (1 - \Delta)^{1-\eta}}{1 - \eta} d\tau + \text{constant}$$

where

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

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

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

where

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

Let

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

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

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

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

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

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

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

$$\tilde{\theta}(\tau) \equiv \theta(\tau) e^{g(1-\eta)\tau} = \beta e^{-\tilde{\gamma}\tau} + (1-\beta) e^{-\tilde{\delta}\tau}, \quad (10)$$

where

$$\tilde{\gamma} \equiv \gamma + g(\eta - 1) \quad \text{and} \quad \tilde{\delta} \equiv \delta + g(\eta - 1). \quad (11)$$

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

6 Equilibria

Different assumptions about commitment ability and about the set of feasible policies lead to different equilibrium sets. If the decisionmaker at time 0 can commit to an arbitrary function $w(t)$ (conditional on the event not having occurred before t), the solution is obtained by solving a standard non-stationary optimal control problem. This “full commitment” solution is time-inconsistent (unless it happens to involve the boundary solution $w(t) \equiv 0$ or $w(t) \equiv 1$, i.e. never begin stabilization, or begin full stabilization immediately). Since “full commitment” over a long period of time is implausible, we do not consider it further and focus instead on Markov Perfect Equilibria (MPE) to a sequential game. The agents in this game consist of a sequence of policymakers. We study the limiting game where each agent acts for an arbitrarily short period of time, leading to a continuous time model (Karp 2007).

In a MPE, the current regulator cannot commit future generations to a specific course of action but she can *influence* successors’ actions by affecting the world they inherit, i.e. by changing the payoff-relevant state variable. The MPE recognizes the difference between influencing future policies and choosing those policies. In a MPE agents condition their actions on (only) the payoff-relevant state variable, and they understand that their successors do likewise. Therefore, an agent’s beliefs about future policies depend on her beliefs about the future trajectory of the state variable. An agent’s action has an immediate effect on her current flow payoff and it also affects the continuation value via its influence on the state variable. We provide the necessary condition for a MPE for the general case and then analyze a binary action specialization. In order to provide a benchmark for the set of MPE in this binary case, we then consider an equilibrium involving “restricted commitment”.

6.1 MPE in the general model

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

The state variable h is standard: at a future time $t + \tau, \tau > 0$, the value of $h(t + \tau)$ depends on the current hazard $h(t)$ and the intervening decisions $w(t + \xi)$, $0 \leq \xi \leq \tau$. The probability of survival until time $t + \tau$, conditional

on $T > t$, is $Pr\{T > t + \tau | T > t\} = e^{-y(t,\tau)}$, which also depends on $h(t)$ and the intervening decisions. However, if the regulator at time t is in a position to make a decision, the event has not yet occurred: $y(t,t) = 0$. Therefore, a *stationary* equilibrium depends only on the current hazard, $h(t)$. Conditional on survival at time t , $h(t)$ is the only payoff-relevant state variable. We restrict attention to stationary pure strategies.

Let $q(h(t + \tau), w(t + \tau))$ denote the right-hand side of (4) and let h and w stand for $h(t)$ and $w(t)$, respectively. The following Lemma gives the necessary condition for a MPE; proofs are in the appendix:

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

$$K(h) + (\tilde{\gamma} + h) V(h) = \max_{w \in \Omega} \{U(w) + q(h, w) V'(h)\}, \quad (12)$$

with the “side condition”

$$K(h) \equiv (\delta - \gamma) (1 - \beta) \int_0^\infty e^{-(\delta\tau + y(t,\tau))} U(\chi(h(t + \tau))) d\tau. \quad (13)$$

Remark 1. *The control rule that maximizes the right-hand side of equation (12) depends on the payoff relevant state h , but not on y . This control rule also depends on the current regulator's beliefs about her successors' policies. Those policies affect the shadow value of the hazard, $V'(h)$.*

Remark 2. *The DPE is “generalized” in the sense that it collapses to the standard model with constant discounting in the two limiting cases $\beta = 1$ and $\beta = 0$. The former case is obvious from equation (13). To demonstrate the latter case, note that for $\beta = 0$,*

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

Substituting this equation into (12) produces the DPE corresponding to the constant discount rate $\tilde{\delta}$.

6.2 A binary action specialization

We focus on the situation where $w(t)$ is limited to either full stabilization ($w = 1$) or BAU ($w = 0$). There are in general multiple MPE because the optimal decision for the current regulator depends on her beliefs about the actions of subsequent regulators. The equilibrium beliefs of the current regulator (i.e. those that turn out to be correct) depend on her beliefs about the beliefs (and thus the actions) of successors. There is an infinite sequence of these higher order beliefs, leading to generic multiplicity of equilibria. However, the equilibrium set has a simple characterization.

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

$$U(1) = \frac{(1 - \Delta X)^{1-\eta} - (1 - \Delta)^{1-\eta}}{1 - \eta} \quad \text{and} \quad U(0) = \frac{1 - (1 - \Delta)^{1-\eta}}{1 - \eta}. \quad (14)$$

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

$$x \equiv 1 - \frac{U(1)}{U(0)} = 1 - \frac{(1 - \Delta X)^{1-\eta} - (1 - \Delta)^{1-\eta}}{1 - (1 - \Delta)^{1-\eta}}. \quad (15)$$

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

$$X = \frac{1}{\Delta} \left[1 - \{1 - x [1 - (1 - \Delta)^{1-\eta}] \}^{\frac{1}{1-\eta}} \right]. \quad (16)$$

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

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

Characterization of the equilibrium set uses two functions of the state variable, h , and model parameters; we denote these functions as $x^U(h)$ and $x^L(h)$. These functions divide the (h, x) plane into three regions, which have the following properties: (i) if $x \leq x^L(h)$ (so that the utility cost of stabilization is small) the unique MPE is perpetual stabilization; (ii) if $x \geq x^U(h)$ (so that the utility cost of stabilization is large) the unique MPE is perpetual BAU; and (iii) if $x^L(h) < x < x^U(h)$ there are MPE with either perpetual stabilization or perpetual BAU. For a subset of $(x^L(h), x^U(h))$ there are additional MPE that involve delayed stabilization, i.e., a BAU policy ($w(t) = 0$) for a while, followed by a perpetual stabilization.

Once we obtain the critical “utility cost of stabilization” values ($x^L(h)$ and $x^U(h)$), we can use equation (16) to obtain the critical income cost of stabilization values, denoted $X^L(h)$ and $X^U(h)$. The latter values are important for our simulations in section 8.

This model shows the potentially offsetting effects of an increase in the elasticity of marginal utility, η . This parameter affects the equilibrium by altering the “effective discount rate” $\rho(t) + g(\eta - 1)$ and it also enters the function $X(x)$ defined in equation (16). For $g > 0$, an increase in η increases

the “effective discount rate” $\rho(t) + g(\eta - 1)$, which in turn decreases the critical values $x^U(h)$, $x^L(h)$; that is, the change makes the decision-maker less willing to sacrifice current utility for future reduction in risk. However, the larger value of η makes the decision-maker more risk averse; it shifts up the graph of $X(x)$, as shown in Figure 1, so the smaller value of the critical x (resulting from the increase in η) might correspond to a larger value of the critical X . Thus, in general the effect of η on critical values of X is ambiguous. For our calibration, an increase in η reduces these critical values.

The following functions are used in the analysis below; superscripts B and S denote functions under perpetual BAU or stabilization, respectively. Under BAU, using equation (4), the probability of disaster by time t is

$$F^B(t) = 1 - \exp\left(\frac{-at\mu + (a - h_0)(1 - e^{-\mu t})}{\mu}\right). \quad (17)$$

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

$$V^B(h_0) \equiv U(0) \int_0^\infty (1 - F^B(t)) \tilde{\theta}(t) dt = U(0) \nu(h_0), \quad (18)$$

where

$$\nu(h) \equiv \int_0^\infty \exp\left(\frac{-a\mu t + (a - h)(1 - e^{-\mu t})}{\mu}\right) \tilde{\theta}(t) dt. \quad (19)$$

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

$$V^S(h_0) \equiv U(1) \int_0^\infty e^{-h_0 t} \tilde{\theta}(t) dt = U(1) \xi(h_0) \quad (20)$$

where

$$\xi(h) \equiv \int_0^\infty e^{-ht} \tilde{\theta}(t) dt = \frac{(1 - \beta)(\gamma + g(\eta - 1)) + h + \beta(\delta + g(\eta - 1))}{(\delta + g(\eta - 1) + h)(h + \gamma + g(\eta - 1))}. \quad (21)$$

6.2.1 Markov Perfect Equilibria

The control space is $w(t) \in \{0, 1\}$, the flow payoffs are given in equation (14) and the hazard evolves according to equation (4). Let $\chi(h)$ be a MPE decision rule. Using the equilibrium condition (12) and the convention that in the event of a tie the regulator chooses stabilization, in the binary setting χ satisfies

$$\chi(h) = \begin{cases} 1 & \text{if } U(1) \geq U(0) + \mu(a - h)V'(h) \\ 0 & \text{if } U(1) < U(0) + \mu(a - h)V'(h) \end{cases}. \quad (22)$$

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, the current regulator must want to stabilize when she believes that all future regulators will stabilize. Under this belief, $V(h) = V^S(h)$ and $V'(h) = V^{S'}(h) = U(1)\xi'(h)$, where $V^S(h)$ and $\xi(h)$ are defined in equations (20) and (21), respectively. Thus, using the equilibrium rule (22), $U(1) \geq U(0) + \mu(a - h)U(1)\xi'(h)$ must hold for stabilization to be a MPE. Defining

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

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 $U(1) < U(0) + \mu(a - h)V^{B'}(h) = U(0) + \mu(a - h)U(0)\nu'(h)$. Defining

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

with $\nu(h)$ given by equation (19), 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 2. *The functions $\pi(h)$ and $\sigma(h)$ are increasing over $(0, a)$ with $\pi(a) = \sigma(a) = 1$, and $\sigma(h)$ is concave.*

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

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

$$\pi(h) < \sigma(h), \quad h \in (0, a). \quad (25)$$

Under inequality (25), there exists a MPE with perpetual stabilization ($w \equiv 1$) if and only if at the initial hazard h the cost of stabilization satisfies

$$x < x^U(h) \equiv 1 - \pi(h); \quad (26)$$

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

$$x > x^L(h) \equiv 1 - \sigma(h). \quad (27)$$

Figure 2 illustrates Proposition 1. 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 MPE.

Because the region between these two curves has positive measure (when inequality (25) is satisfied), the existence of multiple equilibria is generic in this model.⁴ The multiplicity of equilibria stems from the fact that the

⁴Laibson (1994) shows that there are multiple equilibria to this kind of sequential game under non-Markov policies. Krusell and Smith (2003) show the existence of a continuum of MPE when agents use step functions. Elements of this equilibrium set involves an infinite

Figure 2: 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.

optimal action today depends on the shadow value $V'(h)$, which depends on future actions *that the current regulator does not choose*. If future regulators will stabilize, the shadow cost of the state ($-V'(h)$) is high, relative to the shadow cost when future regulators follow BAU. The current regulator has more incentive to stabilize if she believes that future regulators will also stabilize. Actions are “strategic complements”, a circumstance common to coordination games. Our problem resembles the dynamic coordination game familiar from the “history versus expectations” literature (Matsuyama 1991, Krugman 1991). In those coordination games, the optimal decision for (non-atomic) agents in the current period depends on actions that will

sequence of steps, and the step sizes are endogenous. Our setting contains a single, exogenously determined step size. Karp (2005, 2007) shows the existence of multiple *candidates* solving the necessary conditions for MPE, due to an indeterminacy in the steady state conditions. In our setting, the multiplicity arises because of a non-convexity in the game. Section 7 elaborates on this observation, showing the resemblance between the problem under constant discounting and the familiar “Skiba problem” in optimal control.

be taken by agents in the future. The non-convexity in the payoffs in these problems typically leads to multiple rational expectations equilibria for a set of initial conditions of the state variable. These equilibria are in general not Pareto efficient. We show that inter-generational coordination problems in our game can lead to either too little or too much stabilization, relative to a benchmark under restricted commitment.

Proposition 1 characterizes only equilibrium trajectories in which the action never changes. It is clear that a switch from stabilization to BAU is impossible, since the hazard remains constant under stabilization and the decisionmaker uses a pure strategy. However, the proposition does not rule out the possibility of a MPE with delayed stabilization, i.e. an equilibrium beginning with BAU and switching to stabilization once the hazard reaches a threshold. The next proposition shows that such equilibria exist.⁵ We use the following definition

$$\Theta(h) \equiv \frac{\mu(a-h) \left(\frac{\beta}{\gamma+h} + \frac{1-\beta}{\delta+h} \right)}{h + \beta\tilde{\gamma} + \tilde{\delta}(1-\beta) + \mu(a-h) \left(\frac{\beta}{\gamma+h} + \frac{1-\beta}{\delta+h} \right)}. \quad (28)$$

Proposition 2. *Suppose that Condition (25) 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*

$$\Theta(h) < x < 1 - \pi(h). \quad (29)$$

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

Recall that x equals the utility cost of stabilizing the hazard (or the atmospheric GHG concentration) as a fraction of the value-at-risk $U(0)$. Relation (29) defines the lower and upper bounds of x for a delayed stabilization

⁵>From 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.

MPE to exist. We verify in the appendix that

$$1 - \pi(h) - \Theta(h) = \frac{(\tilde{\delta} - \tilde{\gamma})^2(2h + \tilde{\gamma} + \tilde{\delta})}{(h + \tilde{\gamma})^2(h + \tilde{\delta})^2} \beta(1 - \beta). \quad (30)$$

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

6.2.2 Restricted commitment: a benchmark

We saw in the previous section that a class of MPE decision rules leads to either perpetual BAU or perpetual stabilization. Here we consider a “restricted commitment” benchmark in which the decision-maker at time 0 behaves as if she could commit future generations to either perpetual stabilization or perpetual BAU. In contrast, “full commitment” permits switches between BAU and stabilization – or vice-versa.

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

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

Under restricted commitment there exists a critical function $x^C(h)$ such that initial decision-maker chooses perpetual stabilization if $x \leq x^C(h)$ and she chooses perpetual BAU if $x > x^C(h)$. To determine the critical function, we note that 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

$$\lambda(h) \equiv \frac{\nu(h)}{\xi(h)}. \quad (31)$$

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

$$x^C(h) \equiv 1 - \lambda(h). \quad (32)$$

Thus, when future regulators will follow her policy in perpetuity, the current regulator wants to stabilize if and only if, at the current hazard h , the utility cost of stabilization does not exceed $x^C(h) = 1 - \lambda(h)$. If $x > x^C(h)$, she chooses BAU, in which case the hazard h increases. A restricted commitment policy that involves stabilization is obviously time consistent, since under stabilization the hazard does not change. The restricted commitment policy that involves BAU is also time consistent under fairly general circumstances, because stabilization is more valuable when the hazard is lower. (Obviously, there is no point in incurring a cost to stabilize when the hazard is near its steady state.) If the regulator (under restricted commitment) wants to follow BAU for a given initial value of h , all of her successors would make the same choice at the larger values of h that result from earlier BAU. We summarize this discussion in⁷

Proposition 3. *Given the initial hazard $h \in [0, a]$, the optimal restricted-commitment policy is to stabilize if and only if $x \leq x^C(h)$. This policy is*

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

time consistent for all $h \in [0, a]$ and $x \in [0, 1]$ if and only if $\lambda'(h) \geq 0$. A sufficient condition for this inequality is $\mu \geq a + \delta + g(\eta - 1)$.

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

7 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 $\tilde{\delta}$ with $\tilde{\gamma}$.) The constant discount rate is $\tilde{\delta}$, so the distant future is “heavily discounted”. Following the standard procedure to obtain the DPE, or invoking Remark 2, we have the following DPE:

$$(\tilde{\delta} + h) V(h) = \max_{w \in \{0,1\}} \{U(w) + \mu(a - h)(1 - w)V'(h)\}. \quad (33)$$

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

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

The second implication is that $\lambda(h) = \sigma(h)$ under constant discounting. This equality means that the optimal solution when the regulator is restricted to making a commitment (in perpetuity) at time 0, is equal to the solution when the regulator has the opportunity to switch between BAU and stabilization. For abrupt events, the regulator is tempted to delay stabilization (i.e. the “restriction” in restricted commitment binds) only under hyperbolic discounting. The ability to switch between policies is of no value for abrupt events under constant discounting. The economic explanation for this result is simply that BAU is the optimal policy only if the hazard is sufficiently large; under BAU the hazard increases, whereas it remains constant under stabilization.

8 Policy bounds and constant equivalent rates

When $\eta \neq 1$ and $g \neq 0$ this model has one degree of freedom: for given β , the “effective discount rate” depends on $\tilde{\gamma}$ and $\tilde{\delta}$, determined by two equations in three unknowns, δ , γ , and g . These parameters, unlike η , do not enter the function U , defined in equation (7). We normalize by setting $\gamma = 0$.⁸ This normalization implies that the long run pure rate of time preference is 0, i.e. it means that we are unwilling to transfer utility between two agents living in the infinitely distant future at a rate other than one-to-one. It also implies that the long run effective discount rate is $g(\eta - 1)$.

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

⁸When $\eta = 1$ the equilibrium is always independent of g . For $\eta = 1$ or $g = 0$, γ and δ equal $\tilde{\gamma}$ and $\tilde{\delta}$. In this case, setting $\gamma = 0$ is an assumption, not a normalization.

When $g > 0$, the constant defined in equation (5) is finite if and only if $\eta > 1$. In contrast, the maximand in expression (9) is defined even for some values of $\eta < 1$, because the hazard has an effect similar to discounting. For $\eta \leq 1$ we can adopt the “overtaking criterion” to evaluate welfare.

⁹The exact equivalence occurs if the decision rules under both hyperbolic and constant discounting can be characterized by a single parameter. Barro (1999) also obtains a constant equivalent discount rate, because the single parameter in his logarithmic model is the slope of the decision rule. When the decision rules cannot be described by a single parameter, it is possible only to obtain an approximate constant equivalent discount rate. For example, in the linear-quadratic model there exists a linear equilibrium control rule under both constant and hyperbolic discounting. Because this control rule involves two parameters – the slope and the intercept – it is in general not possible to find an exact constant equivalent discount rate for the hyperbolic model (Karp 2005).

8.1 Calibration

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

In order to be able to compare the damage estimates under our calibration with those used by other models, we define $P^B(t) \equiv Pr\{T \leq t|BAU\}$ as the probability that the catastrophe occurs by time t under BAU, and $P^S(t) \equiv Pr\{T \leq t|Stabilization\}$ as the corresponding probability under stabilization. The future (time t) expected increase in damages from following BAU rather than stabilization, as a percentage of future income, is $D(t) = (P^B(t) - P^S(t)) 100\Delta\%$. For all calibrations where $h(0) > 0$, $\lim_{t \rightarrow \infty} D(t) = 0$, because both probabilities converge to 1.¹⁰ Figure 3 shows the graphs of $D(t)$ over the next millennium for $\Delta = 0.05, 0.1$ and 0.2 . The corresponding damages after 100 and 200 years are $D(100) = \{0.72, 1.43, 2.88\}$ and $D(200) = \{2.03, 4.01, 8.11\}$.

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

The Stern Review assumes that climate-related damages are zero after

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

Figure 3: Percentage expected increased loss of income under BAU: $\Delta = 0.05$ = dashed; $\Delta = 0.1$ = solid; $\Delta = 0.2$ = dotted.

200 years, whereas in our calibration damages continue to rise for 800 years and then decrease asymptotically to 0. The maximum level of $D(t)$ equals $91\Delta\%$, i.e. 4.5%, 9.1% and 18.2% for the three values of Δ . In view of the different profiles of damages in the Stern Review and in our calibration, exact matching is not possible. However, our case $\Delta = 0.2$ approximates one of the high (but not the highest) Stern damage scenarios; the value $\Delta = 0.1$ approximates the Stern “market impacts + risk of catastrophe” scenario, and the value $\Delta = 0.05$ corresponds to a much lower damage scenario.

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

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

This parameterization implies that the pure rate of time preference begins at 3% and falls to 1% by 30 years, eventually declining to 0. Our value of $\rho(30)$ is ten times greater than the Stern Review’s constant pure rate of time preference. An ethical concern for generations in the distant future requires a small pure rate of time preference *only in the case of a constant pure rate of time preference*. A declining pure rate of time preference is consistent with

both ethical considerations and a large pure rate of time preference in the near and medium term. This flexibility means that the model is compatible with both a reasonable ethical view and also with market discount rates.

8.2 Results

For a variety of parametric and equilibrium assumptions, we calculated upper and lower bounds on X – the fraction of income-at-risk that society spends to stabilize risk. These values were insensitive to choices of Δ over the interval $(0.1, 0.2)$, so the tables below report only results for $\Delta = 0.2$. We also report the corresponding constant equivalent pure rate of time preference (ρ). We discuss results for $g \in [1\%, 2\%]$ and $\eta \in [1.1, 4]$. An appendix, available on request, shows that our results are sensitive to parameter changes in the neighborhood of $\eta = 1$, and also to values of our calibration parameter ρ (30).

Tables 1 – 3 show the (X) policy bounds and constant-equivalent ρ values for the 6 cases corresponding to $\eta \in \{1.1, 2, 4\}$ and $g \in \{0.01, 0.02\}$. In each case the constant equivalent social discount rate (not shown) equals the constant equivalent value of ρ plus ηg . We emphasize the case where $\eta = 2$ and compare the results for $g = 1\%$ and $g = 2\%$ across the different equilibria.

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

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

We begin with the restricted commitment equilibrium, which is both time consistent and constrained optimal. For $\eta = 2$, the maximum fraction of the

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

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

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

For $g = 1\%$ and $\eta = 2$ the upper and lower bounds of X in a MPE are 17.8% and 9.9%, with corresponding constant equivalent values of ρ of 0.1% and 0.7% (Tables 2 and 3). In this case, for $17\% < X < 17.8\%$ of the value at risk, the optimal policy is to follow BAU, but there are MPE that result in stabilization. For $9.9\% < X < 17\%$ the optimal policy is to stabilize, but there are MPE that result in BAU. Thus, a MPE may result in either excessive or insufficient stabilization (although, in a sense, the latter is more likely). The broad range of values for which there are multiple MPE indicates the importance of establishing commitment devices that enable the current generation to lock in the desired policy trajectory.

For $g = 2\%$ and $\eta = 2$, the upper and lower bounds (5.4% and 4%) are much closer (compared to when $g = 1\%$), and both lie below the upper bound

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

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

under restricted commitment. In this case, for any X such that stabilization is a MPE, stabilization also maximizes welfare. For $5.4\% < X < 6.1\%$ all MPE involve BAU even though stabilization is optimal. With $g = 2\%$ and $\eta = 2$, the constant equivalent ρ in a MPE ranges between 0.5% to 1% (the upper and lower bounds that correspond to the MPE set). As expected, higher growth rates make the current generation less willing to sacrifice for the sake of wealthier future generations, decreasing the X bounds

9 Conclusion

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

In our binary action model, a reduction in current consumption (“sta-

bilization”) reduces the future hazard rate of a random event that causes permanent loss of utility. There are multiple MPE for an interval of stabilization costs. The upper bound of this interval is the maximum cost consistent with a MPE involving stabilization; the lower bound is the minimum cost consistent with a MPE involving BAU. For each of these bounds we calculated a constant equivalent pure rate of time preference, i.e. a constant rate that leads (in the optimal control problem) to the same decision rule as does the time-varying pure rate of time preference (in the sequential game). We compared the set of MPE to a time-consistent (constrained optimal) reference equilibrium. The MPE equilibrium set indicates how much society would be willing to spend to stabilize the risk if it managed to solve the *intragenerational* but not the *intergenerational* collective action problem; the reference equilibrium indicates how much society should be willing to spend, if it solves both the intra- and the inter-generational problems.

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

Across most dimensions, our model is vastly simpler than the integrated assessment models typically used for policy recommendations. However, catastrophic risk is central to our model, and we take seriously the fact that future policies are not chosen at the current time, but will instead be

conditioned on future fundamentals. In addition, our model of the pure rate of time preference provides a reasonable description of ethics while also being consistent with observed market rates. Ethical concern does not require a small pure rate of time preference in the near and medium run; it requires that the pure rate of time preference eventually become small. Our numerical results are broadly consistent with the recommendations in the Stern Review. The simplicity and parsimony of the model make it easy for other researchers to examine the sensitivity of those results.

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. Weyant (eds), *Discounting and Intergenerational Equity*, Resources for the Future, Washington, DC, pp. 13–22.
- Barro, R.: 1999, Ramsey meets Laibson in the neoclassical growth model, *Quarterly Journal of Economics* **114**, 1125–52.
- 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., Aydede, S. and Portney, P.: 1994, Preferences for life saving programs: How the public discounts time and age, *Journal of Risk and Uncertainty* **8**, 243–265.
- Dasgupta, P.: 2007a, Commentary: The Stern review’s economics of climate change, *National Institute Economic Review* **199**(1), 4–7.
- Dasgupta, P.: 2007b, Discounting climate change. University of Cambridge.
- Heal, G. M.: 1998, *Valuing the Future: Economic Theory and Sustainability*, Columbia University Press.
- Heal, G. M.: 2005, Intertemporal welfare economics and the environment, in K.-G. Mäler and J. R. Vincent (eds), *Handbook of environmental economics*, Vol. 3, Elsevier, chapter 21, pp. 1105–1145.

- Intergovernmental Panel on Climate Change: 2007, Climate change 2007, <http://www.ipcc.ch/SPM040507.pdf>.
- 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**(1), 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.
- Laibson, D.: 1994, Self control and saving, Harvard University Working Paper.
- 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.
- Nordhaus, W.: 2007, A review of the Stern Review on the economics of climate change, *Journal of Economic Literature* **45**, 686–702.
- 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.
- Phelps, E. S. and Pollak, R.: 1968, On second-best saving and game-equilibrium growth, *Review of Economic Studies* **35**(2), 185–199.
- Ramsey, F. P.: 1931, *Foundations of Mathematics and Other Logical Essays*, Routledge and Kegan Paul.

- Schelling, T. C.: 2007, Climate change: the uncertainties, the certainties, and what they imply about action, *Economists' Voice* **July**, 1–5.
- Skiba, A.: 1978, Optimal growth with a convex-concave production function, *Econometrica* **46**, 527–39.
- Stern, N.: 2007, *The Economics of Climate Change*, Cambridge University Press.
- Strotz, R. H.: 1956, Myopia and inconsistency in dynamic utility maximization, *Review of Economic Studies* **23**(3), 165–180.
- Tsur, Y. and Zemel, A.: 1996, Accounting for global warming risks: Resource management under event uncertainty, *Journal of Economic Dynamics & Control* **20**, 1289–1305.
- Weitzman, M.: 2007, A review of the Stern Review on the economics of climate change, *Journal of Economic Literature* **45**(3), 703–724.

Appendix: Proofs

To simplify notation we assume that $g(\eta - 1) = 0$. The proofs extend to the general case by substituting $\tilde{\delta} \equiv \delta + g(\eta - 1)$ and $\tilde{\gamma} \equiv \gamma + g(\eta - 1)$ for δ and γ , respectively.

Proof of Lemma 1: 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 Karp (2007) to our setting, and using the hyperbolic discount factor in equation (2), yields the generalized DPE

$$\hat{K}(z) + \gamma W(z) = \max_{w \in \Omega} (e^{-y(t)}U(w(t)) + W_h g + W_y h), \quad (34)$$

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

$$\hat{K}(z) = (\delta - \gamma)(1 - \beta) \int_0^\infty e^{-(\delta t + y(t))} U(\hat{\chi}(z)) dt \quad (35)$$

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 (34), cancelling e^{-y} and rearranging, yields equation (12). 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 (35) as

$$\begin{aligned} K(h(t)) &= (\delta - \gamma)(1 - \beta) e^{y(t)} \int_t^\infty e^{-\delta(s-t)} \exp\left(-\int_t^s h(\tau)d\tau\right) e^{-y(t)} U(\chi(h(s))) ds \\ &= (\delta - \gamma)(1 - \beta) \int_t^\infty e^{-\delta(s-t)} \exp\left(-\int_t^s h(\tau)d\tau\right) U(\chi(h(s))) ds \end{aligned} \quad (36)$$

Setting $t = 0$ in equation (36) produces equation (13). $\square$

Proof of Lemma 2 Define

$$\varpi(h) \equiv \pi(h)^{-1} = 1 - \rho(a - h)\xi'(h). \quad (37)$$

Differentiating, using equation (21), we obtain

$$\varpi'(h) = \rho\xi'(h) - \rho(a - h)\xi''(h) < 0. \quad (38)$$

Thus,

$$\pi'(h) = -\varpi'(h)/\varpi(h)^2 > 0. \quad (39)$$

Differentiating (24), using equation (19), gives

$$\sigma'(h) = -\rho\nu'(h) + \rho(a - h)\nu''(h) > 0. \quad (40)$$

To establish $\sigma''(h) < 0$, use equation (19) and differentiate three times to obtain $\nu'''(h) < 0$. Differentiating equation (40) gives

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

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

Proof of Proposition 1 We first establish sufficiency of inequality (25) using a constructive proof, which also establishes the claims associated with inequalities (26) and (27). We then show necessity of inequality (25) 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 (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 (20), and

$$V^{S'}(h) = U(1)\xi'(h) \tag{41}$$

with $\xi'(h)$ obtained using equation (21).

Using equation (12) (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

$$U(1) \geq U(0) + \rho(a-h)U(1)\xi'(h) \tag{42}$$

or

$$\frac{U(1)}{U(0)} \geq \pi(h). \tag{43}$$

If inequality (43) 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 (43) is not satisfied, then clearly perpetual stabilization is not an equilibrium. We consider below the case where the weak inequality (43) 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 (18). 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

$$\frac{U(1)}{U(0)} < \sigma(h) \equiv 1 + \rho(a - h)\nu'(h), \quad (44)$$

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 2.

Now we return to the case where inequality (43) 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 (43) is satisfied with equality. From equation (22), 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 2, 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 (43) 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 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

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

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 (43) 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 $\frac{U(1)}{U(0)} = \pi(\hat{h})$. $\square$

Proof of Proposition 2 We use the following definition

$$h_\pi(x) \equiv \begin{cases} \pi^{-1}(1-x) & \text{for } x \in [0, 1 - \pi(0)) \\ 0 & \text{for } x \in [1 - \pi(0), 1] \end{cases} \quad (45)$$

Hazard rates that satisfy $h > h_\pi(x)$ lie above the curve $1 - \pi$ in Figure 2.

(i) The stabilization set is absorbing, because if a (pure strategy) MPE calls for a regulator to stabilize, the hazard never changes. By Proposition 1, 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 1 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 1). 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 1, 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 (22), this action is part of an equilibrium if and only if

$$U(0) - U(1) > -\rho(a - h)V_h(h; \tilde{h}). \quad (46)$$

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

$$\tilde{h} = a - (a - h)e^{-\rho T} \Rightarrow \quad (47)$$

$$T(\tilde{h}; \tilde{h}) = 0 \quad \text{and} \quad \frac{dT}{dh} = \frac{-1}{\rho(a - h)}. \quad (48)$$

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

$$\frac{dy(\tau)}{dh} = \frac{d \int_0^\tau h(s) ds}{dh} = \int_0^\tau \frac{dh(s)}{dh} ds = \int_0^\tau e^{-\rho s} ds = \frac{1 - e^{-\rho \tau}}{\rho}. \quad (49)$$

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

$$\begin{aligned} \frac{dy(\tau)}{dh} &= \frac{d(\int_0^T h(s)ds + \tilde{h}(\tau-T))}{dh} = \\ \int_0^T \frac{dh(s)}{dh} ds + \left(h(T) - \tilde{h} \right) \frac{dT}{dh} &= \int_0^T e^{-\rho s} ds \end{aligned} \quad (50)$$

The last equality uses the fact that $h(T) = \tilde{h}$, from the definition of T . Using equation (47) and (48), 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 (49) and the definition of $y(\tau)$, we have

$$\begin{aligned} y(T) &= \int_0^T h(s)ds \Rightarrow \\ \frac{dy(T)}{dT} &= h(T) + \int_0^T \frac{dh(s)}{dh} \frac{dh}{dT} ds \end{aligned} \quad (51)$$

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 (49)) is

$$\begin{aligned} V_h(h; \tilde{h}) &= (U(0) - U(1)) e^{-y(T)} \theta(T) \frac{dT}{dh} + \\ \int_0^T \frac{d(e^{-y(\tau)})}{dh} \theta(\tau) U(0) dt + \int_T^\infty \frac{d(e^{-y(\tau)})}{dh} \theta(\tau) U(1) dt \\ &= \frac{-(U(0) - U(1))}{\rho(a-h)} e^{-y(T)} \theta(T) - \end{aligned} \quad (52)$$

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

Using this expression, we can write the optimality condition (46) as

$$\begin{aligned} U(0) - U(1) &> (U(0) - U(1)) e^{-y(T)} \theta(T) + \\ \rho(a-h) &\left(\int_0^T \left(\frac{1-e^{-\rho\tau}}{\rho} \right) e^{-y(\tau)} \theta(\tau) U(0) dt + \int_T^\infty \left(\frac{1-e^{-\rho T}}{\rho} \right) e^{-y(\tau)} \theta(\tau) U(0) dt \right). \end{aligned} \quad (53)$$

It is convenient to treat T as the independent variable, recognizing that the initial condition h is a function of T (from equation (47)): $h = h(T)$. The existence of a set A with positive measure requires that inequality (53) 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 (53) is

$$(U(0) - U(1)) - (U(0) - U(1))(\tilde{h} + r(0))T + o(T). \quad (54)$$

This expansion uses equations (3) and (51) 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 (53) is

$$\begin{aligned} \rho(a - \tilde{h})T \int_0^\infty e^{-y(\tau)}\theta(\tau)U(1)d\tau + o(T) = \\ \rho(a - \tilde{h})T \int_0^\infty e^{-\tilde{h}\tau}\theta(\tau)U(1)d\tau + o(T) = \\ \rho(a - \tilde{h})T \frac{(1-\beta)\gamma + \beta\delta + \tilde{h}}{(\tilde{h} + \gamma)(\tilde{h} + \delta)}U(1) + o(T). \end{aligned} \quad (55)$$

Substituting expressions (54) and (55) into inequality (53), dividing by T and letting $T \rightarrow 0$ (from above) produces the inequality

$$(U(0) - U(1))(\tilde{h} + r(0)) > \rho(a - \tilde{h}) \frac{(1-\beta)\gamma + \beta\delta + \tilde{h}}{(\tilde{h} + \gamma)(\tilde{h} + \delta)}U(1). \quad (56)$$

Using $x \equiv 1 - \frac{U(1)}{U(0)}$ and $r(0) = \beta\gamma + \delta(1-\beta)$ (from equation (3)), and replacing $\tilde{h}$ with h , inequality (56) can be expressed as

$$\frac{x}{1-x}(h + \beta\gamma + \delta(1-\beta)) > \rho(a - h) \left(\frac{\beta}{h + \gamma} + \frac{1-\beta}{h + \delta} \right) \quad (57)$$

or, equivalently,

$$x > \Theta(h), \quad (58)$$

where $\Theta(h)$ is defined in equation (28), 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 (30). $\square$

Proof of Proposition 3: (i) This claim follows from differentiating the functions $\nu(h)$ and $\xi(h)$ and by inspection. (ii) We begin with

$$y^B(t, h) \equiv \int_0^t (a - (a - h)e^{-\rho\tau}) d\tau = at - (a - h) \frac{1 - e^{-\rho t}}{\rho}, \quad (59)$$

where $y^B(t, h)$ is a specialization of $y(0, t)$, defined in (8), when the hazard process under BAU evolves (following equation (4)) according to $h(t) = a - (a - h_0)e^{-\mu t}$. From equations (19), (21) and (59),

$$\nu(h) - \xi(h) = \int_0^\infty \theta(t) \left(e^{-y^B(t, h)} - e^{-ht} \right) dt. \quad (60)$$

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, $y^B(t, h) > ht$ when $h < a$ and $\rho > 0$, and the right-hand side of equation (60) is negative. (iii) This claim is merely a summary of the derivation in the text above equation (32).

(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 (18) and (20), we express $\lambda(h)$ as

$$\lambda(h) = \frac{\int_0^\infty e^{-y^B(t,h)} \theta(t) dt}{\int_0^\infty e^{-ht} \theta(t) dt}. \quad (61)$$

Using equation (59) we have

$$y_h^B(t, h) \equiv \partial y^B(t, h) / \partial h = \frac{1 - e^{-\rho t}}{\rho}. \quad (62)$$

The argument h in $y^B(t, h)$ is the initial hazard. Differentiating (61) with respect to h , we see that $\lambda'(h) > 0$ if and only if

$$\int_0^\infty e^{-y^B(t,h)} \theta(t) dt \int_0^\infty e^{-ht} t \theta(t) dt > \int_0^\infty e^{-ht} \theta(t) dt \int_0^\infty e^{-y^B(t,h)} y_h^B(t, h) \theta(t) dt. \quad (63)$$

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 (62), we express (63) as

$$\left(\frac{\beta}{(h+\gamma)^2} + \frac{1-\beta}{(h+\delta)^2} \right) \int_0^\infty e^{-y^B(t,h)} \theta(t) dt > \left(\frac{\beta}{h+\gamma} + \frac{1-\beta}{h+\delta} \right) \int_0^\infty e^{-y^B(t,h)} \theta(t) \frac{1-e^{-\rho t}}{\rho} dt. \quad (64)$$

Since $\delta > \gamma$, the right-hand side of inequality (64) is smaller than

$$\left(\frac{\beta}{(h+\gamma)^2} + \frac{1-\beta}{(h+\delta)^2} \right) \int_0^\infty e^{-y^B(t,h)} \theta(t) \frac{(h+\delta)(1-e^{-\rho t})}{\rho} dt. \quad (65)$$

Thus, it suffices to show that the left-hand side of (64) exceeds (65), i.e., that

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

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 4 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 1) 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 (30) that condition (29) is not satisfied in any of these cases, implying, in view of Proposition 2 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 1, the DPE (33) 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 (33) 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 1 imply that for $\frac{U(1)}{U(0)} < \sigma^0(h)$, $V^B(h)$ satisfies

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

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

$$\begin{aligned} V^S(h) &= \frac{1}{\delta+h} \max\{U(1), U(0) + \rho(a-h)V_h^S(h)\} \\ &= \frac{1}{\delta+h} U(1) \geq \frac{1}{\delta+h} (U(0) + \rho(a-h)V_h^S(h)). \end{aligned} \quad (67)$$

>From (66) and (67) 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 1, $V^S(h)$ 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. $\square$