Available online at www.sciencedirect.com
ScienceDirect
ELSEVIER
Journal of Environmental Economics and Management 56 (2008) 83–101
JOURNAL OF ENVIRONMENTAL ECONOMICS AND MANAGEMENT
www.elsevier.com/locate/jeem

Numerical analysis of non-constant pure rate of time preference: A model of climate policy

Tomoki Fujii[a,*], Larry Karp[b]

[a]School of Economics, Singapore Management University, 90 Stamford Road, 178903 Singapore, Singapore
[b]Department of Agricultural and Resource Economics, University of California, 207 Giannini Hall, Berkeley, CA 94720, USA

Received 30 May 2006
Available online 20 February 2008

Abstract

When current decisions affect welfare in the far-distant future, as with climate change, the use of a declining pure rate of time preference (PRTP) provides potentially important modeling flexibility. The difficulty of analyzing models with non-constant PRTP limits their application. We describe and provide software (available online) to implement an algorithm to numerically obtain a Markov perfect equilibrium for an optimal control problem with non-constant PRTP. We apply this software to a simplified version of the numerical climate change model used in the Stern Review. For our calibration, the policy recommendations are less sensitive to the PRTP than widely believed.

JEL classification: C63; Q20

Keywords: Non-constant pure rate of time preference; Hyperbolic discounting; Numerical methods; Non-renewable resources; Observational equivalence; Climate change; Stern Review

1. Introduction

The pure rate of time preference (PRTP) is a component of the discount rate, and as such plays a (potentially) important role in dynamic policy decisions. A model with non-constant pure rate of time preference (NCPRTP) is useful for reconciling observations on medium term market discount rates with ethical considerations relevant to the distant future. The optimal solution to such a model is typically time-inconsistent; a subgame perfect (and therefore time consistent) solution requires the solution to a dynamic game. Numerical methods have proven useful in many areas of economics, both to solve old problems and to suggest new ones. Numerical methods can be similarly useful in NCPRTP models. We introduce and apply a numerical package that solves a fairly general NCPRTP model.

We use the software to study an aspect of the climate change debate that has recently received a great deal of attention. The Stern Review of Climate Change [26] uses a constant 0.1% per annum pure rate of time

*Corresponding author. Fax: +65 68280833.
E-mail addresses: tfujii@smu.edu.sg (T. Fujii), karp@are.berkelely.edu (L. Karp).

0095-0696/$ - see front matter
doi:10.1016/j.jeem.2007.11.004

preference and an elasticity of marginal utility equal to 1. Some commentators, including [4,23,29], think that these values lead to unreasonably low social discount rates, and thus call into question the Stern Review's policy recommendations. The Stern Review defends its discounting parameters on ethical grounds. It is difficult to satisfy both ethical criteria and to obtain a social discount rate that reflects medium term market rates, using a single PRTP. Karp and Tsur [18] discuss the role of NCPRTP in reconciling these two objectives in a model that focuses on low probability catastrophes.

We construct a simple model of climate change that is consistent with the orders of magnitude of costs and benefits in the integrated assessment model used in [26]. Our most surprising finding is that for this calibration, the PRTP is much less important than has previously been thought. The explanation is that a modest level of expenditures produces major gains, but much larger expenditures are worth little, and may even be counterproductive. Therefore, the optimal level of expenditures with a 3% PRTP is not much smaller than the level under a 0.1% discount rate. Of course, this result is model-specific.

The rest of this Introduction explains why NCPRTP may be an important feature in economic problems, and we explain what we mean by the "solution" to such a model. Groom et al. [12] provide a recent review of much of this literature.

The resurgence of interest in NCPRTP is due largely to behavioral economics' use of hyperbolic discounting, a particular form of NCPRTP in which the PRTP decreases. Hyperbolic discounting models have been used to explain anomalies such as apparent reversals in an individual's preferences [25]. These models use a relatively short period of time, such as the life of an individual. However, NCPRTP is also important for the study of long-lived environmental problems, such as greenhouse gasses (GHGs), where it is reasonable to use a very long or infinite horizon.

Our interest in NCPRTP arises from natural resource/environmental problems. Constant discounting at a non-negligible rate makes the possibility of extremely large damages in the far-distant future irrelevant to current actions. Constant discounting at a negligible rate causes current generations to save too much for (possibly richer) future generations (or, for example, to spend too much on GHG abatement). NCPRTP, with a discount rate that approaches a very low level, provides a balance that takes into account legitimate reasons for impatience in the near to middle term, while still giving non-negligible weight to welfare in the distant future [2,21,14].

The PRTP measures our willingness to trade utility between (for example) two successive generations. A constant rate implies that our willingness to exchange utility between the current and the next generation is the same as our willingness to exchange utility between two successive generations in the far-distant future. Introspection and survey data [3] suggest that many people do not have such preferences. We can distinguish between the current and the next generations and therefore might exhibit impatience when thinking about transferring utility between the two, leading to a positive PRTP in the near term. However, two generations in the very distant future are indistinguishable to us, suggesting that the PRTP in the distant future might be close to zero. The longest financial instruments mature within 30 or 40 years, so we cannot rely on markets to reflect a declining discount rate over long spans of time.

Two other circumstances produce models that are formally equivalent to a model with NCPRTP. In the first, there exists a "correct" constant discount rate, but the decision-maker has only a probability distribution for this parameter. Moreover, the decision-maker either obtains no information about this parameter or is unable to act on that information. If the decision-maker maximizes the expected value of a payoff, using the subjective distribution of the discount rate, the resulting maximization problem involves a discount rate that falls over time [1,6,28]. Second, if the decision-maker maximizes a convex combination of the payoffs of two or more agents with constant discount rates, the discount rate to the resulting problem falls over time [11].

Some decision problems, such as those that involve a large sunk cost, can be modeled as consisting of a single choice. Once a nuclear power plant is built, it is unlikely to be de-commissioned before its lifetime has expired. The undiscounted future costs of disposing of spent fuel may be of the same order of magnitude as the current construction costs, so the discount rate(s) are critical in determining the cost–benefit ratio of the construction project. However, once the trajectory of discount rates is chosen, the computation of the cost–benefit ratio is standard.

Other problems require a sequence of decisions, leading to a sequence of costs and rewards. Efforts to control climate change involve abatement costs and possible benefits (from reduced climate-related damages) over many periods. NCPRTP qualitatively changes these kinds of dynamic problems (rather than simply complicating a computation), because of the time consistency issue. With NCPRTP, the current decision-maker's willingness to transfer income (or costs, or utility) between two future time periods depends on how far distant those periods are from the current period. That distance decreases with time, so decision-makers' willingness to make the transfer changes with calendar time: the choice that is optimal from the standpoint of the current time is not optimal at a future point in time.

If the current decision-maker were able to commit to a future sequence of policy rules (the "full commitment equilibrium"), the time-consistency problem would disappear. However, it is not reasonable to think that today's decision-maker can make binding decisions for generations more than a century in the future. Therefore, the full commitment equilibrium is implausible. This equilibrium is also of questionable normative interest, because it gives special status to the first generation, eliminating the ability of subsequent generations to make their own choices. In some circumstances the first generation is more selfish in the full commitment equilibrium, relative to the Markov perfect equilibrium (MPE) [16]. Thus, away from the steady state, some future generations may be better off (i.e. have a higher present-discounted value of future payoffs) in the MPE than in the full commitment equilibrium (because they might inherit a worse state variable in the former equilibrium).

Although unable to directly choose successors' policies, the current decision-maker in our setting is able to influence the environment that future generations inherit, thereby affecting the decisions that they choose to make. By using a game theoretic model and adopting a MPE as the solution concept, we strike this balance: the policy-maker in each period chooses the action or decision rule for that period, understanding the effect their decision has on future policy-makers.

The game theoretic model leads to a continuum of steady-state candidates, making it difficult to identify an equilibrium [17]. A similar feature arises in differential games [27,5]. In an equilibrium or optimization problem with unique (or "isolated") steady states, local analysis (in the neighborhood of a steady state) can be accomplished using only the Euler equation and the equation of motion. This analysis enables us to learn something about the problem without obtaining the full solution to that problem. In contrast, where we have a continuum of candidate steady states, there is no criteria for selecting the specific point around which to conduct local analysis. (It is possible to provide a local analysis around an arbitrarily selected candidate steady state, but the arbitrariness of the selection makes that exercise of dubious value.) This difficulty increases the importance of a numerical solution: without such a solution, we learn little about the model.[1]

There is an enormous literature on the term structure of interest rates. The time-profile of returns on assets with different maturities typically is non-constant, due to uncertainty over future consumption or growth rates [8–10,30]. In the absence of market imperfections, the private interest rate and the social discount rate are equal. Thus, with uncertain future consumption paths, the social discount rate is typically non-constant even when the PRTP is constant. If the social discount rate is non-constant due to this type of uncertainty, the environmental policy problem should be solved as a stochastic control problem; *the issue of time-inconsistency does not arise.* The situation is different when the social discount rate is non-constant because of changing pure rates of time preference, and the PRTP is "stationary", i.e. the discount factor used to compare utility in two contiguous future periods depends on the number of periods in the future—not on calendar time. In this case, the time-inconsistency is central, and the environmental policy problem should (arguably) be modeled as a dynamic game. Both sources of the non-constant social discount rate are important in environmental problems, but in order to have a tractable model, we consider only non-constancy arising from a changing PRTP.

The reader who is interested only in the economic results, not in the methods, can move directly to the application in Section 4, skipping Sections 2 and 3 without significant loss of continuity.

[1]Phase plane analysis leads to some qualitative insights in the case of differential games (where there also exists a continuum of steady-state candidates). The MPE to the problem with NCPRTP has a different structure, either eliminating or severely reducing the possibility of using phase plane analysis. In this sense, numerical methods are arguably more important for studying NCPRTP models than differential games.

2. Numerical methods

This section presents a method to numerically solve a non-stochastic discrete time dynamic programming problem with NCPRTP. We derive a *quasi dynamic programming equation* (QDPE) for the control problem. The modifier "quasi" reminds the reader that in order to obtain a MPE to the control problem under NCPRTP, we need to solve that problem as a dynamic game amongst a sequence of decision-makers. Then we describe the numerical algorithm used to solve the QDPE. For the software to implement this algorithm, see [7] and the online appendix available through the journal's archive for supplementary material, which can be accessed at ⟨http://www.aere.org/journal/index.html⟩.

2.1. The optimization problem

It is instructive to consider first a standard autonomous optimal control problem in a discrete time, continuous state setting. Suppose that a decision-maker wants to maximize an objective functional by choosing a control variable $x_t \in \Omega \subset \mathbf{R}$ for each period t. The payoff in each period is given by the *reward function* $f(x_t, S_t)$, where S_t is the *state variable* for time t and $S_t \in \mathscr{S} = [\underline{S}, \overline{S}]$ is the state space. The state variable obeys the equation of motion $S_{t+1} = g(x_t, S_t)$; g is the *transition function*. Both the reward function and transition function are bounded over $S \times \Omega$.

The decision-maker maximizes the following function by choosing $\{x_t\}_{t=0}^{\infty}$:

$$\sum_{t=0}^{\infty} \theta_t f(x_t, S_t) \quad \text{s.t. } S_t = g(x_t, S_t), \;\; S_0 = S, \tag{1}$$

where θ_t is the *discount factor for utility* in period t. The discount rate associated with θ_t is the PRTP, i.e. the "utility discount rate", for period t. The sequence (θ_t) is exogenously given. In the standard setting, the PRTP is a constant, so that we can write $\theta_t = \delta^t$ for $0<\delta<1$. Defining the maximum value of the above summation as $V(S)$, we write the dynamic programming equation under constant discounting:

$$\begin{aligned} V(S) &\equiv \max_{\{x_t\}_{t=0}^{\infty}} \left[\sum_{t=0}^{\infty} \delta^t f(x_t, S_t) \text{ s.t. } S_t = g(x_t, S_t), \;\; S_0 = S \right] \\ &= \max_{x_0} [f(x_0, S) + \delta V(g(x_0, S))]. \end{aligned} \tag{2}$$

Now we relax the assumption that the discount rate is constant. Set the current time to $t = 0$ and let the one-period discount factor for time $t>0$ be σ_t so that

$$\theta_t = \prod_{\tau=1}^{t} \sigma_\tau.$$

We define $\theta_0 \equiv 1$ and require that the one-period discount factor becomes a positive constant after a finite time T. That is, $\sigma_t = \delta<1$, $\forall t \geqslant T$, so that $\theta_t = \theta_T \delta^{t-T}$, $\forall t \geqslant T$. By choosing T large, this model approximates a model in which the discount rate approaches a constant only asymptotically. The standard problem corresponds to the case where $T = 0$. With quasi-hyperbolic discounting $T = 1$ and $\sigma_1 = \beta\delta$ where $0<\beta<1$; the current generation discounts the next generation's utility by the factor of $\beta\delta$, but each successive generation is discounted at a constant factor δ.

Under NCPRTP, the trajectory of plans that are optimal for a decision-maker in a particular period is not time-consistent in general. The value of control variable x_t for $t>0$ that is optimal for the decision-maker at time 0 does not in general equal the value of the control variable that the decision-maker in a subsequent period would want to choose. This inconsistency stems from the non-constant discounting. In our model, time-inconsistency arises only if the preferences of decision-makers at two different points in time are different. For example, fix the current calendar time to τ. The decision-maker at time τ evaluates a trade-off between utility in periods $\tau + t$ and $\tau + t + 1$ using the discount factor σ_{t+1}; the decision-maker in calendar time $\tau + 1$ evaluates the same trade-off using the one-period discount factor σ_t. When $\sigma_t \neq \sigma_{t+1}$ the two decision-makers have different preferences; in this case, the policy that is optimal for the decision-maker at calendar time τ is typically not time-consistent.

We obtain a MPE by solving this problem as a game amongst a succession of generations of policy-makers. Each generation's payoff depends on the sequence of future rewards, but not on past rewards. Generation t ranks outcomes using the objective function

$$\sum_{s=0}^{\infty} \theta_s f(x_{t+s}, S_{t+s}).$$

No generation can directly choose actions taken in the future, but each generation can influence future actions by changing the value of the state variable that it leaves to future generations. This situation can be viewed as a sequential game. We consider a symmetric Nash equilibrium, in which each generation chooses a control rule that is the best response to future generations' control rules. Generations are symmetric: each takes the current state variable as given and is followed by an infinite sequence of future generations. Since the functions $f(\cdot)$ and $g(\cdot)$ are time-independent, we look for a stationary equilibrium control rule. In the equilibrium, the problem of time-inconsistency is resolved because each generation understands how its action affects future actions, via changes in the state variable. We obtain the necessary equilibrium conditions to this problem using a straightforward generalization of the methods in [13], who considered the case where $T = 1$.

The Introduction explained why we do not explicitly study the full commitment equilibrium. However, we note that the steady state in that equilibrium is the same as the steady state in the model where the decision-maker has a constant PRTP equal to the lowest value under NCPRTP [17]. Our section on climate change calculates this steady state. Thus, our numerical results permit comparison of the MPE and the full commitment equilibrium steady states.

2.2. The algorithm

We search for a differentiable control rule $\chi : \mathscr{S} \to \Omega$. The control rule is an *equilibrium control rule* if no decision-maker in the infinite sequence of decision-makers wants to deviate from it.

Definition 1. An equilibrium control rule χ satisfies the following relationship $\forall S$:

$$\chi(S) = \arg\max_{x_0} \sum_{t=0}^{\infty} \theta_t f(x_t, S_t) \quad \text{s.t. } S_0 = S, \;\; S_{t+1} = g(x_t, S_t), \;\; x_t = \chi(S_t) \quad \forall t \geqslant 1. \tag{3}$$

Eq. (3) states that $\chi(S)$ is the current decision-maker's best response, under the belief that future decision-makers will use the rule $\chi(S)$. This equation embodies the Nash equilibrium assumption.

We write the value function as $W_\chi(S)$; the subscript emphasizes the dependence of the value function on the (possibly non-unique) equilibrium control rule, χ. The following result is the basis for our algorithm.

Proposition 1. *Assume that there exists a differentiable Markov perfect control rule $\chi(S)$ and a differentiable value function $W_\chi(S)$. An equilibrium control rule $\chi(S)$ satisfies the QDPE*

$$\chi(S) = \arg\max_{x_0} f(x_0, S) + \sum_{t=1}^{T} (\theta_t - \delta\theta_{t-1}) f(\chi(S_t), S_t) + \delta W_\chi(S_1), \tag{4}$$

where the value function $W_\chi(S)$ satisfies

$$W_\chi(S) = \max_{x_0} \left[f(x_0, S) + \sum_{t=1}^{T} (\theta_t - \delta\theta_{t-1}) f(\chi(S_t), S_t) + \delta W_\chi(S_1) \right]$$
$$\text{with } S_0 = S \text{ given} \quad \text{and} \quad S_{t+1} = g(\chi(S_t), S_t) \text{ for } t \geqslant 1. \tag{5}$$

Proof. We obtain the equilibrium payoff $W_\chi(S)$ by evaluating the payoff in expression (1) at $x = \chi(S)$:

$$W_\chi(S) \equiv \sum_{t=0}^{\infty} \theta_t f(\chi(S_t), S_t) \quad \text{s.t. } S_0 = S. \tag{6}$$

Treating x_0 as a choice variable, and setting $x_t = \chi(S_t)$ $\forall t \geqslant 1$ we write the maximization problem in Definition 1 as

$$\begin{aligned}
&\max_{x_0}\left[f(x_0, S_0) + \sum_{t=1}^{\infty}\theta_t f(x_t, S_t)\right] \\
&= \max_{x_0}\left[f(x_0, S) + \sum_{t=1}^{T}\theta_t f(\chi(S_t), S_t) + \sum_{t=T+1}^{\infty}\delta\theta_{t-1} f(\chi(S_t), S_t)\right] \\
&= \max_{x_0}\left[f(x_0, S) + \sum_{t=1}^{T}(\theta_t - \delta\theta_{t-1}) f(\chi(S_t), S_t) + \delta\sum_{t=0}^{\infty}\theta_t f(\chi(S_{t+1}), S_{t+1})\right].
\end{aligned}$$

We use the definition in Eq. (6) to replace the third term in the last line of the above equation. Evaluating W_χ at $S = S_1$, we obtain Eq. (5). Eq. (4) is an immediate consequence of Eq. (5) and Definition 1. □

The QDPE, Eq. (5), is similar to Eq. (2) except that the former has an extra term, the summation on the right side. The collocation method and function iteration provide a means of solving the QDPE. To see how this method works, it helps to review the standard setting with constant discounting. We start by guessing the value function $V(S)$, and denote the initial guess by $V^{(0)}(S)$. Consider a particular value of S_0. Solving the maximization problem in Eq. (2) yields the maximizing value of the control variable x_0^* at S_0. We can use this to "update" the value function. That is, we let $V^{(1)}(S_0) \leftarrow f(x_0^*, S_0) + \delta V^{(0)}(g(x_0^*, S_0))$. In principle, if we do this for each possible value of S_0, we would have $V^{(1)}(S)$. Hence, we would be able to use $V^{(1)}(S)$ on the right-hand side of Eq. (2) to again update the value function. We repeat this process until $V^{(r)}(S)$ and $V^{(r+1)}(S)$ are close enough, where r denotes the round of iteration.

In practice, we cannot evaluate the value function at every possible value of S_0. Instead, we apply the collocation method, in which a set of K prescribed points $s^1, \ldots, s^K$ called the *collocation nodes* is used to evaluate $V^{(r)}(s^i)$ for $1 \leqslant i \leqslant K$ in a fixed domain $[\underline{S}, \overline{S}]$ [15,22]. The values of $V^{(r)}$ outside the collocation nodes are approximated by a linear combination of $N \leqslant K$ known *basis functions* $\phi_n(S)$, so that the approximant has the form of

$$\hat{V}^{(r)}(S) \equiv \sum_{n=1}^{N} c_n \phi_n(S);$$

the basis functions must be linearly independent at the collocation nodes. The N coefficients $c_1, c_2, \ldots, c_N$ are determined by minimizing the residual at the collocation nodes using (for example) ordinary least square. When $N = K$, the approximant $\hat{V}^{(r)}$ takes the same value as $V^{(r)}$ at the collocation nodes.

The collocation nodes could be evenly spaced over this domain, but the method works better with Chebyshev nodes; this method places more nodes closer to the boundaries of the domain. Depending on the nature of the function to approximate, we can use different types of approximants. Section 4 uses Chebyshev nodes and cubic splines. Cubic splines tend to perform well when approximating a function that has a portion that may not be smooth, and when the order of the approximant is high. The program allows the user to select Chebyshev polynomials as an alternative.

The collocation method and function iteration described above do not directly apply to Eq. (5), because of the presence of the middle term. We need two approximants in our algorithm, one for $\chi(\cdot)$ and the other for $W(\cdot)$. Our approach starts by guessing the control rule at the collocation nodes. Let s^k be the kth collocation node with $s^k < s^{k+1}$ for all $k < K$. The initial guess of the control rule consists of $x_1^{(0)}, \ldots, x_K^{(0)}$. Given a choice of basis functions for the control rule, we can then find the approximant $\hat{\chi}^{(0)}$ that exactly or approximately satisfies $\hat{\chi}^{(0)}(s^k) = x_k^{(0)}$.

Because we do not know W, we also need to guess this function. However, $W^{(0)}$ must be consistent with $\hat{\chi}^{(0)}$. To obtain $W^{(0)}$ we replace the infinite sum in Eq. (6) by a finite sum from 0 to time T_w, where T_w is a large number (i.e. much greater than T). Letting s_t be the value of the state variable at time t when the initial state is $s_t = s^k$ and the control rule $\hat{\chi}^{(0)}$ is used, our guess of the value function at the collocation nodes is

$$W^{(0)}(s^k) \leftarrow \sum_{t=0}^{T_w} \theta_t f(\hat{\chi}^{(0)}(s_t), s_t) \quad \text{s.t. } s_0 = s^k, \quad s_{t+1} = g(\hat{\chi}^{(0)}(s_t), s_t).$$

Given the choice of basis functions for the value function, we then choose the coefficients of the approximant to (exactly or approximately) satisfy $\hat{W}^{(0)}(s^k) = W^{(0)}(s^k)$ for all k. Having the initial guess $\hat{\chi}^{(0)}$ and $\hat{W}^{(0)}$, we start the iteration. We begin each iteration with $\hat{\chi}^{(r)}$ and $\hat{W}^{(r)}$ and update these functions during the iteration. We can evaluate the control rule at each of the approximation nodes by maximizing the right-hand side of the QDPE, Eq. (5):

$$\begin{aligned} \chi^{(r+1)}(s^k) \leftarrow \arg\max_x f(x, s^k) + \sum_{t=1}^{T} (\theta_t - \delta\theta_{t-1}) f(\hat{\chi}^{(r)}(s_t^k), s_t^k) + \delta \hat{W}^{(r)}(s_1^k) \\ \text{s.t. } s_1^k = g(x, s^k), \quad s_{t+1}^k = g(\hat{\chi}(s_t^k), s_t^k) \quad \text{for } t \geqslant 1. \end{aligned}$$

We choose new coefficients of the approximant of the control rule to obtain the approximant $\hat{\chi}^{(r+1)}$. Likewise, we can get $\hat{W}^{(r+1)}$ by evaluating the value function at the collocation nodes with the following equation, and finding the coefficients for its approximant:

$$\begin{aligned} W^{(r+1)}(s^k) \leftarrow \max_x f(x, s^k) + \sum_{t=1}^{T} (\theta_t - \delta\theta_{t-1}) f(\hat{\chi}^{(r)}(s_t^k), s_t^k) + \delta \hat{W}^{(r)}(s_1^k) \\ \text{s.t. } s_1^k = g(x, s^k), \quad s_{t+1}^k = g(\hat{\chi}(s_t^k), s_t^k) \quad \text{for } t \geqslant 1. \end{aligned}$$

The iteration continues until $\chi^{(r+1)}$ and $\chi^{(r)}$ are close enough, and $W^{(r+1)}$ and $W^{(r)}$ are also close enough. Our convergence criterion is

$$\max_{s^k} \{\chi^{(r+1)} - \chi^{(r)}, W^{(r+1)} - W^{(r)}\} \leqslant \text{tol},$$

where tol is a small positive value, the tolerance level. See [7] and the online appendix for further details.

2.3. The steady state

We are often interested in the characteristics of the steady state S^*. We can numerically find the steady state by solving for S in $g(\hat{\chi}(S), S) = S$, where $\hat{\chi}$ is the converged approximant of the control rule. The value of the control variable at the steady state is simply $x^* = \hat{\chi}(S^*)$. Here, in order to find the steady state we first need to approximate the solution to the entire problem.

In control problems with constant discounting, analysis of the steady state is much simpler. There, we use the transversality condition $\lim_{t\to\infty} \delta^t V'(S_t) = 0$ to conclude that the state variable must converge. Providing that it converges to an interior steady state (as is the case for a broad class of problems), we then use the Euler equation and the equation of motion, both evaluated at a steady state. These are two algebraic equations, so their analysis, using numerical or qualitative methods, is straightforward (and does not require the solution to the control problem). Even if there are multiple solutions to these two algebraic equations, those solutions are "locally unique".

The same approach leads to an interval (or possibly several intervals) of candidate steady states in the case of NCPRTP. With NCPRTP, we show in the Appendix that the Euler equation evaluated at a steady state is a Tth order polynomial in χ'^*:

$$f_x^* + g_x^*(f_x^* \chi'^* + f_s^*) \sum_{t=1}^{T} (\theta_t - \delta\theta_{t-1})(g_x^* \chi'^* + g_s^*)^{t-1} + \delta(f_s^* g_x^* - f_x^* g_s^*) = 0. \quad (7)$$

In Eq. (7) the subscripts x and s denote the partial derivatives, and $*$ denotes the value at the steady state. For example, f_x^* is $\partial f(x,s)/\partial x$ evaluated at $(x,s) = (\chi(S^*), S^*)$.

Eq. (7), together with the steady-state condition of the equation of motion, $S^* = g(x^*, S^*)$, comprise two equations in three unknowns, x^*, S^*, and χ'^*. In contrast, with constant discounting ($T = 0$), the summation in Eq. (7) vanishes, eliminating the unknown value χ'^*. With constant discounting we have the standard Euler equation evaluated at the steady state

$$f_x^* + \delta(f_s^* g_x^* - f_x^* g_s^*) = 0. \tag{8}$$

Eq. (8) and the steady-state condition $S^* = g(x^*, S^*)$ comprise two equations in two unknowns, which can be solved to identify steady-state candidates. Thus, under constant discounting, analysis of the steady state requires only that we analyze two algebraic equations, without actually needing to solve the optimization problem.

The necessary condition for asymptotic stability,

$$|g_x^* \chi'^* + g_s^*| < 1, \tag{9}$$

can be used to restrict the range of candidate steady states under NCPRTP. However, the two algebraic equations and the inequality identify a continuum of candidates, rather than a unique candidate or isolated (i.e. "locally unique") candidates.

The economic explanation for the lack of even "local" uniqueness of the steady state is essentially the same as in the differential game literature. There, the cause of the indeterminacy of the steady state is referred to as an "incomplete transversality condition". In the steady state, the current decision-maker needs to consider how a change in her decision, and the resulting change in the state in subsequent periods, would change subsequent decisions, and the effect that these changes would have on her payoff. The subsequent decisions *do not* maximize the payoff of the current decision-maker, so the envelope theorem cannot be invoked, as is done in control problems with constant discounting. Each decision-maker's optimal choice depends on how that choice will alter the decisions of her successors. This dependence holds at every point, including at the steady state. Thus, the value of the steady state, S^*, depends on $\chi'(S^*)$, as Eq. (7) shows.

We can obtain an analogous equation for the higher-order derivatives of the control rule evaluated at the steady state. The online appendix gives the second derivative of the control rule at the steady state, χ''^*. We use this value to check the performance of the approximant, but it does not assist us in identifying a finite set of candidate steady states. In principle, we can also use these higher-order derivatives to construct a Taylor approximation of the control rule that drives the state to a particular steady state. Krussel et al. [19] implement this procedure in the special case of quasi-hyperbolic discounting.

The approximated control rule may appear much smoother than the function actually is. For example, consider a function $h(z) = z + \varepsilon \sin(z/\varepsilon^2)$ where ε is a small positive number; $\hat{h}(z) = z$ approximates $f(z)$ well because $|h(z) - z|$ never exceeds ε. However, $\hat{h}'(z)$ does not approximate $h'(z)$ well, because $|h'(z) - \hat{h}'(z)| = (1/\varepsilon)|\cos(z)|$ can be very large. The QDPE does not explicitly involve the derivatives of the control rule. However, we can use the Euler equation and its derivatives, evaluated at the steady state, to improve or to validate the approximant, as discussed in the following section.

3. Testing the software

We conducted extensive tests of this software, since we know from earlier work that models with quasi-hyperbolic discounting (a simple form of hyperbolic discounting) sometimes do not converge [20]. For our first tests we chose the parameters to give a constant discount rate, and confirmed that the solution is equivalent to that obtained using the packages provided by Miranda and Fackler.

We then tested the software using a linear-quadratic model with quasi-hyperbolic discounting. This specification is convenient, because it admits a closed form solution for the *linear* equilibrium, given in terms of a solution to a cubic equation [16]. We can compare the solution returned by the program with this closed form solution in order to test the algorithm and the code. To check the robustness of the results, we used Chebyshev polynomials in addition to cubic splines.

Use of the linear-quadratic model also provides insight into the question of multiple equilibria, arising from the indeterminacy of the steady state. We can identify the interval of candidate steady states using the requirement of asymptotic stability, Eq. (9). For general functional forms, it is a simple matter to determine this interval numerically. For the linear-quadratic problem we can determine the interval analytically. The steady state to the linear equilibrium is a point in this interval.

For general functional forms there is an interval of values of the state that satisfy the necessary conditions for optimality at a steady state, and that are asymptotically stable. As noted in Section 2.3, it is possible to use the derivatives of the Euler equation, evaluated at a steady state, to obtain higher-order derivatives of the control rule at the steady state.[2] This procedure provides no information about the domain over which that particular control is defined. We therefore attempted to identify these equilibria directly.

To that end, we wrote the program so that the user has the option of specifying the steady state. The user is able to pick a point S^{u} in the interval of candidate steady states and require that this point be a steady state to the equilibrium problem. By imposing the steady state condition $g(\chi(S^{\mathrm{u}}), S^{\mathrm{u}}) = S^{\mathrm{u}}$, we ensure that S^{u} is the steady state. In addition, the user has an option to impose the steady-state Euler condition, Eq. (7), that specifies the (Markov perfect) equilibrium condition for the first derivative of the control rule. Further, the user can check whether the (Markov perfect) equilibrium condition for the second derivative of the control rule is approximately satisfied using a condition given in the online appendix.

Both the steady-state condition and the steady-state Euler condition can be expressed as a linear constraint on the coefficients $c_1^{\chi}, c_2^{\chi}, \ldots, c_N^{\chi}$ for the approximant of the control rule χ. That is, from the steady-state condition, we can find a value of the control rule x^{u} that satisfies $g(x^{\mathrm{u}}, S^{\mathrm{u}}) = S^{\mathrm{u}}$. Thus, the coefficients must satisfy

$$x^{\mathrm{u}} = \sum_{n=1}^{N} c_n^{\chi} \phi_n(S^{\mathrm{u}}).$$

Likewise, the steady-state Euler equation (Eq. (7)) gives the value of χ' at S^{u}. Hence, the coefficients must satisfy the following linear constraint:

$$\chi'(S^{\mathrm{u}}) = \sum_{n=1}^{N} c_n^{\chi} \phi_n'(S^{\mathrm{u}}).$$

When the user-specified conditions are imposed, the coefficients are found using constrained least squares.[3]

The user also has the option of not specifying the steady state. With this option, the program approximates the equilibrium control rule without special consideration to the steady state. We then find the steady state by solving $g(\hat{\chi}(S), S) = S$. We can then use the steady-state condition and the steady-state Euler condition to validate the approximant.

That is, under both options (with and without a user-specified steady state) we use the steady-state condition and Euler equation, Eq. (7). However, we use the second equation in different ways. That condition is imposed with the user-specified steady-state option, and it is merely a diagnostic tool when the program chooses the steady state. The derivative of the Euler equation, which involves second derivatives of the control rule, is used as a diagnostic tool under both options.

We know that for the linear-quadratic problem with quasi-hyperbolic discounting, the linear equilibrium is defined over the entire real line [16], but we have no information about the domains of other equilibria. Therefore, if the user imposes a steady state other than the value corresponding to the linear equilibrium, it is necessary to experiment with the state space. It might be the case that a particular non-linear equilibrium is defined over only a small interval in the neighborhood of the steady state corresponding to that equilibrium.

[2]It is in theory possible to obtain higher derivatives of the control rule for the entire state space by solving higher-order derivatives of the Euler equation. However, numerical solutions are difficult to obtain because the expression generally involves the value of the state variable for the future periods.

[3]It is also possible to impose the equation involving second derivatives (in the online appendix), but we dropped this option for two reasons. First, the higher-order derivatives directly derived from the approximants are in general not very reliable. Second, the additional linear constraints make it more difficult to obtain reliable numerical results using constrained least squares.

Consequently, in order to try to identify non-linear equilibria, it is necessary to experiment with different (e.g. small) definitions of state space.

To reiterate, the purpose in writing the program with the option of a user-specified steady state was to see if the algorithm returns a unique solution, or whether it can return many solutions, each of which corresponds to a particular steady state. When the user *does not* specify the steady state, the program always returns the linear equilibrium, independently of the starting values that we use in the algorithm, and of the choice of the basis function. If the user specifies the steady state to the linear equilibrium, the program returns the linear equilibrium. If the user specifies a steady state within the candidate set (i.e. the points that are asymptotically stable), but not equal to the steady state corresponding to the linear equilibrium, there are two possibilities. If the user *does not* impose the steady-state Euler condition (Eq. (7)), the program converges to a highly non-linear control rule. However, if the user specifies the steady state and also imposes the steady-state Euler condition (as is appropriate, since this conditions must hold in a MPE), the program fails to converge.

In summary, if the user specifies the steady state, the algorithm does not converge unless the user happens to get the "right" steady state, or unless the user neglects the steady-state Euler condition. The algorithm converges when a steady state is not imposed upon it. That is, for these experiments, we find that our algorithm identifies a unique equilibrium. Perhaps this outcome should not be surprising. First, the linear equilibrium is the only equilibrium function within the class of finite polynomials in S that is defined over the entire real line.[4] Second, recent years have seen the development of algorithms to obtain MPE in dynamic games (e.g. [24]) which are designed to be used to estimate parameters in Industrial Organization models. When authors of these and related papers mention uniqueness, they note that although there may be multiple equilibria to the games they study, they report that for their experiments their algorithm obtains a unique equilibrium. We know of no theory that explains why these kinds of algorithms return unique equilibria.

We repeated searches for multiple equilibria employing the user-specified steady-state option and using other functional forms for which we are not able to identify any closed form equilibrium. We did not find multiple equilibria. Thus, our algorithm has not proved useful for exploring the multiplicity of equilibria. However, the fact that in our experiments it returns a unique equilibrium, and that this equilibrium is the "natural" one (i.e. the linear equilibrium of the linear-quadratic problem) makes the algorithm well suited for policy experiments. Hereafter, we allow the program to identify the steady state; that is, we do not employ the user-specified steady-state option. An earlier version of this paper describes extensive experiments with a fishery model using the Schaefer growth function.

4. The application to climate change

We use our QDPE solver to determine the MPE under NCPRTP in a model of climate change based on the Stern Review [26]. We want to investigate the sensitivity of the optimal policy to different levels and forms (constant versus non-constant) of discounting. We explain how we calibrate the model, then discuss steady states, and then provide comparisons of policy functions and equilibrium trajectories.

4.1. Model calibration

We use runs from the PAGE model, reported in [26], to calibrate a simple deterministic model with one state variable.[5] We select the Stern Review runs that use their baseline climate scenario and economic impacts inclusive of market impacts, risk of catastrophe and non-market impacts. The Stern Review takes into account uncertainty using repeated simulations, but we use only the mean estimates to calibrate our deterministic model.

[4]The proof of this assertion is in the online manual.

[5]Our model cannot reproduce the complicated dynamics that can arise with multi-state variable models, particularly features that capture the inertia in the climate system. We merely use the PAGE results to calibrate a model that captures the "orders of magnitude" implicit in that model.

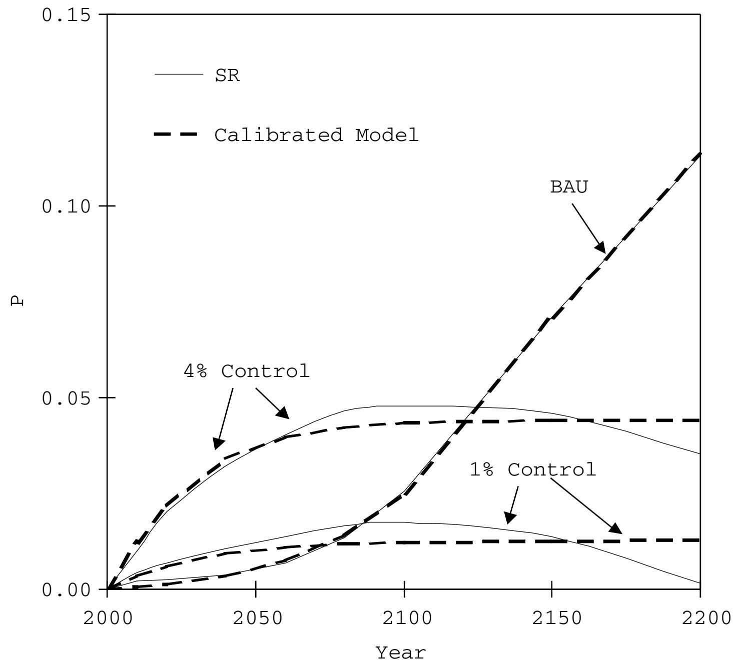

Fig. 1. Time evolution of P_t for the SR and calibrated models under various scenarios.

We assume that in the absence of the climate change (i.e. no damage occurs even under the business-as-usual (BAU) scenario), the output per capita (GWP) y grows at a constant rate g. We also assume that under BAU the saving rate s is constant so that (per capita) consumption at time t is $c_t = (1-s)y_t$. We normalize y so that the base year consumption is equal to unity. Hence, $c_t = (1+g)^t$ in the absence of climate change.

The state variable in our model, $P_t \in (-\infty, 1]$, is the fraction of GWP lost as a result both of climate change and past efforts (e.g. re-directing savings) to reduce that change. With this notation, $(1+g)^t(1-P_t)$ is the fraction of GWP available for consumption. The control variable, x_t, is the fraction of this amount that is diverted from consumption to spend on climate change mitigation and avoidance in the current period; current consumption is $c_t = (1+g)^t(1-P_t)(1-x_t)$.

The solid lines in Fig. 1 show three Stern Review trajectories for P_t (which they call the "output gap") when x is held constant at 0%, 1% and 4%. The Stern Review provides the values of P under these three scenarios for eight years, 2010, 2020, 2040, 2060, 2080, 2100, 2150 and 2200. We use these $8 \times 3 = 24$ "data" points for our calibration.

Initially, the graph of P is higher under 4% expenditures than under either 1% or under BAU, illustrating the possibility that large current expenditures on mitigation and avoidance can actually reduce future GWP. The Stern Review does not explain this feature of the model. One explanation for it follows from the observation that expenditures on climate-related policy are financed by both a reduction in consumption and a reallocation of savings (which in our setting is a constant fraction of output); these two sources are likely to be positively correlated. A larger value of x reduces current consumption. If x is positively correlated with the fraction of the total amount of savings devoted to climate policy, then a large x reduces investment in non-climate-related projects. The loss in future production due to the reduction in future capital may be greater than the gain in future production due to lower climate-related damage. For our purposes, the important consequence of this feature is that it puts an upper bound on the optimal level of x, regardless of the discount rate.

Under BAU, P rises to about 0.13; in this case, climate change reduces consumption by over 13% by the end of the next century. The solid curves in Fig. 1 turn down because of the Stern Review's assumption that damages end at year 2200. We adopt a calibration in which damages eventually level off.

Table 1
Calibrated parameter values

Parameter	Value
a_1	3.8746×10^{-4}
a_2	1.9700
a_3	0.73819
b_1	-1.8241×10^{-2}
b_2	3.1511×10^{-4}
b_3	0.21807

To this end, we use a discrete time version of the standard ''S-shaped'' adoption–diffusion model. In the absence of control, this model implies that the growth equation for P is

$$P_{t+1} = \frac{P_t + a_1}{a_2 P_t + a_3}.$$

Based on convenience and goodness of fit with our calibration points, we chose the following functional form for including the effect of control, x:

$$P_{t+1} = \frac{P_t + a_1 + b_1 x_t P_t/(x_t^2 + b_2) + b_3 x_t}{a_2 P_t + a_3} = g(x_t, P_t). \tag{10}$$

The dashed lines in Fig. 1 show the values of P using our calibration, for the three scenarios (constant $x \in \{0, 0.01, 0.04\}$). Table 1 reports the values of the parameters that minimize the sum of squared residuals at the data points.

Our calibration almost exactly reproduces the evolution of P_t in the Stern Review model (solid line) under BAU. The fit for the mitigation scenarios (1% Control and 4% Control) is not as good, but this is due to our decision to use a model in which damages do not fall (when x is held constant). However, our calibrated trajectories lie well within the 90% confidence interval for the Stern Review estimates of damage (not shown, to avoid clutter).

In a steady state, P_t and x_t are constant, so GWP and consumption grow at the rate of g. The Stern Review assumes that the economy grows at a constant rate of 1.3% per annum after year 2200. We use the same growth rate; taking one period to be a decade, we set $g = 0.13787$.

We define $\Delta = 1 - (1 - P)(1 - x)$, the aggregate consumption loss due to mitigation expenditures and remaining climate-related damage. Hereafter we evaluate this function only at a steady state. Fig. 2 shows the graph of Δ in the steady state. With our calibration, this function reaches a minimum at $x = 0.00845$, where $\Delta = 0.022$. These are the optimal steady-state values in the limiting problem in which the constant discount rate approaches 0. Here, society spends slightly more than 0.8% of its income on climate policies, maintaining a constant level of climate-related damage such that the aggregate consumption loss equals 2.2% of the level absent all climate-related effects. Any problem with a positive discount rate leads to a steady state at a lower level of x and a higher level of Δ. The graph of Δ falls rapidly for very small levels of x, where a small increase in expenditures leads to a large fall in steady-state damages. Thus, our calibration represents a world in which small expenditures yield large future benefits—the situation described in [26].[6]

As in [26], we assume that welfare in period t depends only on current consumption, and the utility function is isoelastic:

$$U_\eta(c_t) = \frac{c_t^{1-\eta} - 1}{1 - \eta} \quad (\eta \neq 1).$$

[6]Our calibration implies that the steady-state cost of doing too little is much greater than the cost of doing too much. If we were to extend this model by including uncertainty, and if (in the interest of simplicity) the goal were to maximize the expected steady-state flow of welfare, it would be optimal to set $x > 0.00845$.

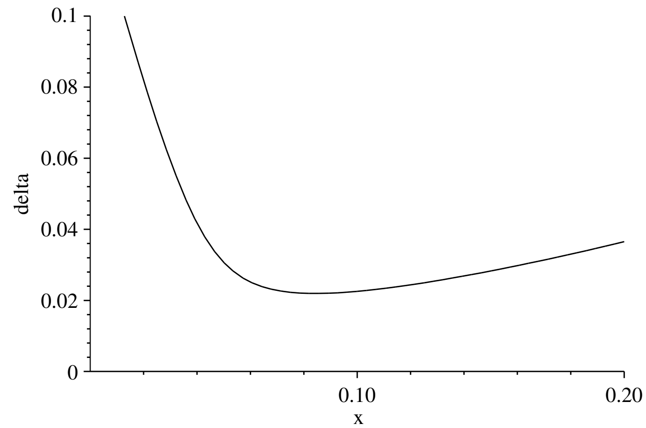

Fig. 2. The graph of Δ in the steady state using parameters in Table 1.

Table 2
The values of the state variable P^*, the control variable x^*, and the loss in consumption Δ in the steady state under different discounting schemes

PRTP	$\eta = 1$			$\eta = 2$		
	P^*	x^*	Δ	P^*	x^*	Δ
Const. 0.1%	0.0137	0.0084	0.0220	0.0147	0.0076	0.0222
Const. 3.0%	0.0165	0.0068	0.0232	0.0185	0.0062	0.0246
NCPRTP	0.0151	0.0074	0.0224	0.0171	0.0066	0.0235

For $\eta = 1$, we define $U_1(c_t) = \ln c_t$. The present-discounted utility W when $\eta \neq 1$ is

$$W = \sum_{t=0}^{\infty} \theta_t U(c_t) = \sum_{t=0}^{\infty} \theta_t \frac{[(1+g)^t (1-P_t)(1-x_t)]^{1-\eta} - 1}{1-\eta} = \sum_{t=0}^{\infty} \tilde{\theta}_t \frac{[(1-P_t)(1-x_t)]^{1-\eta} - 1}{1-\eta},$$

where $\tilde{\theta}_t = \theta_t (1+g)^{t(1-\eta)}$ is the utility discount factor accounting for growth. The PRTP can be variable, with θ_t equal to the utility discount factor (exclusive of growth).

4.2. Steady states

We consider two benchmark cases with constant PRTP equal to either 0.1% per annum ($\theta_t = 0.99005^t$) and 3% per annum ($\theta_t = 0.74409^t$). Ref. [26] uses a PRTP of 0.1% and other integrated assessment models use a 3% rate. (Recall our comment in Section 2.1 that the steady state corresponding to a constant PRTP of 0.1% is also the steady state of the "full commitment" equilibrium to the problem with NCPRTP.) We also take two values of $\eta \in \{1, 2\}$. Ref. [26] used $\eta = 1$; some authors state that $\eta = 2$ is a better choice. For the non-constant discounting (NCPRTP) scenario, we use $3/(100 + 145t)$ for the per annum PRTP, which declines from 3% in $t = 0$ (Year 2000) to the terminal rate of 0.1% in $t = 20 (= T)$ (Year 2200). We use Chebyshev nodes and cubic splines, and set both the number of collocation nodes (K) and the number of basis functions for the approximant (N) equal to 200.[7]

The first two lines of Table 2 show the steady-state values of P and x for various combinations of a (constant) PRTP and η; the third line shows these values under NCPRTP. We solved all models using our

[7]Neither our solver nor the CompEcon toolbox converged even under constant discounting, when we used Chebyshev polynomials (rather than cubic splines) for the basis functions. The polynomial is not able to approximate the kink in the control rule. Our results are not sensitive to the choice of K and N.

QDPE solver. We found that defining carefully appropriate bounds for P_t and x_t was crucial for obtaining convergence, because the approximated values of W and χ are generally not reliable when they are extrapolated to outside the range of the evaluation nodes.

Consider first the steady states under constant discounting. The steady-state levels of expenditure are moderately sensitive to the discount rate. With $\eta = 1$, increasing the constant annual discount rate from 0.1% to 3% decreases the level of expenditures by $((0.0084 - 0.0068)/0.0084)100 = 19\%$. (The levels of the state and control are small fractions of a very large number—GWP; therefore, it is more informative to speak of the percentage changes in these values, rather than the difference in levels.) The steady-state measure of climate-related damage is considerably smaller (in percentage terms) at the lower discount rate. However, the two effects nearly offset each other, so that the aggregate steady-state loss increases by only $((0.0232 - 0.0220)/0.022)100 = 5\%$ when we increase the discount rate. Thus, the steady-state climate-related costs are quite insensitive to the discount rate.

The steady-state costs are also quite insensitive to η, and the direction of the change is as expected. The growth-inclusive discount rate is $\delta^{-1}(1+g)^{(\eta-1)} - 1$, so a larger value of η has an effect similar (but not identical) to an increase in the PRTP. For the 0.1% constant discount rate, the growth-inclusive discount rate rises from 0.1% to 1.40% as η increases from 1 to 2. This change reduces the steady-state level of expenditures by $((0.0084 - 0.0076)/0.0084)100 = 9.5\%$. Due to the offsetting increase in P, the flow of damages increases by $((0.0222 - 0.0220)/0.0022)100 = 9.1\%$.

Not surprisingly, the steady-state levels under NCPRTP lie between the levels with the constant discount rates. Again, there is a moderate difference in the level of the control variable and the state with different types of discounting, and a smaller difference in the steady-state flow of damages.

Fig. 2 provides the key to understanding these results. Moderate levels of expenditure (levels less than 0.5% of income) produce very large gains relative to BAU, so it is optimal to spend at least a moderate amount, even under fairly high discount rates. In the neighborhood of the minimum of Δ, even a significant (percent) change in the level of expenditures leads to a small change in welfare. Therefore, the flows of welfare are quite insensitive to changes in the discount rate (within the range of 0.1–3%).

The Stern Review estimates that spending 1% every year would almost completely eliminate the damage from climate change. Our steady-state estimates suggest that a 0.6–0.7% expenditure in the steady state is optimal. Since our deterministic model considers only a moderate damage scenario, and does not account for (very) low probability catastrophic events, our recommendations are in line with [26]. This relatively low level of expenditures reduces total losses from over 13% (under BAU) to about 2.3–2.6%.

In order to test the sensitivity of these steady-state results with respect to our calibration, we first increased and then decreased each of the six parameters in Eq. (10) by 5%, holding all other parameters constant. The largest change in Δ occurred when we decreased a_3, causing a 12% increase in Δ relative to the baseline (from 0.0232 to 0.0260). Most of the changes were much smaller. The full set of sensitivity results is in the online appendix. These results suggest that the results are reasonably robust to changes in the calibration: the elasticity of Δ with respect to any of the parameters is bounded by $\frac{12}{5} = 2.4$.

4.3. The control rules and trajectories

Fig. 3 shows the graph of the control rule for $\eta = 1$ as P ranges between 0% and 3%, under the three discounting scenarios (discount rates of 0.1% and 3% and a non-constant discount rate ranging between these two levels). The control rule for $\eta = 2$ is qualitatively similar. In all cases it is optimal to wait until the state reaches a threshold level until beginning to spend on climate policy.[8] A lower discount rate lowers the threshold and increases the optimal level of expenditures when these are positive. The control rule under non-constant discounting lies between the two rules for constant discounting.

[8]The use of a single state variable means that our model does not capture the complex dynamics that actually arise with climate change, particularly the role of inertia. The use of a deterministic model means that the model does not adequately deal with (very) low probability catastrophic events. For both of these reasons, we reject the model's prescription to wait until the state has reached a threshold before acting.

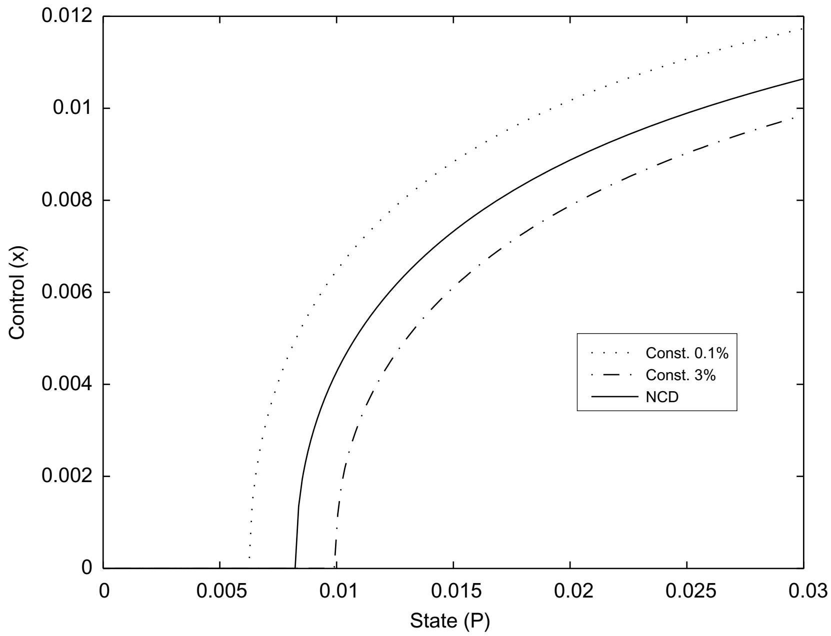

Fig. 3. Control rule under constant (0.1% and 3%) and non-constant discounting when $\eta = 1$.

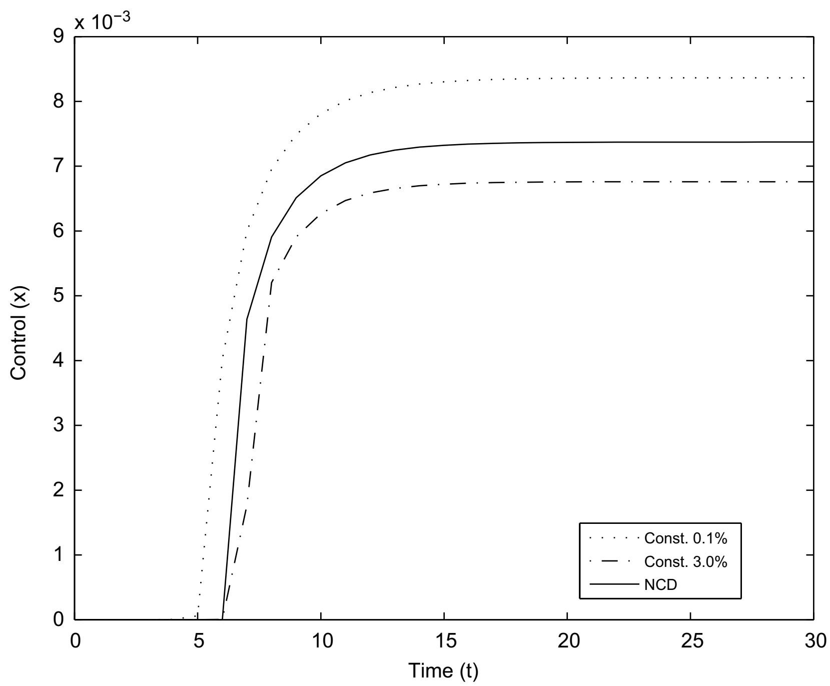

Fig. 4. Time evolution of control variable.

Figs. 4 and 5 show the time evolution of the state and control variable under the three discounting schemes when $\eta = 1$. The qualitative nature of these graphs is similar when $\eta = 2$. The broken line in Fig. 5 corresponds to the case of the BAU (no control). In our model, the optimal (or equilibrium) control policy—just like the steady state—is not very sensitive to the choice of discount rate. In all cases, no control is taken until around Year 2050 ($t = 5$). After that, expenditures are phased in and reach almost the steady-state level

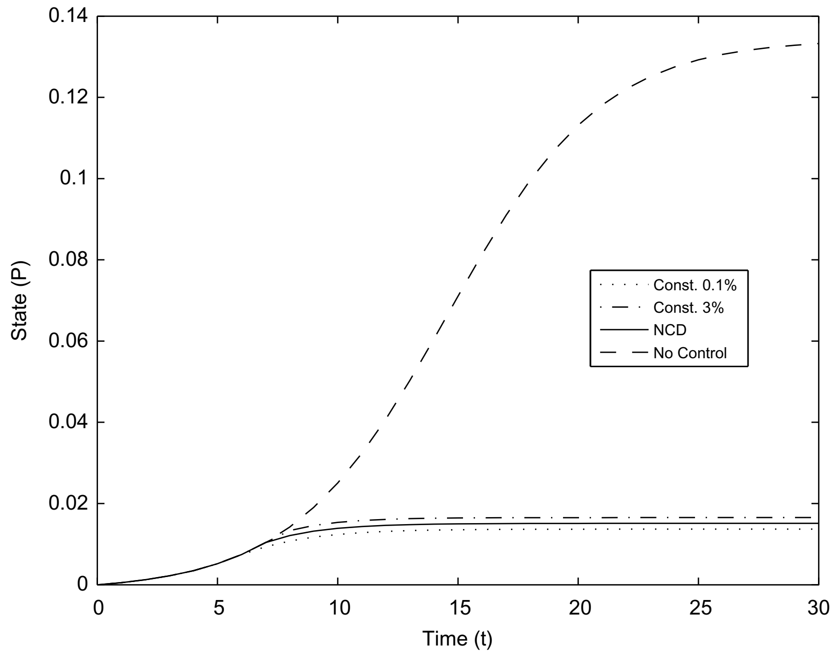

Fig. 5. Time evolution of state variable.

by around Year 2100 ($t = 10$). In the BAU scenario, the output gap (the vertical distance of the broken line) is already sizable by Year 2100, and continues to grow until Year 2300, at which time it is close to its steady-state level. Under the optimal control policy, the state is close to the steady state by 2200. Therefore, our model prediction is consistent with the Stern Review assumption that consumption grows at a constant rate after 2200 in the presence of control.

By substituting the steady-state values of S, x into Eq. (8) and solving for δ we obtain a constant discount rate that supports the steady state obtained under NCPRTP. We refer to this value as the steady-state-equivalent (SSE) discount factor. For our discounting function, the per annum SSE discount rate is 1.79% for $\eta = 1$ and 2.05% for $\eta = 2$. We computed the control rule under the SSE discount factor and compared it to the MPE control rule under NCPRTP. By construction, these rules are equivalent at the steady state, but they are also virtually indistinguishable even outside the steady state. We obtained a similar result for the fishery example used to test the software, although in that case the difference between the two control rules was slightly larger. This evidence suggests that we can obtain a good approximation to the MPE under NCPRTP by using a constant discount rate between the short-run and the long-run rates under NCPRTP. This result was not a foregone conclusion. The introduction of non-constant discounting not only involves a change in the discount rate, but also a change in the nature of the equilibrium—from an optimization problem to a game amongst a succession of regulators.

5. Summary and discussion

Models with non-constant discount rates are potentially useful in studying environmental and resource problems where current actions have long-lived effects. These models make it possible to include non-negligible pure rates of time preference for the near and middle term, while allowing the long-run rate to become small. This flexibility satisfies ethical criteria, giving the future a non-negligible weight. It reflects the view that the current generation does not differentiate (much) between two generations in the distant future—even if it does differentiate between the current and the next generation. At the same time, this flexibility enables us to calibrate a social discount rate that matches observed short and medium term market rates. Despite the importance of this model in environmental and resource problems, it has seldom been applied, owing largely to the difficulty of analyzing subgame perfect equilibria.

We developed and programmed a method to find a MPE when the decision-maker has non-constant PRTP. The procedure uses the standard collocation method and function iteration, but requires that we iterate with both the value function and control rule. An important objective of this paper is to describe and publicize this software. We think that its availability will promote further applied research that uses non-constant PRTP.

The program can be applied using any finite sequence of non-constant PRTP, followed by an infinite sequence of constant rates. In our model, the PRTP approaches a constant in finite time, a parameter in the model; by allowing that parameter to be large, we approximate the model in which the PRTP approaches a constant only asymptotically.

Previous theoretical results show that there is a continuum of candidate steady states that are asymptotically stable and that satisfy the equilibrium conditions at the steady state. In experiments using three functional forms (the linear-quadratic model, the fishery model, and the climate change model) and a range of parameter values our program returns a unique equilibrium. For the linear-quadratic case, this equilibrium is linear in the state. It is the equilibrium to the limiting problem, obtained by taking the limit of the finite horizon problem, letting the horizon approach infinity. In this respect at least, the program returns the "natural" equilibrium.

We used our numerical package to address the current controversy over the role of discounting in determining socially optimal climate change policy. This controversy arises because in some circumstances optimal policy is very sensitive to the choice of the PRTP and the elasticity of marginal utility. Moreover, the choice of these parameters depends on value judgments. In some cases, however, the optimal policy may be quite insensitive to the social discount rate. We constructed a model that is consistent with the orders of magnitude for abatement costs and environmental damage used for the numerical policy recommendations in [26].

In our calibration damages fall very rapidly with small expenditures, and then flatten out and then rise with larger expenditures. In this circumstance it may be optimal to be near the bottom of the damage curve under a broad range of discount rates. Lowering the discount rate lowers the equilibrium level of damages by a small amount in the neighborhood of the minimum of the damage curve. Because it is possible to eliminate most damages with moderate expenditures, and larger expenditures do little good, the optimal policy is quite insensitive to the discount rate.

We solved the model for a constant 0.1% and a 3% per annum PRTP (leading to a standard control problem) and under a PRTP that falls from 3% to 0.1% (requiring the solution to an equilibrium problem, and the use of our software). The MPE with non-constant discounting lies between the optimal solutions with the two constant discount rates. Moreover, it is possible to approximate the MPE using a constant discount rate between the two exogenously given rates. This result means that for applied work there may be little cost in using a constant moderate discount rate to approximate a problem with a non-constant discount rate.

Our results do not imply that "the discount rate does not matter very much" in general. The results do show that in some circumstances the discount rate may not be as important a parameter as is commonly believed. In those circumstances, the choice between constant and non-constant discounting may also be relatively unimportant.

It is simple enough to vary the constant PRTP in complex integrated assessment models. For a model in which the discount rate is very important, a non-constant discount rate is useful for capturing both the short-run and long-run objectives. However, it is not straightforward to obtain a subgame perfect (e.g. Markov perfect) equilibrium to complex models with non-constant discounting. By using a simple model that captures some features of the complex integrated assessment model, we can at least get an idea of the level of the constant discount rate that reflects both the short- and long-run objectives. That constant rate can then be used in the original integrated assessment model, to obtain an approximation of the problem with non-constant discounting.

Acknowledgments

We thank two anonymous referees for their comments on an earlier version of this paper. We also thank Simon Dietz and Dimitri Zenghelis for their information about the PAGE model. Tomoki Fujii thanks the Lee Foundation for financial support. The usual disclaimer applies.

Appendix A. Derivation of Euler equation

Here, we derive the Euler equation corresponding to Eq. (5). First, by differentiating $S_{t+1} = g(\chi(S_t), S_t)$ with respect to S_t, we have

$$\frac{dS_{t+1}}{dS_t} = g_x(\chi(S_t), S_t)\chi'(S_t) + g_s(\chi(S_t), S_t),$$

where the partial derivatives are denoted by subscripts. With a little abuse of notation, we denote $g(t) = g(\chi(S_t), S_t)$. We use similar shorthand notation for f and the derivatives. Also, we denote $\chi(t) = \chi(S_t)$. Then, we have

$$\frac{dS_t}{dS_1} = \frac{dS_t}{dS_{t-1}} \cdot \frac{dS_{t-1}}{dS_{t-2}} \cdots \frac{dS_2}{dS_1} = \prod_{\tau=1}^{t-1}[g_x(\tau)\chi'(\tau) + g_s(\tau)] \quad (t \geqslant 2)$$

and

$$\frac{dS_1}{dx} = g_x(0).$$

Using the notation above, we can write Eq. (5) as

$$W(S) = \max_x \left\{ f(x, S) + \sum_{t=1}^{T}(\theta_t - \delta\theta_{t-1})f(t) + \delta W(S_1) \right\}.$$

The first-order condition is

$$f_x(0) + \left[\sum_{t=1}^{T}(\theta_t - \delta\theta_{t-1})(f_x(t)\chi'(t) + f_s(t))\frac{dS_t}{dS_1} + \delta W'(S_1) \right] g_x(0) = 0. \tag{11}$$

By the envelope theorem,

$$W'(S) = f_s(0) + \left[\sum_{t=1}^{T}(\theta_t - \delta\theta_{t-1})(f_x(t)\chi'(t) + f_s(t))\frac{dS_t}{dS_1} + \delta W'(S_1) \right] g_s(0).$$

Applying the first-order condition, Eq. (11), we obtain

$$W'(S) = f_s(0) - \frac{f_x(0)g_s(0)}{g_x(0)}.$$

Then, advancing one period, we have

$$W'(S_1) = f_s(1) - \frac{f_x(1)g_s(1)}{g_x(1)}.$$

Substituting this expression into the first-order condition Eq. (11), we have

$$f_x(0) + \left[\sum_{t=1}^{T}(\theta_t - \delta\theta_{t-1})(f_x(t)\chi'(t) + f_s(t))\frac{dS_t}{dS_1} + \delta\left(f_s(1) - \frac{f_x(1)g_s(1)}{g_x(1)} \right) \right] g_x(0) = 0. \tag{12}$$

Eq. (12) is the *Euler equation* for the QDPE Eq. (5). In a steady state, we have $g(0) = g(1) = \cdots = g^*$. Using similar notation for other functions and variables, we have

$$\left.\frac{dS_t}{dS_1}\right|_{(\chi(S^*), S^*)} = (g_x^*\chi'^* + g_s^*)^{t-1}.$$

Substituting this expression into Eq. (12) and arranging the terms, we obtain Eq. (7).

References

[1] O. Azar, Rationalizing hyperbolic discounting, J. Econ. Behav. Organ. 38 (1999) 245–252.
[2] G. Chichilnsky, An axiomatic approach to sustainable development, Soc. Choice Welfare 13 (1996) 231–257.
[3] M.L. Cropper, S.K. Ayded, P.R. Portney, Preferences for life-saving programs: how the public discounts time and age, J. Risk Uncertainty 8 (1994) 243–266.
[4] P. Dasgupta, Commentary: the Stern review's economics of climate change, Nat. Inst. Econ. Rev. 199 (1) (2007) 4–7.
[5] E. Dockner, S. Jorgensen, N. Van Long, G. Sorger, Differential Games in Economics and Management Sciences, Cambridge University Press, Cambridge, UK, 2000.
[6] P.H. Dybvig, J.E. Ingersoll, S.A. Ross, Long forward and zero coupon rates can never fall, J. Bus. 69 (1996) 1–25.
[7] T. Fujii, Manual for discrete time continuous state quasi-dynamic programming equation solver, Singapore Management University, 2008.
[8] C. Gollier, Discounting an uncertain future, J. Public Econ. 85 (2002) 149–166.
[9] C. Gollier, Time horizon and the discount rate, J. Econ. Theory 107 (2002) 463–473.
[10] C. Gollier, Consumption-based determinants of the term structure of discount rates, Mathematics Finan. Econ. 1 (2) (2007) 81–101.
[11] C. Gollier, R. Zeckhauser, Aggregation of heterogeneous time preferences, J. Polit. Economy 113 (2005) 878–896.
[12] B. Groom, C. Hepburn, P. Koundouri, D. Pearce, Declining discount rates: the long and the short of it, Environ. Resource Econ. 32 (2005) 445–493.
[13] C. Harris, D. Laibson, Dynamic choices of hyperbolic consumers, Econometrica 69 (5) (2001) 935–957.
[14] G. Heal, Intertemporal welfare economics and the environment, in: K.-G. Maler, J. Vincent (Eds.), The Handbook of Environmental Economics, North-Holland, Amsterdam, 2001.
[15] K.L. Judd, Numerical Methods in Economics, MIT Press, Cambridge, MA, 1998.
[16] L. Karp, Global warming and hyperbolic discounting, J. Public Econ. 89 (2005) 261–282.
[17] L. Karp, Non-constant discounting in continuous time, J. Econ. Theory 132 (2007) 557–568.
[18] L. Karp, Y. Tsur, Time perspective, discounting and climate change policy, 2007, Unpublished working paper ⟨http://are.Berkeley.EDU/karp/⟩.
[19] P. Krusell, B. Kuruscu, A. Smith, Equilibrium welfare and government policy with quasi-geometric discounting, J. Econ. Theory 105 (2002) 42–72.
[20] D. Laibson, A. Repetto, J. Tobacman, Self control and saving for retirement, Brookings Pap. Econ. Act. 1998 (1998) 91–196.
[21] C.Z. Li, K.G. Lofgren, Renewable resources and economic sustainability: a dynamic analysis with heterogeneous time preferences, J. Environ. Econ. Manage. 40 (2000) 236–250.
[22] M.J. Miranda, P.L. Fackler, Applied Computational Economics and Finance, MIT Press, Cambridge, MA, 2002.
[23] W.D. Nordhaus, A review of the Stern review on the economics of climate change, J. Econ. Lit. 45 (3) (2007) 686–702.
[24] A. Pakes, P. McGuire, Computing Markov-perfect Nash equilibria: numerical implications of a dynamic differentiated product model, RAND J. Econ. 25 (4) (1994) 555–589.
[25] M. Rabin, Psychology and economics, J. Econ. Lit. 36 (1998) 11–46.
[26] N. Stern, Stern review on the economics of climate change, Technical Report, HM Treasury, UK, 2006.
[27] S. Tsutsui, K. Mino, Nonlinear strategies in dynamic duopolisitc competition with sticky prices, J. Econ. Theory 52 (1990) 136–161.
[28] M. Weitzman, Gamma discounting, Amer. Econ. Rev. 91 (1) (2001) 260–271.
[29] M. Weitzman, A review of the Stern review on the economics of climate change, J. Econ. Lit. 45 (3) (2007) 703–724.
[30] M. Weitzman, Subjective expectations of asset price puzzles, Amer. Econ. Rev. 97 (4) (2007) 1102–1130.