

Global warming and hyperbolic discounting

Larry Karp

207 Giannini Hall, University of California, Berkeley, CA 94720, USA

Received 17 December 2002; received in revised form 20 January 2004; accepted 10 February 2004

Available online 18 August 2004

Abstract

The use of a constant discount rate to study long-lived environmental problems such as global warming has two disadvantages: the prescribed policy is sensitive to the discount rate, and with moderate discount rates, large future damages have almost no effect on current decisions. Time-consistent quasi-hyperbolic discounting alleviates both of these modeling problems, and is a plausible description of how people think about the future. We analyze the time-consistent Markov Perfect equilibrium in a general model with a stock pollutant. The solution to the linear-quadratic specialization illustrates the role of hyperbolic discounting in a model of global warming.

© 2004 Elsevier B.V. All rights reserved.

JEL classification: D83; L50

Keywords: Stock pollutant; Hyperbolic discounting; Global warming; Time consistency

1. Introduction

There are two important consequences of using a constant discount rate to model the control of long-lived environmental stocks such as greenhouse gasses. First, the optimal program is likely to be sensitive to the discount rate; a parameter about which there is some disagreement. Second, discounting at a non-negligible rate makes the present value of future damages small. The effects of greenhouse gasses might not be felt for a century (if ever). At an annual discount rate of 1%, we would invest 37 cents today to avoid a dollar's worth of damages in a century, and at a discount rate of 4% that amount falls to 1.8 cents. These values differ by a factor of more than 20. The corresponding values if damages occur after 10 years rather than after a century are 90 and 67 cents, numbers that

E-mail address: karp@are.berkeley.edu (L. Karp).

differ by a factor of less than 1.4. The cost-benefit ratio for investments related to global warming may be largely determined by the discount rate.

It seems reasonable to apply a non-negligible discount to the future, but using a constant and non-negligible discount rate make us callous toward the far-distant future. An obvious remedy is to use a declining discount rate. This remedy introduces the problem of time-inconsistency. A time-consistent equilibrium can be studied, but even with strong assumptions such as Markov beliefs and differentiable policies, this equilibrium is typically not unique in an infinite horizon problem. Despite this limitation, discounting is sufficiently important in problems of long-lived pollutants that it is worth considering carefully an alternative to constant discounting. This paper is a step towards providing that analysis.

Section 2 reviews the literature, and explains why hyperbolic discounting is a useful way to model long-lived environmental problems. Section 3 uses [Harris and Laibson \(2001\)](#)'s heuristic method to derive the Euler equation in a stock-pollutant model. This section provides sufficient conditions under which the expected stock either oscillates or changes monotonically. The section then discusses the non-uniqueness and the Pareto ranking of the set of differentiable Markov perfect equilibria (MPE). Section 4 specializes the model to linear-quadratic functions and presents the equations that determine the (linear) equilibrium decision rule. A calibration that represents plausible magnitudes of costs and benefits associated with global warming shows that hyperbolic discounting provides a useful method of studying long-lived environmental problems.

To the extent that a social planner has a declining discount rates, the analysis here is positive. The paper may also contribute to the large normative literature on the optimal control of greenhouse gasses, which assumes a constant discount rate. If a decreasing discount rate provides a better description of society's preferences, and if the time-consistency problem is important, then the model studied here is useful for normative analysis.

2. Literature review

[Arrow et al. \(1996\)](#) suggest that the appropriate discount rate for environmental damages in the distant future depends on whether the modeling exercise is "descriptive" or "prescriptive". They conclude that in the former case, it is appropriate to use a market rate of interest, typically in excess of 7%; in the latter case, a social discount rate no greater than 3% should be used. The collection of papers edited by [Portney and Weyant \(1999\)](#)—in particular, [Dasgupta et al. \(1999\)](#); [Cline \(1999\)](#)—provides a variety of perspectives on this issue. [Heal \(1998, 2001\)](#) examines the effect of different kinds of discounting in environmental contexts. [Frederick et al. \(2002\)](#) review the genesis of models based on discounted utility, and they survey the empirical literature that measures individuals' discount rates.

[Cropper and Laibson \(1999\)](#) suggest using hyperbolic discounting to evaluate payoffs under global warming. They use [Phelps and Pollack \(1968\)](#)'s model in which an individual chooses a time-profile of consumption, subject to a growth rate for capital. We use a similar idea, but modify the model so that it describes a situation in which the accumulation of a pollution stock causes future economic damages.

In a continuous time setting, where $U(c_t)$ is the social utility of consumption and $\phi(t)$ is the discount factor for consumption, the payoff at time t is $\int_0^\infty \phi(\tau)U(c_{t+\tau})d\tau$. The present value today of \$1 additional consumption τ units of time in the future is $\phi(\tau)U'(c_{t+\tau})$. The *social discount rate*, $r(t)$, equals the negative of the rate of change of the present value of future marginal utility of consumption:

$$r(t) = \frac{-d\ln(\phi(t)U'(c_t))}{dt} = \xi(t) + v(c_t) \frac{\dot{c}(t)}{c(t)}, \quad (1)$$

where $\xi(t)$ is the pure rate of time preference and $v(c)$ is the elasticity of marginal utility of consumption. In standard usage, hyperbolic discounting refers to a falling pure rate of time preference. This paper interprets hyperbolic discounting as a declining social discount rate.

Rabin (1998) describes the psychological basis for a declining pure rate of time preference. Ainslie (1991); Cropper et al. (1994); Kirby and Herrnstein (1995) provide empirical evidence that individuals actually discount the future in this manner. Read (2001); Rubinstein (2003) offer other interpretations of this evidence. Rubinstein presents experimental evidence that is not consistent with either constant or hyperbolic discounting, but is consistent with a decision-making procedure based on “similarity relations”. This procedure assumes that individuals ignore small differences and focus on large differences when comparing two alternatives.

Hyperbolic discounting and “similarity relations” models have important differences, but they have in common the idea that a decision-maker’s ability to distinguish between the levels of characteristics of alternatives is important in making a choice. Hyperbolic discounting assumes that the ability to make distinctions diminishes for more distant events. For example, an individual might prefer \$1 today to \$2 tomorrow; in a short time frame, a single day is an appreciable delay. The same individual might prefer to receive \$2 in 10 years and 1 day rather than \$1 in 10 years; over a long time frame, the elapse of a single day is nearly irrelevant.

This idea is compelling when considering long-lived environmental problems. We may feel appreciably closer to our children than to our grandchildren, and therefore be willing to discount the welfare of the second generation. It seems plausible that there is a smaller difference in our emotional attachment to the 10th relative to the 11th future generation. In that case, our future rate of time preference is lower. Two successive generation in the distant future appear more similar to the current generation, compared to two successive generations in the near future.

Eq. (1) shows that even if the pure rate of time preference is constant the social discount rate could change, if for example, the growth rate of consumption changes or the elasticity of marginal utility changes with consumption. In a stationary model with a constant rate of time preference, Gollier (2002) provides sufficient conditions for a declining yield curve; this implies a falling social discount rate.

Weitzman (2001) suggests an additional rationale for using a decreasing social discount rate. Suppose that there actually exists a constant discount rate, r , that is “correct” for the specific modeling objective; the value of r is unknown, so it makes sense to treat it as a random variable. Define $\pi(t) \equiv E_r e^{-rt}$ as the (subjective) expectation of the social discount factor, and $\vartheta(t) = -\frac{d\ln\pi(t)}{dt}$ as the corresponding

social discount rate. Weitzman shows that $\vartheta(t)$ is decreasing when r has a gamma distribution. A previous draft of this paper shows that $\vartheta(t)$ is decreasing when r has an *arbitrary* discrete distribution. The intuition for this result is that as t increases, smaller values of r in the support of the distribution are relatively more important in determining the expectation of e^{-rt} . An alternative, but formally equivalent interpretation of Weitzman's model is that there are agents with different but constant discount rates. A social planner chooses the time path of a public good in order to maximize a convex combination of the present discounted value of the utility of these different agents.

Hyperbolic discounting implies that optimal policies are time-inconsistent (Strotz, 1956).¹ This time-inconsistency arises because the marginal rate of substitution between consumption at two points in the future depends on the ratio of the discount factors. With hyperbolic discounting, this ratio changes as the gap between the current period and the two future points diminishes with the elapse of time.

Chichilnsky (1996) proposes a variation of hyperbolic discounting as a means of modeling sustainable development. Li and Lofgren (2000) build on this proposal to study the sustainable use of natural resource stocks. This modeling approach allows the current regulator to commit to future actions, thereby avoiding (by assumption) the time-consistency problem.

Cropper and Laibson (1999) show (in a particular setting) that a 1-period ahead interest rate subsidy, together with the ability to choose current consumption, provides a substitute for commitment.² This result might suggest that the time-consistency issue should be ignored, since it can be resolved given a sufficiently rich policy menu. A different interpretation is that the impracticality of writing and enforcing sufficiently detailed contingent contracts, and the limitations of the policy menu in the real world eliminate the kinds of remedies that arise in simple models. If we accept that time-consistency problems put us in a second-best world, it is worth trying to understand the resulting equilibrium.

Here we assume that, for one of the reasons suggested above, the regulator has a declining discount rate. She is unable to commit to future actions, and does not have a commitment device that solves the time-consistency problem. The regulator makes the current decision with the understanding of how this will influence the environment and thereby influence future decisions. Equivalently, there are a succession of regulators; each regulator's tenure is limited, perhaps due to a political cycle. The current regulator can

¹ If the declining discount rate depends on calendar time or the calendar values of other state variables, optimal policy is not time-inconsistent; see page 71–72 of Blanchard and Fischer (1990). In Newell and Pizer (2003) the discount rate follows an ARMA process. The current discount rate is known and the regulator learns about future values over time. Here also the changing discount rate does not lead to a time-consistency problem.

² This result holds in a model with quasi-hyperbolic discounting where there are only two discount rates. In a sense, there is a single "distortion", so it is perhaps not surprising that a single interest rate subsidy can achieve the first best outcome. More complicated policies (perhaps the use of a stream of future subsidies) would be needed in a model with a more complicated discount function. The policies or institutional change needed to eliminate the time-consistency problem are sensitive to the details of the model.

influence her successors' decisions by means of influencing the environment that they inherit, but cannot directly choose her successors' decisions. Regulators are identified by the time at which they act. Each regulator cares about current and future payoffs, but treats bygoners as bygoners. The equilibrium is Markov perfect.³

3. Hyperbolic discounting with a stock pollutant

The first subsection describes the model and derives the necessary condition for a differentiable MPE. The second subsection analyzes the necessary condition for a MPE. The third discusses the non-uniqueness and Pareto ranking of the equilibria.

3.1. Model description and derivation of equilibrium condition

Let S_t and z_t be the stock and the flow of emissions in period t , and η the fraction of the stock that persists until the next period⁴. Using the convention that the flow in the current period contributes only to next period's stock, the equation of motion for the stock is

$$S_{t+1} = \eta S_t + z_t. \quad (2)$$

The payoff in the current period is $h(S_t, z_t)$. This function is concave in both arguments; it is decreasing in its first argument and increasing in the second argument. A higher stock causes environmental damages; a higher level of emissions is associated with increased GNP or lower abatement costs.

In period t the regulator's present discounted value of the payoff is

$$h(S_t, z_t) + \beta \sum_{\tau=1}^{\infty} \delta^\tau (h(S_{t+\tau}, z_{t+\tau})). \quad (3)$$

At time t the discount factor used to compare payoffs in periods s and $s+1$, for $s \geq t+1$, is the constant $0 < \delta < 1$; the discount factor used to compare payoffs in periods t and $t+1$ is $\beta\delta$, with $0 \leq \beta \leq 1$. The value $\beta=1$ produces the standard model of constant discounting, and if $0 < \beta < 1$ there is quasi-hyperbolic discounting. In this case, the regulator at time t discounts the payoff in the subsequent period ($t+1$) at a higher

³ The time-inconsistency issue arises not only because of the nature of the problem—agents' objectives and constraints—but also because of the assumption that decisions are Markovian, i.e. they depend on the payoff-relevant state variable (in this case, the stock of pollution). Allowing agents to use history dependent controls—i.e. to have history dependent beliefs—typically leads to a multiplicity of equilibria, some of which may be approximately first-best, as in Ausebel and Deneckere (1989); Chari and Kehoe (1990).

⁴ A previous version of this paper considers the slightly more general case in which Eq. (2) is replaced by the stochastic equation $S_{t+1} = \eta S_t + z_t + \theta_t$, where θ_t is an iid random variable. This generalization accounts for the possibility that the regulator can control emissions only imperfectly, or other sources of random changes in the stock.

rate than she uses to compare payoffs in two contiguous future periods. For example, the regulator at period t compares the payoffs in periods $t+1$ and $t+2$ using the discount factor δ . However, in the next period, at time $t+1$, the regulator compares the payoffs at time $t+1$ and $t+2$ using the discount factor $\beta\delta \leq \delta$. Matters appear different at time $t+1$ than they did at time t .

The regulator is able to choose the level of emissions in the current period, but cannot commit to decision rules that will be followed in the future. It is as if the regulator plays a dynamic game with her future selves; thus, we speak of “Regulator t ” as being the regulator who chooses z_t .

We want to find a differentiable equilibrium Markov control rule, $\chi(S_t)$, such that (from the standpoint of the regulator at time t) the optimal level of z_t is $z_t = \chi(S_t)$, given that the regulator knows that her “future selves” will choose emissions according to the rule $z_{t+\tau} = \chi(S_{t+\tau})$. We find a symmetric Nash equilibrium in the sequential game, using [Harris and Laibson \(2001\)](#)’s heuristic derivation. (Their method can be extended to a general model of non-constant discounting, as in [Karp \(2004\)](#).)

Regulator t ’s payoff is given by expression (3) and the constraint is given by Eq. (2). The single period payoff in equilibrium is

$$H(S_t) \equiv h(S_t, \chi(S_t)). \tag{4}$$

The dynamic programming equation used to generate the MPE in this game is

$$W(S_t) = \max_z \{h(S_t, z) + \delta[W(S_{t+1}) - H(S_{t+1})(1 - \beta)]\}. \tag{5}$$

(Details are in Appendix A.) In solving this problem, Regulator t takes the function $H(\cdot)$ as given. A symmetric equilibrium requires that the solution to this problem, the control rule χ , is the same as the function that appears in Eq. (4).

For $\beta = 1$, the control problem is identical to the standard problem with constant discounting. For the other extreme case, $\beta = 0$, the regulator at time t puts no value on future payoffs. In that case she maximizes the single period payoff in each period, leading to the control rule $\chi = \arg \max h(S, z)$.

In the more interesting case where $0 < \beta < 1$, hyperbolic discounting changes the nature of the control problem. The necessary condition for the problem in Eq. (5) is

$$h_z(S_t, z) + \delta[W'(S_{t+1}) - H_s(S_{t+1})(1 - \beta)] = 0. \tag{6}$$

The stock of pollution creates damages, so the shadow cost of pollution (the negative shadow value of pollution) is positive, $-W' > 0$. In the problem with constant discounting ($\beta = 1$), the first order condition requires equality between the marginal benefit of current emissions (h_z) and the discounted shadow cost of pollution. With hyperbolic discounting, the shadow cost of pollution is reduced by the constant factor $(1 - \beta)$ times the single period marginal equilibrium cost, H_s . For $\beta < 1$, the “effective shadow cost” of pollution falls from $-W'$ to $-W' + (1 - \beta) H_s$. Since a value $\beta < 1$

reduces the effective shadow cost of the stock, we expect it to lead to a larger level of emissions at a given stock.

Using standard manipulations (given in Appendix A) we can write the “Generalized Euler equation” corresponding to the DPE (Eq. (5)) as

$$h_z(t) = -\delta\{\beta h_s(t+1) - h_z(t+1)(\eta + (1-\beta)\chi'(S_{t+1}))\}, \quad (7)$$

using the notation that $h_y(\tau)$ is the partial derivative of h with respect to y evaluated at $h(S_\tau, z_\tau)$.

3.2. Analysis of the equilibrium condition

This subsection presents the intuition for the Euler equation, discusses the monotonicity of the trajectory, and considers the nature of the strategic interaction amongst different generations of the regulator.

3.2.1. Intuition

For $0 < \beta < 1$ the outcome is an equilibrium to a game, rather than the solution to an optimization problem. In this case, the intuition from the standard Euler equation (associated with $\beta = 1$) is not directly applicable. However, it helps to recall the standard case, to see how matters are different here.

If $\beta = 1$, the Euler equation has the following familiar interpretation. Consider a perturbation of a reference path; this perturbation marginally increases emissions in period t and makes an offsetting reduction in the following period, so that the stock inherited in period $t+2$ is the same as in the reference path. If the reference path is optimal, the gain from this perturbation must equal the cost. The left side of Eq. (7) gives the gain of a slight increase in emissions in the current period. A unit increase in emissions in period t leads to a unit increase in stock in the next period. The first term on the right side is the discounted cost due to this higher stock. An additional unit of stock in period $t+1$ results in η additional units in period $t+2$. In order for the perturbation to return the stock to the reference level, it is necessary to reduce emissions in period $t+1$ by η , incurring a cost of $h_z(t+1)\eta$.

If $\beta < 1$ the regulator at t cannot choose emissions in period $t+1$. Nevertheless, the costs and benefits of the perturbation are as described above. In addition, there is an “automatic” equilibrium change of period $t+1$ emissions, due to Regulator $t+1$ ’s response to the changed stock. This change equals χ' and costs $h_z(t+1)\chi'(S_{t+1})$ in that period. The last term in Eq. (7) accounts for this cost.

3.2.2. Monotonicity

At time t , the equilibrium stock in the next period is $S_{t+1} = \eta S_t + \chi(S_t)$. The next period stock is a monotonically increasing function of the *current stock* if and only if $\eta + \chi'(S) > 0$. In this case, the trajectory of the stock is a monotonic function of time. If the inequality is reversed, the next period stock is a monotonically decreasing function of the current stock. In this case, the stock trajectory oscillates over time. The following proposition provides sufficient conditions for these two cases.

Proposition 1. *A sufficient condition for the next period stock to be non-decreasing function of the current stock is*

$$h_{S_z} - \eta h_{zz} \geq 0 \quad (8)$$

evaluated at $z = \chi(S)$. A sufficient condition for the next period stock to be everywhere non-increasing in the current stock is

$$h_{S_z} - \eta h_{zz} \leq 0 \quad (9)$$

evaluated at $z = \chi(S)$.

This proposition holds for $0 < \beta \leq 1$, that is, it also holds for the case of exponential discounting. However, the equilibrium decisions rule changes with β . Thus, we cannot rule out the possibility that one of the two inequalities (8) or (9) holds for one value of β but not for some other value. Of course, if either of these inequalities holds for all (S, z) (not only for the equilibrium z), the next period stock is monotonic in the current stock.

In view of the concavity of $h(\cdot)$ in z , a sufficient condition for the next period stock to be monotonically increasing in the current stock is $h_{S_z} \geq 0$. Thus, additive separability in abatement costs and environmental damages ($h_{S_z} = 0$) is sufficient for monotonicity; we use this fact in Section 3.3. A large value of η (as with global warming) or a large absolute value of h_{zz} also make it “more likely” that Eq. (8) holds. In this case, the pollution stock is a monotonic function of time.

When Eq. (9) holds, the stock oscillates, a high value of S is followed by a low value, and vice-versa. If the absolute value of h_{zz} is small, the regulator is not particularly concerned with smoothing emissions. (For example, emissions may be positively correlated with GNP, and the regulator is not concerned with smoothing GNP.) If η is small, emissions in period t have little effect on the stock in periods $t+j$, $j \geq 2$. If in addition, $h_{S_z} < 0$, so that the marginal utility of emissions is small when stocks are high, the regulator wants to alternate periods of high and of low emissions, causing the stock to oscillate. Although this outcome is possible, the more natural case seems to be where Eq. (8) holds.

3.2.3. Strategic substitutes and complements

Since the stock is a bad and the flow is a good in this setting, it might seem that any “reasonable” equilibrium decision rule would satisfy $\chi' < 0$. This inequality implies that actions are “strategic substitutes”; that is, when the stock increases, the regulator responds by decreasing emissions. If this inequality holds, the presence of the last term in Eq. (7) reduces the right side of the equation. Since $h_{zz} < 0$, this reduction requires an increase in period t emissions in order to maintain the equality. In this case, reducing β leads to an increase in emissions for any stock level. However, the inequality $\chi' < 0$ might not hold.

There are two types of effects of reducing β . First, there is the obvious fact that discounting the future more heavily encourages higher emissions in the current period. However, a reduction in β not only means that Regulator t values the current payoff more highly relative to future payoffs. It also means that her valuation of moving benefits from period $t+1$ to period $t+2$ is higher than Regulator $t+1$'s valuation of the same transfer. The discount factor between these 2 periods is δ for Regulator t , and it is $\delta\beta$ for Regulator $t+1$.

As a consequence of a reduction in β , Regulator t not only wants to emit more in the current period rather than the future, but she also would like to see a reallocation of emissions from period $t+1$ to subsequent periods. An increase in period t emissions, leading to an increase in S_{t+1} , reduces period $t+1$ emissions provided that actions are strategic substitutes ($\chi' < 0$). Regulator t 's desire to influence the decision of Regulator $t+1$ encourages the former to emit more when actions are strategic substitutes.

3.3. Non-uniqueness and welfare

This subsection explains why the equilibrium is not unique;⁵ it shows how to Pareto rank the equilibria, and it considers the equilibrium under full commitment.

3.3.1. Non-uniqueness

Asymptotic stability of the steady state requires

$$-(1 + \eta) < \chi'(S_\infty) < 1 - \eta \quad (10)$$

where S_∞ is a steady state. Inequality (10) is consistent with either a monotonic or oscillatory state trajectory. It is also consistent with actions being strategic substitutes or complements in the neighborhood of the steady state.

In this model, the necessary equilibrium conditions are consistent with a continuum of steady states when $0 < \beta < 1$. Using Eq. (2), the steady state stock and flow satisfy $S_\infty(1 - \eta) = z_\infty$. This restriction and the Euler equation (7) evaluated at the steady state comprise two algebraic equations involving the three variables z_∞ , S_∞ and $\chi'(S_\infty)$. Since the function χ is unknown (and therefore $\chi'(S_\infty)$ is unknown), the equilibrium steady state conditions are under-determined, even with the assumption of local stability. In other words, the requirement that a trajectory satisfy the Euler equation, and the assumption that it approaches a steady state, do not determine a (locally) unique steady state. This circumstance is analogous to the situation noted by [Tsutsui and Mino \(1990\)](#) in differential games.^{6,7}

⁵ [Krusell and Smith \(2003\)](#) show that the equilibrium in a model of quasi-hyperbolic discounting is not unique when the equilibrium decision rule is a step function, and therefore not everywhere differentiable. We rule out this source of non-uniqueness by requiring the decision rule to be everywhere differentiable, an assumption used in our derivation of the Euler equation.

⁶ [Tsutsui and Mino \(1990\)](#) refer to this circumstance as an “incomplete transversality condition”. The transversality condition is $\lim_{t \rightarrow \infty} \delta^t W'(S_t) = 0$. This condition implies the steady state condition $S_\infty = \eta S_\infty + z_\infty$. With constant discounting, the Euler equation evaluated at the steady state and the steady state condition comprise two algebraic equations in two unknowns. Their solution yields (locally, but perhaps not globally) unique values of S_∞ and z_∞ . Under hyperbolic discounting the transversality condition also implies the steady state condition, again yielding two algebraic equations. However, with hyperbolic discounting there is a third unknown variable, $\chi'(S_\infty)$. The transversality condition is “incomplete” because it does not enable us to identify even a locally unique steady state.

⁷ [Karp \(1996\)](#) notes that the same circumstance can arise when a monopolist sells a slowly depreciating durable good, or more generally where a decision-maker who is confronted with a time-inconsistency problem uses a stationary Markov decision rule. Our model of a decision-maker with hyperbolic discounting is an example of this kind of problem.

The non-uniqueness can be illustrated graphically when $h_{S_z}=0$ (i.e. the function h is additively separable in the stock and the flow). We noted that in this case the stock trajectory is monotonic, so Eq. (10) is strengthened to

$$-\eta < \chi'(S_\infty) < 1 - \eta. \tag{11}$$

For this case, define $A(z)=h_z$, the marginal benefit of emissions, and $D(S)=-h_S$, the marginal damage of the pollution stock. By concavity $A' < 0$ and $D' > 0$. The Euler equation (7) evaluated at the steady state can be written as $\omega A = \delta \beta D$, where $\omega \equiv 1 - \delta (\eta + (1 - \beta) \chi')$. Eq. (11) implies

$$1 - \delta(1 - \beta + \beta \eta) \equiv \omega_1 < \omega < \omega_2 \equiv 1 - \delta \beta \eta.$$

For fixed β , with $0 < \beta < 1$, Fig. 1 graphs $\delta \beta D$, $\omega_1 A$ and $\omega_2 A$ (evaluated at $z=(1 - \eta)S$), shows the intersection points S_1 and S_2 . The set of candidate steady states under quasi-hyperbolic discounting is the open interval between S_1 and S_2 . (We use this notation below.) The values of S_1 and S_2 depend on β and the other parameters of the model. For $\beta=0$, the steady state is given by $A(S(1 - \eta))=0$; for $\beta=1$, $\omega_1 = \omega_2 = \omega^* \equiv 1 - \delta \eta$. Thus, the interval (S_1, S_2) collapses to a point in the extreme cases where $\beta=0$ or $\beta=1$. The equilibrium steady state is unique in these limiting cases. Fig. 1 also shows the dashed curves $\omega^* A$ and δD whose intersection S^* is the steady state under constant discounting ($\beta=1$). When h is additively separable, the set of candidate steady states corresponding to $\beta < 1$ lies strictly above the unique steady states under constant discounting.

Thus far, we have used only the necessary conditions for equilibrium. There is no guarantee that the candidates steady states (i.e. those in the interval $S_1 < S < S_2$ that is identified in Fig. 1) are globally asymptotically stable or that they are actual equilibria. That is, we do not know whether the function $\chi(S)$ that drives the state to a particular steady state exists for all S , or that it induces functions $W(S)$ and $H(S)$ such that the maximand in Eq. (5) is concave for all values of S (i.e. for all initial conditions).

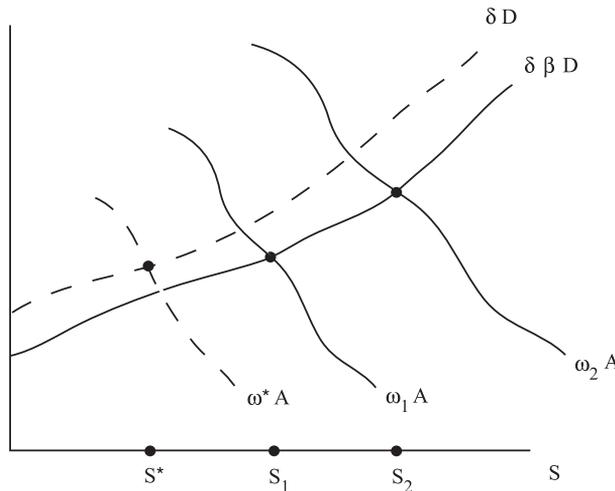


Fig. 1. The steady state conditions under hyperbolic and constant discounting.

However, all values of S satisfying $S_1 < S < S_2$ can be supported as MPE steady states given initial conditions in the neighborhood of that candidate. To confirm this assertion, pick an arbitrary candidate steady state S_∞ . When it is important to emphasize the dependence of the policy function on the steady state (toward which that policy function drives the state), we write the policy function as $\tilde{\chi}(S; S_\infty)$ (instead of $\chi(S)$). This function satisfies $\tilde{\chi}(S_\infty; S_\infty) = (1 - \eta)S_\infty$ and inequality (10).

Concavity of the maximand of (5), evaluated at the steady state, requires that $\frac{\partial^2 \tilde{\chi}(S_\infty; S_\infty)}{\partial S^2}$ satisfy an inequality.⁸ Stability imposes bounds on the first derivative of the policy function (as shown by inequality (10)). Concavity of the maximand imposes bounds on the second derivative of the policy function, without further restricting the candidate steady states. The assumption of concavity implies an additional inequality, but that inequality involves an additional choice variable, the second derivative of the policy function.

The multiplicity of (at least “local”) MPE raises the issue of equilibrium selection. One alternative is to take the limiting equilibrium of the finite horizon game, as the horizon goes to infinity (Driskill, 2002). A second alternative is to admit only equilibria that are defined over the entire state space and that induce a concave problem for all of state space, i.e. to introduce “global” criteria. This alternative could be implemented numerically.

3.3.2. Welfare

A third alternative is to Pareto rank the MPE. We will also use a Pareto criterion to compare a MPE and a non-Markov equilibrium, e.g. one that involves some degree of commitment. To this end, we first compare emissions (as distinct from welfare) under a MPE and in the equilibrium where the initial regulator is able to choose the entire trajectory of emissions (the full commitment equilibrium). We noted in Section 2 that given a sufficiently rich policy menu or a different institutional structure, it might be possible to support the full commitment equilibrium.

If the initial regulator had a commitment device, the Euler equation for the first period is

$$h_z(t) = -\beta\delta\{h_S(t+1) - \eta h_z(t+1)\}. \tag{12}$$

The difference between the functions on the right sides of Eqs. (7) and (12) is

$$RHS(7) - RHS(12) \equiv F(S, \chi(S)) = \delta(1 - \beta)h_z(t+1)[\eta + \chi'(S_{t+1})]. \tag{13}$$

A necessary and sufficient condition for the first period level of emissions to be greater under full commitment is $F(\cdot) > 0$. In view of the inequality $h_z > 0$, a sufficient condition for $F(\cdot) > 0$ is $\eta + \chi'(S_{t+1}) > 0$. The discussion of Proposition 1 notes that a sufficient condition for this inequality is $h_{S_z}(S, z) \geq 0$.

⁸ The derivation of that inequality is straightforward, but the inequality is not informative so we do not present it. To obtain the inequality, substitute the equilibrium control $\tilde{\chi}(S; S_\infty)$ into the dynamic programming Eq. (5) and differentiate the resulting equation twice with respect to S (using the envelope theorem). Evaluate the result at the steady state, and solve to obtain an expression for $W''(S)$. Use this expression to eliminate W'' from the inequality that is necessary and sufficient for concavity of the maximand of Eq. (5). The result is an inequality involving the first and second derivatives of χ and the primitive functions.

In the full commitment equilibrium, the steady state is equal to the steady state in a control problem with constant discount factor δ (since the effect of the higher discounting in the first period eventually wears off). We noted in Section 3.3.1 that at least in the case where $h(S, z)$ is additively separable, the steady state under constant discounting (S^*) is strictly below the infimum of the set of MPE steady state (S_1).

These observations imply

Proposition 2. *For additively separable $h(S, z)$, the regulator who can make full commitments begins with a higher flow of emissions and eventually drives the stock to a lower level (with correspondingly lower steady state emissions), relative to all MPE.*

The ability to make commitments means that future stocks will be relatively low, implying that the shadow cost of the stock in the first period is relatively low, encouraging the regulator to have high emissions in the first period.

We now turn to welfare comparisons. A policy rule $C(S)$ “locally” Pareto dominates a different rule $B(S)$ if for initial conditions \tilde{S} in the neighborhood of the steady state corresponding to $B(S)$ the payoff (on the equilibrium trajectory emanating from \tilde{S}) of the current and all successive regulators is at least as high under $C(S)$ as under $B(S)$, and the payoff is strictly higher for at least one regulator. To evaluate these payoffs we use the expression in (3): each regulator discounts utility τ periods in the future by $\beta\delta^\tau$. The qualifier “locally” in our definition emphasizes that we consider only initial conditions near the steady state corresponding to the rule $B(S)$. If we are near the steady state of $B(S)$, the current and all future regulators would be willing to switch from the rule $B(S)$ to a rule $C(S)$ that locally Pareto dominates $B(S)$.

Consider an arbitrary “reference” MPE rule $\tilde{\chi}(S; S_\infty)$, i.e. a rule that drives the state to a particular steady state S_∞ . The previous subsection establishes that at least in the neighborhood S_∞ there is an equilibrium rule that supports S_∞ . There is also a rule that supports a neighboring steady state; we denote this neighboring rule as $\tilde{\chi}(S; S_\infty - \epsilon)$ for small ϵ .

Since both the reference rule and the neighboring rule are equilibria, each of these is a best response if the current regulator believes that future regulators will use that particular rule. We do not have an explanation for which of the infinitely many equilibria is actually selected. However, we can Pareto rank these equilibria “locally”, i.e. in the neighborhood of the steady state. Under the reference rule, for initial condition $S = S_\infty$, the current regulator’s equilibrium action is to set $z = (1 - \eta) S_\infty$, maintaining the state at the current level. It is feasible for the current regulator to deviate from that action, but it is not optimal if the current regulator believes that her successors will use the rule $\tilde{\chi}(S; S_\infty)$.

One feasible deviation is for the current regulator to reduce emissions slightly, setting $z = (1 - \eta) S_\infty - \epsilon$ with $\epsilon > 0$, so that the state in the next period is $S_\infty - \epsilon$. Since $\tilde{\chi}(S; S_\infty - \epsilon)$ is an equilibrium rule, $S_\infty - \epsilon$ can be maintained as a steady state in equilibrium. We therefore consider the deviation in which the current regulator drives next period stock to $S_\infty - \epsilon$ and future regulators maintain the stock at that level. The question is: does this deviation benefit the current regulator and all her successors? If the answer is “yes”, then the rule $\tilde{\chi}(S; S_\infty - \epsilon)$ locally Pareto dominates the rule $\tilde{\chi}(S; S_\infty)$.

Denote the value of the deviation for the current regulator as $J(S, \epsilon)$ and denote the value for all successive regulators as $K(S, \epsilon)$. For $\epsilon > 0$, it is clear that $J(S_\infty, \epsilon) < K(S_\infty, \epsilon)$ since the current regulator makes a larger decrease in emissions $((1 - \eta) S_\infty - \epsilon < (1 - \eta) (S_\infty - \epsilon))$ than do future regulators, but does not enjoy the reduced stock in the current period. Consequently, for $\epsilon > 0$, the deviation benefits the current and all future regulators if and only if it benefits the current regulator. A necessary and sufficient condition for the current regulator to benefit from a small deviation is $\frac{\partial J(S_\infty, 0)}{\partial \epsilon} > 0$. The function $J(S_\infty, \epsilon)$ is

$$J(S_\infty, \epsilon) \equiv h(S_\infty, (1 - \eta)S_\infty - \epsilon) + \beta \sum_{\tau=1}^{\infty} \delta^\tau h(S_\infty - \epsilon, (1 - \eta)(S_\infty - \epsilon)).$$

A straightforward computation implies

$$\begin{aligned} \frac{\partial J(S_\infty, 0)}{\partial \epsilon} &= -h_z - \frac{\beta \delta}{1 - \delta} (h_s + (1 - \eta)h_z) \\ &= \frac{\delta h_z}{1 - \delta} ((1 - \beta)(1 - \eta) - (1 - \beta)\chi'), \end{aligned} \tag{14}$$

where the second equality uses Eq. (7) evaluated at the steady state. Eq. (14), the fact that $h_z > 0$, the stability condition (10), and the definition of S_1 as the infimum of the set of stable MPE steady states imply

$$\frac{\partial J(S_\infty, 0)}{\partial \epsilon} > 0 \Leftrightarrow S_\infty > S_1. \tag{15}$$

This equivalence relation implies

Proposition 3. (i) *More conservative MPE policy rules—those that lead to a lower steady state pollution stock—locally Pareto dominate less conservative rules. That is, the equilibrium policy function $\tilde{\chi}(S; S_\infty - \epsilon)$ locally Pareto dominates the neighboring policy rule $\tilde{\chi}(S; S_\infty)$ for $\epsilon > 0$.* (ii) *A (non-Markov) policy rule that leads to a steady state strictly lower than S_1 does not Pareto dominate a MPE that drives the state close to the lower bound S_1 .*

The inability to make commitments results in a higher steady state stock, at least when $h(S, z)$ is additively separable (Proposition 2). Therefore, it is not surprising that more conservative MPE rules (locally) Pareto dominate less conservative rules. Part (ii) of Proposition 3 states that if we changed the game, e.g. by allowing the regulator some commitment ability or by introducing additional policies that substitute for commitment ability, the resulting policy rule would not (locally) Pareto dominate a sufficiently conservative MPE rule (one that drives the state close to S_1).

For example, compare a conservative MPE rule that maintains the state slightly above S_1 , $\tilde{\chi}(S; S_1 + \epsilon_1)$ and an alternative non-Markov rule $C(S; S_1 - \epsilon_2)$ that maintains the state slightly below S_1 (with $\epsilon_2 > 0$ and small). For fixed ϵ_2 and sufficiently small ϵ_1 , the rule $C(S; S_1 - \epsilon_2)$ does not locally Pareto dominate $\tilde{\chi}(S; S_1 + \epsilon_1)$ since the *current* regulator would want to switch from $C(S; S_1 - \epsilon_2)$ to $\tilde{\chi}(S; S_1 + \epsilon_1)$ in view of the relation (15). In addition, the rule $\tilde{\chi}(S; S_1 + \epsilon_1)$ does not locally Pareto dominate the rule $C(S; S_1 - \epsilon_2)$: if the current state is at $S_1 - \epsilon_2$, a switch to $\tilde{\chi}(S; S_1 + \epsilon_1)$ drives the state to a higher steady state level, lowering the payoff of *future* regulators.

4. An application to global warming

The linear equilibrium of a linear-quadratic control problem illustrates the effect of hyperbolic discounting in modeling the regulation of a stock pollutant. The linear equilibrium is defined for all values of the state and it is also the limit of the finite horizon model. For our numerical example, the linear equilibrium drives the state close to the lower bound of the set of feasible steady states, S_1 ; in view of Proposition 3, the linear equilibrium therefore Pareto dominates “most” MPE. The first subsection presents the system of algebraic equations that determine the linear equilibrium control rule. The second subsection discusses numerical results.

4.1. The linear-quadratic model

Abatement costs are $\frac{b}{2}(\bar{x} - z)^2$, where b and \bar{x} are positive parameters; the former is the slope of marginal abatement costs and the latter is the cost-minimizing level of emissions; \bar{x} is the business as usual (BAU) level of emissions. The benefits of emissions equal the reduction in abatement costs. Environmental damages are $\frac{g}{2}(\bar{S} - S)^2$ where g and \bar{S} are positive parameters; the former is the slope of marginal damages and the latter is the damage-minimizing level of stocks.

Using these two functions, the single period payoff (benefits minus damages) is

$$h(S, z) = f + az - \frac{b}{2}z^2 - cS - \frac{g}{2}S^2.$$

This equation uses the definitions $f \equiv -\frac{b}{2}\bar{x}^2 - \frac{g}{2}\bar{S}^2$, $a \equiv b\bar{x}$, and $c \equiv -g\bar{S}$. The dynamic programming equation is

$$W(S) = \max_z f + az - \frac{bz^2}{2} - cS - \frac{g}{2}S^2 + \delta[W(S_{t+1}) - H(S_{t+1})(1 - \beta)]. \tag{16}$$

A linear-quadratic equilibrium involves a quadratic value function, $W(S) = \lambda + \mu S + \frac{\rho}{2}S^2$, and a linear control rule, $\chi(S) = A + BS$, where λ , μ , ρ , A and B are constants to be determined. Appendix A shows that the constant B is a root of the cubic

$$\Psi(1 - \beta) - \delta b \eta B^2 + (\delta g + b - b \eta^2 \delta) B + \delta g \eta = 0 \tag{17}$$

with $\Psi \equiv -(\delta b B^3 + B \delta g + \delta b B^2 \eta + \delta g \eta)$. The intercept of the control rule is

$$A = \frac{-(\delta a B - \delta c)(1 - \beta) - \delta a \eta - \delta c + a}{-(\delta b B^2 + \delta g + \delta b B)(1 - \beta) - \delta b \eta B - \eta b \delta + b + \delta g}. \tag{18}$$

When $\beta = 1$ the unique negative root of Eq. (17) is the correct root, since the positive root violates the transversality condition $\lim_{t \rightarrow \infty} \delta^t W'(S_t) = 0$. For $\beta < 1$, there are two negative roots (or two complex roots with negative real parts). We can show analytically that only the larger of these negative roots (the one that is near the unique negative root when $\beta = 1$) satisfies the stability condition (10); in addition, the linear policy function associated with this root induces a globally concave problem.

For purpose of comparison, we present the bounds on the MPE steady states for general (non-linear) rules, previously illustrated in Fig. 1. For the linear-quadratic functional forms, these bounds are

$$S_1 \equiv b\bar{x}\zeta < S_\infty < b\bar{x}\psi \equiv S_2, \tag{19}$$

$$\zeta \equiv \frac{1 - (1 - \beta + \eta\beta)\delta}{\delta g\beta + b(1 - \eta)(1 - \delta + \delta\beta - \delta\eta\beta)}, \quad \psi \equiv \frac{1 - \delta\eta + \delta(1 + \eta)(1 - \beta)}{\delta g\beta + b(1 - \eta)(1 - \delta\eta\beta + \delta - \delta\beta)}.$$

The steady state in the linear equilibrium lies in this interval.

4.2. Numerical results

The numerical results are based on a calibration taken from Karp and Zhang (2002), where its relation to previous literature is explained in detail. That calibration fits the linear-quadratic model to data and estimates used in previous simulation models. It relies heavily on information from IPCC (Intergovernmental Panel on Climate Change, 1996) and Nordhaus (1994). Using a period of 10 years, the parameter values are given in Table 1.

In order to be able to use the formulae in the preceding section, define a new state, $s_t \equiv S_t - 590$. The equation of motion for this state is $s_{t+1} = \eta s_t + z_t$ and damages are $\frac{g}{2}(s_t - \bar{s})^2$, with $\bar{s} \equiv \bar{S} - 590 = 0$. The equilibrium z is given by $A + Bs$.

The parameter values in Table 1 and Eqs. (19), (17) and (18) enable us to compute the boundaries of the set of candidate steady states and the linear emissions rule for different combinations of δ and β . To describe the results, define d as the continuous annual discount rate for future periods, so $\delta \equiv \exp(-10d)$ (because a period lasts for 10 years); r is the additional yearly discount rate for the first period, so $\beta = \exp(-10r)$. The annual discount rate during the current period is $r + d$, and the annual rate at which the current regulator discounts subsequent payoffs is d .

4.2.1. Non-uniqueness

Fig. 2 is constructed using the parameter values in Table 1 and $\delta = e^{-0.3}$ (a yearly discount rate of 3%). The lens-shaped area contains the interval of Markov Perfect steady states that satisfy the stability and monotonicity constraint, Eq. (11). These steady states are shown as a fraction of the unregulated steady state, and graphed as a function of β . For

Table 1
Base-line parameters

Parameter	Note	Value
g	Slope of the marginal damage, billion \$/(billion tons of carbon) ²	0.0223
S_0	Initial stock, billion tons of carbon	781
\bar{S}	Zero damage stock	590
a	Intercept of the marginal benefit, \$(/ton of carbon)	224.26
b	Slope of the marginal benefit, billion \$/(billion tons of carbon) ²	1.9212
\bar{x}	Unregulated emissions	116.7
η	An annual decay rate of 0.0083	0.9204

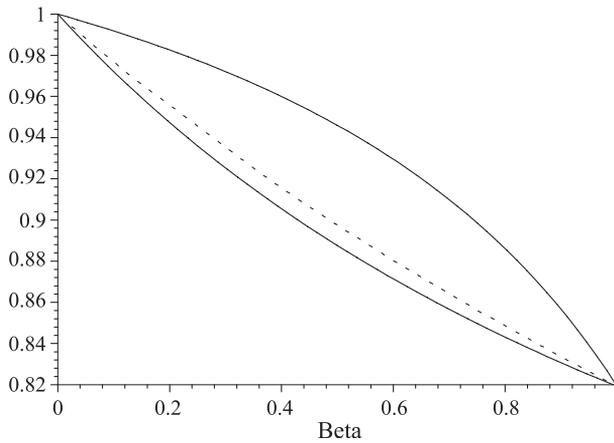


Fig. 2. Ratio of regulated to unregulated steady states. The lens is the set of (stable, monotonic) MPE and the dotted curve is the linear MPE.

example, for $\beta = e^{-0.2} = 0.82$ (a yearly discount rate of 2%), the ratio between the MPE steady state and the unregulated steady state ranges from 0.84 to 0.88. For $\beta = 0$ there is no regulation, and for $\beta = 1$ there is constant discounting. For both of those cases, there is a unique equilibrium.

The dotted curve shows the ratio between the steady state in the linear equilibrium and the unregulated steady state as a function of β . The linear equilibrium achieves nearly the lowest steady state stock that is feasible in a MPE. In view of Proposition 3, this fact means that the linear equilibrium is “close to” the Pareto dominant MPE. For example, for $\beta = 0.82$, the ratio of steady states in the linear and unregulated equilibria is 0.846, only 0.6% higher than in the lowest MPE steady state (and 4% smaller than the largest MPE steady state). In addition, the linear equilibrium is defined over the entire real line; as noted in Section 3.3.1, the domain of other (non-linear) equilibria is unknown.

4.2.2. The short and medium run

Table 2 shows the short and medium term effects of different combinations of d (the columns) and r (the rows) in the linear equilibrium. The first element of each entry is the percentage reduction in emissions during the first period, relative to the BAU emissions. The second element is the percentage reduction of the stock after 10 periods (100 years), relative to the BAU level.

Table 2

First element of each entry: the first period percentage abatement; second element: percentage reduction in stock after 10 periods

$r \backslash d$	0.01	0.03	0.05	0.07
0	(25.1, 18.4)	(9.8, 8.8)	(5.2, 5.2)	(3.3, 3.5)
0.02	(20.8, 15.6)	(8.1, 7.3)	(4.3, 4.3)	(2.7, 2.9)
0.04	(17.4, 13.2)	(6.7, 6.1)	(3.5, 3.6)	(2.2, 2.4)

The first row of the table shows these two values for four levels of d when the regulator has a constant discount rate ($r=0$). Higher discounting leads to a drop in abatement since the costs of abatement are borne in the current period and the benefits arise from lower environmental damages in the future. Beginning with $d=0.01$, an increase in the discount rate causes a substantial fall in abatement. For example, with $r=0$, an increase in the constant discount rate from $d=0.01$ to $d=0.03$ reduces first period abatement from about 25% of the BAU level to approximately 10%, a large change. A further increase in the discount rate to $d=0.05$ causes abatement to fall by an additional 50%.

It is instructive to compare the sensitivity of the results to different parameter changes. For example, if we fix $d=0.03$, $r=0$ but triple the estimate of damages (multiply g by 3), the first period abatement increases to 25.5% and the stock reduction (relative to the BAU level) after 100 years increases to 20.7%. In other words, beginning with our baseline parameters and $d=0.03$, $r=0$, a reduction in the annual discount rate to 0.01 has approximately the same effect on policies as does a tripling of the estimate of damages.

These results illustrate the two problems associated with constant discounting in determining policies to control long-lived environmental problems: The optimal policies tend to be very sensitive to the discount rate, and for reasonable discount rates the regulator may be unwilling to bear moderate costs today in order to prevent substantial damages in the distant future.

The last two rows show the effect of hyperbolic discounting. As expected, this discounting reduces abatement. Perhaps the most interesting result is that the magnitude of this change is moderate. When $d=0.01$ and $r=0.02$, the annual discount rate during the first 10 year period is 0.03; subsequent payoffs are discounted at the rate of 0.01 per year. Holding $d=0.01$ and increasing r from 0 to 0.02 causes first period abatement to fall from approximately 25% to 21%, a moderate change.

Compare the following two changes. In the first, we change the parameters from $d=0.01$, $r=0$ to $d=0.03$, $r=0$ and in the second, we change the parameters from $d=0.01$, $r=0$ to $d=0.01$, $r=0.02$. These two changes have the same effect on the discount rate during the first 10 year period—it changes from 1 to 3%—but they obviously have different effects on the discounting applied to subsequent periods. The first change causes a large reduction in the level of abatement, and the second causes a moderate reduction.

These numerical experiments suggest that the optimal control of greenhouse gasses may be relatively insensitive to the discount rate for the near future, holding fixed the discounting between periods in the distant future. In addition, the regulator may be willing to incur substantial abatement costs even if the short-term discount factor is non-negligible, provided that the long-term discount factor is small.

5. Conclusion

There is a strong argument for discounting the future, but the use of a constant discount rate has unfortunate implications for models of long-lived environmental problems. The optimal policy is likely to be sensitive to the choice of the discount rate, and moderate discounting makes us unwilling to incur even moderate costs today to avoid large damages in the distant future.

Hyperbolic discounting is a plausible description of how people think about trading-off costs and benefits in the distant future. It also may ameliorate some of the modeling defects of constant discounting. Numerical examples show that if the inter-period discount rate used for distant events is held constant, equilibrium policies are relatively insensitive to the discount rate applied to events in the near future. In addition, in equilibrium, a planner with a relatively large near-term discount rate may be willing to incur substantial costs to protect the future.

The time-inconsistency problem is an integral aspect of hyperbolic discounting. Dynamic environmental models can incorporate this feature, rather than assuming it away by allowing the current regulator to commit. Although the resulting equilibria are non-unique, they have a simple Pareto ranking, at least in the neighborhood of the steady state. The linear-quadratic model is particularly useful in this context, because the linear equilibrium exists for all state space and it can be analyzed so easily; in addition, examples show that it is close to the Pareto dominant MPE.

Optimal control methods have been used to study a wide range of stock-related environmental issues. Many of these same issues can also be studied in the (arguably) more realistic situation where the regulator uses hyperbolic discounting and cannot commit to future actions.

Acknowledgements

This paper benefited from comments from Partha Dasgupta, Geoffrey Heal, Cameron Hepburn, Martin Weitzman, seminar participants at Arizona State University, the participants at the 2003 EAERE conference in Bilbao and four anonymous referees. The usual disclaimer applies.

Appendix A. Derivations and Proof

Derivation of Eq. (5).

Suppose that “Regulator t ” believes that all subsequent regulators will use the control rule $z = \chi(S)$. In this case, the present value of Regulator t ’s equilibrium continuation payoff from time $t + 1$ onwards is a function $V(S_{t+1})$ that satisfies the recursive relation

$$V(S_{t+1}) = [h(S_{t+1}, \chi(S_{t+1})) + \delta V(\eta S_{t+1} + \chi(S_{t+1}))]. \quad (20)$$

Since Regulator t treats the function χ as given, she takes the function V as given. Regulator t solves the following dynamic optimization problem

$$W(S_t) = \max_z [h(S_t, z) + \beta \delta V(\eta S_t + z)]. \quad (21)$$

Recall that Regulator t discounts next period’s payoff using the factor $\beta \delta$. A necessary condition for the function χ to be a stationary Markov Perfect Nash equilibrium in this game is that it solves the dynamic programming problem in (21); χ must maximize the

right side of Eq. (21). We obtain the equilibrium value function by substituting the equilibrium control rule into Eq. (21), giving

$$W(S_t) = [h(S_t, \chi(S_t)) + \beta\delta V(\eta S_t + \chi(S_t))]. \tag{22}$$

Regulator t understands that Regulator $t+1$ solves an analogous control problem, possibly with a different value of the initial state, S . Thus, the value function $W(S_{t+1})$ also satisfies Eq. (22) with t replaced by $t+1$. Using Eqs. (20) and (22), and defining $H(S_t) \equiv h(S_t, \chi(S_t))$ we have

$$\beta\delta V(S_{t+1}) = \delta[W(S_{t+1}) - H(S_{t+1})(1 - \beta)]. \tag{23}$$

Substituting Eq. (23) into Eq. (21) we obtain the dynamic programming Eq. (5). Derivation of Eq. (7).

We first use the envelope theorem and Eq. (5) and rearrange to obtain

$$W'(S_t) - \{h_S(S_t, z_t) - \delta(1 - \beta)\eta H_S(S_{t+1})\} = \delta\eta W'(S_{t+1}). \tag{24}$$

Rearranging Eq. (6) and multiplying both sides of the resulting equation by η gives

$$-\{h_z(S_t, z_t) - \delta H_S(S_{t+1})(1 - \beta)\}\eta = \delta\eta W'(S_{t+1}). \tag{25}$$

These two equations imply $W'(S_t) = \{h_S(S_t, z_t) - h_z(S_t, z_t)\eta\}$. Advancing this equation by 1 period and substituting the result into Eq. (6) implies

$$h_z(S_t, z) = -\{\delta[h_S(S_{t+1}, z_{t+1}) - h_z(S_{t+1}, z_{t+1})\eta - H_S(S_{t+1})(1 - \beta)]\}. \tag{26}$$

Use the definition of $H(\cdot)$ to write $H_S(S) = h_S(S, \chi(S)) + h_z(S, \chi(S))\chi'(S)$. Substituting this expression into Eq. (26) and simplifying yields Eq. (7).

Proof of Proposition 1. We begin with some definitions to ease the notation and then prove the proposition. Define the value of the next period stock, given current stock S and current emissions z , as $y \equiv \eta S + z$. By Eq. (2) $S_{t+1} = y_t$. With this definition, the continuation payoff in the maximand of the DPE (5) can be written as $\{\delta [W(S_{t+1}) - H(S_{t+1})(1 - \beta)]\} = \{\delta [W(y_t) - H(y_t)(1 - \beta)]\} \equiv U(y_t)$.

The single period payoff, written in terms of y , is $k(S, y) \equiv h(S, z)$ from which we obtain $k_S + \eta k_y = h_S$, $k_{S_y} + \eta k_{yy} = h_{S_z}$, $k_{y_y} = h_{z_z}$, which implies $k_{S_y} = (h_{S_z} - \eta h_{z_z})$. Thus, we have the following relation

$$k_{S_y} \geq 0 \Leftrightarrow (h_{S_z} - \eta h_{z_z}) \geq 0. \tag{27}$$

Define the equilibrium value of $y(S)$ as $\psi(S) \equiv S + \chi(S)$. The stock is non-decreasing if $\psi'(S) \geq 0$; the stock is non-increasing if $\psi'(S) \leq 0$.

Proof. (Proposition 1) Consider two arbitrary stock levels, $S^* > S^{**}$, and let $y^{**} = \psi(S^{**})$ be the corresponding optimal levels of y . By optimality,

$$k(S^*, y^*) + U(y^*) \geq k(S^*, y^{**}) + U(y^{**})$$

$$k(S^{**}, y^{**}) + U(y^{**}) \geq k(S^{**}, y^*) + U(y^*).$$

Adding these two equations implies

$$0 \leq k(S^*, y^*) - k(S^*, y^{**}) + k(S^{**}, y^{**}) - k(S^{**}, y^*) = \int_{S^{**}}^{S^*} \int_{y^{**}}^{y^*} \frac{\partial^2 k(S, y)}{\partial S \partial y} dy dS \tag{28}$$

If Eq. (8) holds, then $k_{Sy} \geq 0$ by Eq. (27), so $y^* \geq y^{**}$ by Eq. (28). If Eq. (9) holds, the same argument implies that $y^* \leq y^{**}$. \square

Derivation of Eqs. (17) and (18)

Substitute $A + BS'$ into the expression for $h(S', z')$ and use the equation of motion, $S' = \eta S + z$, to write the resulting expression as a function of the current stock and emissions. The single period payoff in the next period, as a function of the current stock and control is

$$H(S, z) \equiv f + a(A + B(\eta S + z)) - \frac{1}{2}b(A + B(\eta S + z))^2 - c(\eta S + z) - \frac{1}{2}g(\eta S + z)^2.$$

Using the quadratic value function, the value of W in the next period is

$$\lambda + \mu(\eta S + z) + \frac{\rho}{2}(\eta S + z)^2$$

Using the definition $\epsilon = 1 - \beta$, the dynamic programming Eq. (5) specializes to

$$\begin{aligned} \lambda + \mu S + \frac{\rho}{2}S^2 = \max_z f + az - \frac{b}{2}z^2 - cS - \frac{g}{2}S^2 + \delta(\lambda + \mu(\eta S + z)) \\ + \frac{\rho}{2}(\eta S + z)^2 - \epsilon H(S, z). \end{aligned} \tag{29}$$

The first order condition implies the control rule $z = A + BS$, with

$$A = \frac{-(-\delta c + \delta a B - \delta b B A)\epsilon + a + \delta \mu}{-(\delta b B^2 + \delta g)\epsilon + b - \delta \rho} \tag{30a}$$

$$B = -\delta \eta \frac{(b B^2 + g)\epsilon + \rho}{(\delta b B^2 + \delta g)\epsilon - b + \delta \rho}. \tag{30b}$$

Solving Eq. (30b) for ρ gives

$$\rho = \frac{-(\delta b B^3 + B \delta g + \delta b B^2 \eta + \delta g \eta)\epsilon + B b}{\delta(B + \eta)}. \tag{31}$$

Substituting the control rule into equation (29) produces the maximized DPE. Equating coefficients in orders of S implies

$$\mu = -\frac{1}{2} \frac{\Omega \epsilon + 2bc - 2a\delta\rho\eta - 2\eta b\delta\mu - 2\delta c\rho}{-(\delta bB^2 + \delta g)\epsilon + b - \delta\rho} \quad (32)$$

$$\begin{aligned} \Omega &\equiv -(-2\eta baB + 2cg + 2\eta b^2 BA + 2abB^2\eta + 2\eta bc + 2B^2bc + 2ag\eta)\delta \\ \rho &= -\frac{(B^2\delta\eta^2b^2 + \delta bg\eta^2 + \delta g^2 + B^2\delta dg)\epsilon + \delta g\rho + \delta b\rho\eta^2 - bg}{(\delta bB^2 + \delta g)\epsilon - b + \delta\rho}. \end{aligned} \quad (33)$$

Substituting the expression for ρ in Eq. (31) into the right side of Eq. (33) and simplifying implies that $\rho = b\eta B - g$. Setting this value of ρ equal to the right side of Eq. (31) implies Eq. (17). To obtain Eq. (18) we use $\rho = b\eta B - g$ in Eq. (32) to obtain an expression for μ . Using this expression in Eq. (30a) implies Eq. (18).

References

- Ainslie, G., 1991. Derivation of ‘rational’ economic behavior from hyperbolic discount curves. *American Economic Review* 81 (2), 334–340.
- Arrow, K.W., Cline, K.-G., Maler, M., Munasinghe, M., Squitieri, R., Stiglitz, J., 1996. Intertemporal equity, discounting, and economic efficiency. In: Bruce, H.L.J.P., Haites, E. (Eds.), *Climate Change 1995: Economic and Social Dimensions of Climate Change*. Cambridge University Press, Cambridge, pp. 128–144.
- Ausebel, L., Deneckere, R., 1989. Reputation and bargaining in the durable goods monopoly. *Econometrica* 57 (3), 511–532.
- Blanchard, O., Fischer, S., 1990. *Lectures on Macroeconomics* MIT Press, Cambridge, MA.
- Chari, V., Kehoe, P., 1990. Sustainable plans. *Journal of Political Economy* 98 (6), 783–802.
- Chichilnisky, G., 1996. An axiomatic approach to sustainable development. *Social Choice and Welfare* 13, 231–257.
- Cline, W.R., 1999. Discounting for the very long term. In: Portney, P.R., Weyant, J. (Eds.), *Discounting and Intergenerational Equity (Resources for the Future)*.
- Cropper, M., Laibson, D., 1999. The implications of hyperbolic discounting for project evaluation. In: Portney, P.R., Weyant, J. (Eds.), *Discounting and Intergenerational Equity (Resources for the Future)*.
- Cropper, M.L., Ayded, S.K., Portney, P.R., 1994. Preferences for life-saving programs: how the public discounts time and age. *Journal of Risk and Uncertainty* 8, 243–266.
- Dasgupta, P., Maler, K.-G., Barrett, S., 1999. Intergenerational equity, social discount rates and global warming. In: Portney, P.R., Weyant, J. (Eds.), *Discounting and Intergenerational Equity (Resources for the Future)*.
- Driskill, R., 2002. A Proposal for a Selection Criterion in a Class of Dynamic Rational Expectations Models with Multiple Equilibria Department of Economics, Vanderbilt University.
- Frederick, S., Loewenstein, F., O’Donoghue, T., 2002. Time discounting and time preference: a critical review. *Journal of Economic Literature* XL, 351–401.
- Gollier, C., 2002. Time horizon and the discount rate. *Journal of Economic Theory* 75, 229–253.
- Harris, C., Laibson, D., 2001. Dynamic choices of hyperbolic consumers. *Econometrica* 69 (5), 935–957.
- Heal, G., 1998. *Valuing the Future: Economic Theory and Sustainability* Columbia University Press, New York.
- Heal, G., 2001. Intertemporal welfare economics and the environment. In: Maler, K.-G., Vincent, J. (Eds.), *The Handbook of Environmental Economics*. North Holland, Amsterdam.

- Intergovernmental Panel on Climate Change, 1995. In: Houghton, J.T., Meira Filho, L.G., Callander, B.A., Harris, N., Kattenberg, A., Maskell K. (Eds.). *Climate Change 1995: The Science of Climate Change IPCC* (Cambridge, UK).
- Karp, L., 1996. Depreciation erodes the coase conjecture. *European Economic Review* 40 (2), 473–490.
- Karp, L., 2004. Non-constant discounting in continuous time Department of Ag and Resource Economics, University of California, Berkeley (<http://are.Berkeley.EDU/karp/>).
- Karp, L., Zhang, J., 2002. Regulation with Anticipated Learning about Environmental Damage Giannini Foundation of Agricultural Economics, University of California, Berkeley (Working Paper No. 926; <http://are.Berkeley.EDU/karp/>).
- Kirby, K., Herrnstein, R., 1995. Preference reversals due to myopic discounting and delayed reward. *Psychology Science* 6 (2), 83–89.
- Krusell, P., Smith, A., 2003. Consumption–saving decisions with quasi-geometric discounting. *Econometrica* 71 (1), 365–375.
- Li, C.Z., Lofgren, K.G., 2000. Renewable resources and economic sustainability: a dynamic analysis with heterogeneous time preferences. *Journal of Environmental Economics and Management* 40, 236–250.
- Newell, R., Pizer, W., 2003. Discounting the distant future: how much do uncertain rates increase valuations. *Journal of Environmental Economics and Management* 46 (1), 52–71.
- Nordhaus, W., 1994. *Managing the Global Commons: The Economics of the Greenhouse Effect* MIT Press, Cambridge, MA.
- Phelps, E., Pollack, R., 1968. On second-best national savings and game: equilibrium growth. *Review of Economic Studies* 35, 185–199.
- Rabin, M., 1998. Psychology and economics. *Journal of Economic Literature* 36, 11–46.
- Read, D., 2001. Is time-discounting hyperbolic or subadditive? *Journal of Risk and Uncertainty* 23 (1), 5–32.
- Rubinstein, A., 2003. Economics and psychology? The case of hyperbolic discounting. *International Economic Review* 44 (4), 1207–1216.
- Strotz, R., 1956. Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies* 23, 165–180.
- Tsutsui, S., Mino, K., 1990. Nonlinear strategies in dynamic duopolistic competition with sticky prices. *Journal of Economic Theory* 52, 136–161.
- Weitzman, M., 2001. Gamma discounting. *American-Economic-Review* 91 (1), 260–271.