Numerical Analysis of Non-constant Discounting with an Application to Renewable Resource Management

Tomoki Fujii* Larry Karp†

May 31, 2006

Abstract

The possibility of non-constant discounting is important in environmental and resource management problems where current decisions affect welfare in the far-distant future, as with climate change. The difficulty of analyzing models with non-constant discounting limits their application. We describe and provide software to implement an algorithm to numerically obtain a Markov Perfect Equilibrium for an optimal control problem with non-constant discounting. The software is available online. We illustrate the approach by studying welfare and observational equivalence for a particular renewable resource management problem.

Keywords: Non-constant discounting, numerical methods, non-renewable resources, observational equivalence.

JEL classification numbers: C63, Q20

*School of Economics and Social Sciences, Singapore Management University, 90 Stamford Road, 178903 Singapore email:fujii@smu.edu.sg

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

1 Introduction

Recent research on models with non-constant discount rates explores the causes of non-constant discounting, examines how agents with non-constant discount rates behave, and attempts to determine empirically whether discount rates change with the planning horizon. Non-constant discounting (hereafter "NCD") increases the complexity of dynamic models, making their analysis more difficult. 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 NCD models. Here we introduce and illustrate a numerical package that solves a fairly general NCD model. Our model is stationary; in particular, it has an infinite horizon. This kind of model has an "incomplete transversality condition", a feature that also occurs in some differential games, but not in standard optimal control problems. Our numerical approach must confront this feature. We illustrate our methods by examining the extent to which a decision rule induced by NCD is observationally equivalent to a decision rule associated with a constant discount rate. We also calculate the loss in steady state welfare resulting from the inability to make binding commitments.

The rest of this Introduction explains why NCD may be an important feature in economic problems, and we explain what we mean by the "solution" to such a model. We review the reason for the incomplete transversality condition, and discuss how this feature complicates the analysis of NCD problems. We then explain the importance of the question of observational equivalence between models with constant versus non-constant discounting. In the process, we discuss some of the relevant literature; Groom, Hepburn, Koundouri, and Pearce (2005) provide a recent review of much of this literature.

The resurgence of interest in NCD in recent years is due largely to its application in behavioral economics, where it has been used to explain anomalies such as apparent reversals in an individual's preferences (Rabin 1998). This context involves a relatively short, or at least finite, period of time, such as the life of an individual. However, NCD is also important for the study of long-lived environmental problems, such as greenhouse gasses (GHGs), where it is reasonable to use an infinite horizon.

Our interest in NCD arises from these kinds of 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). NCD, 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 ((Chichilnsky 1996), (Li and Lofgren 2000), (Heal 2001)).

There are several reasons why NCD might be reasonable. Introspection provides a justification for NCD in the context of long-lived environmental/natural resource problems We are able to make a distinction between our children and our grandchildren. Ignoring for the moment possible differences in wealth, we might have a pure rate of time preference that favors the former over the latter. It is less likely that we would distinguish between the welfare of the 10'th and the 11'th future generation, suggesting that our pure rate of time preference falls over long time spans. 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.

A second justification for NCD is that there exists a "correct" constant discount rate, but the decision-maker has only a probability distribution for this parameter. 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. ((Azar 1999), (Dybvig, Ingersoll, and Ross 1996), (Weitzman 2001)). Equivalently, 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 (Gollier and Zeckhauser 2005). A fourth justification for NCD involves uncertain growth and risk aversion (Gollier 2002a), (Gollier 2002b).

Some decision problems, such as those that involve a large sunk cost, can be modelled as consisting of a single choice. Once a nuclear power plant is built, it is unlikely to be decommissioned 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. NCD qualitatively changes these kinds of dynamic problems (rather than simply complicating a computation), because of the time consistency issue. With NCD, our willingness to transfers 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 our 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.

In using a dynamic model to make policy recommendations, we could simply wish away the time-consistency problem, by assuming that the decision-maker today can commit to a future sequence of policy rules (or actions). 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. The decision-maker can, however, influence the environment that future generations inherit, thereby affecting the decisions that they choose to make. By adopting a Markov Perfect Equilibrium (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.

In a finite horizon model, we would compute the MPE using backward recursion. When each step of this recursion has a unique solution (as is the case in a broad class of problems), the MPE is unique. Matters are more complicated when the horizon is infinite, since this dynamic problem has an "incomplete transversality condition". In (stationary deterministic) dynamic models with constant discounting, the "transversality condition at infinity" in many cases implies that the system asymptotically approaches a steady state. We are able to find candidates for this steady state by finding the steady state of the system consisting of the Euler Equation and the equation of motion for the state variable. For example, in the case of a model with one state variable, this system, when evaluated at a steady state, comprises two algebraic equations that jointly determine a (candidate) steady state and the value of the control variable(s) at the steady state.

The MPE to the model with NCD satisfies a generalized Euler Equation. The two-equation algebraic system that is obtained by finding the steady state of the (generalized) Euler Equation and the equation of motion contains three unknowns: the steady state value of the state variable, the control variable, and the derivative of the equilibrium policy rule with respect to the state. The system is under-determined. Together with the requirement of local stability, it can be used only to find an interval of candidate steady states.

This indeterminacy increases the importance of the numerical tool that we provide. Only in very special cases can we analytically solve dynamic optimization problems, even when these use constant discounting. However, we can often learn something about the solution to those problems by local analysis of the steady state; this analysis involves algebraic equations, and does not require the solution to the control problem. In some differential games, the incomplete transversality condition (resulting in the inability to identify a unique steady state) precludes this kind of local analysis ((Tsutsui and Mino 1990), (Dockner, Steffen, Long, and Sorger 2000)). However, in these single state variable differential games, as with optimal control problems with constant discounting, we can use phase plane analysis to study the problem. For optimal control problems with NCD (as in the differential game), the incomplete transversality precludes the simple identification and subsequent local analysis of the steady state. Moreover, the presence of a "side condition" to the MPE under NCD precludes the use of phase plane analysis (unlike in a differential game model) ((Karp 2006)). Thus, numerical tools are arguably more essential for studying optimal control problems with non-constant discounting, relative to their importance in studying constant discounting problems or symmetric differential games.

We illustrate our numerical package by studying welfare under NCD, and examining the observational (non-)equivalence between problems with NCD and constant discounting. In some cases, e.g. when the payoff function is logarithmic in the control variable and the equation of motion is linear, the affine MPE decision rule for the NCD problem is identical to the optimal decision rule for some constant discount rate (Barro 1999). Observational equivalence means that the qualitative properties of the problem with NCD can be studied by examining the simpler control problem with constant discounting. The equivalence also means that it may not be possible to empirically detect NCD. (The estimation of even a constant discount rate is a challenging empirical problem.[1]) In another example, with quadratic payoff and linear equation of motion, there is a linear MPE, and of course the solution to the control problem with constant discounting is linear. In this case, there is *not* observational equivalence when the NCD is "quasi-hyperbolic" (i.e. the sequence of discount factors is $\beta\delta, \beta\delta^2, \beta\delta^3...$) (Karp 2005). Even when observational equivalence does not hold exactly, it may hold approximately. Simulations offer a means of assessing the the extent to the solution to a control problem under constant discounting approximates a MPE to a problem under non-constant discounting.

[1]There is a burgeoning literature on estimating the parameters of dynamic optimization problems ((Rust 1994), (Aguirregabiria and Mira 2002)). Early papers in this tradition (e.g. (Hansen and Sargent 1980)) discussed the difficulty of estimating the discount factor in dynamic models. Most subsequent papers estimate other parameters of the optimization problem, conditional on an assumed discount rate.

The next section describes the optimization problem and the algorithm, and discusses the role of the incomplete transversality condition. Section 3 explains how we tested the program. It then considers a familiar optimal control problem, modified to include NCD. There we discuss welfare under NCD, and the degree of observational equivalence between NCD and constant discounting.

2 Numerical Methods

In this section, we develop a method to numerically solve a non-stochastic discrete time dynamic programming problem with NCD. 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 NCD, 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. The software to implement this algorithm is available at http://www.mysmu.edu/faculty/tfujii/ncd/ncd.html. We then explain the problem of the incomplete transversality condition.

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 a certain objective function over time by choosing a control variable $x_t \in \Omega \subset \mathbf{R}$ in each period of time t. The payoff she gets 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 $S_t \in \mathcal{S} = \left[\underline{S}, \overline{S}\right]$ is the state space. The state variable follows 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 within the domain of interest.

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

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

where θ_t is the *discount factor* for period t. In the standard setting, the discount rate 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:

$$
\begin{aligned}
V(S) &\equiv \max_{\{x_t\}_{t=0}^{\infty}} \left[\sum_{t=0}^{\infty} \delta^t f(x_t, S_t) \quad s.t. \quad S_t = g(x_t, S_t), S_0 = S \right] \\
&= \max_{x_0} \left[f(x_0, S) + V(g(x_0, S)) \right]. \qquad (1)
\end{aligned}
$$

Now we relax the assumption that the discount rate is constant. Let the one-period discount factor for time t be σ_t so that $\theta_t = \prod_{\tau=1}^{t} \sigma_\tau$. We also define $\theta_0 \equiv 1$. We assume that the one-period discount factor becomes a positive constant after a finite time T. That is, $\sigma_t = \delta < 1$ for $\forall t \geq T$ so that $\theta_t = \theta_T \delta^{t-T}$ for $\forall t \geq T$. (By choosing T large, this model approximates a model in which the discount rate approaches a constant only asymptotically.) The standard problems 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 utility by the factor of $\beta\delta$, but each successive generation is discounted at a constant factor δ.

Under NCD, the decision-maker's decision is not time-consistent in general. The value of control variable x_t for $t > 0$ which is optimal for the decision-maker at time 0 does not in general equal the value of the control variable that she would want to choose at a later time. This inconsistency stems from the non-constant discounting. For example, in the current period ($t = 0$), the decision-maker compares the calendar time $t + 1$ and $t + 2$ using σ_{t+1}, but she uses σ_t to make the same comparison in the next period (at $t = 1$).

We can obtain a MPE by solving this problem as a game amongst a succession of generations of policymakers. Each generation cares about future generations but not about past generations. 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 insofar as each takes the current state variable as given and each 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 Harris and Laibson (2001), who considered the case where $T = 1$.

2.1.1 The algorithm

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

$$\chi(S) = \arg\max_{x_0} \sum_{t=0}^{\infty} \theta_t f(x_t, S_t) \quad s.t. \quad S_0 = S, S_{t+1} = g(x_t, S_t), x_t = \chi(S_t) \quad for \quad \forall t \geq 1. \tag{2}$$

EQ (2) 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.

To emphasize the dependence of the value function on the (possibly non-unique) equilibrium control rule, χ, we write the value function as $W_\chi(S)$. Using the constraint $S_{t+1} = g(\chi(S_t), S_t)$, we have the following relationship:

$$\begin{aligned} W_\chi(S) &\equiv \sum_{t=0}^{\infty} \theta_t f(\chi(S_t), S_t) \qquad s.t. \quad S_0 = S & (3) \\ &= f(\chi(S), S) + \sum_{t=1}^{T} \theta_t f(\chi(S_t), S_t) + \sum_{t=T+1}^{\infty} \delta\theta_{t-1} f(\chi(S_t), S_t) \\ &= f(\chi(S), S) + \sum_{t=1}^{T} (\theta_t - \delta\theta_{t-1}) f(\chi(S_t), S_t) + \delta \sum_{t=0}^{\infty} \theta_t f(\chi(S_{t+1}), S_{t+1}) \\ &= f(\chi(S), S) + \sum_{t=1}^{T} (\theta_t - \delta\theta_{t-1}) f(\chi(S_t), S_t) + \delta W_\chi(S_1) \end{aligned}$$

Hence, we obtain the QDPE as follows:

$$\begin{aligned} W_\chi(S) &= \max_x f(x, S) + \sum_{t=1}^{T} (\theta_t - \delta\theta_{t-1}) f(\chi(S_t), S_t) + \delta W_\chi(S_1) \\ & \quad s.t. \quad S_1 = g(x, S), S_{t+1} = g(\chi(S_t), S_t) \quad for \quad t \geq 1 \qquad (4) \end{aligned}$$

This equation is similar to EQ (1) except that EQ (4) has an extra term, the summation on the right side. One way to solve EQ (1) is to use the collocation method and function iteration. To see how our method works, it is helpful to review the standard setting with constant discounting.

We start by guessing the value function $V(S)$. We denote the initial guess by $V^{(0)}(S)$. Consider a particular value of S_0. Solving the maximization problem in EQ (1) 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 $V^{(1)}(S_0) \leftarrow f(x_0^*, S_0) + V^{(0)}(g(x_0^*, S_0))$. 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 (1) 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 $s_1, \cdots, s_K$ called the *collocation nodes* is used to evaluate $V^{(r)}(s_i)$ for $1 \leq i \leq K$ in a fixed domain $[\underline{S}, \overline{S}]$ ((Judd 1998), (Miranda and Fackler 2002)). The values of $V^{(r)}$ outside the collocation nodes are approximated by a linear combination of $N \leq K$ known *basis functions* $\phi_n(S)$, so that the approximant has the form of $\hat{V}^{(r)}(S) \equiv \sum_{n=1}^{N} c_n \phi_n(S)$, where 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. In the program, the minimization is done by the ordinary least square method. 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 it is known that certain approximants work better with Chebyshev nodes, which place 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. In the following section we use 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 numerical results that we report below are very similar to those obtained using Chebyshev polynomials.

The collocation method and function iteration described above do not directly apply to EQ (4), because of the extra (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 k-th collocation node with $s^k < s^{k+1}$ for all $k < K$. The initial guess of the control rule consists of $x_1^{(0)}, \cdots, x_K^{(0)}$. Given a choice of basis functions for the control rule, we can then find the approximant $\hat{\chi}^{(0)}$ that exactly or approximately satisfy $\hat{\chi}^{(0)}(s^k) = x_k^{(0)}$.

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 the first line of EQ (3) 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^k and the control rule $\hat{\chi}^{(0)}$ is used, our guess of the value function at the collocation nodes is

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

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 (4), so that

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

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:

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

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

$$\max_{s^k} \left\{ \chi^{(r+1)} - \chi^{(r)}, W^{(r+1)} - W^{(r)} \right\} \leq tol$$

where tol is a small positive value, the tolerance level. Our Matlab implementation of the algorithm uses the CompEcon toolbox that accompanies Miranda and Fackler (2002). Fujii (2006) contains a detailed discussion of our program.

2.2 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 merely evaluate the Euler equation and the equation of motion at a steady state. These are two algebraic equation, so their analysis, using numerical or qualitative methods, is straightforward. We are not able to apply this approach to the model with NCD because of the "incomplete transversality condition" to this problem.

The transversality condition, for the class of control problems that we are interested in, requires that the state variable asymptotically approach a steady state. This requirement explains why it makes sense to examine the steady state of the Euler equation and the equation of motion. With NCD, we show in Appendix A that the Euler equation evaluated at a steady state is a T-th order polynomial in χ'^*:

$$f_x^* + g_x^*(f_x^*\chi'^* + f_s^*)\sum_{t=1}^{T}(\theta_t - \delta\theta_{t-1})(g_x^*\chi'^* + g_s^*)^{t-1} + \delta\left(f_s^*g_x^* - f_x^*g_s^*\right) = 0. \tag{5}$$

In EQ (5) the subscripts x and s denote the partial derivatives, and $*$ denotes the value at the steady state. For example, f_x^* is $\frac{\partial f(x,s)}{\partial x}$ evaluated at $(x, s) = (\chi(S^*), S^*)$.

EQ (5), 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 (5) vanishes, eliminating the unknown value χ'^*. With constant discounting we have the standard Euler equation evaluated at the steady state,

$$f_x^* + \delta\left(f_s^*g_x^* - f_x^*g_s^*\right) = 0. \tag{6}$$

EQ (6) and the steady state condition $S^* = g(x^*, S^*)$ comprise two equations in two unknowns, which can be solved to identify steady state candidates with constant discounting.

With NCD, the necessary condition for asymptotic stability,

$$\mid g_x^*\chi'^* + g_s^* \mid < 1, \tag{7}$$

can be used to restrict the range of candidate steady states. 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. 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 (5) shows.

We can obtain an analogous equation for the higher-order derivatives of the control rule evaluated at the steady state. EQ (12) in Appendix B 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.

The approximated control rule may appear much smoother than the function actually is. For example, consider a function $h(z) = z + \epsilon \sin(\frac{z}{\epsilon^2})$ 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 $|h'(z) - \hat{h}'(z)| = \frac{1}{\epsilon}|\cos(z)|$ 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 Illustration of the software

This section illustrates potential uses of the software. We first explain how we tested the program, and discuss the fact that it always returns a unique solution (in our experiments). We then apply the software to a canonical renewable resource problem. We use this problem as a basis for discussing the observational (non-)equivalence between the model with NCD and constant discounting.

3.1 Testing the program

We tested the program 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 (Karp 2005). 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 the Chebyshev polynomials in addition to the cubic splines.

Use of the linear-quadratic model also provides insight into the question of multiple equilibria. The incomplete transversality condition means that there exists an interval of candidate steady states. We can identify this interval using the requirement of asymptotic stability, EQ (7). 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.

In general 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 (2.2), 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] Those derivatives can be used to obtain a Taylor approximation of the control rule that drives the state to a particular steady state. This procedure would tell us nothing 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 candidates and require that this point be a steady state to the control problem. By imposing the steady state condition $g(\chi(S^u), S^u) = S^u$, we ensure that S^u is the steady state. In addition, the user has an option to impose the steady state Euler condition Eq (5) 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 EQ (12).

Both the steady state condition and the steady state Euler condition can be expressed as a

[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.

linear constraint on the coefficients $c_1^\chi, c_2^\chi, \cdots, c_N^\chi$ for the approximant of the control rule χ. That is, from the steady state condition, we can find such a value of the control rule x^u that satisfies $g(x^u, S^u) = S^u$. Thus, the coefficients must satisfy $x^u = \sum_{n=1}^N c_n^\chi \phi_n(S^u)$. Likewise, the steady state Euler Equation (EQ (5)) gives the value of χ' at S^u. Hence, the coefficients must satisfy the following linear constraint: $\chi'(S^u) = \sum_{n=1}^N c_n^\chi \phi_n'(S^u)$. When the user-specified conditions are imposed, the coefficients are found by the constrained least squares method.[3]

The user also has the option of not specifying the steady state. When this option is used, the program approximates the equilibrium control rule without special consideration to the steady state. We then find the steady state using the approximant, i.e. by solving $g(\hat{\chi}(S), S) = S$. We can then use the steady state condition and the steady state Euler condition (EQ (5) 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 (5). However, we use these two equations in different ways. These conditions can be imposed with the user-specified steady state option, and they are merely a validation tool when the program chooses the steady state. EQ (12) is always used as a validation tool.

We know that for the linear-quadratic problem with quasi-hyperbolic discounting, the linear equilibrium is defined over the entire real line (Karp 2005), but we have no information about the domains of other equilibria.[4] 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. 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

[3] It is also possible to impose EQ (12), 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, as there are more linear constraints, it becomes more difficult to obtain reliable numerical results using constrained least squares.

[4] In the linear-quadratic symmetric differential game with one state variable, the linear equilibrium is defined over the entire real line, but each non-linear equilibrium is defined over only a subset of the real line (Tsutsui and Mino 1990).

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 (not surprisingly) 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 (5), 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 conditions. 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. This finding has no implications for the existence of multiple equilibria; it merely reports that our algorithm returns 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.[5] Second, recent years have seen the development of algorithms to obtain MPE in dynamic games (e.g. Pakes and McGuire (1994)) 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) 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 it returns a unique equilibrium (in our experiments), and the fact 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. The results discussed below allow the program to identify the steady state. (We do not employ the user-specified steady state option.)

[5]The proof of this assertion, available on request, uses the Euler equation evaluated at the steady state, EQ (4).

3.2 The renewable resource problem

This section considers a problem of renewable resource management taken from Miranda and Fackler (2002) [page 250]. We use a problem whose solution (under constant discounting) is readily available, in order to emphasize the role of NCD. A social planner manages a renewable resource, a stock of fish. The social planner determines the harvest, x, at the beginning of the period, and the fish stock grows according to the Schaefer growth function. The transition function is

$$g(x, S) = \alpha(S - x)(1 - \frac{S - x}{\kappa}),$$

where α is the intrinsic growth rate of the resource stock and κ is the carrying capacity. The planner faces the inverse demand function $x^{-\gamma}$, and the unit cost of harvest is a constant c. The planner's reward function is

$$f(x, S) = \frac{x^{1-\gamma}}{1-\gamma} - cx.$$

The parameter values are: $\alpha = 4.0$, $\kappa = 8.0$, $\gamma = 0.5$, $c = 0.2$ and $\delta = 0.9$. Since costs are stock-independent in this example, the optimal steady state under constant discounting depends only on the discount factor and on the parameters of the growth equation, α and κ.

The maximum sustainable yield in this model is $x_{MSY} = \frac{(\alpha-1)^2\kappa}{4\alpha} = 4.5$, and the corresponding stock level is $S_{MSY} = \frac{(\alpha^2-1)\kappa}{4\alpha} = 7.5$. With constant discounting, the optimal steady-state stock level S^* and the corresponding harvest x^* (when positive) are

$$S^* = \frac{\kappa(\alpha^2 - \delta^{-2})}{4\alpha} = 7.3827 \qquad (8)$$

$$x^* = S^* - \frac{\kappa(\alpha - \delta^{-1})}{2\alpha} = 4.4938. \qquad (9)$$

When the discount factor is sufficiently small ($\delta \leq \alpha^{-1} = 0.25$) the decision maker harvests until the population is driven to extinction.

To solve the optimization problem/dynamic game, we use cubic splines and Chebyshev nodes with $N = K = 20$, and implemented the program in Matlab. We let $\underline{S} = 2$ and $\bar{S} = 9$. The program stops when $|\chi^{(r+1)} - \chi^{(r)}| < 10^{-5}$ and $|W^{(r+1)} - W^{(r)}| < 10^{-5}$ are satisfied at all the collocation nodes. Convergence was deemed achieved at that point. Figure 1 shows the "benchmark control rule" for the benchmark case when the decision-maker has a constant discount factor of $\delta = 0.9$.

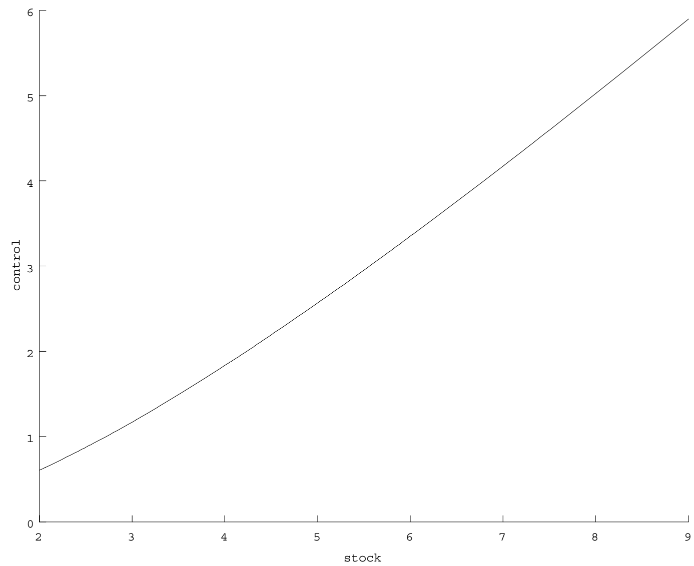

Figure 1: The control rule under constant discounting

We illustrate the effect of NCD on the steady state values S^* and x^* using two one-parameter sequences of single-period discount factors. The first sequence, $S1(\tau)$, sets $T = \tau$ and $\sigma_t = \delta^2$ for $\forall t \leq \tau$. The second sequence, $S2(\tau)$, sets $T = 1$ and $\sigma_1 = \delta^{\tau+1}$. For example, for $\tau = 3$, $S1(3) = \{1, \delta^2, \delta^4, \delta^6, \delta^7, \delta^8, ...\}$ and $S2(3) = \{1, \delta^4, \delta^5, \delta^6, \delta^7, \delta^8, ...\}$. The sequences $S1(\tau)$ and $S2(\tau)$ are identical for $t \geq \tau$, and the long run discount factor is always δ. For $\tau = 1$ the two (identical) sequences are a special case of quasi-hyperbolic discounting, with $\beta = \delta$. The two sequences thus provide a parsimonious extension of quasi-hyperbolic discounting. Under both sequences, a larger value of τ implies more impatience.

For each discounting scheme, we solved the model for $\tau = 1, 2, 3, 4$ and 5. In each case, we obtained convergence. We also checked the values of χ' and χ'' at the steady state. The left-hand-side of EQ (5) and EQ (12) were on the order of 10^{-6} and 10^{-1}, respectively. Thus, when the program is free to determine the steady state, the Euler equation is almost exactly satisfied at the steady state, and the derivative of the Euler equation is reasonably close to its theoretical value.

Table 1 shows the steady state values for various discounting schemes. The first column shows the discounting scheme. The second column shows the steady-state stock level; the third column shows the steady-state control variable; and the fourth column shows the one-period steady-state reward $f(S^*, x^*)$. The fifth column χ'^* shows the derivative of the control rule at the steady state. This value is used to confirm that the Euler equation holds at the steady state. It is also useful for determining the effect of the state, on the control rule, in the neighborhood of the steady state. The sixth column gives the constant discount factor, denoted δ^* for which the steady state under constant discounting equals the steady state under the discounting scheme $Sj(\tau)$, $j = 1, 2$. The seventh column provides the derivative of the control rule, at the steady state, under a constant discount factor δ^*.

Constant and non-constant discounting are observationally equivalent in the steady state. If an econometrician knew (or estimated) the values of α and κ and observed the steady state value S^*, it would be possible to use EQ (8) to calculate δ^*, a constant discount rate that supports S^* as a steady state. We refer to δ^* as the "observationally equivalent" discount rate, with the understanding that "observational equivalence" holds at the steady state, but perhaps nowhere else. Comparison of columns 5 and 7 show that in the neighborhood of the steady state, it would be difficult to distinguish between constant and non-constant discounting, since the first derivatives of the respective control rules are nearly the same.

The steady state reward is useful in studying the welfare effects of the inability to make commitments under NCD. Recall that in our model, the discount factor becomes constant after T periods, so in every case the long run one-period discount factor is δ. If the decision-maker at time 0 was able to commit to future policies, she would solve a non-stationary control problem with a changing discount rate. The *long run* effect of non-constant discounting during the first T periods would be 0, so the steady state in this non-stationary control problem (*with commitment*) equals the steady state in the control problem under the constant discount factor δ. In the MPE the decision-maker is not able to commit to future policies, but she still discounts far-distant payoffs using a factor δ. Therefore, the steady state welfare effect of the inability to make commitments is equal to $\frac{1}{1-\delta}$ times the difference in the steady state payoff flow under commitment and under the MPE.

The largest difference between constant discounting and NCD occur for $S2(5)$. For the experiments reported in Table 1, the steady state stocks under NCD differ from the steady state under constant discounting by (approximately) 10% or less ($1 - \frac{6.60}{7.39} = 0.106\,9$) and the steady state payoff flows differ by less than 2% ($1 - \frac{3.28}{3.34} = 0.018$). For these experiments, the steady state stock effect of being unable to make binding commitments is moderate, and the steady state welfare effect is small. However, column 6 of Table 1 shows that the observationally equivalent discount factor, δ^*, is quite sensitive to NCD; there is a 33% difference between the

1	2	3	4	5	6	7
Discounting	S^*	x^*	$f(x^*, S^*)$	$\chi'(S^*)$	δ^*	$\tilde{\chi}'(S^*)$
Constant	7.3827	4.4938	3.3410	0.8487	0.9000	0.8487
$S1(1), S2(1)$	7.2641	4.4773	3.3365	0.8411	0.8243	0.8406
$S1(2)$	7.2423	4.4733	3.3354	0.8395	0.8124	0.8393
$S1(3)$	7.2386	4.4726	3.3352	0.8392	0.8104	0.8391
$S1(4)$	7.2380	4.4725	3.3352	0.8391	0.8101	0.8391
$S1(5)$	7.2380	4.4725	3.3352	0.8391	0.8100	0.8391
$S2(2)$	7.1282	4.4487	3.3286	0.8344	0.7573	0.8331
$S2(3)$	6.9738	4.4064	3.3170	0.8281	0.6980	0.8262
$S2(4)$	6.7997	4.3491	3.3011	0.8225	0.6454	0.8196
$S2(5)$	6.6049	4.2753	3.2803	0.8176	0.5987	0.8173

Table 1: Steady state values for various sequences of discount factors

largest and smallest values.

To appreciate the importance of these differing magnitudes, consider the situation in which we observe the steady state behavior of two decision-makers facing the problem in our example. The first has a constant discount factor and the second uses NCD with $S2(5)$. Suppose that we incorrectly assume that they both have constant discount factors. If we knew all of the parameters of the control problem except for the discount factor, we could use the steady state behavior to infer the constant discount factor. For the first decision-maker (who actually does have constant discounting), we would correctly infer the discount factor 0.9 and would conclude that the steady state value of the fishery is $\frac{3.341}{1-0.9} = 33.41$. Under the incorrect assumption that the second decision-maker also has a constant discount factor, we would infer that the steady state value of the fishery to her is $\frac{3.2803}{1-0.5987} = 8.1742$. These two values differ by a factor of $\frac{33.41}{8.17} = 4.09$

Since the flow payoffs are nearly the same in the two settings, almost all of the difference in the present discounted values is due to the difference in implicit discount factors ($\frac{1-0.5987}{1-0.9} = 4.01$). Consequently, if we used the estimated discount factors to estimate the shadow value of the stock at the steady state, the two would also differ by a factor of approximately 4. Here, by incorrectly imputing a constant discount factor to a decision-maker who has NCD, we seriously underestimate the shadow value of the resource. This underestimation occurs even though the observed (steady state) behavior of the agent is close to that of a different agent who does have a constant discount rate. The stock's shadow value (or some other function related to the solution to the optimization problem) can be important in policy analysis. For example, it can be used to assess the efficiency of an investment that protects the fishery. Our example here shows that incorrectly attributing constant discounting to a decision-maker can lead to substantial mistakes.

As τ increases, the future is discounted more heavily, causing the social planner to harvest more in the current period, leading to a lower steady-state stock. $S2$ discounts the near future more heavily than $S1$. Consequently, the steady-state stock is lower for $S2\,(\tau)$ than for $S1\,(\tau)$. With $S1$, the impact (on the steady state) of increasing τ diminishes quickly, because the change (resulting from higher τ) occurs ever further in the distant future. In contrast, with $S2$ an increase in τ causes a significant decrease in the short run discount factor; in this case, the impact of increasing τ does not diminish quickly.

Figure 2 compares the control rules for various discounting schemes, relative to the benchmark. The horizontal axis is the stock level and the vertical axis is the difference between the control rule for each discounting scheme and the benchmark control rule. The figure shows that the largest difference in behavior, under different discounting schemes, occurs far from the steady state. This result is reasonable, because the long run discount rate is the same under all schemes, but the short run discount rate is much higher under NCD (in our examples). In addition, an increase in τ has a small effect under $S1$, but a substantial effect under $S2$.

Constant and non-constant discounting are observationally equivalent in the steady state. (Even in that case, we noted above that welfare and policy implications can be very different under constant and non-constant discounting.) Outside the steady state, observational equivalence holds in special cases, but it is unlikely to hold generally. For example, in the linear-quadratic model, the linear control rule under NCD requires two parameters, the slope and the intercept. In general, it is not possible to choose a constant discount rate that leads to a control rule with the same values of the slope and intercept.

For any empirical problem, model parameters are estimated with noise. It may therefore be more interesting to know whether the control rule under NCD is close enough to a control rule under constant discounting, that the difference between the two cannot be distinguished from estimation error. Our software provides a means of assessing this possibility.

To illustrate how we might perform this assessment, for each of nine NCD trajectories in Table 1 we solved the control problem using the constant "observationally equivalent discount

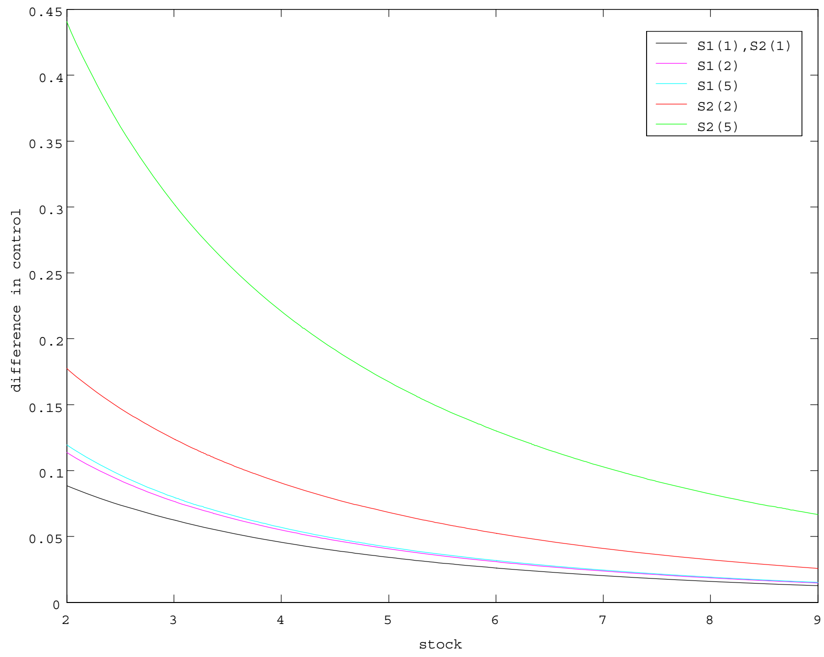

Figure 2: Control rule for the various discounting schemes. Vertical axis measures the difference from the benchmark control rule (with a constant discount factor of $\delta = 0.9$

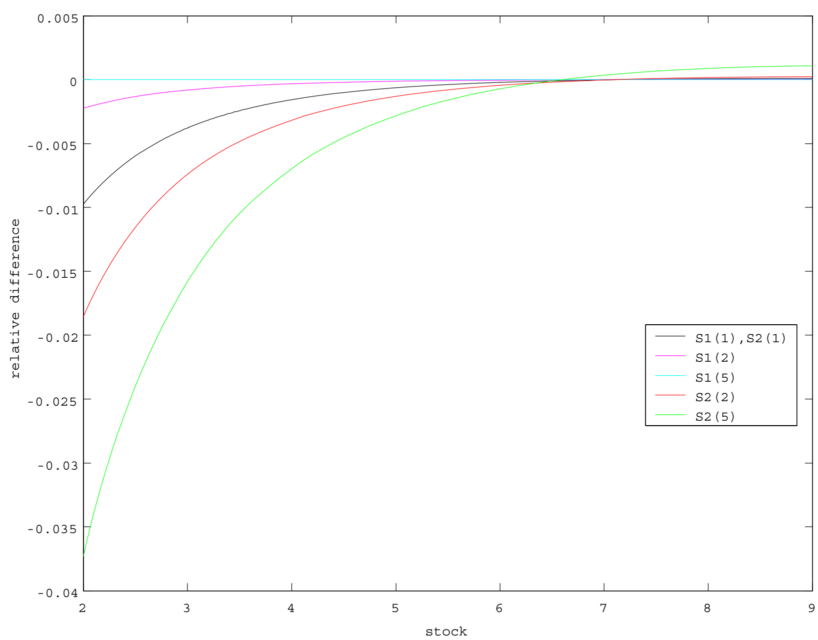

Figure 3: Relative difference in the control rule between non-constant discounting and the “observationally equivalent” discount rate.

rate" (given in column 6 of the table). We took the difference between the value of the control rule under this observationally equivalent discount rate and the original problem under NCD, and divided it by the value of the control rule under NCD in order to make the measure unit-free.

Figure 3 shows the graphs of these ratios (the "bias") as a function of the state variable. By construction, the value of the ratio equals 0 at the respective steady state. The figure shows that the graphs are asymmetric; the bias is larger below than above the steady state. In addition, the bias is increasing in τ for $S2$, but decreasing in τ for $S1$. The explanation for this reversal is that under $S1$ a larger value of τ makes the discounting scheme more similar to constant discounting (with a discount factor of δ^2). In contrast, under $S2$ a larger value of τ makes the discounting scheme less similar to constant discounting. The control rule under $S1(5)$ and the constant discount factor $\delta = 0.81$ are "nearly" observationally equivalent, although they are not identical. Even where the bias is greatest (for $S2(5)$) the largest difference in values between the control rules (under constant discounting and NCD) is less than 4% of the value under NCD. For our example, it appears that it would be extremely difficult to detect NCD. Even though NCD and constant discounting are not observationally equivalent (except at a single point, such as the steady state), the difference between the two could easily be mistaken for estimation error.

4 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 discount rates for the near and middle term, while allowing the long run discount rate to become small. With this flexibility, we can incorporate "reasonable" rates of short and medium run time preference, while recognizing that our time preference in the very distant future is likely to be small. Despite the importance of this model, it has seldom been applied to environmental and resource problems, owing largely to the difficulty of analyzing subgame perfect equilibria.

We developed and programmed a method to find a Markov Perfect equilibrium when the decision-maker has non-constant discounting. 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 discounting.

We used two parsimonious extensions of quasi-hyperbolic discounting, but the program can applied using any finite sequence of non-constant discount rates, followed by an infinite sequence of constant discounting. In our model, the discount rate approaches a constant in finite time; by allowing that time (a parameter in the model) to be large, we approximate the model in which the discount rate 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. Experiments with different functional forms show that our program returns a unique equilibrium. For the linear-quadratic case, this equilibrium is linear in the state. It is defined over the entire real line, and is the equilibrium to the limiting problem, obtained by taking the limit of the finite horizon problem (letting the horizon approach infinity).

We illustrated the methods by comparing equilibria under constant and non-constant discounting using a familiar renewable resource problem. Since the long run discount rate in our model is constant, the steady state of a decision-maker (with NCD) who is able to make binding commitments equals the steady state of planner with the constant discount rate equal to the long-run rate. Using this fact, we can find the steady state flow payoff of the regulator (with NCD) who is able to make commitments, by solving a pair of algebraic equations. We find the steady state payoff of the regulator (with NCD) who cannot make commitments using our program. Comparing these two values provides a measure of the steady state cost of the inability to make binding commitments (since the discount rates are the same in the two cases). In our experiments this loss was very small.

However, the (steady state) "observationally equivalent" discount factor is substantially smaller than the true long run discount factor. If we used steady state behavior to estimate the discount factor under the mistaken hypothesis that the planner has constant discounting, we would significantly underestimate the value of the resource. Other experiments showed that although observational equivalence does not hold over the entire state space, in our examples it would in practice be difficult to distinguish non-constant discounting from estimation error.

A Appendix A: Derivation of Euler Equation

Here, we derive the Euler equation corresponding to EQ (4). 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

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

and

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

Using the notation above, we can write EQ (4) 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,

$$f_x(0) + \left[\sum_{t=1}^{T} (\theta_t - \delta\theta_{t-1})(f_x(t)\chi'(t) + f_s(t)) \frac{dS_t}{dS_1} + \delta W'(S_1) \right] g_x(0) = 0. \tag{10}$$

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 (10), 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 (10), we have

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

EQ (11) is the *Euler Equation* for the QDPE EQ (4). 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 (11) and arranging the terms, we obtain EQ (5).

B Appendix B: Condition on χ''^*

The Euler Equation evaluated at the steady state EQ (5) provides a condition on χ'^*. We can also find a condition that χ''^* has to satisfy. This is helpful to see if the approximant is good at the higher order.

First notice that the following holds for $t\geq 2$.

$$\begin{aligned}
\frac{d^2S_t}{dS_1^2} &= \frac{d}{dS_1}\prod_{\tau=1}^{t-1}(g_x(\tau)\chi'(\tau)+g_s(\tau))\\
&= \left[\prod_{\tau=1}^{t-1}(g_x(\tau)\chi'(\tau)+g_s(\tau))\right]\cdot\left[\sum_{\tau=1}^{t-1}\frac{1}{g_x(\tau)\chi'(\tau)+g_s(\tau)}\frac{d(g_x(\tau)\chi'(\tau)+g_s(\tau))}{dS_\tau}\frac{dS_\tau}{dS_1}\right]\\
&= \left[\prod_{\tau=1}^{t-1}(g_x(\tau)\chi'(\tau)+g_s(\tau))\right]\cdot\left[\sum_{\tau=1}^{t-1}\frac{g_{xx}(\tau)(\chi'(\tau))^2+2g_{xs}(\tau)\chi'(\tau)+g_x(\tau)\chi''(\tau)+g_{ss}(\tau)}{g_x(\tau)\chi'(\tau)+g_s(\tau)}\frac{dS_\tau}{dS_1}\right]
\end{aligned}$$

Because EQ (11) has to hold for all S, we can totally differentiate the equation with respect to S. Applying the chain rule and product rule multiple times and using EQ (11), we have

$$
\begin{aligned}
& (f_{xx}(0)\chi'(0) + f_{xs}(0)) + g_x(0)\frac{dS_1}{dS_0}\Bigg[\sum_{t=1}^{T}(\theta_t - \delta\theta_{t-1})\cdot\Big\{(f_x(t)\chi'(t) + f_s(t))\frac{d^2 S_t}{dS_1^2} \\
+ \quad & (f_{xx}(t)(\chi'^2 + 2f_{xs}(t)\chi'(t) + f_x(t)\chi''(t) + f_{ss}(t))\left(\frac{dS_t}{dS_1}\right)^2\Big\} \\
+ \quad & \delta\Big\{(f_{xs}(1)\chi'(1) + f_{ss}(1)) + \frac{f_x(1)g_s(1)(g_{xx}(1)\chi'(1) + g_{xs}(1))}{(g_x(1))^2} \\
- \quad & \frac{(f_{xx}(1)\chi'(1) + f_{xs}(1))g_s(1) + f_x(1)(g_{xs}(1)\chi'(1) + g_{ss}(1))}{g_x(1)}\Big\}\Bigg] \\
- \quad & \frac{f_x(0)}{g_x(0)}(g_{xx}(0)\chi'(0) + g_{xs}(0)) = 0
\end{aligned}
$$

In a steady state, we have

$$
\begin{aligned}
\left.\frac{d^2 S_t}{dS_1^2}\right|_{(\chi(S^*),S^*)} &= (g_x^*\chi'^* + g_s^*)^{t-2}(g_{xx}^*(\chi'^*)^2 + 2g_{xs}^*\chi'^* + g_x^*\chi''^* + g_{ss}^*)\sum_{\tau=1}^{t-1}(g_x^*\chi'^* + g_s^*)^{\tau-1} \\
&= (g_{xx}^*(\chi'^*)^2 + 2g_{xs}^*\chi'^* + g_x^*\chi''^* + g_{ss}^*)\frac{(g_x^*\chi'^* + g_s^*)^{2t-3} - (g_x^*\chi'^* + g_s^*)^{t-2}}{g_x^*\chi'^* + g_s^* - 1}
\end{aligned}
$$

Therefore, χ'' must satisfy the following equation linear in χ''.

$$
\begin{aligned}
& (f_{xx}^*\chi'^* + f_{xs}^*) + g_x^*(g_x^*\chi'^* + g_s^*)\cdot\Bigg[\sum_{t=1}^{T}(\theta_t - \delta\theta_{t-1}) \\
\cdot \quad & \Big\{(f_x^*\chi'^* + f_s^*)(g_{xx}^*(\chi'^*)^2 + 2g_{xs}^*\chi'^* + g_x^*\chi''^* + g_{ss}^*)\frac{(g_x^*\chi'^* + g_s^*)^{2t-3} - (g_x^*\chi'^* + g_s^*)^{t-2}}{g_x^*\chi'^* + g_s^* - 1} \\
+ \quad & (f_{xx}^*(\chi'^*)^2 + 2f_{xs}^*\chi'^* + f_x^*\chi''^* + f_{ss}^*)(g_x^*\chi'^* + g_s^*)^{2t-2}\Big\} \\
+ \quad & \delta\Big\{(f_{xs}^*\chi'^* + f_{ss}^*) + \frac{f_x^* g_s^*(g_{xx}^*\chi'^* + g_{xs}^*)}{(g_x^*)^2} - \frac{(f_{xx}^*\chi'^* + f_{xs}^*)g_s^* + f_x^*(g_{xs}^*\chi'^* + g_{ss}^*)}{g_x^*}\Big\}\Bigg] \\
- \quad & \frac{f_x^*}{g_x^*}(g_{xx}^*\chi'^* + g_{xs}^*) = 0
\end{aligned} \tag{12}
$$

References

AGUIRREGABIRIA, V., AND P. MIRA (2002): "Swapping the Nested Fixed Point Algorithm: a Class of Estimators for discrete Markov decision models," *Econometrica*, 70(4), 1519–43.

AZAR, O. (1999): "Rationalizing hyperbolic discounting," *Journal of Economic Behavior and Organization*, 38, 245–52.

BARRO, R. (1999): "Ramsey meets Laibson in the neoclassical growth model," *Quarterly Journal of Economics*, 114, 1125–52.

CHICHILNSKY, G. (1996): "An Axiomatic Approach to Sustainable Development," *Social Choice and Welfare*, 13, 231–257.

DOCKNER, E., J. STEFFEN, N. V. LONG, AND G. SORGER (2000): *Differential Games in Economics and Management Sciences*. Cambridge University Press, Cambridge, U.K.

DYBVIG, P., J. INGERSOLL, AND S. ROSS (1996): "Long forward and zero coupon rates can never fall," *Journal of Business*, 69, 1–25.

FUJII, T. (2006): "Manualfor discrete time continuous state quasi-dynamic programming equation solver," Singapore Management University.

GOLLIER, C. (2002a): "Discounting an uncertain future," *Journal of Public Economics*, 85, 149–166.

——— (2002b): "Time horizon and the discount rate," *Journal of Economic Theory*, 107, 463–73.

GOLLIER, C., AND R. ZECKHAUSER (2005): "Aggregation of heterogeous time preferences," *Journal of Political Economy*, 113, 878–896.

GROOM, B., C. HEPBURN, P. KOUNDOURI, AND D. PEARCE (2005): "Declining discount rates: The long and the short of it," *Environment and Resource Economics*, 32, 445–93.

HANSEN, L. P., AND T. J. SARGENT (1980): "Formulating and Estimating dynamic linear rational expectations models," *Journal of Economic Dynamics and Control*, 2(1), 7–46.

HARRIS, C., AND D. LAIBSON (2001): "Dynamic Choices of Hyperbolic Consumers," *Econometrica*, 69(5), 935–957.

HEAL, G. (2001): "Intertemporal Welfare Economics and the Environment," in *The Handbook of Environmental Economics*, ed. by K.-G. Maler, and J. Vincent. North Holland, Amsterdam.

JUDD, K. (1998): *Numerical Methods in Economics*. MIT Press, Cambridge, Massachusetts.

KARP, L. (2005): "Global Warming and hyperbolic discounting," *Journal of Public Economics*, 8, 261–282.

——— (2006): "Non-constant discounting in continuous time," *Journal of Economic Theory, forthcoming*.

LI, C. Z., AND K. G. LOFGREN (2000): "Renewable Resources and Economic Sustainability: A Dynamic Analysis with Heterogeneous Time Preferences," *Journal of Environmental Economics and Management*, 40, 236–250.

MIRANDA, M. J., AND P. L. FACKLER (2002): *Applied Computational Economics and Finance*. MIT Press, Cambridge, Massachusetts.

PAKES, A., AND P. MCGUIRE (1994): "Computing Markov-perfect Nash equilibria: numerical implications of a dynamic differentiated product model," *Rand Journal of Economics*, 25(4), 555–589.

RABIN, M. (1998): "Psychology and Economics," *Journal of Economic Literature*, 36, 11–46.

RUST, J. (1994): "Structural Estimation of Markov Decision Processes," in *Handbook of Econometrics, Vol 10*, ed. by D. L. McFadden. Elsevier Science, Amsterdam.

TSUTSUI, S., AND K. MINO (1990): "Nonlinear Strategies in Dynamic Duopolisitc Competition with Sticky Prices," *Journal of Economic Theory*, 52, 136–161.

WEITZMAN, M. (2001): "Gamma Discounting," *American-Economic-Review*, 91(1), 260–271.