# A Comparison of Discrete and Parametric Approximation Methods for Continuous-State Dynamic Programming Problems[†]

Hugo Benítez-Silva, *SUNY at Stony Brook*
George Hall, *Yale University*
Günter J. Hitsch, *Yale University*
Giorgio Pauletto, *University of Geneva*
John Rust[‡] *Yale University*

Preliminary and Incomplete. Comments Welcome

October 30, 2000

[†] Benítez-Silva and Hitsch are grateful for the financial support of the Cowles Foundation for Research in Economics through a Carl Arvid Anderson Dissertation Fellowship. Benítez-Silva is also grateful to the John Perry Miller Fund for additional financial support. Hall and Rust gratefully acknowledge financial support from a National Science Foundation grant, SES-9905145. Pauletto gratefully acknowledges financial support from the Swiss National Science Foundation grant 8210–50418.

[‡] **Corresponding author:** John Rust, Department of Economics, Yale University, 37 Hillhouse Avenue, New Haven CT 06520-8264, phone: (203) 432-3569, fax: (203) 432-6323, e-mail: `jrust@gemini.econ.yale.edu`

**Abstract**

We compare alternative numerical methods for approximating solutions to continuous-state dynamic programming (DP) problems. We distinguish two approaches: *discrete approximation* and *parametric approximation*. In the former, the continuous state space is discretized into a finite number of points $N$, and the resulting finite-state DP problem is solved numerically. In the latter, a function associated with the DP problem such as the *value function*, the *policy function*, or some other related function is approximated by a smooth function of $K$ unknown parameters. Values of the parameters are chosen so that the parametric function approximates the true function as closely as possible. We focus on approximations that are linear in parameters, i.e. where the parametric approximation is a linear combination of $K$ *basis functions*. We also focus on methods that approximate the value function $V$ as the solution to the *Bellman equation* associated with the DP problem. In finite state DP problems the method of *policy iteration* is an effective iterative method for solving the Bellman equation that converges to $V$ in a finite number of steps. Each iteration involves a *policy valuation step* that computes the value function $V_\alpha$ corresponding to a trial policy $\alpha$. We show how policy iteration can be extended to continuous-state DP problems. For discrete approximation, we refer to the resulting algorithm as *discrete policy iteration* (DPI). Each policy valuation step requires the solution of a system of linear equations with $N$ variables. For parametric approximation, we refer to the resulting algorithm as *parametric policy iteration* (PPI). Each policy valuation step requires the solution of a linear regression with $K$ unknown parameters. The advantage of PPI is that it is generally much faster than DPI, particularly when $V$ can be well-approximated with small $K$. The disadvantage is that the PPI algorithm may either fail to converge or may converge to an incorrect solution. We compare DPI and PPI to parameteric methods applied to the Euler equation for several test problems with closed-form solutions. We also compare the performance of these methods in several “real” applications, including a life-cycle consumption problem, an inventory investment problem, and a problem of optimal pricing, advertising, and exit decisions for newly introduced products.

**Keywords:** Dynamic Programming, Numerical Methods, Policy Iteration, Linear-Quadratic problems, Consumption/Saving problems, Stochastic growth problems, Inventory control problems, Product advertising and pricing problems.
**JEL classification:** C0, D0

# 1 Introduction

Despite the rapid growth in computing power and new developments in the literature on numerical dynamic programming in economics (for recent surveys see Rust 1996, Santos 1999, the text by Judd 1998, and the collection of essays edited by Marimon and Scott 1999), multi-dimensional infinite-horizon continuous-state dynamic programming (DP) problems are still quite challenging to solve. Most economists are aware of the "curse of dimensionality" and the limits it places on our ability to solve high-dimensional DP problems. Despite recent theoretical results that suggest that it is possible to break the curse of dimensionality under certain conditions (see Rust 1997a and Rust, Traub and Woźniakowski 2000), solutions to most high-dimensional DP problems are still beyond our grasp even using the best algorithms and the fastest workstations and supercomputers.

There is considerable disagreement in the literature about the most efficient algorithms to solve high-dimensional DP problems. The debate is roughly whether it is better to solve DP problems by *discrete approximation* or by *parametric approximation.* In the former approach, the continuous state space is discretized into a finite number of *grid points,* $N$, and the resulting finite-state DP problem is solved numerically. The value function and policy function can be computed at points in the state space that are not elements of the predefined grid via interpolation. In parametric approximation, the value or policy function (or some other related function) is approximated by a smooth parametric function with $K$ unknown parameters. These parameters are chosen in such a way that the resulting function "best fits" the true solution according to some metric. The argument for the superiority of the parametric approximation approach is roughly that in many cases, one can obtain a good global approximation to a function in question using a small number of parameters $K$, whereas in high-dimensional problems discretization requires very large values of $N$ to obtain a comparably accurate approximation.

It is true that naive discretization of multidimensional DP problems leads directly to the curse of dimensionality, since in a $d$-dimensional problem one can show that $O(1/\varepsilon)$ points in each dimension, or a total of $N = O(1/\varepsilon^d)$ grid points, are required in order to obtain an $\varepsilon$-approximation to the value or policy function. Since $N$ increases exponentially fast in the dimension $d$, it follows that naive discretization results in a curse of dimensionality. However the fact that naive discretization leads to a curse of dimensionality does not imply that *all* ways of discretizing the problem necessarily produce a curse of dimensionality. Rust's (1997a) "random multigrid algorithm" breaks the curse of dimensionality using a random discretization of the state space. This algorithm results in approximate solution to the DP problem with an *expected error* of $\varepsilon$ using only $N = O(1/\varepsilon^2)$ points. However the regularity conditions for Rust's result require a Lipschitz-continuous transition probability for the state variables, and in some economic applications this condition will not be satisfied. In addition, Rust's result applies to DP problems where the control variable takes on only a finite number of possible values: we do not know whether Rust's result can be extended to problems where the control variables are continuous.

The appeal of parametric approximation methods is that a potentially infinite-dimensional problem (e.g. finding the solution $V$ to the Bellman functional equation) is reduced to a finite-dimensional problem with a relatively small number $K$ of unknown parameters. To illustrate this approach, suppose we are interested in approximating the value function $V(s)$, which is the unique solution to *Bellman's equation*

$$V(s) = \Gamma(V)(s) \equiv \max_{a \in A(s)} \left[ u(s,a) + \beta \int V(s') p(s'|s,a) ds' \right]. \tag{1}$$

Suppose we conjecture that $V$ can be approximated as a linear combination of a relatively small number $K$ of "basis functions" $\{\rho_1(s), \ldots, \rho_K(s)\}$

$$V_\theta(s) = \sum_{k=1}^{K} \theta_k \rho_k(s). \tag{2}$$

If the true $V$ is not too irregular, and if we have chosen a "good basis" we will be able to find a good approximation to $V$ for a relatively small value of $K$. The goal is to find a particular parameter value $\hat{\theta}$ such that the approximate value function $V_{\hat{\theta}}$ "best fits" the true value function $V$. Since $V$ is not known, this can't be done directly. However since $V$ is the zero to a certain *residual function*, $\Psi(V) = V - \Gamma(V)$, this suggests that there should be ways of solving for $\hat{\theta}$ so that the resulting function $V_{\hat{\theta}}$ should be a good approximation to $V$.

Consider the case where the state space $S$ is a compact subset of $R^d$, where $u(s,a)$ is a bounded, continuous function of $(s,a)$, and the conditional expectation operator is weakly continuous (i.e. where $Eh(s,a) = \int h(s') p(s'|s,a)$ is a bounded, continuous bounded function of $(s,a)$ for each continuous, bounded function $h$). In this case we know the value function $V$ will be an element of $B(S)$, the Banach space of bounded, continuous functions of $S$. It will be the unique solution Bellman's equation, or alternatively the zero to the residual operator $\Psi$. This can be expressed as

$$V = \underset{\{W \in B(S)\}}{argmin} \|W - \Gamma(W)\|. \tag{3}$$

where $\|W\|$ is the usual sup-norm, i.e. $\|W\| = \sup_{s \in S} |W(s)|$. This representation of the problem suggests that we should choose $\hat{\theta}$ as the corresponding solution to the finite-dimensional minimization problem:

$$V_{\hat{\theta}} = \underset{\{V_\theta | \theta \in R^K\}}{argmin} \|V_\theta - \Gamma(V_\theta)\|. \tag{4}$$

Using fact that $\Gamma$ is a contraction mapping, a simple application of the triangle inequality yields the following error bound:

$$\|V_{\hat{\theta}} - V\| \leq \frac{\|V_{\hat{\theta}} - \Gamma(V_{\hat{\theta}})\|}{(1-\beta)}. \tag{5}$$

Thus, to the extent that we can find a "good basis" $\{\rho_1, \ldots, \rho_K\}$ with a relatively small number of elements $K$ such that the quantity $\|V_{\hat{\theta}} - \Gamma(V_{\hat{\theta}})\|$ is small, we can be guaranteed that $V_{\hat{\theta}}$ is a good global approximation to the true solution $V$. Further, to the extent that it does not take too many evaluations of the error function $g(\theta) = \|V_\theta - \Gamma(V_\theta)\|$ to find the minimizing parameter vector $\hat{\theta}$, the parametric approximation approach could be much faster than discrete approximation of $V$.

However we are not aware of a formal proof that parametric approximation methods similar to the one outlined above succeed in breaking the curse of dimensionality. Indeed there are several reasons why we would expect parametric approximation to be subject to an unavoidable curse of dimensionality. First, in the absence of some sort of "special structure", the number of basis functions required to provide a uniform approximation to a smooth function of $d$ variables increases exponentially in $d$ (see, e.g. Traub, Wasilkowski, and Woźniakowski 1988, and Traub and Werschulz 1998). Second, the objective function $g(\theta) = \|V_\theta - \Gamma(V_\theta)\|$ is generally not concave in $\theta$ (and may not even be smooth in $\theta$), and there is a well known curse of dimensionality associated with solving non-concave minimization problems, regardless of whether deterministic or random algorithms are allowed (see Nemirovsky and Yudin 1978). Indeed, we are not aware of any formal analysis of the complexity or parametric approximation methods, or even a derivation of error bounds or proofs of convergence that account for the fact that the function $g(\theta)$ cannot generally be evaluated exactly. Instead both the Bellman operator $\Gamma(V_\theta)$ and the sup-norm, $\|V_\theta - \Gamma(V_\theta)\|$ must be approximated, and it can be costly to approximate these objects to a sufficient level of accuracy to insure that $V_{\hat{\theta}}$ does in fact provide a good approximation to $V$.

Practical applications of parametric approximation methods (see e.g. Taylor and Uhlig 1990, Deaton and Laroque 1992, Gaspar and Judd 1997, Santos 1999, Miranda 1998 and Christiano and Fisher 2000) have yielded mixed results. In some cases the nonlinear optimization problem can be solved quickly and reliably; but others have been plagued by problems multiple optima and have experienced considerable difficulty in getting the minimization problem (4) to converge, especially when the underlying function being approximated has kinks or discontinuities. In this paper we propose an alternative parametric approximation strategy based on iterative solution of a sequence of parametric minimization problems each of which can be solved by the method of *ordinary least squares* (OLS). This method is motivated by the iterative *policy iteration algorithm* for solving finite and infinite-dimensional DP problems (see Howard 1960, and Puterman and Brumelle 1979). Under mild regularity conditions it can be proved that policy iteration results in a monotonically improving sequence of approximate value functions that converge to $V$ in a finite number of iterations.

Policy iteration will be described more formally in section 2, but briefly, it consists of an alternating sequence of *policy improvement* and *policy valuation* steps. The policy valuation step results in a linear functional equation for the value function $V_\alpha$ corresponding to policy $\alpha$:

$$V_\alpha(s) = u(s,\alpha(s)) + \beta \int V_\alpha(s')p(s'|s,\alpha(s))ds' \tag{6}$$

Discrete approximation methods involve solving an approximate finite state DP problem defined over a grid of $N$ points $\{s_1,\ldots,s_N\}$ in the state space. Discretization converts the infinite-dimensional linear functional equation into a system of $N$ linear equations in the $N$ unknowns $\{V_\alpha(s_1),\ldots,V_\alpha(s_N)\}$. The amount of work required to solve this system is bounded by $O(N^3)$, the time required by standard linear equation solvers (e.g. LU factorization and back-substitution) for dense systems.

Now consider solving the policy valuation step (6) via a linear parametric approximation to $V_\alpha$ such as in equation (2). Substituting the parametric approximation of $V_\alpha$ into equation (6) we obtain

$$\sum_{k=1}^{K} \theta_k \rho_k(s) = u(s,\alpha(s)) + \beta \sum_{k=1}^{K} \int \theta_k \rho_k(s')p(s'|s,\alpha(s)). \tag{7}$$

If we evaluate this equation at $M$ points $\{s_1,\ldots,s_M\}$ where $M \geq K$, we can solve for the value $\hat{\theta}$ that approximately solves equation (7) by the method of *ordinary least squares* (OLS). In fact, if $M = K$ and the points $\{s_1,\ldots,s_K\}$ are chosen so that the $K \times K$ matrix $X$ whose $(i,j)$ element is given by

$$x_{ij} = \rho_i(s_j) - \beta \int p_i(s')p(s'|s_j,\alpha(s_j))ds' \tag{8}$$

has full rank, then we can find an exact solution to the system (7). The solution is given by $\hat{\theta} = (X'X)^{-1}X'y$ where $y_j = u(s_j,\alpha(s_j))$. The hope is that if we have chosen a “good basis” that enables us to find a good approximation to $V_\alpha$ for small $K$, then *parametric policy iteration* (PPI) will be far faster than a *discrete policy iteration* (DPI), since the value of $K$ necessary to obtain an $\varepsilon$-approximation to $V_\alpha$ will be far smaller than the value $N$ that would be required by discrete approximation methods. The other key advantage of PPI is that unlike problem (4) the implied global minimization problem has an explicit solution and can be carried out in $O(K^3)$ time in the worst case. If PPI also shares the rapid, global convergence rate of ordinary policy iteration, then it could be quite promising for solving high-dimensional DP problems.

We compare the performance of DPI and PPI in a number of “test problems” that admit closed-form solutions for the value and policy functions. These test problems include the infinite-horizon consumption/saving problem studied by Phelps (1962) and Hakansson (1970), the finite-horizon consumption/savings problem, the linear-quadratic optimal control problem studied by Holt, Modigliani, Muth and Simon (1960) and Hansen and Sargent (1999), an optimal replacement model studied by Rust (1985, 1986, 1987), and an stochastic growth model analyzed in Santos (1999). We also compare the performance of parametric and discrete approximation methods in several "real" applications including a model of optimal consumption and labor supply at the end of the life-cycle studied by Benítez-Silva (2000), a model optimal inventory investment and commodity price speculation studied in Hall and Rust (1999a,b), and a model of optimal pricing and advertising and product exit decisions for newly introduced products studied by Hitsch (2000).

For each of these test problems, we compare DPI and PPI to Euler-based parametric algorithms, often referred to as *projection* or *minimum weighted residual* methods. These methods are widely throughout the economics to study issues such as commodity storage, asset pricing, and optimal fiscal policy.[1] We do not attempt to survey all the algorithms that fall under this broad class of strategies, but instead study particular projection methods that involve parameterizing either directly or indirectly a certain conditional expectation function that enters the *Euler equation.* The Euler equation is derived from the first order condition to DP problems with continuous control variables. Thus the methods we study can be considered *parametrized expectations algorithms* (PEA).[2] We implement our PEAs guided by the recommendations of Judd (1992, 1994, 1998) and Christiano and Fisher (2000).

Section 2 reviews the finite- and infinite-dimensional versions of the Howard (1960) policy iteration algorithm. However since it is not feasible to exactly solve infinite-dimensional linear equations (linear functional equations), we describe ways of forming feasible approximations to the systems that must be solved when using policy iteration in DP problems with continuous state spaces. This leads us to formally define the DPI and PPI variants of policy iteration. We provide a taxonomy of different variants of of these algorithms corresponding to different ways of discretizing the state space, different basis functions for parametric approximations, different quadrature methods for computing integrals underlying conditional expectations, different optimization algorithms for approximating the max operator in the Bellman equation, and so forth. We also provide a general description of the particular Euler-based projection methods we use as comparative solution techniques. Section 3 introduces the test problems used in our study and

[1] From the commodity storage literature, examples of these techniques can be found in Miranda and Helmberger (1988) and Miranda (1998). Marshall (1992) uses these methods to study asset returns. For examples in the optimal fiscal policy literature, see Braun and McGrattan (1993), Chari, Christiano and Kehoe (1994) and Marcet, Sargent and Seppala (2000).

[2] The term, parameterized expectations algorithm, was coined by Marcet (1988) and den Hann and Marcet (1990); however, the first use of PEA appears to be Wright and Williams (1982a, 1982b, 1984).

presents their analytical solutions. Section 4 presents results for a stochastic growth problem with and without leisure. Section 5 presents the results of our numerical comparisons for the finite and infinite-horizon consumption/savings problem. Section 6 presents results for an optimal replacement problem, which unlike the previous problems is one with a discrete control variable and where there is a kink in the optimal value function. Section 7 presents results for the linear-quadratic-gaussian (LQG) control problem, our last test problem. Section 8 introduces the more realistic models presenting results for Hall and Rust's (1999a,b) model of optimal inventory investment and commodity price speculation. Section 9 presents results for Hitsch's (2000) model of pricing and advertising and product exit decisions for newly introduced consumer products. Section 10 shows the results for Benítez-Silva's (2000) analysis of consumption/savings and labor supply decisions at the end of the life cycle. Section 11 presents our conclusions about the performance of the various algorithms, and our recommendations for future research in this area.

We realize that the large number of methods and algorithms available for solving DP problems actually presents a daunting burden to non-experts who are interested in solving a specific problem. Our hope is that by studying a larger range of problems, practioners interested in solving a specific DP problem will find their problem to be sufficiently similar to one of the problems analyzed here that they might be able to use this analysis to help them select the algorithm that is likely to be best for their particular problem. We have attempted to provide a clear summary of the strengths and weaknesses of various methods and to present our "bottom line" recommendations about the algorithms that work best for various problems. Most importantly, we also provide (via the web site `http://gemini.econ.yale.edu/jrust/sdp`) fully documented source code in Gauss, Matlab, and C that implement all the methods and will recreate all the results presented in this paper. Our hope that providing this software library to the economics community will accelerate the use of these methods and enable the profession to get further practical experience with these methods and hopefully, expand the range of interesting applied problems that can be solved in practice.

## 2 Algorithms

This section reviews some basic facts about infinite horizon DP problems and provides a brief description of policy iteration, DPI and PPI algorithms, and the parameterized expectations algorithm.

## 2.1 Review of the DP Problem

Consider an infinite horizon dynamic programming problem where the state $s \in S \subset R^d$. Bellman's equation is

$$V(s) = \max_{a \in A(s)} [u(s,a) + \beta \int V(s')p(s'|s,a)ds']. \tag{9}$$

The *optimal policy* $\alpha(s)$ is the solution to:

$$\alpha(s) = \underset{a \in A(s)}{argmax} [u(s,a) + \beta \int V(s')p(s'|s,a)ds']. \tag{10}$$

In abstract terms $V = \Gamma(V)$ is the unique fixed point to the *Bellman operator* $\Gamma : B \to B$, where $B$ is a Banach space of functions from $S$ to $R$ and the Bellman operator is given by

$$\Gamma\Big(V\Big)(s) \equiv \max_{a \in A(s)} [u(s,a) + \beta \int V(s')p(s'|s,a)ds']. \tag{11}$$

Most previous work has focused on proving approximation theorems based on some type of *Discretized Bellman operator:*

$$\Gamma_N\Big(V\Big)(s) \equiv \max_{a \in A(s)} \left[ u(s,a) + \beta \sum_{i=1}^{N} V(s_i) p_N(s_i|s,a) \right], \tag{12}$$

where $p_N$ is a discrete probability distribution over a finite *grid* $\{s_1, \ldots, s_N\}$ in $S$, where $S = [0,1]^d$ for simplicity. In that case $\Gamma_N$ has a dual interpretation, it can be regarded as a contraction mapping on $R^N$, (this is where the computation is done, resulting in a fixed point $V_N \in R^N$), but it is also a valid contraction mapping $\Gamma_N : B \to B$. This latter feature makes it easy to prove approximation bounds since the function $\Gamma(V_N)$ can be regarded as an element of $B$, and thus is a natural candidate as an approximation to $V = \Gamma(V)$.

Now consider an alternative way of approximating $V$, namely as a linear combination of a set of *basis functions* $\{\rho_1(s), \rho_2(s), \ldots, \rho_K(s)\}$. These functions may not literally be a basis for $B$, but should have the property that the sequence is *ultimately dense* in $B$ in the sense that for any $V \in B$ we have:

$$\lim_{K \to \infty} \inf_{\theta_1, \ldots, \theta_k} \sup_{s \in S} |V(s) - \sum_{i=1}^{K} \theta_i \rho_i(s)| = 0. \tag{13}$$

It is possible that families of functions that are nonlinear in the parameters $\theta$ could be considered also, such as neural net and wavelet bases. We restrict attention to bases which are linear in parameters for simplicity, since as we will see below it vastly simplifies the problem of determining the optimal values of $\theta$: the optimal $\hat{\theta}$ will be the solution to a simple ordinary least squares problem which is trivial to compute. If the basis is a nonlinear function of $\theta$ then we will have to solve a nonlinear least squares problem, which could be more time consuming and it may be difficult or impossible to prove that the algorithm breaks the curse of dimensionality.[3]

## 2.2 Policy Iteration for Continuous and Discrete MDPs

To understand the PPI algorithm we first review the infinite dimensional version of policy iteration. Puterman and Shin (1978) proved the convergence of this algorithm, showing that it is basically equivalent to an infinite-dimensional version of Newton's method for solving the nonlinear equation

$$(I-\Gamma)(V)=0.$$

The algorithm consists of alternating *policy valuation* and *policy improvement* steps.

**Policy Iteration (Infinite-Dimensional Version)**

1. **Policy Valuation Step:** given an initial guess of policy $\alpha$ compute $V_\alpha$, the value function implied by policy $\alpha$:

$$V_\alpha(s)=u(s,\alpha(s))+\beta E_\alpha V_\alpha(s), \tag{14}$$

   where $E_\alpha$ is the *Markov operator* corresponding to $\alpha$:

$$E_\alpha V(s)=\int V(s')p(s'|s,\alpha(s))ds'. \tag{15}$$

   There is a unique solution to the linear operator equation defining $V_\alpha$ (Fredholm integral equation of the second kind):

$$V_\alpha=u_\alpha+\beta E_\alpha V=(I-\beta E_\alpha)^{-1}u_\alpha, \tag{16}$$

   where $(I-\beta E_\alpha)$ exists and has the following geometric series or *Neumann series* representation:

$$(I-\beta E_\alpha)^{-1}=\sum_{t=0}^{\infty}[\beta E_\alpha]^t. \tag{17}$$

2. **Policy Improvement Step:** Compute improved policy $\alpha'$ using $V_\alpha$:

$$\alpha'(s)=\underset{a\in A(s)}{argmax}[u(s,a)+\beta\int V_\alpha(s')p(s'|s,a)ds']. \tag{18}$$

[3] Barron's 1993 result on the properties of neural nets as a means of breaking the curse of dimensionality of approximating certain classes of functions notwithstanding, there is the computational problem of finding a globally minimizing $\theta$ vector and this is where a curse of dimensionality could arise.

If the Policy iteration algorithm converges, it is easy to see that the policy $\alpha^*$ that it converges to, and the corresponding value function $V_{\alpha^*}$ are solutions to Bellman's equation, and thus are the solution to the DP problem. It is well known that policy iteration always converges in a finite number of steps from any starting point if the state space $S$ and action sets $A(s)$, $s \in S$ are finite sets. Puterman and Shin provided sufficient conditions for policy iteration to converge when $S$ and $A(s)$ contain a continuum of points.

One strategy for approximating the solutions to continuous state DP problems is via discretation that results on an approximate MDP problem on a finite state space $S_N = \{s_1, \dots, s_N\}$, and the use of policy iteration for the finite MDP on $S_N$. This results in the finite-dimensional version described below.

**Policy Iteration (Finite-Dimensional Version)**

1. **Policy Valuation Step:** given an initial guess of policy $\alpha$ compute $V_{\alpha,N}$, the value function (in $R^N$) implied by policy $\alpha$:

$$V_{\alpha,N}(s) = u(s,\alpha(s)) + \beta E_{\alpha,N} V_{\alpha,N}(s), \tag{19}$$

where $E_{\alpha,N}$ is the *discrete Markov operator* corresponding to $\alpha$:

$$E_{\alpha,N} V(s) = \sum_{i=1}^{N} V(s_i) p_N(s_i|s,\alpha(s)). \tag{20}$$

There is a unique solution to the linear system of equations defining $V_{\alpha,N} \in R^N$

$$V_{\alpha,N} = u_\alpha + \beta E_{\alpha,N} V_{\alpha,N} = (I - \beta E_{\alpha,N})^{-1} u_\alpha, \tag{21}$$

where $(I - \beta E_{\alpha,n})$ exists and has the following geometric series representation:

$$(I - \beta E_{\alpha,N})^{-1} = \sum_{t=0}^{\infty} [\beta E_{\alpha,N}]^t. \tag{22}$$

where $I$ is the $N \times N$ identity matrix and $E_{\alpha,N}$ is the $N \times N$ Markov transition matrix with $(i,j)$ entry given by:

$$E_{\alpha,N}[i,j] = [p_N(s_i|s_j,\alpha(s_j))]. \tag{23}$$

2. **Policy Improvement Step:** Compute improved policy $\alpha'$ using $V_\alpha$:

$$\alpha'(s) = \underset{a \in A(s)}{argmax} [u(s,a) + \beta \sum_{i=1}^{N} V_{\alpha,N}(s_i) p_N(s_i|s,a)]. \tag{24}$$

### 2.3 Parametric Policy Iteration

This algorithm is basically the same as the infinite dimensional version of policy iteration, except that we approximately solve each policy valuation step by approximating the solution $V_\alpha$ as a linear combination of $k$ basis functions $\{\rho_1,\dots,\rho_k\}$. Thus, suppose we set

$$V_\alpha(s) \simeq \sum_{i=1}^{k} \theta_i \rho_i(s). \tag{25}$$

Then the equation for $V_\alpha$

$$V_\alpha(s) = u(s,\alpha(s)) + \beta \int V_\alpha(s')p(s'|s,\alpha(s)ds' \tag{26}$$

is transformed into a linear equation with $k$ unknown parameters $\theta \equiv \{\theta_1,\dots,\theta_k\}$:

$$\sum_{i=1}^{K} \theta_i \rho_i(s) = u(s,\alpha(s)) + \beta \int \sum_{i=1}^{K} \theta_i \rho_i(s')p(s'|s,\alpha(s))ds' \tag{27}$$

Suppose we evaluate the above equation at a set of $M$ points in $S$, with $M \geq K$. Then define the $(M \times K$ matrices $P$ and $EP$ with elements $P_{j,k}$ and $EP_{j,k}$ given by

$$P_{j,k} = \rho_k(s_j) \tag{28}$$

$$EP_{j,k} = \int \rho_k(s')p(s'|s_j,\alpha(s_j)). \tag{29}$$

Define the $(M \times 1)$ vector $y$ with $j-th$ element $y_j$ given by

$$y_j = u(s_j,\alpha(s_j)). \tag{30}$$

and let the $(M \times K)$ matrix $X$ be given by

$$X = (P - \beta EP) \tag{31}$$

Then the system of equations (27) can be written in matrix form as

$$y = X\theta. \tag{32}$$

If $M = K$ and $X$ is invertible the solution for $\theta$ is simply

$$\hat{\theta} = y/X = X^{-1}y. \tag{33}$$

If $M > K$ we have an over-determined system and in general there is no $\theta \in R^K$ that allows us to exactly solve $y = X\theta$. However we can form an approximate solution using the ordinary least squares estimator (OLS), i.e. the value $\hat{\theta}$ that minimizes the distance $\|y - X\theta\|^2$, is given by

$$\hat{\theta} = y/X = (X'X)^{-1}X'y \tag{34}$$

In general we will not be able to exactly integrate the basis functions and must use a quadrature rule to approximate the elements of $EP$ given in equation (29). Thus, the PPI algorithm requires the following choices:

1. The quadrature rule for computing the elements of $EP$.

2. The sample points $\{s_1, \dots, s_M\}$ at which $P$ and $EP$ are evaluated.

3. The set of basis functions $\{\rho_1, \dots, \rho_K\}$.

Note that the policy improvement step would only be done on the same $M$ points $\{s_1, \dots, s_M\}$. Thus only $M \times K$ numerical integrations and $M$ maximizations are required for each policy valuation and policy improvement step, so if $M$ and $K$ can be chosen to be small, it is possible to find an approximation solution to the DP problem with amazingly few computations, provided the basis functions $\{\rho_1(s), \dots, \rho_K(s)\}$ are sufficiently easy to evaluate at each $s \in S$. The resulting solution is defined by a parameter vector $\hat{\theta}$ that enables us to evaluate $V_{\hat{\theta}}(s) = \sum_{k=1}^{K} \hat{\theta}_k \rho_k(s)$ very rapidly at any $s \in S$. Evaluating the corresponding decision $\alpha(s)$ at that point would require an approximate solution to

$$\alpha(s) = \max_{a \in A(s)} \left[ u(s,a) + \beta \int \sum_{k=1}^{K} \hat{\theta}_k \rho_i(s') p(s'|s,a) ds' \right]. \tag{35}$$

In many cases this can be done quite rapidly, or alternatively, using the values $\{\alpha(s_i)\}$, $j = 1, \dots, M$ from the last step of policy iteration, it might be possible to interpolate values of $\alpha(s)$ for $s \in \{s_1, \dots, s_M\}$ if the decision rule is sufficiently smooth.

## 2.4 Euler-based Projection Methods

In the subsection we sketch the basic strategy behind applying projection methods to solve numerically for decision rules that satisfy a set of first-order equations. This class of methods is broad and encompasses a wide variety of algorithms. We will not attempt to survey the entire class. Instead we refer interested readers to Judd (1998) and McGrattan (1999). Guided by the insights of Judd (1992, 1994, 1998) and Christiano and Fisher (2000) we focus on a small number of algorithms that parameterize either directly or indirectly the conditional expectation term in the Euler equation; hence these approaches are types of *parameterized expectations algorithms* (PEA).

Usually projection or minimum weighted residual methods are not applied directly to the Bellman equation, (9), but instead to a set of first order conditions or an Euler equation.[4] To get the stochastic Euler equation for the general problem described in (9) we take first-order conditions and applying the envelop theorem. This yields:

$$\frac{\partial u(s,a)}{\partial a} + \beta \int \frac{\partial u(s',a')}{\partial s} \frac{\partial p(s'|s,a)}{\partial a} ds' = 0. \tag{36}$$

The basic strategy of the projection methods we employ is to find a parametric approximation to $\int \frac{\partial u(s',a')}{\partial s} \frac{\partial p(s'|s,a)}{\partial a} ds'$ which depends only on a finite-dimensional $(K \times 1)$ vector of parameters $\theta$. This expectation is approximated by a finite linear combination of $K$ known basis functions, $\rho(s)$. Thus we can write

$$\int \frac{\partial u(s',a')}{\partial s} \frac{\partial p(s'|s,a)}{\partial a} ds' \approx \sum_{i=1}^{K} \theta_i \rho_i(s).$$

It is often convenient to parameterize this conditional expected marginal return function indirectly such as by parameterizing the decision rule(s) for $a$. In this case:

$$\int \frac{\partial u(s',a')}{\partial s} \frac{\partial p(s'|s,a)}{\partial a} ds' \approx f\left(\sum_{i=1}^{K} \theta_i \rho_i(s)\right).$$

Either way, the basic strategy remains the same. In sections 4.1 and 4.2, we solve the stochastic growth using this method: first by parameterizing a decision rule, second by directly parameterizing the conditional expectation function.

Given an approximation to the conditional expectation, we find the vector $\theta$ which which set a weighted sum of a residual function, $R(s,a)$ as close as possible to 0 for all $s$. Mathematically this means choosing $\theta$ such that

$$\int w(s) R(s, a(s)|\theta)\ ds = 0 \tag{37}$$

where $w(s)$ is a weighting function. In this paper, the Euler equation itself will be the residual function.

Since it is rarely the case that the integrals in (37) can be evaluated directly, we will need to approximate the integrals via a quadrature method over a discrete grid of $M$ points in $S$. Thus we find $\theta$ such that

$$\frac{\partial u(s,a)}{\partial a} + \beta \sum_{i=1}^{M} \left( w(s_i) \frac{\partial p(s_i|s,a)}{\partial a} \sum_{j=1}^{K} \theta_j \rho_j(s_i) \right) = 0 \tag{38}$$

[4] However, it is possible to formulate the parametric policy iteration algorithm as a projection problem.

where $p(s_i|s,a)$ is a discretized approximation to the transition probability density.

As with the PPI algorithm, this method requires the following choices:

1. The quadrature rule for computing $\int w(s)R(s,a|\theta)\,ds$.

2. The sample points $\{s_1,\ldots,s_M\}$ at which the residual function is evaluated.

3. The set of basis functions $\{\rho_1,\ldots,\rho_K\}$.

### 2.5 Choices, choices, choices

From the above discussion it may seem as if there are only a couple of choices one must make when picking a solution technique: discrete or parametric approximation? if discrete: value function iteration or discrete policy function iteration? if parametric: parameterize the decision rule or parameterize the value function? But in fact these choices are just the start. There are a seemingly unlimited number of choices a researcher must make before finalizing any decision about solution methods. Furthermore, each of these choice can be made a la carte.

In Table 1 we outline the primary choices a research must make if s/he wishes to implement any of the posed algorithms. Of course the first decision is whether to use a discrete or parametric approach. If one decides to use a discrete approach, the first choice is what grid to use. As discussed in the introduction, naive discretization of multidimensional DP problems leads directly to the curse of dimensionality, since in a $d$-dimensional problem one can show that $O(1/\varepsilon)$ points in each dimension, or a total of $N = O(1/\varepsilon^d)$ grid points, are required in order to obtain an $\varepsilon$-approximation to the value or policy function. Since $N$ increases exponentially fast in the dimension $d$, it follows that naive discretization results in a curse of dimensionality.

Although naive discretization leads to a curse of dimensionality, there exist ways of discretizing the problem that avoid this curse. For example, Rust's (1997a) "random multigrid algorithm" breaks the curse of dimensionality by using a random discretization of the state space. This algorithm results in approximate solution to the DP problem with an expected error of $\varepsilon$ using only $N = O(1/\varepsilon^2)$ points. However, the regularity conditions for Rust's result require a Lipschitz-continuous transition probability for the state variables, and in some economic applications this condition will not be satisfied. In addition, Rust's result applies to DP problems where the control variable takes on only a finite number of possible values: we do not know whether Rust's result can be extended to problems where the control variables are continuous.

As can be seen from the Bellman equation (9) and the definition of the optimal policy (10), the policy improvement step requires the solution of a constrained optimization problem involving the conditional expectation of the value function. Since in general no analytic solutions to this conditional expectations will

| Methodological Choices for Solving Dynamic Economic Models | |
|---|---|
| DISCRETE APPROXIMATION | PARAMETRIC APPROXIMATION |
| GRIDS | FUNCTION TO PARAMETERIZE |
| uniform | value function |
| random | decision rule/expected value function |
| low discrepancy | BASIS |
| INTEGRATION | Chebyshev |
| Monte Carlo | poly-log |
| quadrature | neural networks |
| low discrepancy | FITTING TECHNIQUE |
| SMOOTHING | linear regression |
| bilinear interpolation/simplicial | non-linear regression |
| linear regression | INTEGRATION |
| probability weighted/Tauchen-Hussey | Monte Carlo |
| | quadrature |
| | low discrepancy |

Table 1: Methodological Choices

exist, we usually must resort to numerical integration. The most common approach to numerical integration is quadrature. The quadrature approach approximates the integral by a probability weighted sum:

$$\int V(s')p(s'|s,a)ds' = \frac{1}{N}\sum_{i=1}^{N}\hat{V}(s_i)p(s_i|s,a)$$

where the quadrature points and weights are selected in such a way that finite-order polynomials can be integrated exactly using quadrature formulae. The weights used have the natural interpretation of probabilities associated with intervals around the quadrature points. In this case, $p(s_i|s,a)$ is a discretized approximation to the transition probability density $p(s'|s,a)$.[5]

A second approximation method this integral is the "Monte Carlo" method given by

$$\int V(s)p(s'|s,a)ds' = \frac{1}{N}\sum_{i=1}^{N}\hat{V}(\tilde{s}_i) \tag{39}$$

where $\tilde{s}_i$ are draws from the density $p(s'|s,a)$ computed from uniformly distributed draws $\tilde{u}_i$ from the unit interval via the probability integral transform.

Instead of using pseudo-random random draws for $\{\tilde{u}_i\}$ one can obtain acceleration using *Generalized Faure sequences* (also known as *Tezuka sequences*). Using number theoretic methods (see, e.g. Neiderreiter 1992, or Tezuka 1995), one can prove that for certain classes of integrands, the convergence of Monte Carlo methods based on deterministic *low discrepancy sequences* is $O(\log(N)^d/N)$ (where $d$ is the dimension of the integrand and $N$ is the number of points), whereas traditional Monte Carlo methods converge at rate $O_p(1/\sqrt{N})$. These favorable rates of convergence have been observed in practice (see e.g. Papageorgiou and Traub 1996 and 1997).

It is critical to use numerical integration methods that provide accurate approximations of both the levels and the derivatives of the value function, since the latter determine the first order conditions for a constrained optimum for $a$. In regions where the value function is nearly flat in $a$, small inaccuracies in the estimated derivatives can create large instabilities in the estimated value of $a$. In our own experimentation with these methods, we have found these two methods can also be sensitive to the discretization of the $s$ and $a$ axes and the number of points used in the discretization. We find it useful to experiment with different integration methods, and different choices for grids

Interpolation is any method that construct a smooth function that satisfies a predetermined set of conditions. We use interpolation in one and two dimensions extensively in the numerical solutions to the test and real problems. When running DPI solutions and finite horizon problems we use linear and bilinear interpolation but we have experimented with other smoothing methods.

[5] For a detailed characterization of quadrature methods we refer the reader to Tauchen and Hussey (1991), Judd (1998), and Burnside (1999).

If a parametric method such as PPI or PEA is chosen, one must first choose which function to parameterize. It is often the case that any one function may be parametrized in multiple way either directly or indirectly. For example if one parameterizes the decision rule, often the conditional expectation function is implicitly parameterized.

Having chosen a function to parameterize one must chose which class of basis function should be used. These function are typically quite simple. Examples of commonly used basis function include polynomials or piece-wise linear functions. Judd (1992, 1998) and Christiano and Fisher (2000) advocate the use of Chebyshev polynomials as basis function. For several of the models we study the results from PPI and Euler-based projection methods are sensitive to basis chosen. For example in the simple consumption and saving problem described below, the value function is linear in the logarithms of $s$. If we parameterize the value function using a poly-log basis, PPI solves the model exactly. However if we parameterize the value function using Chebyshev polynomials, PPI provides an inaccurate solution.

Finally as with discrete methods, finding a method of integration that provides accurate approximation to both the levels and slopes of the function being integrated is critical to the success of any solution algorithm.

Within the two broad classes of solution methods, discrete or parametric, there are numerous variations of techniques. In this paper we do not attempt to survey every possible combination of these choices. Instead the methods we employ we have made are based on our own experimentation and biases. The computer code we are making available is designed with numerous "switches" which allow user to experiment with different choices. We invite users to try different combinations of these methods, and let us know if we have overlooked a particular accurate and/or fast set of choices.

In the following section we describe four test problems and present their analytical solutions. In sections 4 to 11, we will apply the algorithms described above to these test problems as well as to three more complicated, "real world" problems. In each section we will compare and contrast the different solution methods in terms of speed and accuracy.

## 3 Test problems used in the numerical experiments

This section describes several "test problems" used in the numerical experiments in sections 4-7. These test problems have closed-form solutions which are extremely useful in enabling us to judge the accuracy of alternative algorithms. We will defer a description of the three "real" applications until they are introduced in sections 8, 9 and 10, respectively.

## 3.1 The Stochastic Growth Model with Leisure

We study the one-sector stochastic growth model as formulated by Santos (1999). In this model the representative agent wishes to:

$$\max_{l_t,c_t,k_{t+1}} E_0 \sum_{t=0}^{\infty} \beta^t \lambda \ln c_t + (1-\lambda) \ln l_t$$

subject to:

$$\begin{aligned} c_t + i_t &= z_t A k_t^{\alpha} l_t^{1-\alpha} \\ k_{t+1} &= (1-\delta) k_t + i_t. \end{aligned}$$

We assume $z_t$ evolves according to an exogenous first order Markov process with a transition density $g(z', z)$. In particular:

$$\ln z_{t+1} = \rho \ln z_t + \varepsilon_{t+1} \quad \varepsilon \sim N(0, \sigma^2).$$

There is a single consumption good which is produced each period via a Cobb-Douglas production technology. At each date $t$, a single agent starts the period with a given level of capital $k_t$ and learns the value of $\varepsilon_t$. The agent then decides how much labor,$1 - l_t$, to provide, how much of the good to consume, $c_t$ and how much of the good to save $k_{t+1}$. We assume $\varepsilon_t$ is an identically distributed random variable with normal distribution with mean zero and standard deviation $\sigma$. We assume leisure, $l_t$ is bounded by (0,1); and we assume the depreciation rate on capital, $\delta$, and discount factor $\beta$ are between zero and one.

The Bellman equation can be written as:

$$V(k,z) = \max_{l,c,k'} \left\{ \lambda \ln c + (1-\lambda) \ln l + E\beta V(k', z') \right\}$$

subject to:

$$c + k' = zAk^{\alpha} l^{1-\alpha} + (1-\delta)k.$$

We derive decision rules for both $c$ and $k'$ as functions of $l$. From the first-order conditions of the problem we get

$$c = \frac{\lambda z A k^{\alpha} l (1-\alpha)}{(1-\lambda)(1-l)^{\alpha}}, \tag{40}$$

and

$$k' = zAk^{\alpha}(1-l)^{1-\alpha} + (1-\delta)k - \frac{\lambda}{1-\lambda} l (1-\alpha) z A k^{\alpha} (1-l)^{-\alpha} \tag{41}$$

So the problem reduces to a unidimensional choice problem in $l$.

In the special case of $\delta = 1$, it is well known that an analytical solution to the Bellman equation exists and takes the form:

$$V(k,z) = F + G \ln k + H \ln z$$

where

$$\begin{aligned} G &= \frac{\lambda\alpha}{1-\alpha\beta}, \text{ and} \\ H &= \frac{\lambda}{(1-\rho\beta)(1-\alpha\beta)}. \end{aligned}$$

The decision rules involve working a constant fraction of the time endowment regardless of the state:

$$l = \frac{(1-\lambda)(1-\alpha\beta)}{\lambda(1-\alpha)+(1-\lambda)(1-\alpha\beta)},$$

and consuming a constant fraction of current output:

$$c = (1-\alpha\beta) z A k^{\alpha} (1-l)^{1-\alpha},$$
$$k' = \alpha\beta z A k^{\alpha} (1-l)^{1-\alpha}.$$

We also solve this model for the case without leisure, $l_t = 1\ \forall t$. In the case, the model is similar to the one studied by Taylor and Uhlig (1990) and the accompanying papers; in those papers, $delta = 0$.

### 3.2 The Consumption/Saving Problem

We now consider the problem of optimal consumption and saving first analyzed by Phelps (1962). The state variable $s$ denotes a consumer's current wealth, and the decision $d$ is how much to consume in the current period. Since consumption is a continuous decision, we will use $c_t$ rather than $d_t$ to denote the values of the control variable, and let $w_t$ to denote the state variable wealth.

The consumer is allowed to save, but is not allowed to borrow against future income. Thus, the constraint set is $D(w) = \{c \mid 0 \le c \le w\}$. The consumer can invest his savings in a single risky asset with random rate of return $\tilde{R}$ which is *IID* with distribution $F$. Thus, $p(dw_{t+1} \mid w_t, c_t) = F\left((dw_{t+1}/(w_t - c_t)\right)$. Let the consumer's utility function be given by $u(w,c) = \ln(c)$. Then Bellman's equation for this problem is given by:

$$V^*(w) = \max_{0 \le c \le w} \left[\ln(c) + \beta \int_0^{\infty} V^*\big(R(w-c)\big) F(dR)\right]. \tag{42}$$

As in the previous example, $V_t$ has the form, $V = F + G\ln(w)$ for constants $A_t$ and $B_t$. Thus, it is reasonable to conjecture that this form holds in the limit as well. Inserting the conjectured functional form $V^*(w) = A_\infty \ln(w) + B_\infty$ into (42) and solving for the unknown coefficients $A_\infty$ and $B_\infty$ we find:

$$\begin{aligned} A_\infty &= 1/(1-\beta) \\ B_\infty &= \ln(1-\beta)/(1-\beta) + \beta\ln(\beta)/(1-\beta)^2 + \beta E\{\ln(\tilde{R})\}/(1-\beta)^2, \end{aligned} \tag{43}$$

and the optimal decision rule or *consumption function* is given by:

$$\alpha(w) = (1-\beta)w, \tag{44}$$

as shown in Phelps (1962) and Hakansson (1970). Thus, the logarithmic specification implies that a strong form of the *permanent income hypothesis* holds in which optimal consumption is independent of the distribution $F$ of investment returns.

Section 5 shows the closed form solutions using other utility functions. It also shows the closed form solutions of the finite horizon case with the different utility functions and compares all these results with those of numerical solutions of the problems.

### 3.3 The Optimal Replacement Problem

Sometimes one can derive a differential equation for $V$ and in certain cases one can derive analytical solutions to this differential equation and use it to characterize the optimal decision rule. Consider, for example, the problem of optimal replacement of durable assets analyzed in Rust (1985, 1986, and 1987). In this case the state space $S = R_+$, where $s_t$ is interpreted as a measure of the accumulated utilization of the durable (such as the odometer reading on a car). Thus $s_t = 0$ denotes a brand new durable good. At each time $t$ there are two possible decisions {keep,replace} corresponding to the binary constraint set $D(s) = \{0,1\}$ where $d_t = 1$ corresponds to selling the existing durable for scrap price $\underline{P}$ and replacing it with a new durable at cost $\overline{P}$. Suppose the level of utilitization of the asset each period has an exogenous exponential distribution. This corresponds to a transition probability $p$ given by:

$$p(ds_{t+1}|s_t,d_t) = \begin{cases} 1-\exp\{-\lambda(ds_{t+1}-s_t)\} & \text{if } d_t = 0 \text{ and } s_{t+1} \geq s_t \\ 1-\exp\{-\lambda(ds_{t+1}-0)\} & \text{if } d_t = 1 \text{ and } s_{t+1} \geq 0 \\ 0 & \textit{otherwise.} \end{cases} \tag{45}$$

Assume the per-period cost of operating the asset in state $s$ is given by a function $c(s)$ and that the objective is to find an optimal replacement policy to minimize the expected discounted costs of owning the durable over an infinite horizon. Since minimizing a function is equivalent to maximizing its negative, we can define the utility function by:

$$u(s_t,d_t) = \begin{cases} -c(s_t) & \text{if } d_t = 0 \\ -\left[\overline{P} - \underline{P}\right] - c(0) & \text{if } d_t = 1. \end{cases} \tag{46}$$

Bellman's equation takes the form:

$$\begin{aligned} V^*(s) = \max\Big[ \quad & -c(s) + \beta \int_s^\infty V^*(s')\lambda \exp\{-\lambda(s'-s)\}ds', \\ & -[\overline{P} - \underline{P}] - c(0) + \beta \int_0^\infty V^*(s')\lambda \exp\{-\lambda(s')\}ds'\Big]. \end{aligned} \tag{47}$$

Observe that $V^*$ is a non-increasing, continuous function of $s$ and that the second term on the right hand side of (47), the value of replacing the durable, is a constant independent of $s$. Note also that $\overline{P} > \underline{P}$ implies that it is never optimal to replace a brand-new durable $s = 0$. Let $\gamma$ be the smallest value of $s$ such that the agent is indifferent between keeping and replacing. Differentiating Bellman's equation (47), it follows that on the *continuation region*, $[0,\gamma)$, $V^*$ satisfies the differential equation:

$$V^{*'}(s) = -c'(s) + \lambda c(s) + \lambda(1-\beta)V^*(s). \tag{48}$$

This is known as a *free boundary value problem* since the boundary condition:

$$V^*(\gamma) = [\overline{P} - \underline{P}] + V^*(0) = -c(\gamma) + \beta V^*(\gamma) = \frac{-c(\gamma)}{1-\beta}, \tag{49}$$

is determined endogenously. Equation (48) is a linear first order differential equation which can be integrated to yield the following closed-form solution for $V^*$:

$$V^*(s) = \max\Big[\frac{-c(\gamma)}{1-\beta}, \frac{-c(\gamma)}{1-\beta} + \int_s^\gamma \frac{c'(y)}{1-\beta}[1 - \beta e^{-\lambda(1-\beta)(y-s)}]dy\Big], \tag{50}$$

where $\gamma$ is the unique solution to:

$$[\overline{P} - \underline{P}] = \int_0^\gamma \frac{c'(y)}{1-\beta}[1 - \beta e^{-\lambda(1-\beta)y}]dy. \tag{51}$$

It follows that the optimal decision rule is given by:

$$\alpha^*(s) = \begin{cases} 0 & \text{if } s \in [0,\gamma] \\ 1 & \text{if } s > \gamma. \end{cases} \tag{52}$$

### 3.4 Linear-Quadratic Control Problems

Consider the following linear-quadratic-gaussian (LQG) control problem whose solution is given in the following theorem.

**Theorem:** *Let $S = R$ and $A(s) = S$, $\forall s \in S$. Consider a DP with the following utility function and transition density:*

$$u(s,a) = [\lambda_2 a^2 + \lambda_1 a + \lambda_0] + s[\rho_0 + \rho_1 a] + \mu s^2 \tag{53}$$

$$\begin{aligned} p(s'|s,a) &= \frac{1}{\sqrt{2\pi}\sigma} \exp\left\{-(s' - \kappa_0 - \kappa_1 a - \kappa_2 s)^2/(2\sigma^2)\right\} \\ \sigma &= [\eta_0 + \eta_1 a + \eta_2 s] \end{aligned} \tag{54}$$

where

$$\mu < 0, \quad \lambda_2 < 0, \quad \textit{and} \quad \rho_1^2 - 4\mu\lambda_2 < 0. \tag{55}$$

*Then $V(s)$ is given by:*

$$\begin{aligned} V(s) &= \max_{a \in A(s)} \left[ u(s,a) + \beta \int V(s') p(s'|s,a) ds' \right] \\ &= \gamma_0 + \gamma_1 s + \gamma_2 s^2 \end{aligned} \tag{56}$$

*and the optimal decision rule $\alpha(s)$ is given by:*

$$\alpha(s) = f_0 + f_1 s \tag{57}$$

*where:*

$$\begin{aligned} f_1 &= -\frac{\rho_1 + 2\beta\gamma_2(\kappa_1\kappa_2 + \eta_1\eta_2)}{2\left[\lambda_2 + \beta\gamma_2(\eta_1^2 + \kappa_1^2)\right]} \\ f_0 &= -\frac{\lambda_1 + \beta\gamma_1\kappa_1 + 2\beta\gamma_2(\eta_0\eta_1 + \kappa_0\kappa_1)}{2\left[\lambda_2 + \beta\gamma_2(\eta_1^2 + \kappa_1^2)\right]} \end{aligned} \tag{58}$$

*where:*

$$\begin{aligned} \gamma_2 &= \frac{-k_1 - \sqrt{k_1^2 - 4k_2 k_0}}{2k_2} \\ \gamma_1 &= \frac{\rho_0 + 2\beta\gamma_2[\kappa_0\kappa_2 + \eta_0\eta_2] + f_1\left[\lambda_1 + 2\beta\gamma_2(\eta_0\eta_1 + \kappa_0\kappa_1)\right]}{1 - \beta(\kappa_1 f_1 + \kappa_2)} \\ \gamma_0 &= \frac{\lambda_2 f_0^2 + \lambda_1 f_0 + \lambda_0 + \beta\gamma_1(\kappa_0 + \kappa_1 f_0) + \beta\gamma_2\left[(\eta_0 + \eta_1 f_0)^2 + (\kappa_0 + \kappa_1 f_0)^2\right]}{(1 - \beta)} \end{aligned} \tag{59}$$

*where:*

$$\begin{aligned} k_0 &= \rho_1^2 - 4\mu\lambda_2 \\ k_1 &= 4\left[\lambda_2[1-\beta(\kappa_2^2+\eta_2^2)] - \mu\beta(\eta_1^2+\kappa_1^2) + \rho_1\beta(\kappa_1\kappa_2+\eta_1\eta_2)\right] \\ k_2 &= 4\left[\beta(\kappa_1^2+\eta_1^2)[1-\beta(\kappa_2^2+\eta_2^2)] + \beta^2(\kappa_1\kappa_2+\eta_1\eta_2)^2\right]. \end{aligned} \tag{60}$$

If $\eta_1$ and $\eta_2$ are set to zero, the DP given in equation (53) can be formulated as an optimal linear regular problem (OLRP) and solved recursively.[6] In particular, this DP can rewritten as:

$$\max E_0 \sum_{t=0}^{\infty} \beta^t \{x_t' R x_t + u_t' Q u_t + 2u_t' W x_t\} \tag{61}$$

subject to:

$$x_{t+1} = Ax_t + Bu_t + \varepsilon_{t+1} \tag{62}$$

where the matrices $R$ and $Q$ are symmetric, negative definite matrices. Setting $\eta_1$ is set to zero, we assume $\varepsilon_{t+1}$ is a $2 \times 1$ vector of random variables that is independently and identically distributed through time with a mean vector zero and a covariance matrix:

$$E\varepsilon_{t+1}\varepsilon_{t+1}' = \Sigma.$$

For the problem at hand, we set $x_t = [1 \;\; s_t]'$, $u_t = a_t$,

$$R = \begin{bmatrix} \lambda_0 & \frac{1}{2}\rho_0 \\ \frac{1}{2}\rho_0 & \mu \end{bmatrix}, \quad Q = [\lambda_2], W = \frac{1}{2}[\lambda_1 \quad \rho_1],$$

$$A = \begin{bmatrix} 1 & 0 \\ \kappa_0 & b \end{bmatrix}, B = \begin{bmatrix} 0 \\ \kappa_1 \end{bmatrix}, \text{ and } \Sigma = \begin{bmatrix} 0 & 0 \\ 0 & \eta_0 \end{bmatrix}.$$

Bertsekas (1995) and Hansen and Sargent (1999) show the value function can then written as $V(x) = x_t' P x_t + d$ where $P$ solves the algebraic matrix Riccati equation:

$$P = R + \beta A'PA - (\beta A'PB + W')(Q + \beta B'PB)^{-1}(\beta B'PA + W)$$

and the optimal decision rule is:

$$u_t = -(Q + \beta B'PB)^{-1}(\beta A'PB + W)x_t \quad \text{or } u_t = -Fx_t.$$

[6] See Bertsekas (1995) section 4.1, and Hansen and Sargent (1999).

The scalar $d$ is given by $-\frac{\beta}{(1-\beta)} trace P\Sigma$.

For the example presented here, iterating on the Riccati equation yields the analytical solution to the value function and the decision rules presented in the theorem above. Indeed if one writes out the matrices $P$ and $F$ and solves the fixed point problem via brute force (as done in equations 58 -60), one can show:

$$P = \begin{bmatrix} \gamma_0 + d & \frac{1}{2}\gamma_1 \\ \frac{1}{2}\gamma_1 & \gamma_2 \end{bmatrix}, \quad F = \begin{bmatrix} -f_0 & -f1 \end{bmatrix}.$$

# 4 The Stochastic Growth Model

In this section we solve the stochastic growth model (SGM) presented in section 3.1 using Euler-based projection methods, discrete policy iteration, and parametric policy iteration.

We first solve the stochastic growth with leisure as studied by Santos (1999) by parameterizing the decision rule for leisure (thus parameterizing the conditional expectation indirectly) and finding coefficients that satisfy the Euler equation in an average sense. We then solve the stochastic growth model without leisure. When solving this model, we follow Christiano and Fisher's (2000) Galerkin-Chebyshev PEA discussed in section 4.2.3 of their paper. The model we study differs from the one studied by Christiano and Fisher in two ways: in our case, investment is fully reversible and the shock is continuous. For the case without leisure, the model is almost identical the one studied by the Taylor and Uhlig (1990) symposium; the only exception is we allow the depreciation rate on capital to be non-zero. In section 4.1 and 4.2, we draw heavily on the work of Judd (1992, 1994 and 1998) and Christiano and Fisher (2000).

## 4.1 Solving the SGM with leisure

Consider the model presented in section 3.1:

$$\max_{l_t, c_t, k_{t+1}} E_0 \sum_{t=0}^{\infty} \beta^t \lambda \ln c_t + (1-\lambda) \ln l_t$$

subject to:

$$\begin{aligned} c_t + i_t &= z_t A k_t^{\alpha} l_t^{1-\alpha} \\ k_{t+1} &= (1-\delta) k_t + i_t \\ \ln z_{t+1} &= \rho \ln z_t + \varepsilon_{t+1} \quad \varepsilon \sim N(0, \sigma). \end{aligned}$$

The Euler equation is:

$$\frac{\lambda}{z_t A k_t^{\alpha}(1-l_t)^{\alpha} + (1-\delta)k_t - k_{t+1}} = \lambda\beta E_t \left( \frac{\alpha z_{t+1} A k_{t+1}^{\alpha-1}(1-l_{t+1})^{1-\alpha} + 1 - \delta}{z_{t+1} A k_{t+1}^{\alpha}(1-l_{t+1})^{1-\alpha} + (1-\delta)k_{t+1} - k_{t+2}} \right). \quad (63)$$

As discussed above, we can write both $c_t$ and $k_{t+1}$ as functions of $l_t$:

$$c_t = \frac{\lambda z_t A k^{\alpha} l_t (1-\alpha)}{(1-\lambda)(1-l_t)^{\alpha}}, \tag{64}$$

and

$$k_{t+1} = z_t A k^{\alpha}(1-l_t)^{1-\alpha} + (1-\delta)k_t - \frac{\lambda}{1-\lambda} l_t (1-\alpha) z_t A k_t^{\alpha}(1-l_t)^{-\alpha}. \tag{65}$$

So the problem reduces to a unidimensional choice problem.

Now we have to make a decision about which function to parameterize within the Euler equation. As with any weighted residual method, we assume $f(k,z)$ is a finite linear combination of known basis functions. In this case we use Chebyshev polynomials as the basis functions. Chebyshev polynomials are defined on $[-1,1]$ and the $i^{th}$ polynomial is given by $T^i = \cos(i(\arccos(x)))$. Since the domain $(k,z)$ is not given by $[-1,1]$, let $\phi(x) = 2(x-a)/(b-a)-1$ where $a$ and $b$ denotes the lower and upper bounds of the variable.

If we parameterize the conditional expectation function by parameterizing the marginal utility of consumption (i.e. the left-hand side of equation 63), we must use a non-linear equation solver to back out the decision rule for leisure (and thus the decision rules for consumption and next period's capital stock). Therefore we parameterize the conditional expectation function indirectly by parameterizing the leisure function:

$$l_t \approx l(k,z) \equiv \frac{1}{1+\exp(\theta' \mathcal{T}(\phi(k),\phi(z)))}.$$

We let $\beta \exp(f(k,z)) = \theta' \mathcal{T}(\phi(k),\phi(z))$ where $\theta$ is a vector of $N \times 1$ vector of polynomials, $\mathcal{T}$ is a $N \times 1$ vector of complete degree $j$ Chebyshev polynomials in 2 variables. This parameterization forces $l$ to take values between 0 and 1.

Thus we can define the residual function as:

$$R(k,z) \equiv \ln\left(\frac{1}{zAk^{\alpha}(1-l(k,z|\theta))^{1-\alpha} + (1-\delta)k - k'(l(k,z|\theta))}\right) -$$
$$\ln\beta \int \frac{\alpha z' A k'(l(k,z|\theta))^{\alpha-1}(1-l(k,z|\theta))^{1-\alpha} + 1 - \delta}{z' A k'(l(k,z|\theta))^{\alpha}(1-l(k',z'|\theta))^{1-\alpha} + (1-\delta)k'(l(k,z|\theta)) - k''(l(k'(l(k,z|\theta)),z'|\theta))} g(z'|z)dz' \tag{66}$$

where, using equation (65), $k'$ is written as a function of $l(k,z)$. As in the previous problem, we discretized the state space using with the Chebyshev zeros of $k$ and $z$, and approximated the integral in equation (66) with Gaussian quadrature.

Thus the strategy for solving the model involves finding a vector of parameters $\theta$ which set a weighted sum of $R(k,z)$ as close as possible to 0 for all $k$ and $z$. Mathematically this means choosing $\theta$ such that

$$\int\int w(k,z) R(k,z|\theta) \, dk \, dz = 0 \tag{67}$$

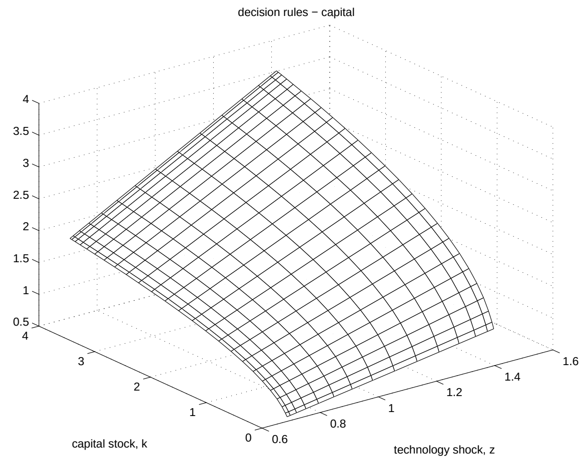


Figure 1: Stochastic growth model with leisure: the $k'(k,z|\theta)$ function

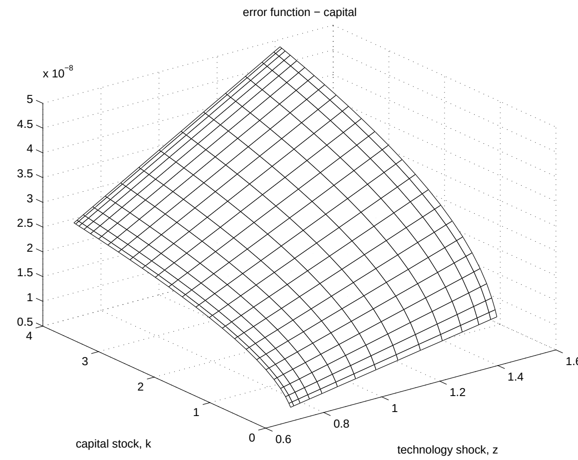


Figure 2: The difference between the parameterized decision rule, $k'(k,z|\theta)$ and the analytical solution.

where $w(k,z)$ is a weighting function.

The problem of finding a $\theta$ that solves (67) can be approximated by Gauss-Chebyshev quadrature. Let $X$ denote the $M \times N$ matrix of $N$ Chebyshev polynomials evaluated at each of the $M$ $(k,z)$ grid points. Hence (67) can be approximated by

$$X'R(k,z|\theta) = 0. \tag{68}$$

We evaluated (68) at the Chebyshev zeros of $k$ and $z$. Let $M_k$ and $M_z$ denote the number of grid points choose for $k$ and $z$ respectively. Thus $M = M_k \times M_z$. If $M > N$ this method is often called a Galerkin method. If $M = N$ (so $X$ is a square), this method is referred to as a collocation method. We set the following parameter values: $\alpha = 0.34$, $A = 10$, $\beta = 0.95$, $\delta = 1$, $\lambda = 1/3$, $\rho = 0.90$, and $\sigma = 0.008$. For the grid of the log of the technology shock $\ln z$ we choose the 15 $(M_z)$ Chebyshev zeros between -0.35 and 0.35. For the capital stock grid $k$ I chose the 25 Chebyshev zeros $(M_k)$ and bound the grid between .1 and 4.0.

This method seems to work well only if one starts off with good initial guess for $\theta$. We parameterized the decision rule for leisure as

$$\begin{aligned}
\ln\left(\frac{1}{l(k,z)} - 1\right) &= \theta'\mathcal{T}(\phi(k),\phi(z)) \\
&= \begin{bmatrix} \theta_1 & \theta_2 & \theta_3 \end{bmatrix} \begin{bmatrix} 1 \\ T^1(\phi(k)) \\ T^1(\phi(z)) \end{bmatrix}.
\end{aligned}$$

Since the analytical solution for leisure is constant, the algorithm should set $\theta_2 = \theta_3 = 0$. Indeed any polynomial approximation of the leisure decision rule (including a Chebyshev polynomial) should nail the solution exactly.

If we initialize θ at [ -1; 0 ; 0 ], this algorithm leads to the correct solution: [ -.72; 0; 0 ]. The Matlab program converges in 132 seconds on a 266 Mhz computer. Figures 1 and 2 display the numerical decision rule for capital and the difference between the numerical decision rule for capital and the exact analytical solution. The difference between the numerical and analytical solutions are tiny. However if we initialize θ at a starting value away from the correct value, (e.g. θ = [ -1; .5; -.5 ]) the non-linear equation solver (which uses a least squares method) fails to find the correct solution.

### 4.2 Solving the SGM without leisure

Consider a special case of the stochastic growth model described above with $\lambda = 1$. For these special case to make sense, we want to assume $(1-\lambda)\ln 0 = 0$ when $\lambda = 1$. The Euler equation becomes:

$$\frac{1}{z_t A k_t^\alpha + (1-\delta)k_t - k_{t+1}} - \beta E_t\left(\frac{\alpha z_{t+1} A k_{t+1}^{\alpha-1} + 1 - \delta}{z_{t+1} A k_{t+1}^\alpha + (1-\delta)k_{t+1} - k_{t+2}}\right) = 0$$

In this special case, we parameterize the marginal utility of consumption:

$$\frac{1}{zAk^\alpha + (1-\delta)k - k'(k,z)} \approx \beta \exp(f(k,z)).$$

Solving for $k'(k,z)$ yields:

$$k'(k,z) = \frac{\beta \exp(f(k,z))(zAk^\alpha + (1-\delta)k) - 1}{\beta \exp(f(k,z))}.$$

So implicitly we have parametrized the decision rule for next period's capital stock.

We let $\beta \exp(f(k,z)) = \theta' \mathcal{T}(\phi(k), \phi(z))$. As before, $\mathcal{T}$ is a $N \times 1$ vector Chebyshev polynomials. Having parametrized the decision rule for capital, we define the residual function $R(k,z|\theta)$:

$$R(k,z|\theta) \equiv f(k,z|\theta) - \ln\left(\int \frac{\alpha z' A k'(k,z|\theta)^{\alpha-1} + 1 - \delta}{z' A k'(k,z|\theta)^\alpha + (1-\delta)k'(k,z|\theta) - k'(k'(k,z|\theta), z'|\theta)}\, g(z'|z)dz'\right). \qquad (69)$$

Unlike the model with leisure there are (at least) three alternative strategies for finding a vector θ such this residual function is set to zero in an average sense.

1. **Direct Gauss-Chebyshev quadrature:** Simply find the vector θ which sets $X'R(k,z|\theta) = 0$. We employed this method for stochastic growth model with leisure. However in the case without leisure, as Christiano and Fisher (2000) point out, it is convenient to exploit the special structure of this problem. This leads to two other approaches.

2. **As an iterative linear regression problem:** To see this let

$$Y(k,z) \equiv \ln \int \left( \frac{\alpha z' A k'(k,z)^{\alpha-1} + 1 - \delta}{z' A k'(k,z)^{\alpha} + (1-\delta) k'(k,z) - k'(k'(k,z),z')} \right) g(z'|z) dz'.$$

   The function $X'R(k,z) = 0$ can then be rewritten as:

$$X'(X\theta - Y(k,z)) = 0.$$

   Premultipling both sides of this equation by $(X'X)^{-1}$ yields

$$\theta - (X'X)^{-1} X' Y(k,z) = 0 \quad (70)$$

$$\theta = (X'X)^{-1} X' Y(k,z). \quad (71)$$

   Note that since Chebyshev polynomials are orthogonal to each other $(X'X)$ is a diagonal matrix; so taking its inverse is trivial.

   So the following algorithm could be followed:

   (a) Guess an initial $N \times 1$ vector of $\theta$.

   (b) Compute $Y(k,z|\theta)$.

   (c) Regress $Y(k,z|\theta)$ on $X$ to obtain a new value of $\theta$.

   (d) Repeat steps (b) and (c) until convergence.

3. **As a simple non-linear equation problem:** Instead of finding the $\theta$ vector that solves $X'R(k,z) = 0$, one can find the $\theta$ vector that solves equation (70).

In practice, we solved this problem via method 3 with a non-linear equation solver. We parameterized $f(k,z)$ to be a complete polynomial of degree three in two variables, $k$ and $z$. Thus $N = 10$. We obtained a starting value for $\theta$ by arbitrarily setting the initial $\theta$ vector and then iterating via the "regression" strategy (method 2) a couple of times.

This method works well. We set the following parameter values: $\alpha = 0.34$, $A = 10$, $\beta = 0.95$, $\delta = 1$, $\rho = 0.90$, and $\sigma = 0.008$. For the grids of the log of the technology shock $\ln z$ we choose the 15 (thus $M_z$=15) Chebyshev zeros between -0.35 and 0.35. For the capital stock $k$ I set $M_k = 25$ and bound the grid between 1.1 times the steady-state value of the capital stock evaluated at $z = .35$ and .9 times the steady-state value of the capital evaluated at $z = -.35$. Again, the Chebyshev zeros were used. To evaluate the integral in (69) we used Gaussian quadrature at 100 nodes. Using Matlab on 266Mhz machine, the model took 457 seconds to solve. As discussed in section 3.1, since $\delta = 1$, there is an analytical solution to this model. In figure 3 we plot the parametrized decision rule for capital. In figure 4 we plot the difference between the parametrized decision rule and the analytical solution.

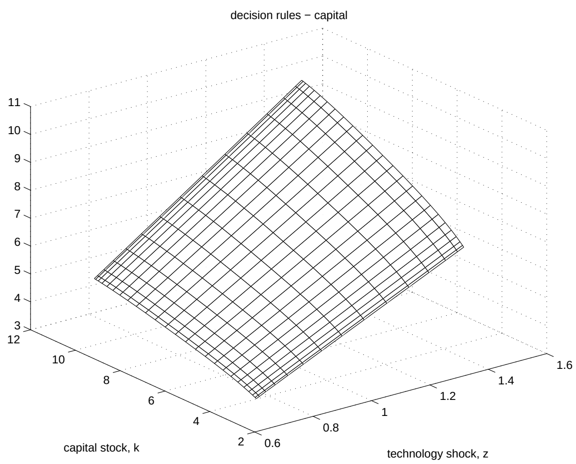


Figure 3: The $k'(k,z|\theta)$ function

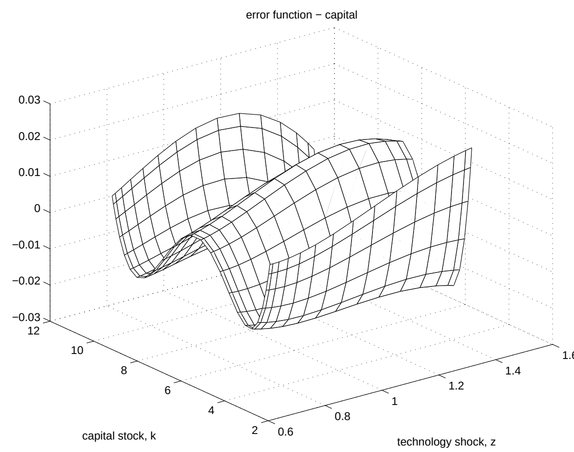


Figure 4: The difference between the parameterized decision rule, $k'(k,z|\theta)$ and the analytical solution.

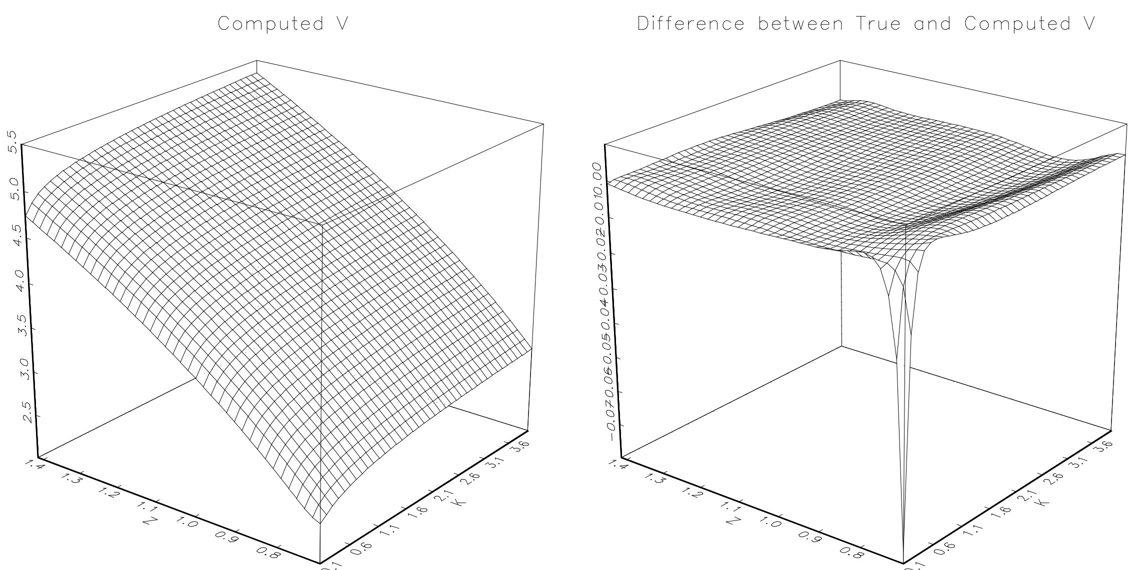


Figure 5: PPI Value Function

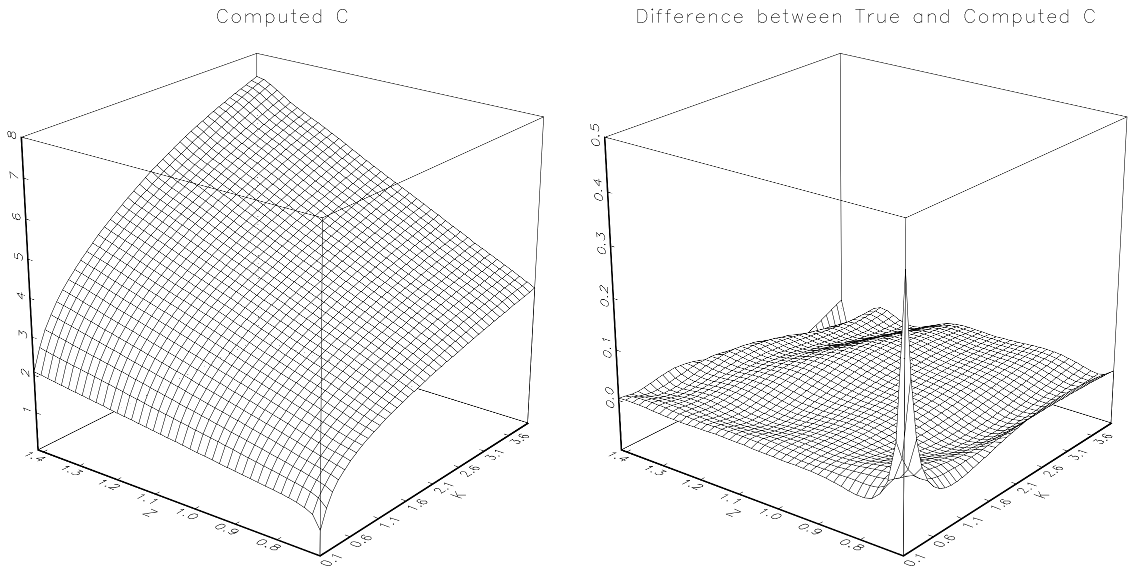


Figure 6: PPI Consumption Function

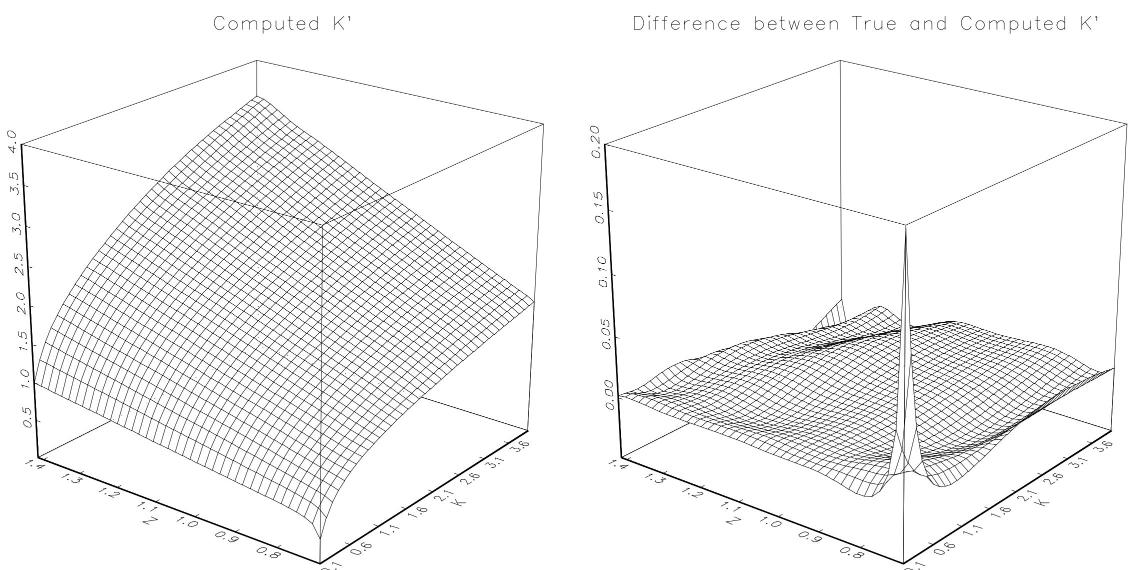


Figure 7: PPI Investment Function

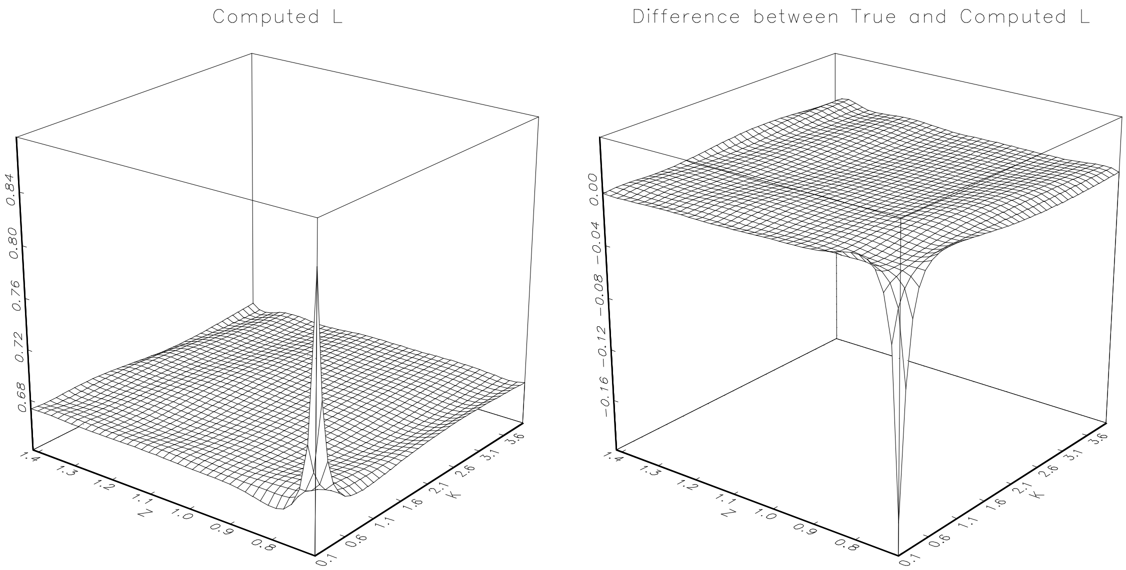


Figure 8: PPI: Optimal Leisure Function

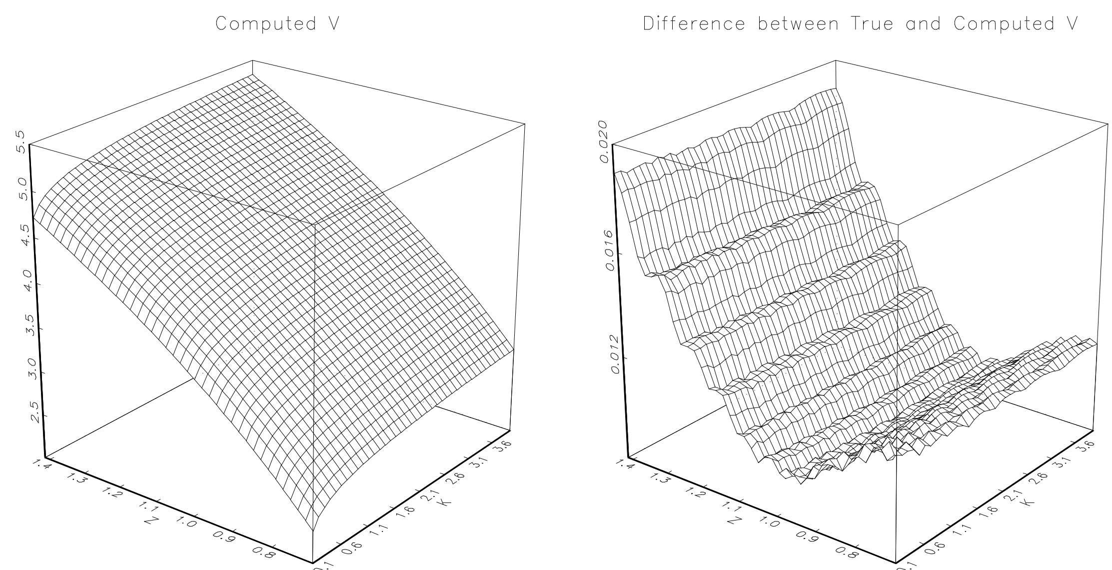


Figure 9: DPI Value Function

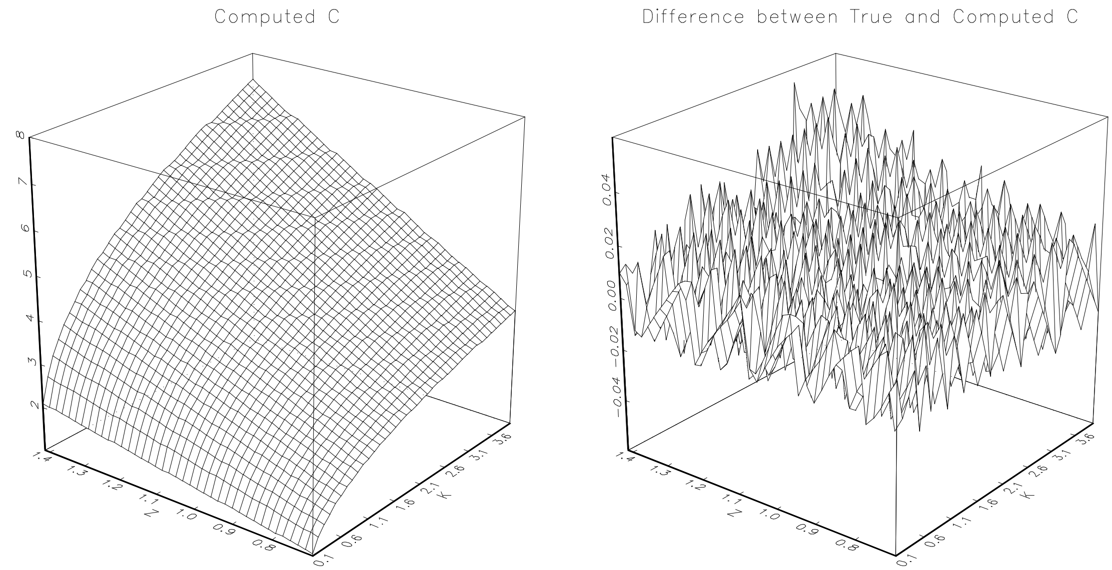


Figure 10: DPI Consumption Function

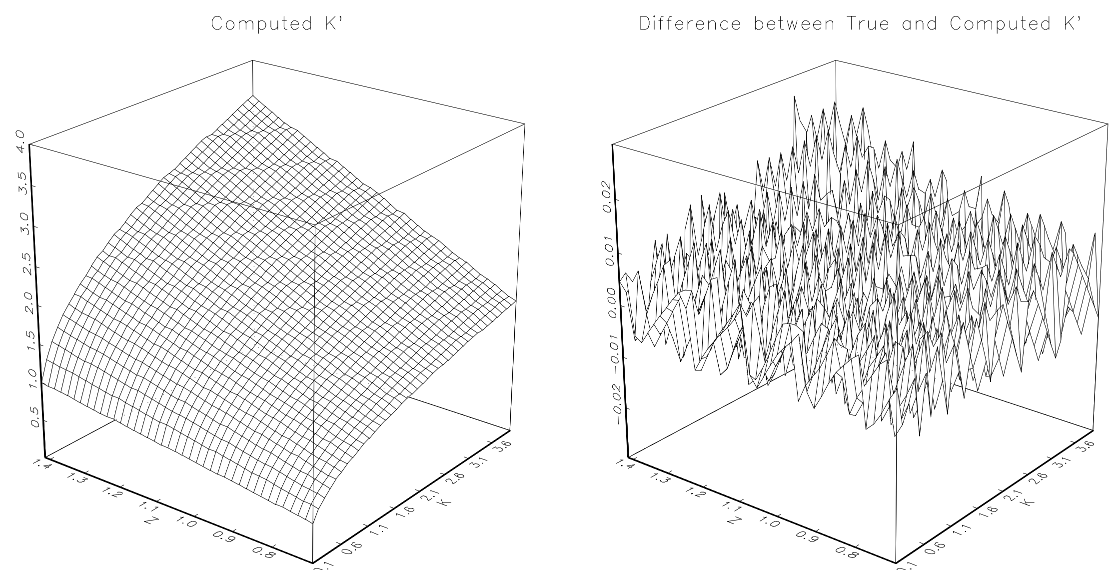


Figure 11: DPI Investment Function

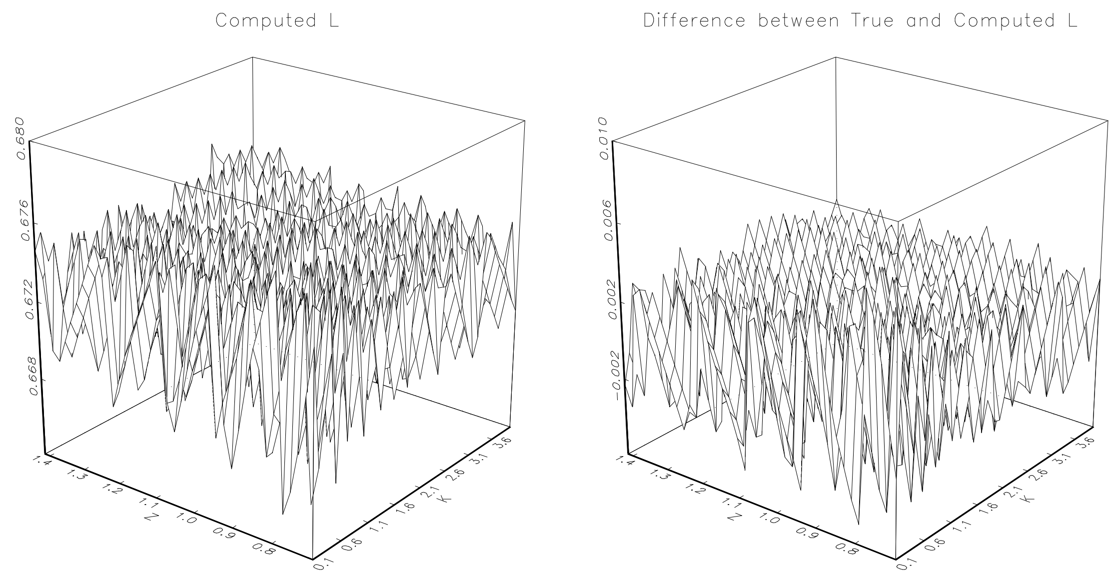


Figure 12: DPI: Optimal Leisure Function

# 5 The Consumption/Saving Model

In this section we show the closed form solutions to the classical consumption/saving problem presented in Section 3.2, using other utility functions. We also solve the problem in the finite horizon and compare all the solutions with those of our numerical computations.

## 5.1 Infinite Horizon: Closed Form Solutions

In solving this classical problem we do not need to restrict our attention to the logarithmic utility case. We can also consider the very same problem but assuming that the utility function is of the CRRA type. That is,

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma}, \tag{72}$$

where $\gamma > 0$ is the parameter of relative risk aversion. We can again find a closed form solution for the value function and the decision rule for this model, as shown in Phelps (1962), Levhari and Srinivasan (1969), and Hakansson (1970). They are obtained by the same procedure outlined in Section 3.2. We replicate these solutions below,

$$V(w) = \frac{\left(1-\beta^{\frac{1}{\gamma}}\left[E\left(R^{1-\gamma}\right)\right]^{\frac{1}{\gamma}}\right)^{-\gamma}}{1-\gamma}, \tag{73}$$

and

$$c(w) = \left[1 - \beta^{\frac{1}{\gamma}} \left[E\left(R^{1-\gamma}\right)\right]^{\frac{1}{\gamma}}\right] w \,. \tag{74}$$

The other interesting utility function we can use is the constant absolute risk aversion (CARA) utility. That is,

$$u(c) = -e^{-\gamma c}, \tag{75}$$

where $\gamma > 0$ is the parameter of absolute risk aversion. Unfortunately, we have not been able to find, so far, a closed form solution for the value function and the decision rule, using the same distributional assumptions for the interest rates as in the other cases.[7] We present here the derivation for the certainty case, and further below the numerical solution of the uncertainty case using log-normal returns to capital.

Using Hakansson's (1970) presentation of the problem we can conjecture that the value function has the following form:

$$V(w) = -N\, e^{-\lambda\, w}\,, \tag{76}$$

where $N$ and $\lambda$ are positive constants. Then we can write,

$$-N\, e^{-\lambda\, w} \;\; = \;\; \max_{0 \le c \le w} \left[-e^{\gamma c} - \beta N\, e^{-\lambda R(w-c)}\right]\,, \tag{77}$$

where $R = (1+r)$, with $r$ as the fixed interest rate on capital investments, and $\gamma > 0$ is the parameter of absolute risk aversion. Taking $f.o.c.$ we reach a solution for the decision rule in terms of the parameters of the problem,

$$c(w) \;\; = \;\; \frac{\lambda R w}{\lambda R + \gamma} - \frac{\ln\left(\frac{\beta N \lambda R}{\gamma}\right)}{\lambda R + \gamma}\,. \tag{78}$$

We can then calculate $w - c$ and rewrite the value function in such a way that we can start matching unknown coefficients,

$$-N\, e^{-\lambda w} \;\; = \;\; -\left[e^{-\gamma\left(\frac{\lambda R w}{\lambda R+\gamma}\right)} e^{\left(\frac{\gamma \ln\left(\frac{\beta N \lambda R}{\gamma}\right)}{\lambda R+\gamma}\right)}\right] - \left[\beta N e^{-\lambda R\left(\frac{\gamma w}{\lambda R+\gamma}\right)} e^{\left(\frac{-\lambda R \ln\left(\frac{\beta N \lambda R}{\gamma}\right)}{\lambda R+\gamma}\right)}\right]\,. \tag{79}$$

[7] Hakansson (1970) presents, in fact, closed form solutions using the CARA utility for a model that extends Phelps' (1962) by allowing borrowing, and appropriately treating non-labor income. We have not been able to use his results to find the solution to our problem given that the assumptions under which his results are valid seem unclear and are not discussed in his paper.

Then from this equation we can find the value of $\lambda$,

$$\lambda = \frac{\gamma(R-1)}{R} \,. \tag{80}$$

The value function can then be rewritten as,

$$V(w) = -N\, e^{\gamma(1-\frac{1}{R})w} \,. \tag{81}$$

Next we can find the solution for the remaining unknown coefficient, $N$,

$$N \;=\; e^{\left(\frac{\gamma \ln\left(\frac{\beta N \lambda R}{\gamma}\right)}{\lambda R+\gamma}\right)} + \beta N\, e^{\frac{\left(-\lambda R \ln\left(\frac{\beta N \lambda R}{\gamma}\right)\right)}{\lambda R+\gamma}} \,, \tag{82}$$

using the solution we obtained for $\lambda$ this expression simplifies to the following equation,

$$N\left(1-\beta\,[\beta N(R-1)]^{\frac{1-R}{R}}\right) \;=\; [\beta N(R-1)]^{\frac{1}{R}} \,. \tag{83}$$

A trivial and uninteresting solution for the this equation is $N = 0$. Assuming that $N > 0$ we can derive a closed form solution for $N$ and obtain the closed form solutions for the value function and consumption rule,

$$N \;=\; \left[[\beta(R-1)]^{\frac{1}{R}} + \beta\,[\beta(R-1)]^{\frac{1-R}{R}}\right]^{\frac{R}{R-1}} \,. \tag{84}$$

Then we can write,

$$V(w) \;=\; \left[[\beta(R-1)]^{\frac{1}{R}} + \beta\,[\beta(R-1)]^{\frac{1-R}{R}}\right]^{\frac{R}{R-1}} \; e^{-\gamma(1-\frac{1}{R})w} \,, \tag{85}$$

and

$$c(w) \;= \tfrac{R-1}{R} w - \frac{\ln\left(\beta(R-1)\left[[\beta(R-1)]^{\frac{1}{R}} + \beta[\beta(R-1)]^{\frac{1-R}{R}}\right]^{\frac{R}{R-1}}\right)}{\gamma R} \tag{86}$$

## 5.2 Infinite Horizon: Numerical vs. Closed Form Solutions

To solve these problems numerically we first use the policy iteration algorithm, and in particular Discrete Policy Iteration (DPI).

Figure 13 shows the difference between the true infinite horizon decision rule of the Phelps' problem with logarithmic utility, and the computed solution using a discrete uniform grid of 200 points and integration using probability weights and low discrepancy sequences.

It is worth mentioning how we approximate the conditional expectation operator via the Probability Integral Transform method:

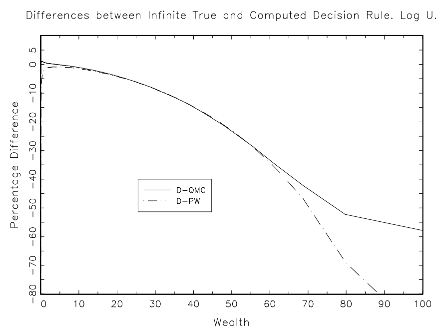


Figure 13: Infinite vs. Computed Decision Rule. Log U.

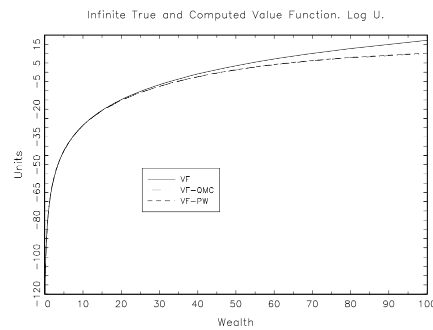


Figure 14: Infinite vs. Computed Value Function. Log U.

$$\hat{EV}(w,c) = \frac{1}{S}\sum_{s=1}^{S} V(F^{-1}(u_s)(w-c)) \tag{87}$$

where $F(r) = \int_{-\infty}^{r} f(x)dx$ and $\{u_1, \ldots, u_S\}$ are *IID* draws from $U(0,1)$, or alternatively, draws from a low discrepancy sequence such as a *Generalized Faure sequence*.

In this case the numerical methods do not perform too well. The differences are large in percentage terms and they are increasing in wealth. In part the differences are the product of the method of extrapolation that penalizes investments once you have a high level of wealth, what leads to overconsumption by the agents. We will see below that in the finite horizon case once we extrapolate linearly the opposite effect appears, that is, underconsumption, to take advantage of the better investing opportunities as wealth increases beyond the space of the grid. Figure 14 shows the value functions resulting from solving the problem numerically compared with the true solution. We can see that they are fairly similar except for higher levels of wealth when the numerical solutions consistently underpredicts the true solution.

We will also present the solutions of PPI and PEA methods for this problem, along with the discussion of the performance of the model using other utility functions.

### 5.3 Finite Horizon: Closed Form Solutions

In this subsection we solve a finite horizon version of the consumption/saving problem. Agents choose consumption according to the following utility maximizing framework:

$$\max_{0 \le c_s \le w} E_t \left[ \sum_{s=t}^{T} \beta^{s-t} u(c_s) \right] , \tag{88}$$

where $\beta$ is the discount factor, which includes the mortality probabilities, $c$ represents consumption, and $w$ is wealth at the beginning of the period. Savings accumulate at an uncertain interest rate of return $\tilde{R}$ such that $w_{t+1} = \tilde{R}(w_t - c_t)$, as in the infinite horizon case. Utility still depends only on consumption.

We can again solve this problem using Dynamic Programming and Bellman's principle of optimality. We solve it by backward induction starting in the last period of life, in which the individual solves

$$V_T(w) = \max_{0 \le c \le w} \ln(c) + K \ \ln(w-c) , \tag{89}$$

assuming a logarithmic utility function where $K \in (0,1)$ is the bequest factor.[8] By deriving the first order condition with respect to consumption we find that

$$c_T = \frac{w}{1+K} , \tag{90}$$

and from this we can write the analytical expression for the last period value function:

$$V_T(w) = \ln(\frac{w}{1+K}) + K \ \ln(\frac{wK}{1+K}) . \tag{91}$$

We can then iterate by backward induction and write the next to last period value function as:

$$V_{T-1}(w) = \max_{0 \le c \le w} \ln(c) + \beta E \ V_T(w-c) , \tag{92}$$

where the second term in the right hand side can be written as

$$E \ V_T(w-c) = \int_0^{Rmax} V_T(\tilde{R}(w-c)) f(\tilde{R}) \ d\tilde{R} , \tag{93}$$

where $\tilde{R}$ is the stochastic return on capital accumulation, and *Rmax* is the truncation point of the log-normal distribution of returns. Then we can write

$$V_{T-1}(w) = \max_{0 \le c \le w} \ln(c) + \beta E \ \ln(\tilde{R} \ (\frac{w-c}{1+K})) + \beta K \ E \ \ln(\tilde{R} \ (\frac{(w-c)K}{1+K})) . \tag{94}$$

[8] Agents in this model care only about the absolute size of their bequests, leading to its been called the "egoistic" model of bequests. A bequest factor of one would correspond to valuing bequest in the utility function as much as current consumption. The importance of bequest motives is still an open issue in the literature. Here we take the position of acknowledging that bequests do exist and explore the implications of changing the importance of the bequest motive in the utility function. Hurd (1987, 1989), Bernheim (1991), Modigliani (1988), Wilhem (1996) and Laitner and Juster (1996) are some of the main references on the debate over the significance of bequests and altruism in the life cycle model. Kotlikoff and Summers (1981) stress the importance of intergenerational transfers in aggregate capital accumulation.

Here the logarithmic utility simplifies the problem. Again taking first order conditions with respect to consumption, we obtain an expression for the consumption rule in the next to last period of life:

$$c_{T-1} = \frac{w}{1+\beta+\beta K} \,. \tag{95}$$

We then have an expression for $V_{T-1}$ in the following form:

$$V_{T-1}(w) = \ln(\frac{w}{1+\beta+\beta K}) + \beta \ln(\frac{w\beta}{1+\beta+\beta K}) + \beta K \ln(\frac{w\beta K}{1+\beta+\beta K}) + \Upsilon \,, \tag{96}$$

where $\Upsilon$ gathers all the terms that do not depend on $w$. From here we can write $V_{T-2}$ and again derive first order conditions, resulting in

$$c_{T-2} = \frac{w}{1+\beta+\beta^2+\beta^2 K} \,. \tag{97}$$

Through backward induction, we continue iterating to find $c_{T-k}$

$$c_{T-k} = \frac{w}{1+\beta+\beta^2+\beta^3+\ldots+\beta^k+\beta^k K} \,. \tag{98}$$

for any $k < T$. From these decision rules, we can observe that as T grows large, the finite horizon solution with bequests converges to the infinite horizon solution, already shown, since the influence of the bequest parameter becomes less important as the time horizon increases.

The derivation of the decision rules in the case of the CRRA utility function is similar though somewhat more involved. We show below only the optimal decision rule for the last period of life and the recursive formula to obtain the optimal consumption in all other periods, and present the full derivation in the Appendix. We now use the utility function specified in (72). In the last period of life agents consume

$$c_T = \frac{w}{1+K^{\frac{1}{\gamma}}} \,, \tag{99}$$

where $w$ is wealth at the beginning of that last period, and $K$ is the bequest factor. Then we can write the general closed form solution for the decision rule as

$$c_{T-k} = \frac{w}{1+\beta^{\frac{1}{\gamma}} E\,(\tilde{R}^{1-\gamma}) + \beta^{\frac{2}{\gamma}} E\,(\tilde{R}^{1-\gamma}) + \ldots + \beta^{\frac{k}{\gamma}} K^{\frac{1}{\gamma}} E\,(\tilde{R}^{1-\gamma})} \,, \tag{100}$$

where $\beta$ is the discount factor, and the interest rate, $\tilde{R}$, follows a log-normal distribution with mean $\mu$ and variance $\sigma^2$, then given that $E(\tilde{R}) = e^{\mu+\frac{\sigma^2}{2}}$ and denoting $E(\tilde{R})$ as $\overline{R}$ we can write

$$E\,(\tilde{R}^{1-\gamma}) = \overline{R}^{1-\gamma} e^{-\gamma(1-\gamma)\frac{\sigma^2}{2}} \,. \tag{101}$$

We can also see that if $\gamma$ is equal to 1 we are back to the logarithmic utility case. It is also important to emphasize that this expression is the finite horizon counterpart to the one obtained in Levhari and Srinivasan

(1969), and also replicated in the previous section, once a bequest motive is introduced, and that their results regarding the effects of uncertainty (decreasing proportion of wealth consumed as the uncertainty grows if $\gamma > 1$) go through in this case.

We next can assume a constant absolute risk aversion (CARA) utility function. Similarly to the infinite horizon case, we have not found a closed form solution for this problem under uncertain returns that follow a log-normal distribution as in the cases above. We therefore solve the finite horizon problem under certainty. The utility function used is the one presented in (75). We again present below only the optimal decision rule for the last period of life and the general solution for the rest of the periods, the full derivation is presented in the Appendix.

In the last period of life the optimal consumption rule is

$$c_T \quad = \quad \min\left(max\left(0, \frac{w}{2} - \frac{1}{2}\,\frac{\ln K}{\gamma}\right), w\right) \;, \tag{102}$$

where K is again the bequest factor.

We can then characterize the decision rule for any other period up to the first period of life

$$c_{T-k} \quad = \quad \min\left(\max\left(0, \frac{R^k w}{2+R+\ldots+R^k} - \left[\frac{2+R+\ldots+R^{k-1}}{2+R+\ldots+R^k}\right]\frac{\ln K_{T-k}}{\gamma}\right), w\right) \;, \tag{103}$$

where $R = 1+r$ and $r$ is the fixed rate of interest, and $K_{T-k}$ is shown below, and it is a function of some of the parameters of the problem and the previous constants, and where we can write the expression $2+R+R^2+\ldots+R^k$ as $1+\left[\frac{1-R^{k+1}}{1-R}\right]$ and similarly for the other series,

$$K_{T-k} \quad = \quad \frac{\beta\, R^k}{1+\left[\frac{1-R^k}{1-R}\right]}\left[e^{\frac{1+\left[\frac{1-R^{k-1}}{1-R}\right]}{1+\left[\frac{1-R^k}{1-R}\right]}\ln K_{T-k+1}} + \beta\, e^{-\frac{R^{k-1}\ln K_{T-k+1}}{1+\left[\frac{1-R^k}{1-R}\right]}}\left[\frac{K_{T-k+1}}{\frac{\beta R^{k-1}}{1+\left[\frac{1-R^{k-1}}{1-R}\right]}}\right]\right] \tag{104}$$

We can see from these solutions the different effect of risk given the utility functions, as the theory tells us. The $\gamma$ coefficient of risk aversion only has an absolute effect for the CARA utility, regardless of the wealth level. In the case of the CRRA utility the effect is relative to the level of resources.

### 5.4 Finite Horizon: Numerical vs. Closed Form Solutions

Our ability to derive an analytical solution for these finite horizon models allows us to evaluate the effectiveness of our numerical methods, which are all that we have available in more complicated models. The exercise of solving the model numerically is also interesting on its own given that the infinite horizon version of this model has been shown to be quite difficult to replicate using numerical methods, even with the logarithmic utility function, as discussed in the previous section and in Rust (1999).

The numerical procedure is by nature very similar to the analytical approach, involving backward recursion starting in the last period of life. We discretize wealth and compute the optimal value of consumption for all those wealth levels using bisection. Bisection is an iterative algorithm with all the components of a nonlinear equation solver. It makes a guess, computes the iterative value, checks if the value is an acceptable solution, and if not, iterates again. The stopping rule depends on the desired precision given that the solution is bracketed by the nature of the algorithm and that the round-off errors will probably not allow us to increase the precision beyond a certain limit. In each iteration of the numerical solution, except for the final one where all uncertainty has been eliminated, we have to compute the expectation in equation (93), which is potentially the most computationally demanding step. For this we use Gaussian-Legendre quadrature. We also compute the derivative of this expectation using numerical differentiation, also requiring quadrature as part of its routine. Here the analytical derivatives are simple to compute, but this is not always the case for more complicated models. We therefore wish to evaluate the accuracy of the numerical strategy. As explained in Section 2.5, Gaussian quadrature approximates the integral through sums using rules to choose points and weights based on the properties of orthogonal polynomials corresponding to the density function of the variable over which we are integrating, in this case the draws of the interest rates following a log-normal distribution.

At this point we are considering a one dimensional problem, for which quadrature methods have been shown to be very accurate compared with other techniques of computing expectations (integrals) such as Monte Carlo integration and weighted sums.[9]

This all amounts to manipulating (93) through a change of variables such that we can write it as an integral in the $(0,1)$ interval and then approximate it by a series of sums depending on the quadrature weights and quadrature abscissae which we compute recursively, following readily available routines (e.g. Press et al. 1992).[10]

An additional numerical technique that we use to solve the model completely is function approximation by interpolation. Since savings in a given period are accumulated at a stochastic interest rate, next period's wealth will not necessarily fall in one of the grid points for which we have the value of the function already calculated. Ideally we would solve the next period's problem for any wealth level, but this is computationally infeasible. Therefore, we use linear interpolation to find the corresponding value of the function given the values in the nearest grid points.[11]

---

[9] For an analysis of how these different techniques perform in other applied problems see Rust (1997b).

[10] We can write $\int_r V(r)\, fr\, dr$ after a change of variables as $\int_0^1 V(F^{-1}) du$, which can then be approximated by $\sum_{i=1}^N w_i V(F^{-1}(u_i))$, where $w_i$ are the quadrature weights and $u_i$ are the quadrature abscissae.

[11] More sophisticated interpolation procedures can be used such as splines or Chebyshev interpolation but they are not considered

The bisection algorithm that uses the quadrature and interpolation procedures eventually converges to a maximum of the lifetime consumption problem for a given value of wealth in a given period (or reaches the pre-decided tolerance level). This procedure is repeated until the solution of the first-period problem is obtained.

Once we have solved the model, we have a decision rule for every level of wealth in our initial grid. Here case we have chosen a grid space of 500 points; to gain accuracy more of these points are concentrated at low wealth levels where the function is changing rapidly. Figures 15-17 show the decision rule of the consumption/saving problem for wealth ranging from 0 to 100 units. For expositional purposes we have solved a 10-period model.

Figure 15 plots several decision rules given logarithmic utility. It first plots the numerical solutions for different time periods, denoted $C_1$, $C_2$, and so on. It also plots the solution of the infinite horizon problem, denoted by $CINF$ in the figure. We have chosen a discount factor of 0.95 and a bequest parameter of 0.6. Figure 16 plots the decision rule when we consider a CRRA, with risk aversion parameter equal to 1.5, $\beta = 0.95$, and bequest parameter equal to 0.6, we also plot the analytical solution of the infinite horizon problem. For both types of utility function we observe that the consumption rules increase with wealth and time and that in very few periods we are fairly close to the solution of the infinite problems. Figure 17 plots the consumption rules using a CARA utility function, with the same underlying parameters as the other functional forms. For every level of wealth, consumption is now lower than in the other two cases. We also plot the infinite horizon solution of the model with certainty we derived above. For comparison purposes notice that we have to compare the early periods of life with the infinite horizon solution since these periods are a better approximation to the infinite horizon because the agent is making a decision with more and more periods remaining in his or her lifetime.

Figures 18-20 are concerned with comparing the numerical solutions with the true analytical solutions derived above. We plot in all figures the percentage difference between the two solutions in terms of the value of the true solution, for a sample of time periods. The numerical technique performs quite well for the logarithmic and CRRA utility. For about half of the range of values, the numerical solution is very accurate with deviations below 1%, for both types of utility functions. After that, errors are a bit larger, especially for early time periods. For the first period and for high levels of wealth the error reaches 12% to 13%, depending on the utility function. The underconsumption resulting from our numerical technique seem to be an artifact of the linear extrapolation for values of wealth outside the chosen grid. The implied return from the extrapolation is higher than the one agents were facing before, so the normal reaction is to under-

for this problem.

Figure 15: Consumption Decision Rule. Log Utility

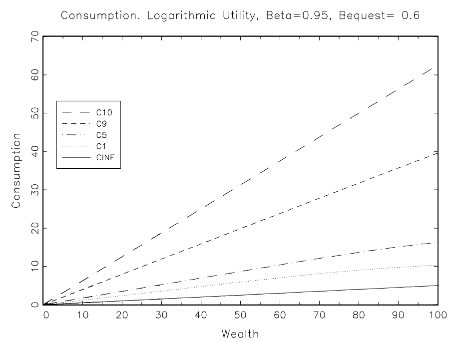


Figure 16: Consumption Decision Rule. CRRA Utility

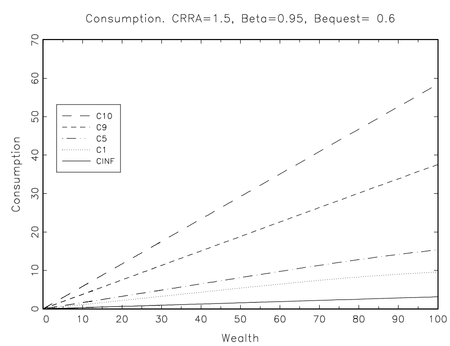


Figure 17: Consumption Decision Rule. CARA Utility

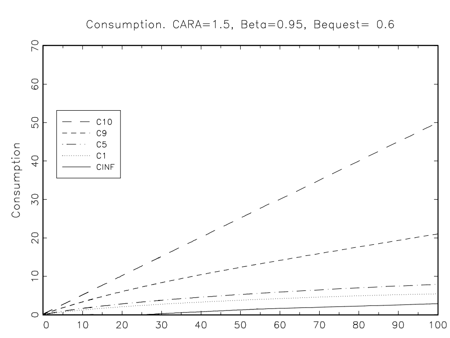

consume with respect to the true solution for levels of wealth approaching the upper bound value for wealth. For the CARA utility we already mentioned that we do not have a closed form solution for the model with uncertainty, but we did derive the certainty case. In Figure 20 we compare the closed form solutions under certainty with the numerical solutions under uncertainty, but with a distribution of returns with a very low variance. The numerical methods perform quite well again, and in this case the differences for most of the range of values do not seem to increase as we move towards the first period of life.

In Figure 21 we plot the decision rules of the limiting finite horizon of our numerical model with uncertainty vs. the closed form solution of the infinite horizon we derived before. For the finite horizon case we solve a 100 and a 200 period model. The limiting finite horizon seems to approximate the infinite horizon but even with 200 periods it stills delivers a significantly higher decision rule for consumption. Another issue to notice is that the true infinite horizon decision rule is zero for a significant portion of the wealth space and then is linear increasing in wealth, the computed limiting finite horizon is positive and higher than the infinite solution for all values of wealth.

In Figures 22 and 23 we simulate this model using the numerical solution for the CRRA utility function, to show the behavior implied by the decision rules shown above. We report the results of 5,000 simulations of an 11-period model with 500 grid points for wealth in the 0 to 200,000 range. We plot consumption and wealth paths with an initial wealth level of 10,000.[12] We also consider several values for the parameters of interest. In the first specification, $\gamma$ is taken to be 1.5 (the parameter of relative risk aversion), and it is increased to 2.5 in the second specification (*hg* lines in the plots). We then increase the bequest parameter to 0.6, leaving $\gamma = 1.5$ (*bq* lines in the plots), and finally, we decrease the relative risk aversion parameter to 0.7 (*lg* lines in the figures).

We observe that people consume less at the beginning of their lives, with increased consumption in the final periods of life, given uncertain interest rates represented by draws from a truncated log-normal distribution. Consumption does, however, decrease if the risk aversion parameter is less than 1. Focusing on the pattern of wealth accumulation, we observe that individuals deaccumulate their wealth gradually. We also see that increasing the relative risk aversion parameter has the effect of making consumption less smooth (with higher wealth accumulation), while decreasing the parameter from the benchmark value of 1.5 leads to more smoothing (with lower wealth accumulation). We can also observe the expected effect of the bequest parameter: those with a higher concern for their offspring, represented by a higher valuation of bequests in the utility function, consume uniformly less over the life cycle than do those with a lower

[12] This is approximately the net worth reported by Poterba (1998), using the Survey of Consumer Finances, for individuals at the beginning of their working lives.

Figure 18: Computed vs. True Decision Rule. Log Utility

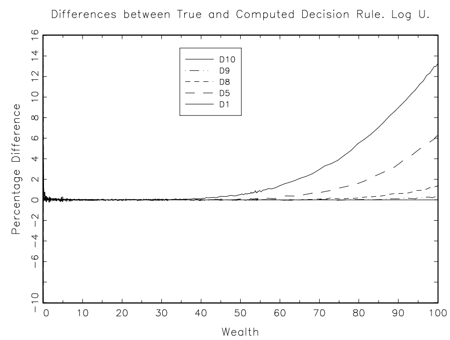


Figure 19: Computed vs. True Decision Rule. CRRA Utility

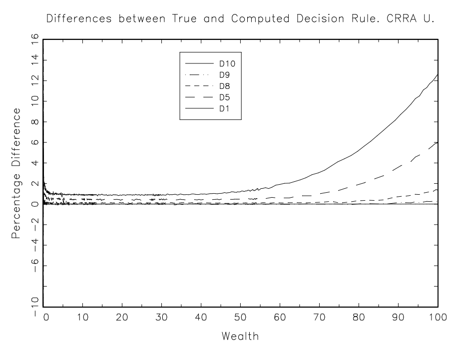


Figure 20: Computed vs. True Decision Rule. CARA Utility

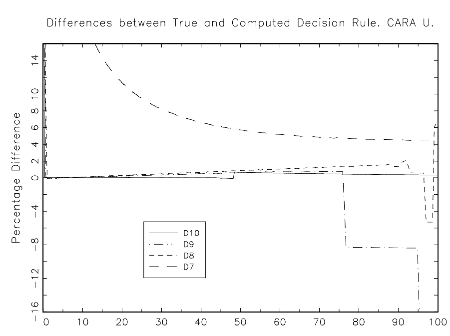

Figure 21: Limiting Finite vs. True Infinite. CARA U.

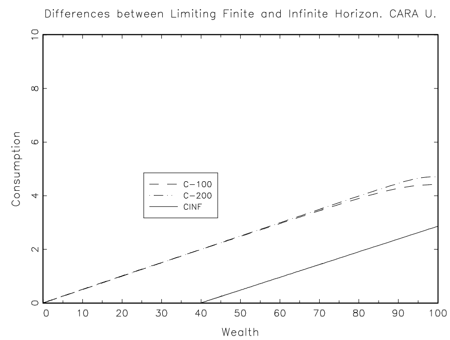


bequest parameter. This former population also accumulates more and for a longer period. These results regarding the effect of the bequest motives are consistent with, and in fact extend, the theoretical model of Hurd (1987) to the case of agents with various levels of bequest.

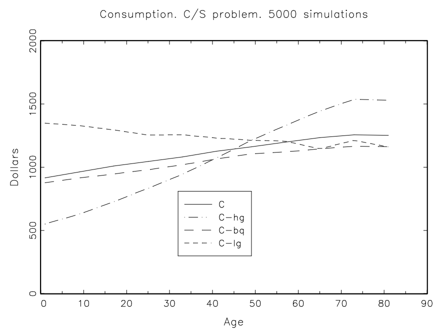


Figure 22: Simulated Consumption. CRRA Utility

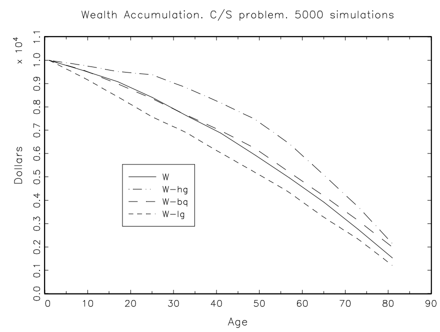


Figure 23: Simulated Wealth Accumulation. CRRA Utility

# 6 Optimal Replacement Problem

This section has yet to be included.

# 7 The Linear-Quadratic Model

This section has yet to be included.

# 8 An Inventory Investment Model

Consider a model of an intermediary studied in Hall and Rust (1999a, 1999b, and 2000) and Scarf (2000). The intermediary does not not undertake any physical production processing: its main function is to buy a durable good at spot prices, store it, and sell it subsequently at a markup.

We model the intermediary as making decisions about buying and selling a durable commodity in discrete time. The state variables for the firm are $(p_t, q_t)$ where $q_t$ denotes the inventory on hand at the start of day $t$, and $p_t$ denotes the per unit spot price at which the intermediary can purchase the commodity at day $t$. We assume $\{p_t\}$ evolves according to an exogenous Markov process with transition density $g(p_{t+1}|p_t)$. At the start of day $t$ the intermediary observes $(p_t, q_t)$ and places an order $q_t^o \geq 0$ for immediate delivery of the commodity at the current spot price $p_t$. The intermediary sets a uniform sales price to its customers, $p_t^s$, via an exogenously specified markup rule over the current spot price $p_t$:

$$p_t^s = f(p_t) = \alpha_0 + \alpha_1 p_t, \tag{105}$$

where $\alpha_0$ and $\alpha_1$ are positive constants.

After receiving $q_t^o$, the intermediary observes the quantity demanded of the commodity by the intermediary's customers, $q_t^d$. Let $H(q_t^d|p_t)$ denote the distribution of realized customer demand. We assume that $H$ has support on $[0,\infty)$ with a mass point at $q^d = 0$, reflecting the event that the intermediary receives no customer orders on a given day $t$. Let $h(q^d|p)$ be the conditional density of sales given that $q^d > 0$. This is a density with support on the interval $(0,\infty)$. Let $\eta(p) = H(0|p)$ be the probability that $q^d = 0$. Then we can write $H$ as follows:

$$H(q^d|p) = \eta(p) + (1-\eta(p)) \int_0^{q^d} h(q'|p)dq'. \tag{106}$$

We assume $h(q|p) > 0$ for all $q > 0$.

There are no delivery lags and unfilled orders are not backlogged. We assume that the intermediary meets the entire demand for its product in day $t$ subject to the constraint that it can not sell more that the quantity it has on hand, the sum of beginning period inventory $q_t$ and new orders $q_t^o$, $q_t + q_t^o$. Thus the intermediary's realized sales to customers in day $t$, $q_t^s$, is given by

$$q_t^s = \min\left[q_t + q_t^o, q_t^d\right]. \tag{107}$$

We assume the durable commodity is not subject to physical depreciation. Therefore the law of motion for start of period inventory holdings $\{q_t\}$ is given by:

$$q_{t+1} = q_t + q_t^o - q_t^s. \tag{108}$$

Since the quantity demanded has support on the $[0,\infty)$ interval, equation (107) implies that there is always a positive probability of unfilled demand $q_t^s < q_t^d$. We let $\delta(p, q+q^o)$ denote the probability of this event:

$$\delta(p, q+q^o) = 1 - H(q+q^o|p). \tag{109}$$

Whenever $q_t^d > q_t^s$, equations (107) and (108) imply that a *stockout* occurs, i.e. $q_{t+1} = 0$.

We define the intermediary's expected sales revenue $ES(p^s, q, q^o)$ by:

$$\begin{aligned} ES(p,q,q^o) &= E\{p^s q^s|q,q^o\} \\ &= p^s E\{q^s|q,q^o\} \end{aligned} \tag{110}$$

where:

$$E\{q^s|p,q,q^o\} = [1-\eta(p)]\left[\int_o^{q+q^o} q^d h(q^d|p)dq^d + \delta(p,q+q^o)[q+q^o]\right]. \tag{111}$$

$ES(p,q,q^o)$ is a strictly monotonically increasing and concave function of $q$.

We assume the intermediary incurs a cost of ordering inventory given by a function $c^o(p,q^o)$ given by:

$$c^o(p,q^o) = \begin{cases} K + pq^o & \text{if} \quad q^o > 0 \\ 0 & \text{otherwise.} \end{cases} \tag{112}$$

The intermediary's single-period profit $\pi$ is given by:

$$\pi(p_t,q_t^s,q_t,q_t^o) = p^s q_t^s - c^o(p_t,q_t^o). \tag{113}$$

The value function $V(p,q)$ is given by the unique solution to Bellman's equation:

$$V(p,q) = \max_{0 \le q^o \le \overline{q}-q}\left[W(p,q+q^o) - c^o(p,q^o)\right], \tag{114}$$

where:

$$W(p,q) \equiv \left[ES(p,q) + \frac{1}{1+r}EV(p,q)\right], \tag{115}$$

and $EV$ denotes the conditional expectation of $V$ given by:

$$\begin{aligned} EV(p,q) &= E\{V(\tilde{p},\max[0,q+q^o-\tilde{q}^d])|p,q) \\ &= \eta(p)\int_{p'} V(p',q)g(p'|p)dp' + [1-\eta(p)]\delta(p,q)\int_{p'} V(p',0)g(p'|p)dp' \\ &+ [1-\eta(p)]\int_{p'}\int_0^q V(p',q-q')h(q')g(p'|p)dq'dp'. \end{aligned} \tag{116}$$

The optimal purchasing rule $q^o(p,q)$ is given by:

$$q^o(p,q) = \inf \mathop{argmax}_{0 \le q^o \le \overline{q}-q}\left[W(p,q+q^o) - c^o(p,q^o)\right]. \tag{117}$$

Hall and Rust (2000) prove that the optimal procurement rule is in the class of generalized $(S,s)$ policies. We define such a policy now.

**Definition:** *A generalized $(S,s)$ policy is a decision rule of the form:*

$$q^o(p,q) = \begin{cases} 0 & \text{if} \quad q \ge s(p) \\ S(p) - q & \text{otherwise} \end{cases} \tag{118}$$

*where $S$ and $s$ are functions satisfying $S(p) \ge s(p)$ for all $p$.*

Hall and Rust (2000) go on to prove the value function $V$ is linear with slope $p$ on the interval $[0,s(p)]$:

$$V(p,q) = \begin{cases} W(p,S(p)) - p[S(p)-q] - K & \text{if} \quad q \in [0,s(p)] \\ W(p,q) & \text{if} \quad q \in (s(p),\overline{q}]. \end{cases} \tag{119}$$

Note that equation (119) shows that the "shadow price" of an extra unit of inventory is $p$ when $q < s(p)$. However for $q > s(p)$ the shadow price is generally not equal to $p$ except at the target inventory level $S(p)$. For $q \in [s(p),S(p))$ we have $\partial W(p,q)/\partial q > p$ and for $q \in (S(p),\overline{q}]$ we have $\partial W(p,q)/\partial q < p$. Thus, the value function is not concave in $q$, but is $K$-concave in $q$. The intuition for this simple result is straightforward: if the firm has an extra unit of $q$ when $q \leq s(p)$ then it needs to order one fewer unit in order to attain its target inventory level $S(p)$. The savings from ordering one fewer unit of inventory is simply the current spot price of the commodity, $p$. When $q > s(p)$ it is not optimal to order and and the shadow price of an extra unit of inventory is no longer equal to $p$. We do know that since $q = S(p)$ maximizes $W(p,q) - pq$, we must have $\partial W(p,q)/\partial q = p$ when $q = S(p)$. If we assume for the moment $W$ is strictly concave, this implies that $\partial W(p,q)/\partial q > p$ when $q \in (s(p),S(p)]$ and $\partial W(p,q)/\partial q < p$ when $q \in (S(p),\overline{q}]$. Thus, there is a kink in $V$ function at the inventory order threshold, $q = s(p)$. As we can see from formula (115) this kink is also present in the expected value function $W(p,q)$. However in our numerical examples below we find that there is only a small discontinuity in the partial derivative $\partial W(p,q)/\partial q$ at $q = s(p)$, so that $W(p,q)$ is approximately strictly concave in $q$.

The first order condition of of the value function determining the target inventory level $S(p)$ can be viewed as an Euler equation. This condition can be written as

$$p = \frac{\partial ES}{\partial q}(p,S(p)) + \frac{1}{1+r}\frac{\partial EV}{\partial q}(p,S(p)). \tag{120}$$

The first terms $\partial ES(p,S(p))/\partial q$ constitute the "convenience yield" net of holding costs of adding an extra unit of inventory. In our case, the convenience yield equals the increase in expected sales of having an extra unit of inventory. The second term, $\frac{1}{1+r}\partial EV(p,S(p))/\partial q$, is the expected discounted shadow price of an extra unit of inventory. As noted above, $\partial V(p,q)/\partial q = p$ for $q < s(p)$, and at $q = S(p)$. Even though $\partial V(p,q) > p$ for $p \in [s(p),S(p))$, we also have $\partial V(p,q)/\partial q < p$ for $q \in [S(p),\overline{q}]$.

We now solve a discrete approximation of (114) numerically using parameters that match the selected first and second moments of the daily dataset presented in Hall and Rust (1999a,b). The daily interest rate, $r$ is equal to $2 \times 10^{-4}$. The firm uses the sales price markup rule $p_t^s = 0.9 + 1.06 p_t$ and spot prices $\{p_t\}$ evolve according to a truncated lognormal $AR(1)$ process:

$$\log(p_{t+1}) = \mu_p + \lambda_p \log(p_t) + \varepsilon_t \tag{121}$$

where $\mu_p = .06$, $\lambda_p = .98$, and $\{\varepsilon_t\}$ is an *IID* $N(0,\sigma_p^2)$ sequence, with $\sigma_p^2 = 3.94 \times 10^{-4}$. The upper and lower truncation bounds on this process were chosen to be $(13,29)$ which are beyond the minimum and maximum spot purchase prices observed in our sample or in long run simulations of the untruncated version of this process. These values yield a order price process with an invariant distribution with mean of 20.0 cents per pound and a standard deviation of 2.00 cents per pound. Given the markup, the mean and standard deviation of the sell price process are 22.1 and 1.11, respectively.

We assumed that quantity demanded, $q_t^d$, is a mixed truncated lognormal distribution conditional on $p_t$. That is, with probability .5 $q_t^d = 0$, and with probability .5 $q_t^d$ is a draw from a truncated lognormal distribution with location parameter $\mu_q(p) = 5.50 - .7\log(p_t)$ and standard deviation parameter $\sigma_q = 1.4$. These parameters yield a stationary distribution for $q_t^d$ (conditional on $q_t^d > 0$) with conditional mean equal to 25.0 and conditional standard deviation equal to 25.0. The units of the quantity variables are in 1,000's of pounds. We assumed that goodwill costs of stockouts $\gamma$ is \$100, the physical holding costs are zero, $c^h(q_t) = 0$, and that the fixed order cost is equal to \$75, i.e. $c^o(0) = 0$ and $c^o(q^o) = \$75$ if $q^o > 0$.

We first solved for the optimal inventory investment rule by the method of parameterized policy iteration (PPI). This PPI algorithm amounts to the following iterative procedure:

1. Approximate the value function $V(p,q)$ with a finite linear combination of basis functions.

2. Discretize the state space into a finite number of $(p,q)$ pairs.

3. Using equation (117), compute the optimal decision rule $q_i^o(p,q)$ at each of the discretized $(p,q)$ pairs. Note that although we discretized the state variables, we treat the control variable $q_i^o$ as a continuous variable subject to the constraint that $0 \le q_i^o \le \overline{q} - q$.

4. Perform a policy iteration step. That is compute

$$V_i(p,q) = E\left\{ \sum_{j=0}^{\infty} \left(\frac{1}{1+r}\right)^j \pi(p_j, p_j^s, q_j^s, q_i^o(p_j,q_j) + q_j^s) \Big| p,q \right\}. \tag{122}$$

5. Regress the updated value function, $V_i(p,q)$, on the discrete set of $p$ and $q$'s to compute a new parameterized approximation of the value function.

6. Iterate over $i$ on steps 3–5 until the coefficients on the parameterized approximation of the value function converge.

We approximated the value function by a complete set of Chebychev polynomials of degree 3 in $p$ and $q$. We discretized the state space with 225 grid points (15 in the $p$ dimension and 15 in the $q$ dimension).

The grid points are fixed at the Chebychev zeros, so the grid points are more heavily weighted toward the boundaries of the state space. Policy iteration is not guaranteed to converge in continuous choice problems such as this one; but for this example, this algorithm converged in 29 iterations.

As can be seen from Bellman's equation (114), the policy improvement step requires the solution of a constrained optimization problem involving the two functions $ES(p,q)$ and $EV(p,q)$, each of which is a conditional expectation of functions of two continuous variables (sales, $p^s q^s$, and the value function, $V(p,q)$). Since no analytic solutions to these conditional expectations exist, we resorted to numerical integration.

Figures 24-27 present the optimal decision rule $q^o$ as a function of $p$ and $q$ and the associated expected profit function, value function and $(S,s)$ bands. Note that neither the PPI nor the DPI solution algorithms exploit our prior knowledge about the form of the decision rule. The computed value function appear to be nearly linearly increasing in current inventory $q$. At low inventory levels (in regions the firm is expecting to purchase new inventory), $V(p,q)$ is decreasing in $p$, whereas at high values of $q$, (in regions the firm is expecting to not buy but just inventory) $V$ is increasing in $p$. The optimal decision rule is decreasing in both $p$ and $q$, although it generally decreases faster in $p$ than in $q$. In particular when $q^o(p,q) > 0$, $\partial q^o(p,q)/\partial q = -1$ which is consistent with the prediction of the generalized $(S,s)$ rule that $q^o(p,q) = S(p) - q$.

Figure 26 shows the generalized $(S(p),s(p))$ bands implied by our model. The set of order limit points, $s(p)$, is the curve on the $(q,p)$ plane where the $q^o(p,q)$ surface intersects the plane at $q^o = 0$. The set of target inventory points, $S(p)$, is the curve on the $(q_o,p)$ plane where the $q^o(q,p)$ surface intersects the plane at $q = 0$. These bands are plotted in figure 27. Due to the fixed costs of ordering (\$75), the $S(p)$ band is strictly above the $s(p)$ band although the difference between the two bands decreases as the price increases. In other words, the order size at $s$ is a decreasing function of the price.

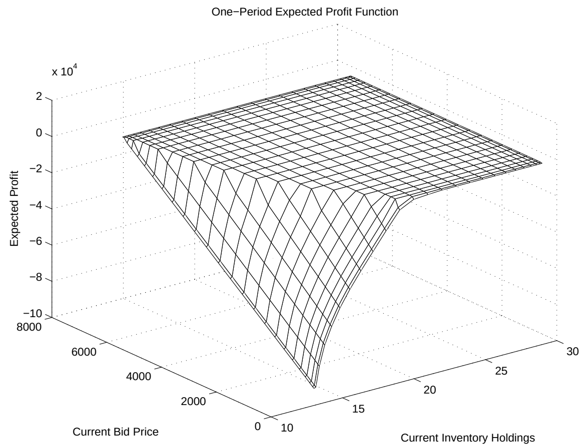


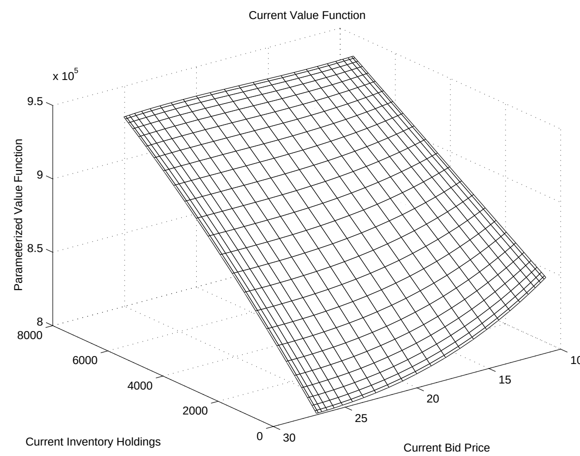


Figure 24: Expected one-period profit function for the inventory problem using PPI.

Figure 25: The value function, $V(q, p)$ for the inventory problem using PPI.

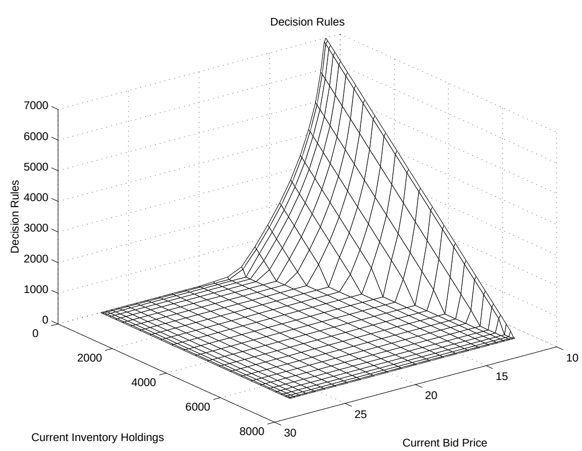


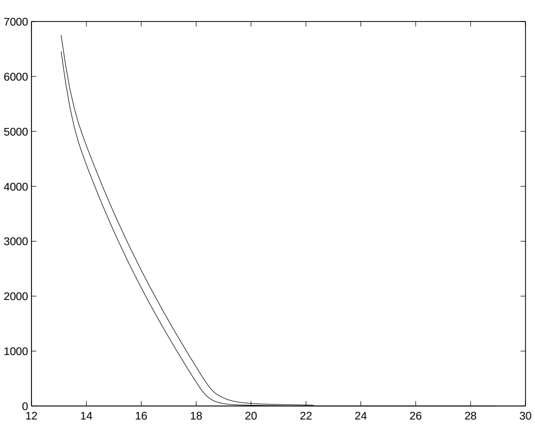

Figure 26: Decision rule, $q^o(q, p)$, for the inventory problem using PPI.

Figure 27: $S(p)$ and $s(p)$ for the inventory problem using PPI.

# 9 A Model of Product Introduction

The model described below is motivated by an empirical phenomenon we observe in many industries: Firms frequently introduce new products, but a majority of these products stays in the market for only a brief time. A specific example is the US ready-to-eat breakfast cereal industry, where the high rate of new product introductions has been noted by economists already a long time ago.[13] The model outlined in this section is used by Hitsch (2000) in a dynamic structural estimation framework to uncover the operation and rationale for new product introductions in the ready-to-eat breakfast cereal industry. However, it can be applied to other industries, and furthermore, it can be used to investigate dynamic aspects of firm entry.

The model is designed to capture the behavior of a firm which introduces a new product in a market, and is uncertain about the demand the product will generate. Eventually, the firm learns about the level of sales it can expect from its new brand, and decides whether to keep the product, or drop it from its product line.

The uncertainty which the firm faces arises because some characteristics of the new product, and the way these characteristics influence consumer choice are unobservable to the firm. The effect of these characteristics is summarized in a parameter $\lambda$, which enters the product demand function, and will be referred to as *product quality*. The firm has an initial prior on $\lambda$, and through time, it learns about the product quality from observing sales.[14]

The decisions the firm takes are as follows: At the beginning of each period the firm decides to keep the product, or drop it from its product line. If the firm decides to keep the product and stay in the market, it sets a product price and spends a certain amount on advertising.[15] At the end of the period, the firm observes sales, and updates its prior on the product quality.

An appropriate framework to describe the problem of the firm necessarily has to be dynamic. In a static framework, the objective function of the firm would be specified as current-period expected profits. If these expected profits were negative, the firm would discontinue producing the product. However, such a static framework would give the wrong prediction on the optimal exit/stay decision of a forward-looking firm, because it cannot recognize the value of staying in the market in order to obtain a more precise estimate of the product quality parameter. Hence, there will be situations where the firm rationally stays in the market

[13] See Schmalensee (1978).

[14] The model resembles Jovanovich (1982), where firms learn about their productivity. However, this model is only concerned with the decisions of a single agent, while Jovanovich's paper is about the evolution of a whole industry.

For recent empirical work including models of learning see Ackerberg (1998), Ching (2000), Crawford and Shum (1999), and Erdem and Keane (1996).

[15] Currently, other marketing mix variables like price promotions or couponing are not included in the model.

to learn more about the product, even though the expectation of current-period profits is negative.

The model outlined below is a model of *market experimentation*, where firms experiment using the exit/stay decision. Furthermore, if some marketing variable, for example advertising, influences the speed at which the firm learns about the true product quality, the firm will optimally alter this variable from the value which maximizes current profits, and hence experiment using that decision.[16]

Beyond learning, there are other potential reasons why the firm's problem has to be treated in a dynamic framework. Current marketing activities may have long-lasting effects in the form of advertising carryover and purchase reinforcement. Advertising carryover occurs if current advertising expenditures affect demand in the future, and purchase reinforcement occurs if a current purchase incidence changes the preferences of a household, such that the household is more likely the buy the currently consumed product again. In both cases, current advertising expenditures, or the current price, will change demand in the future. In the specific model discussed below, only advertising carryover is incorporated to keep the problem as simple as possible. A related intertemporal demand effect is variety-seeking, which is a preference to consume new products, or products whose characteristics differ from the products recently consumed. In the model, we treat variety-seeking in a rather reduced form by allowing the product to yield extra utility to consumers during the first periods after product introduction, i.e. we assume that all extra consumption due to variety-seeking motives occurs shortly after product introduction.

Out model describes the actions of a single agent, and does not take strategic interaction into account. A model allowing for competition among firms would be computationally very intensive, and on currently available hardware (maybe with the exception of advanced supercomputers) it would not be possible to estimate such a model in a reasonably short amount of time.

In the following, the different elements of the model are explained.

**States and decisions.** The state vector $x_t = (\chi_t, b_t, g_t, h_t)$ contains the following components:

1. $\chi_t$ is an indicator variable which equals 1 if the product is in the market, and 0 if it has been dropped from the firm's product line.

2. $b_t$ is the firm's belief about the product quality $\lambda$. $b_t$ could be any arbitrary probability distribution, but to make the model solvable on a computer, the model is restricted in such a way that the firm's belief is always normal, and can therefore be described by the parameter $\mu_t$, the conditional expectation of the product quality, and $\sigma_t^2$, the variance of the belief. Hence $b_t = (\mu_t, \sigma_t^2)$.

[16] See Aghion et al. (1991).

3. $g_t$ is the beginning-of-period goodwill stock, which represents the accumulated effect of past advertising.

4. $h_t$ records the time elapsed since the product has been introduced, where only the first $T$ periods after product introduction are of relevance, i.e. $h$ stays at $h_t = T$ after $T$ periods. $h_t$ accounts for systematic demand effects which occur only shortly after product introduction, and can be due to variety-seeking behavior.

The decisions taken by the firm in each period are $d_t = (\chi_{t+1}, p_t, a_t)$, where $\chi_{t+1}$ is the exit/stay decision, $p_t$ is the product price, and $a_t$ is the dollar amount spent on advertising.[17]

**Advertising and goodwill.** Advertising has intertemporal effects through an accumulated advertising stock called *goodwill*, denoted by $g_t$. At the beginning of each period, the firm decides to spend a certain dollar amount on advertising, which increases beginning-of-period goodwill $g_t$ and yields a quantity called *added goodwill*, denoted by $g^a = \Phi(g,a)$. A larger quantity of added goodwill increases the demand for the firm's product through a function $\Psi(g^a)$. The functional forms chosen are

$$\Phi(g,a) = g + F \cdot (1 - e^{-\phi a}), \tag{123}$$

$$\Psi(g^a) = G \cdot (1 - e^{-\psi g^a}). \tag{124}$$

Both the increase of goodwill through advertising, and the increase of utility through added goodwill are bounded by $F$ and $G$. The parameters $\phi$ and $\psi$ determine the speed at which goodwill and utility can be increased.

The law of motion for the goodwill stock is specified as

$$g_{t+1} = \exp(\nu_{t+1}) \cdot g_t^a, \tag{125}$$

where $\nu$ is i.i.d. normal with mean $\mu_\nu$ and variance $\sigma_\nu^2$. The expectation of the log-normally distributed random variable $\exp(\nu_{t+1})$ is $\mathbb{E}(\exp(\nu_{t+1})) = \exp(\mu_\nu + \sigma_\nu^2/2)$, and hence we restrict its parameters such that $\mu_\nu + \sigma_\nu^2/2 < 0$, in which case goodwill will stochastically decay from one period to the next.

**Demand.** The demand for the firm's product is given by the logit formula

$$Q_t = Ms_t = M \frac{\exp(\delta_t)}{z + \exp(\delta_t)}. \tag{126}$$

[17] Using the same symbol '$\chi$' for both a state and decision variable is not quite concise, but it highlights that in period $t$, the firm decides whether the product will be in the market at the beginning of period $t+1$, indicated by $\chi_{t+1}$.

where $M$ is the market size, and $s_t$ is the market share. $z$ is an indicator of the competitive strength of all rival products.[18] The term $\delta_t$, which in the context of a logit model has the interpretation of mean (across households) utility, is specified as

$$\delta_t = \lambda - \alpha p_t + \Psi(g_t^a) + \tau_{h_t} + \varepsilon_t. \tag{127}$$

$p_t$ is the product price, and $\Psi(g_t^a)$ is the effect of added goodwill. $h_t$ records the time elapsed since product introduction, and indexes one of the time dummies $\tau_0, \ldots, \tau_T$, where $\tau_{T+t} = \tau_T$ for all $t \geq 0$. Systematic differences in demand which arise only during the first periods after product introduction, for example variety-seeking effects, are accounted for by these terms. Finally, $\varepsilon_t \sim N(0, \sigma_\varepsilon^2)$ is an i.i.d. shock to demand, which can be interpreted as a random component of advertising, as the effectiveness of a given dollar amount spent on advertising will generally not be exactly known in advance.

**Learning.** At the beginning of each period, the belief of the firm about the product quality is described by $b_t = (\mu_t, \sigma_t^2)$, which indexes a specific normal distribution. The firm receives a normally distributed signal $\omega_t$, and hence also its posterior will be normal.

At the end of the period the firm can calculate the exact value of the unknown components $\lambda + \varepsilon_t$ by observing demand, which can be seen from the relationship

$$\log(s_t) - \log(s_{0t}) = \delta_t, \tag{128}$$

where $s_{0t}$ is the 'outside' market share. We actually assume that the firm observes only a noisy signal of the sum of these two unknown components:

$$\omega_t = \lambda + \varepsilon_t + \eta_t. \tag{129}$$

The term $\eta_t$ can be interpreted as observation error, for example because exact data on market shares are not immediately available. Alternatively, $\eta_t$ can be thought of as introducing some mild form of bounded rationality into the problem, in the sense that the managers of the firm make slight errors when updating their posteriors.[19] The component $\eta_t$ is i.i.d. normal with mean 0 and variance $\sigma_\eta^2$, and hence also $\omega_t$ is normally distributed: $\omega_t \sim N(\lambda, \sigma_\varepsilon^2 + \sigma_\eta^2)$.

[18] The logit demand system can be derived from the aggregation of brand choices across households. Each household $h$ chooses one of the brands $1, \ldots, J$ or the outside alternative 0. The utility of alternative $i$ is $U_i^h = \delta_i + e_i^h$, where $\delta_i$ is the mean utility of alternative $i$ across households, and $e_i^h$ is an household-specific utility component which has the extreme value distribution. Aggregating across households, one finds that the market share of brand $i$ is $s_i = \exp(\delta_i) / \sum_k \exp(\delta_k)$. In the context of our model, if we consider the demand for product $i$, the indicator of competition is $z = \sum_{k \neq i} \exp(\delta_k)$.

[19] Without the error component $\eta_t$, the model will generally be rejected by the data. See Hitsch (2000) for the details.

The firm updates its belief using Bayes rule, and then, given the beginning-of-period belief $b_t = (\mu_t, \sigma_t^2)$, the belief in the next period is

$$\mu_{t+1} = \mu_t + q_t(\omega_t - \mu_t), \tag{130}$$

$$\sigma_{t+1}^2 = (1 - q_t)\sigma_t^2, \tag{131}$$

where the coefficient $q_t$, which can be interpreted as the speed of learning, is given by the formula

$$q_t = \frac{\sigma_t^2}{\sigma_t^2 + \sigma_\varepsilon^2 + \sigma_\eta^2}. \tag{132}$$

Note that the evolution of the variance $\sigma_t^2$ is deterministic, which means that even though the firm does not know its conditional expectation of the product quality next period, it knows how precise its belief will be. The conditional expectation of the product quality next period, conditional on the information at the beginning of period $t$, is normally distributed:

$$\mu_{t+1} \sim N(\mu_t, q_t^2(\sigma_\varepsilon^2 + \sigma_\eta^2)). \tag{133}$$

**The objective function of the firm.** Current period profits are denoted by $\pi_t$ and given by

$$\pi_t = Q_t(p_t - c) - a_t - k. \tag{134}$$

$c$ is the unit cost of production, which is assumed to be a constant parameter. Realistically, having a product included in the product line incurs a cost which will be positive even if sales are tiny. This per period fixed cost is denoted by $k$, and includes the value of managerial time devoted to managing the project, as well as the opportunity cost of shelf space both in the supermarket where the product is sold and in the warehouse where it is stored intermittently. All uncertainty about current-period profits is due to the uncertainty about current-period demand.

The firm chooses its actions to maximize the objective function

$$V(x_t) = \mathbb{E}\left(\sum_{\tau=t}^{\infty} \beta^{\tau-t} \chi_{\tau+1} \left(Q_\tau(p_\tau - c) - a_\tau - k\right) | x_t\right). \tag{135}$$

The discount rate $\beta$ is constant, and the firm is risk-neutral.[20]

[20]Typically, a new brand accounts only for a small portion of the profits of a cereal firm, and hence does not change the variability of the firm's performance in a significant way.

**Solution of the model.** In the previous subsections all elements of the decision process were outlined. Under the assumptions made, in particular given that all variables follow Markov processes, the optimal policies will be time-invariant functions $d(x) = (\chi(x), p(x), a(x))$. These policies can be recovered from the value function, which satisfies the Bellman equation

$$V(x_t) = \max(0, \sup_{\chi_{t+1}=1, p_t, a_t} \mathbb{E}(\pi_t + \beta V(x_{t+1})|x_t, d_t)). \tag{136}$$

As the firm has the option to drop the product from its product line, the value of the product is always non-negative.

To calculate the expectation of the current-period profit flow, we note that all uncertainty about $\pi$ comes in through the uncertainty about the product quality $\lambda$, and through the demand shock $\varepsilon$. $\varepsilon$ is an i.i.d. random variable with a normal distribution, $\varepsilon_t \sim N(0, \sigma_\varepsilon^2)$. From the point of view of the firm, the product quality $\lambda$ is also normal with mean $\mu_t$ and variance $\sigma_t^2$, $\lambda \sim N(\mu_t, \sigma_t^2)$. Given a state $x_t$, and marketing choices $p_t$ and $a_t$, all randomness in profits is due to randomness in the mean utility level $\delta_t$, and randomness in $\delta_t$ is due to the term $\lambda + \varepsilon_t \equiv \xi_t$. $\xi$ also has a normal distribution, $\xi_t \sim N(\mu_t, \sigma_t^2 + \sigma_\varepsilon^2)$, and we denote its density by $f_\xi(\cdot|x, d)$. The conditional expectation of the current period profit flow can then be calculated as

$$\mathbb{E}(\pi_t|x_t, d_t) = \int_{-\infty}^{\infty} Q_t(p_t - c)dF_\xi(\xi|x_t, d_t) - a_t - k. \tag{137}$$

As regards the expected future value of the product, note that among the state variables only the conditional expectation $\mu$ and the goodwill stock $g$ do not evolve deterministically. We have seen before that $\mu_{t+1}$ has a normal distribution, and $g_{t+1}$ is log-normally distributed. Then, if the product is not withdrawn from the market today, i.e. if $\chi_{t+1} = 1$, the expected future value is

$$\mathbb{E}(V(x_{t+1})|x_t, d_t) = \int_{-\infty}^{\infty} \int_0^{\infty} V(1, \mu, \sigma_{t+1}^2, g, h_{t+1})dF_g dF_\mu. \tag{138}$$

If the product is withdrawn, its value is 0.

The maximization involved in the calculation of the Bellman operator can be made faster by solving for the optimal price $p^*$ before starting the policy iteration algorithm. This is possible because the product price affects only the current period profit flow, but not the future value of the product. In this case, $p^*$ has to be defined as a function of $\mu$, $\sigma^2$, $h$, and $g^a$ instead of $g$, and the profit flow can be redefined as implicitly incorporating the optimal choice of the product price. The only continuous control remaining is the advertising expenditure $a$.

Furthermore, note that $T$ periods after product introduction, the state variable $h$ remains constant at $h = T$. One can then solve the dynamic programming problem by first computing the value function for

$h = T$. The value functions for $h < T$ can then be recursively computed by

$$V^{\mathcal{T}}(x_t) = \max\left(0, \sup_{d_t} \mathbb{E}(\pi_t + \beta V^{\mathcal{T}+1}(x_{t+1})|x_t, d_t)\right), \qquad \mathcal{T} = T-1, ..., 0, \tag{139}$$

where one has to include the appropriate time dummy in the profit flow $\pi_t$. This is much more efficient than iterating on the full state space, as the solution on the region of the state space where $h < T$ is not needed in updating the value function on the region where $h = T$.

**Computational details** The value function is represented by a discrete approximation with uniform grids. The size of the array which stores the value function is $N_\mu \cdot N_{\sigma^2} \cdot N_g \cdot (T+1)$.[21]

As noted before, the current pricing decision does not affect the expect future value of the product, and hence the optimal price and corresponding profit flow is calculated as a function of the current state, where goodwill is replaced by added goodwill. This calculation takes only a small fraction of the total time needed to find a solution to the Bellman equation.

The stationary part of the value function, i.e. the value function at least $T$ periods after product introduction is found by using policy iteration. The optimal advertising decision is calculated using Brent's method.[22] We use Gauss-Hermite quadrature to calculate the expected future value of the product. Because $\sigma^2_{t+1}$ is known at time $t$, the integral has to be taken only over $\mu$ and $g$. The value function is evaluated using three-dimensional linear interpolation, both for points within be boundaries of the grid and for points outside the grid, i.e. the value function is extrapolated from its two closest values at the boundary. The exit/stay decision is then made by checking whether the optimal value from staying in the market is non-negative. Having computed the stationary part, we then compute the value function recursively at time $h = T-1, ..., 0$ after product introduction.

We experimented using Chebyshev polynomials to approximate $V$, but failed to obtain a satisfactory solution. The problem that arises is due to the kink of the value function at the boundary between the exit/stay regions of the state space (see Figure 29). Knowing where the value function hits 0 is essential to determining the exit rule. We tried to solve this problem by approximating

$$W(x_t) = \sup_{\chi_{t+1}=1, p_t, a_t} \mathbb{E}(\pi_t + \beta V(x_{t+1})|x_t, d_t)), \tag{140}$$

instead of directly approximating the value function $V$. $W$ is the value of the product if the firm has to keep it in the market this period, but can withdraw it at any point in the future. Note that $V = \max(0, W)$, and

[21]The indicator $\chi_t$ adds no additional dimension to the computational complexity, as we know that $V = 0$ whenever the product is no longer in the market.

[22]See Press et al. (1992), p. 402.

therefore, we can infer $V$ immediately if we know $W$. As opposed to $V$, $W$ is a smooth function, and hence it can be more easily approximated by polynomials than the value function. In fact, it turns out that $\max(0,W)$ approximates $V$ quite well overall, but often has its kink at the wrong point and therefore yields an inaccurate exit rule. This is due to the fact that the Chebyshev approximation routine does not take into account the importance of finding a very concise approximation close to the exit region. If the model were used only for illustrative purposes, this might not be a major issue, however, the exit/stay decision is of importance when calculating the likelihood function, and in this sense the Chebyshev approximation did not turn out to work well.

As an example, we approximate the continuous part of the state space using $N_\mu = 21$, $N_{\sigma^2} = 11$, and $N_g = 26$ points, and there are 3 initial time periods, i.e. $T = 3$. The stationary part of the state space contains 6006 points. The discount rate $\beta$ is set to 0.975. The solution of the model is found on a computer using an Intel Pentium III 500 MHz processor. All parts of the program are coded in C/C++[23], and the program is compiled using the MS Visual C/C++ compiler using all optimization flags. The solution is found in 35 seconds. Figure 28 shows the value function, and Figures 29-31 show the pricing, advertising, and exit policy.

[23]To be precise, certain parts of the program are coded using C++ language elements, however, we do not make use of the distinguishing object-oriented programming features of C++, and hence the whole program could as well be written using standard C.

Figure 28: Value Function

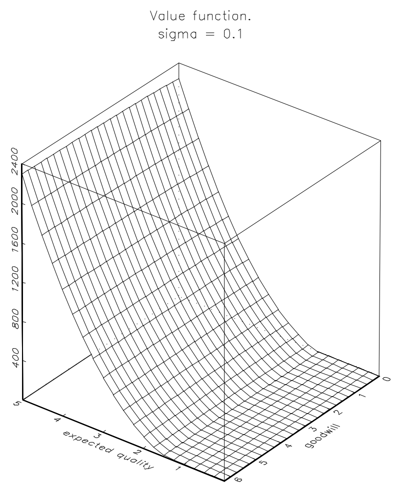


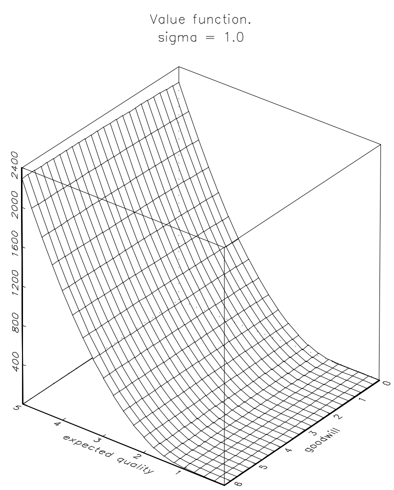

Figure 29: Pricing Policy

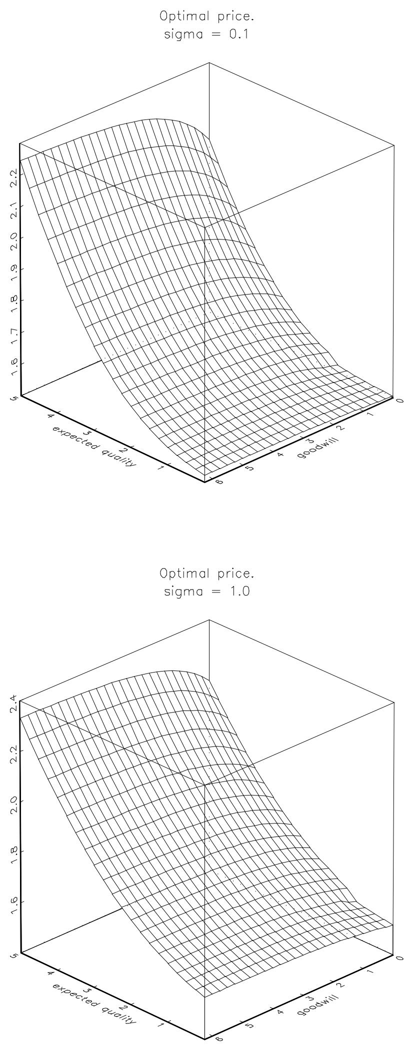

Figure 30: Advertising Policy

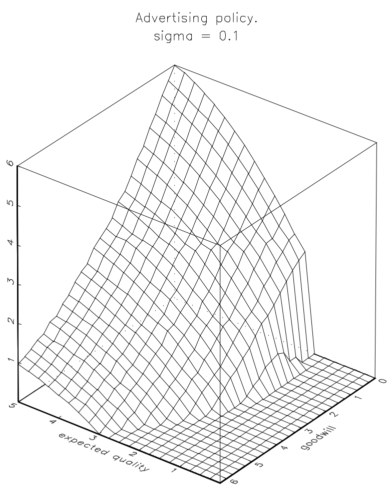


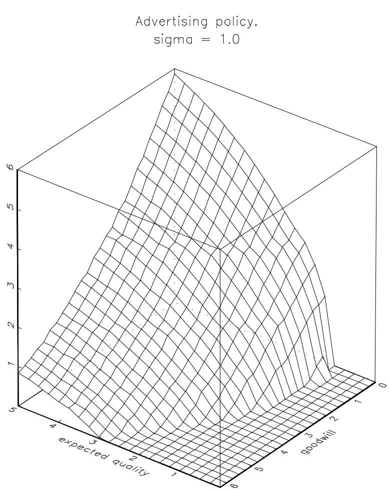

Figure 31: Exit Policy

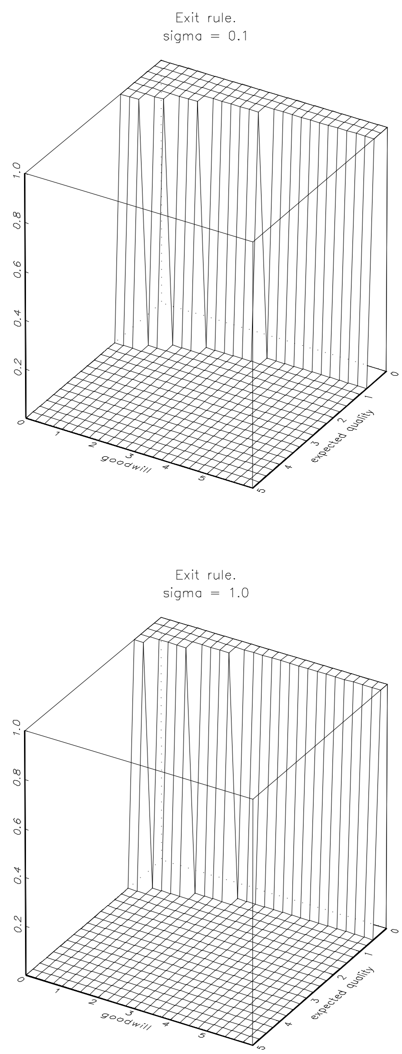

## 10 Consumption/Saving and Labor/Leisure Choices

In this section we present an extension to the classical consumption/saving model under uncertainty discussed in previous sections by introducing the labor/leisure decision as endogenous.[24] Utility is now a function of consumption and leisure, and agents will optimally choose both in every period of their lives. They solve

$$\max_{c_s, l_s} E_t \left[ \sum_{s=t}^{T} \beta^{s-t} u(c_s, l_s) \right] , \tag{141}$$

again in finite horizon. The within-period utility function is assumed to be Isoelastic and Cobb-Douglas between consumption and leisure in time $t$:

$$u(c_t, l_t) = \frac{(c_t^{\eta} l_t^{1-\eta})^{1-\gamma}}{1-\gamma} , \tag{142}$$

where $\gamma$ is the coefficient of *relative risk aversion* and $\eta$ is the valuation of consumption versus leisure.[25] Consumption and leisure are substitutes or complements depending on the value of $\gamma$ as discussed in Heckman (1974) and Low (1998), with the cutoff approximately equal to 1.[26] In our analysis below we will assume values of $\gamma$ larger than 1, implicitly assuming substitutability between consumption and leisure. We will also assume that the agent has only three choices with respect to the labor decision: part-time, full-time, or out of the labor force.[27] It is also important to emphasize that for computational convenience we have chosen a lower bound on leisure equal to 20% of the available time during a given period.[28] Given that we allow for consumption and leisure to influence each other using a CRRA utility function, and considering that we are concerned with corner solutions for the labor decision, the model can only be solved numerically. To do so we employ the techniques presented throughout the paper.

The model introduces, on top of the capital uncertainty we had in the previous models, income uncertainty, and allows for the labor/leisure decision to be endogenous.[29] This feature complicates the model because the value functions now depend on the uncertain wage realizations. We introduce serially correlated

[24] This subsection borrows from Section 3 in Benítez-Silva (2000) and also from Benítez-Silva et al. (2000).

[25] See Browning and Meghir (1991) for evidence on non-separability of consumption and leisure within periods.

[26] Heckman presents a model of perfect foresight and shows that by introducing the labor supply decision it is possible to reconcile the empirical evidence on consumption paths with the life cycle framework, without resorting to credit market restrictions or uncertainty. Low's (1998, 1999) work is fairly close in nature to the model summarize here, but he abstracts from capital uncertainty and allows for borrowing. French's (2000) model is also close to this extended model, although it focuses on the retirement decision and assumes separability between consumption and leisure in the utility function.

[27] We solve in this case an 80-period model, with agents making decision between age 20 and 100.

[28] Different values of this parameter have essentially no effect on the solutions presented below.

[29] We do not allow here for nonzero correlation between income shocks and asset returns. For a discussion of this possibility at the micro level see Davis and Willen (2000).

wages such that,

$$ln\, \omega_t \;=\; (1-\rho)\,\alpha_t + \rho\, ln\, \omega_{t-1} \;+\; \varepsilon_t\,, \tag{143}$$

where $\alpha(t)$ is a quadratic trend that mimics a concave profile of a representative individual. The $\varepsilon_t$ are *i.i.d.* draws from a normal distribution with mean 0 and variance $\sigma_t^2$.

We write the problem solved by the agents in the last period of life as

$$V_T(w,\omega) = \max_{(0\leq c\leq w+\omega(1-l),l)} U(c,l) + K\,U(w+\omega(1-l)-c)]\,, \tag{144}$$

where labor is again chosen among the three possible states. Once we obtain the decision rules numerically we can write the value function in the next to last period:

$$V_{T-1}(w,\omega) = \max_{(0\leq c\leq w+\omega(1-l),l)} U(c,l) + \beta\, E\; V_T(w+\omega(1-l)-c,\omega)\,. \tag{145}$$

The functions for the earlier periods are again obtained recursively. The expectation $E\,V_t(\omega(1-l)+w-c,\omega)$ appearing in the value functions for the different periods can be written as follows:

$$\int_0^{\overline{R}}\int_0^{\overline{\omega}} V(\tilde{R}(w+\tilde{\omega}(1-l)-c),\tilde{\omega})\; f(\tilde{\omega})d\tilde{\omega}\; f(\tilde{R})d\tilde{R}\,. \tag{146}$$

The interpolation of the values of the next period value function has to be carried out in two dimensions, a slightly more cumbersome and slower procedure, we use bilinear interpolation using C to speed up the calculations.[30] The double integrals are again solved by Gaussian-Legendre quadrature, but we use iterated integration since we are assuming independence of wages and interest rates.[31]

Figures 32-34 show the averages of 5,000 simulations of the paths of the relevant variables. Our results show that consumption profiles track income paths very closely up to age 45, when wealth accumulation starts in meaningful amounts. Wealth accumulation then continues up to quite late in life when deaccumulation starts to occur. These two results are quite important since show that the classical life cycle model of consumption can be reconciled with empirical evidence quite closely once we take into account labor supply endogenously, and in the presence of capital and wage uncertainty. The labor supply profile shows full-time work during most of the individuals' life, with part-time work at the very beginning and very end of the life

[30] We also interpolated our functions using bi-simplicial interpolation as suggested by Judd (1998) but found that it was not as accurate as the more standard bilinear and it was not necessarily faster once we wrote the routine in C.

[31] Given that the value function depends on wealth and wages, we needed to discretize both variables in order to approximate the integrals, using 50 points for wealth and 50 points for wages. We found that using fewer points significantly affected the accuracy of the calculations, leading to possible erroneous conclusions.

cycle.[32] We plot the case of individuals starting with wealth of 10,000 units, initial wages of 30,000 units, and serial correlation parameter equal to 0.9.

From the solution and simulation of these models we can conclude that a life cycle model with endogenized labor supply behaves quite consistently with the empirical data on wealth accumulation and consumption profiles and that wealth accumulation seems to start only in mid-life. Additionally, such a model endogenously captures the exiting from the labor force by older individuals who face lower wages. We consider these results as encouraging examples of the interesting models that can be solved with the techniques highlighted in this paper.

[32] Benítez-Silva (2000) shows that once we introduce Social Security in this model labor supply reacts dropping right at the age in which individuals start receiving benefits. Wealth accumulation and welfare are also negatively affected. The author also extends this model to account for an endogenous annuity decision, and presents a possible solution to the “annuity puzzle,” the question as to why the annuity market is so narrow.

Figure 32: Simulated Consumption. Serial Stochastic Wages.


Figure 33: Simulated Labor Supply. Serial Stochastic Wages


Figure 34: Simulated Wealth Path. Serial Stochastic Wages

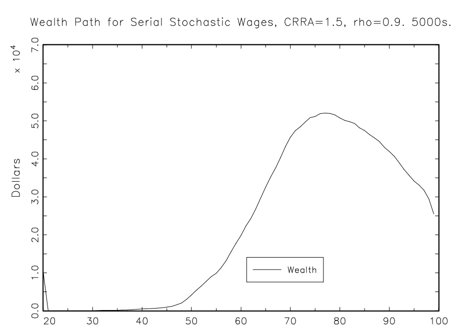

## 11 Conclusion

This section has yet to be written.

## Appendix

In this Appendix we show the details of the derivations of closed form solutions of the decision rules for the finite horizon version of the Phelps'(1962) problem for the CRRA and CARA utility functions.

The derivation of the decision rules in the case of the CRRA utility function is also close in nature to the one performed in Levhari and Srinivasan (1969) for the infinite horizon case. We can again solve this problem relying on Dynamic Programming and Bellman's principle of optimality, using backward induction. In the last period of life agents solve

$$V_T(w) \quad = \quad \max_{0\leq c\leq w} \frac{c^{1-\gamma}}{1-\gamma} + K\,\frac{(w-c)^{1-\gamma}}{1-\gamma}\,, \tag{147}$$

where $\gamma$ is the coefficient of relative risk aversion and $K$ is the bequest factor, characterized as a number between zero and one.[33] By deriving the first order condition with respect to consumption it is straightforward to show that

$$c_T \quad = \quad \frac{w}{1+K^{\frac{1}{\gamma}}}\,, \tag{148}$$

we can then write the analytical expression for the last period value function:

$$V_T(w) \quad = \quad \frac{\left(\frac{w}{1+K^{\frac{1}{\gamma}}}\right)^{1-\gamma}}{1-\gamma} + K\,\frac{\left(\frac{wK^{\frac{1}{\gamma}}}{1+K^{\frac{1}{\gamma}}}\right)^{1-\gamma}}{1-\gamma}\,. \tag{149}$$

Then the problem that agents solve in the next to last period of life is:

$$V_{T-1}(w) \quad = \quad \max_{0\leq c\leq w} \frac{c^{1-\gamma}}{1-\gamma} + \beta E\, V_T(w-c)\,. \tag{150}$$

Using the previous results we can write

$$V_{T-1}(w) \quad = \quad \max_{0\leq c\leq w} \frac{c^{1-\gamma}}{1-\gamma} + \beta E\left[\frac{\left(\frac{\tilde{R}(w-c)}{1+K^{\frac{1}{\gamma}}}\right)^{1-\gamma}}{1-\gamma} + K\left[\frac{\left(\frac{\tilde{R}(w-c)K^{\frac{1}{\gamma}}}{1+K^{\frac{1}{\gamma}}}\right)^{1-\gamma}}{1-\gamma}\right]\right]. \tag{151}$$

Here in order to derive the first order condition with respect to consumption we assume, as in Levhari and Srinivasan (1969), that the value function is differentiable and that the differential and expected value operators can be interchanged. The $f.o.c.$ is then,

$$c^{-\gamma} - \beta E\,(\tilde{R}^{1-\gamma})\left[\left(\frac{(w-c)}{1+K^{\frac{1}{\gamma}}}\right)^{-\gamma}\frac{1}{1+K^{\frac{1}{\gamma}}} + K\left[\left(\frac{(w-c)K^{\frac{1}{\gamma}}}{1+K^{\frac{1}{\gamma}}}\right)^{-\gamma}\frac{K^{\frac{1}{\gamma}}}{1+K^{\frac{1}{\gamma}}}\right]\right] \quad = \quad 0\,. \tag{152}$$

[33] We also follow in this case the "egoistic" model of bequests.

Then some algebraic manipulation allows us to write the $f.o.c.$ as

$$c^{-\gamma} \;=\; \beta E\,(\tilde{R}^{1-\gamma})\,\left(\frac{(w-c)}{1+K^{\frac{1}{\gamma}}}\right)^{-\gamma}. \tag{153}$$

Some more tedious algebra leads to the following expression for the decision rule in the next to last period

$$c_{T-1} \;=\; \frac{w}{1+\beta^{\frac{1}{\gamma}}\,\left[E\,(\tilde{R}^{1-\gamma})\right]^{\frac{1}{\gamma}}\left[1+K^{\frac{1}{\gamma}}\right]}, \tag{154}$$

that can be rewritten as

$$c_{T-1} \;=\; \frac{w}{1+\beta^{\frac{1}{\gamma}}\,\left[E\,(\tilde{R}^{1-\gamma})\right]^{\frac{1}{\gamma}}+\beta^{\frac{1}{\gamma}}\,\left[E\,(\tilde{R}^{1-\gamma})\right]^{\frac{1}{\gamma}}K^{\frac{1}{\gamma}}}. \tag{155}$$

Assuming next that the interest rate, $\tilde{R}$, follows a log-normal distribution with mean $\mu$ and variance $\sigma^2$, then given that $E(\tilde{R}) = e^{\mu+\frac{\sigma^2}{2}}$ and denoting $E(\tilde{R})$ as $\overline{R}$ we can write

$$E\,(\tilde{R}^{1-\gamma}) \;=\; \overline{R}^{1-\gamma}e^{-\gamma(1-\gamma)\frac{\sigma^2}{2}}. \tag{156}$$

We then substitute back in the formula for $c_{T-1}$ and obtain

$$c_{T-1} \;=\; \frac{w}{1+\beta^{\frac{1}{\gamma}}\left(\overline{R}^{1-\gamma}e^{-\gamma(1-\gamma)\frac{\sigma^2}{2}}\right)^{\frac{1}{\gamma}}+\beta^{\frac{1}{\gamma}}K^{\frac{1}{\gamma}}\left(\overline{R}^{1-\gamma}e^{-\gamma(1-\gamma)\frac{\sigma^2}{2}}\right)^{\frac{1}{\gamma}}}, \tag{157}$$

given the similarity with expression (95) it is easy to see how backward induction would lead us to the decision rules for the rest of the periods, for example we can write $c_{T-k}$ as

$$c_{T-k} \;=\; \frac{w}{1+\beta^{\frac{1}{\gamma}}E\,(\tilde{R}^{1-\gamma})+\beta^{\frac{2}{\gamma}}E\,(\tilde{R}^{1-\gamma})+\ldots+\beta^{\frac{k}{\gamma}}K^{\frac{1}{\gamma}}E\,(\tilde{R}^{1-\gamma})}, \tag{158}$$

where we come back to the compact notation for $E\,(\tilde{R}^{1-\gamma})$.

We can also see that if $\gamma$ is equal to 1 we are back to the logarithmic utility case and the expression for $c_{T-1}$ above is equivalent to (95), which is a special case of the expression above. It is also important to emphasize that this expression is the finite horizon counterpart to the one obtained in Levhari and Srinivasan (1969), and also replicated in Section 5.1, once a bequest motive is introduced, and that their results regarding the effects of uncertainty (decreasing proportion of wealth consumed as the uncertainty grows if $\gamma > 1$) go through in this case.

We next assume a constant absolute risk aversion (CARA) utility function. As discussed in the text we have not found a closed form solution for this problem under uncertain returns that follow a log-normal distribution as in the cases above. We therefore solve the finite horizon problem under certainty. We can again solve this problem using backward induction. In the last period of life agents solve

$$V_T(w) \;=\; \max_{0\le c\le w} -e^{-\gamma c}-K\,\left[e^{-\gamma(w-c)}\right], \tag{159}$$

where $\gamma$ is the coefficient of absolute risk aversion and $K$ is the bequest factor. We assume there is no capital accumulation in the last period of life. By deriving the first order condition with respect to consumption it can be shown that

$$c_T \;=\; \min\left(max\left(0,\frac{w}{2}-\frac{1}{2}\,\frac{\ln K}{\gamma}\right),w\right), \tag{160}$$

we can then write the analytical expression for the last period value function:

$$V_T(w) \;=\; -e^{-\gamma(\frac{\gamma w - \ln K}{2\gamma})} - K\left[e^{-\gamma\left(\frac{\gamma w + \ln K}{2\gamma}\right)}\right]. \tag{161}$$

Then the problem that agents solve in the next to last period of life is:

$$V_{T-1}(w) \;=\; \max_{0\le c\le w} -e^{-\gamma c} + \beta V_T(w-c). \tag{162}$$

Using the previous results we can write

$$V_{T-1}(w) \;=\; \max_{0\le c\le w} -e^{-\gamma c} + \beta\left[\left(-e^{-\gamma\left(\frac{\gamma(R(w-c))-\ln K}{2\gamma}\right)}\right) - K\left(e^{-\gamma\left(\frac{\gamma(R(w-c))+\ln K}{2\gamma}\right)}\right)\right]. \tag{163}$$

where $R = 1 + r$ and $r$ is the fixed rate of interest. In order to derive the first order condition with respect to consumption we again assume that the value function is differentiable and that the differential and expected value operators can be interchanged. The $f.o.c.$ is then,

$$\gamma e^{-\gamma c} - \frac{\beta R \gamma}{2}\left[e^{\frac{-\gamma R(w-c)}{2}}\left[e^{\frac{\ln K}{2}} + K\, e^{\frac{-\ln K}{2}}\right]\right] \;=\; 0. \tag{164}$$

Then some algebraic manipulation allows us to write the $f.o.c.$ as

$$\frac{1}{\frac{\beta R}{2}\left[e^{\frac{\ln K}{2}} + K\, e^{\frac{-\ln K}{2}}\right]} \;=\; e^{-\gamma\left[\frac{w-(2+R)c}{2}\right]}. \tag{165}$$

Some more tedious algebra leads to the following expression for the decision rule in the next to last period

$$c_{T-1} \;=\; \min\left(\max\left(0, \frac{Rw}{2+R} - \frac{2}{2+R}\frac{\ln K_{T-1}}{\gamma}\right), w\right), \tag{166}$$

where

$$K_{T-1} \;=\; \frac{\beta R}{2}\left[e^{\frac{\ln K}{2}} + K\, e^{\frac{-\ln K}{2}}\right]. \tag{167}$$

We can then proceed recursively, finding $w - c$, and substituting in $V_{T-1}$, and then writing the problem solved in period $T-2$. After finding the $f.o.c.$ we can again find a closed form solution for $c_{T-2}$,

$$c_{T-2} \;=\; \min\left(\max\left(0, \frac{R^2 w}{2+R+R^2} - \frac{2+R}{2+R+R^2}\,\frac{\ln K_{T-2}}{\gamma}\right), w\right), \tag{168}$$

where

$$K_{T-2} \;=\; \frac{\beta R^2}{2+R}\left[e^{\frac{2}{2+R}\ln K_{T-1}} + \beta\, e^{\frac{-R\ln K_{T-1}}{2+R}}\left[e^{\frac{\ln K}{2}} + K\, e^{\frac{-\ln K}{2}}\right]\right]. \tag{169}$$

Which in fact, can be rewritten as follows:

$$K_{T-2} \;=\; \frac{\beta R^2}{2+R}\left[e^{\frac{2}{2+R}\ln K_{T-1}} + \beta\, e^{\frac{-R\ln K_{T-1}}{2+R}}\left[\frac{K_{T-1}}{\frac{\beta R}{2}}\right]\right]. \tag{170}$$

Similarly, in recursive fashion and after a bit more algebra we find that

$$c_{T-3} \quad = \quad \min\left(\max\left(0, \frac{R^3 w}{2+R+R^2+R^3} - \frac{2+R+R^2}{2+R+R^2+R^3}\ \frac{\ln K_{T-3}}{\gamma}\right), w\right)\ , \tag{171}$$

where in this case

$$K_{T-3} \quad = \quad \frac{\beta\, R^3}{2+R+R^2}\left[e^{\frac{2+R}{2+R+R^2}\ \ln K_{T-2}} + \beta\, e^{-\frac{R^2 \ln K_{T-2}}{2+R+R^2}}\left[\frac{K_{T-2}}{\frac{\beta\, R^2}{2+R}}\right]\right]\ . \tag{172}$$

From this we can characterize the decision rule for any other period up to the first period of life

$$c_{T-k} \quad = \quad \min\left(\max\left(0, \frac{R^k w}{2+R+\ldots+R^k} - \left[\frac{2+R+\ldots+R^{k-1}}{2+R+\ldots+R^k}\right] \frac{\ln K_{T-k}}{\gamma}\right), w\right)\ , \tag{173}$$

where $K_{T-k}$ is shown below and where we can write the expression $2+R+R^2+\ldots+R^k$ as $1+\left[\frac{1-R^{k+1}}{1-R}\right]$ and similarly for the other series,

$$K_{T-k} \quad = \quad \frac{\beta\, R^k}{1+\left[\frac{1-R^k}{1-R}\right]}\left[e^{\frac{1+\left[\frac{1-R^{k-1}}{1-R}\right]}{1+\left[\frac{1-R^k}{1-R}\right]} \ln K_{T-k+1}} + \beta\, e^{-\frac{R^{k-1} \ln K_{T-k+1}}{1+\left[\frac{1-R^k}{1-R}\right]}}\left[\frac{K_{T-k+1}}{\frac{\beta R^{k-1}}{1+\left[\frac{1-R^{k-1}}{1-R}\right]}}\right]\right]\ . \tag{174}$$

## References

Ackerberg, D.A. (1998): “Advertising, Learning, and Consumer Choice in Experience Good Markets: A Structural Empirical Examination,” manuscript, Boston University.

Aghion, P., P. Bolton, C. Harris, and B. Jullien (1991): “Optimal Learning by Experimentation,” *Review of Economic Studies*, **58** 621–654.

Alessie, R., and A. Lusardi (1997): “Consumption, Saving and Habit Formation,” *Economic Letters*, **55** 103–108.

Barron, A.R. (1993); “Universal approximation bounds for superpositions of a sigmoidal function,” *IEEE Transactions on Information Theory*, **39-3** 930–945.

Benítez-Silva, H., M. Buchinsky, S. Cheidvasser, and J. Rust (2000): “Social Insurance at the End of the Life Cycle: A Structural Model of Retirement, Disability and Medicare Benefits,” manuscript. Yale University.

Benítez-Silva, H. (2000): “A Dynamic Model of Labor Supply, Consumption/Savings, and Annuity Decisions Under Uncertainty,” manuscript. Yale University.

Bernheim, B.D. (1991): “How Strong are Bequest Motives? Evidence Based on Estimates of the Demand for Life Insurance and Annuities,” *Journal of Political Economy*, **99-5** 899–927.

Bertsekas, D.P. (1995): *Dynamic Programming and Optimal Control*, Vol 1 and 2. Belmont, MA: Athena Scientific.

Braun, R. A. and E. McGrattan (1993): “The Macroeconomics of War and Peace,” *Macroeconomic Annual* National Bureau of Economic Research, Cambridge, MA.

Browning, M., and C. Meghir (1991): “The Effects of Male and Female Labor Supply on Commodity Demands,” *Econometrica*, **59-4** 925–951.

Burnside, C. (1999): “Discrete State-Space Methods for the Study of Dynamic Economies,” in R. Marimon and A. Scott (eds.), *Computational Methods for the Study of Dynamic Economies*, Oxford University Press, 95–113.

Ching, A. (2000): “Dynamic Equilibrium in the U.S. Prescription Drug Market After Patent Expiration,” manuscript, University of Minnesota.

Crawford, G.S., and M. Shum (1999): “Uncertainty and Learning in Pharmaceutical Demand,” manuscript, Duke University and University of Toronto.

Chari, V.V., L. Christiano, and P. Kehoe (1994): “Optimal Fiscal Policy in a Business Cycle Model,” *Journal of Political Economy*, **102** 617-652.

Christiano, L., and J.D.M. Fisher (2000): “Algorithms for Solving Dynamic Models with Ocasionally Binding Constraints,” *Journal of Economic Dynamics and Control*, **24** 1179–1232.

Davis, S.J., and P. Willen (2000): “Uncertain Labor Income and Portfolio Choice: Covariance and Its Implications,” manuscript, University of Chicago and Princeton University.

Deaton, A. (1992): *Understanding Consumption*, Clarendon Lectures in Economics, Oxford University Press.

Deaton, A., and G. Laroque (1992): "On the Behavior of Commodity Prices," *Review of Economic Studies,* **59** 1–23.

den Haan, W.J., and A. Marcet (1990): "Solving the stochastic growth model by parameterizing expecations," *Journal of Business and Economic Statistics,* **8** 31–34.

Erdem, T., and M.P. Keane (1996): "Decision-making Under Uncertainty: Capturing Brand Choice Processes in Turbulent Consumer Goods Markets," *Marketing Science,* **15 (1)** 1–20.

French, E. (2000): "The Effects of Health, Wealth, and Wages on Labor Supply and Retirement Behavior," manuscript. University of Wisconsin and Federal Reserve Bank of Chicago.

Gaspar, J., and K. Judd (1997): "Solving Large-Scale Rational-Expectations Models," *Macroeconomic Dynamics,* **XX** 45-75.

Hakansson, N.H. (1970): "Optimal Investment and Consumption Strategies Under Risk for a Class of Utility Functions," *Econometrica,* **38** 587–607.

Hall, G., and J. Rust (1999a): "An Empirical Model of Inventory Investment by Durable Commodity Intermediaries," forthcoming *Carnegie-Rochester Conference Series on Public Policy.*

Hall, G., and J. Rust (1999b): "A (S,s) Model of Commodity Price Speculation," manuscript, Yale University.

Hansen, L.P., and T.J. Sargent (1999): *Recursive Methods of Dynamic Linear Economies.* Princeton University Press, Princeton, N.J. Forthcoming.

Heckman, J.J. (1974): "Life Cycle Consumption and Labor Supply: an Explanation of the Relationship between Income and Consumption over the Life Cycle," *American Economic Review,* **64-1** 188–194.

Hitsch, Günter J. (2000): "Estimating a Model of Product Introduction," manuscript, Yale University.

Holt, C., Modigliani, F., Muth, J., and Simon, H. (1960): *Planning Production, Inventories and Work Force.* Englewood Cliffs, N.J.: Prentice-Hall.

Howard, R. (1960): *Dynamic Programming and Markov Processes.* J. Wiley, New York.

Hurd, M. (1987): "Savings of the Elderly and Desired Bequests," *American Economic Review,* **77-3** 298–312.

Hurd, M. (1989): "Mortality Risk and Bequests," *Econometrica,* **57-4** 779–813.

Jovanovich, B. (1982): "Selection and the Evolution of Industry," *Econometrica,* **50-3** 649–670.

Judd, K.L. (1992): "Projection Methods for Solving Aggregate Growth Models," *Journal of Economic Theory,* **58** 410-452.

Judd, K.L. (1994): "Comments on Marcet, Rust and Pakes," *Advances in Econometrics. Sixth World Congress, Vol. II* C. Sims (ed) Cambridge: Cambridge University Press.

Judd, K.L. (1998): *Numerical Methods in Economics,* The MIT Press.

Kotlikoff, L.J. (1979): "Testing the Theory of Social Security and Life Cycle Accumulation," *American Economic Review,* **69-3** 396–410.

Kotlikoff, L.J., and L.H. Summers (1981): “The Role of Intergenerational Transfers in Aggregate Capital Accumulation,” *Journal of Political Economy*, **89-4** 706–732.

Laitner, J., and T. Juster (1996): “New Evidence on Altruism: A Study of TIAA-CREF Retirees,” *American Economic Review*, **86-4** 893–908.

Levhari, D., and T.N. Srinivasan (1969): “Optimal Savings under Uncertainty,” *Review of Economic Studies*, **36-2** 153–163.

Low, H.W. (1998): “Simulation Methods and Economic Analysis,” manuscript, University College London, University of London.

Low, H.W. (1999): “Self Insurance and Unemployment Benefit in a Life-Cycle Model of Labour Supply and Savings,” manuscript, University of Cambridge.

Marimon, R. and A, Scott (1999): *Computational Methods for the Study of Dynamic Economies*, Oxford: The Oxford University Press.

Marcet, A. (1988): “Solving nonlinear stochastic growth models by parameterizing expectations,” manuscript. Carnegie-Mellon University.

Marcet, A., Sargent, T., and J. Seppala (2000): “Optimal Taxation without State-Contingent Debt,” manuscript, Stanford University.

Marshall, D. (1992): “Inflation and Asset Returns in a Monetary Economy,” *Journal of Finance* **47** 1315-1342.

McGrattan, E.R. (1996): “Solving the Stochastic Growth Model with a Finite Element Method,” *Journal of Economic Dynamics and Control* **20** 19–42.

McGrattan, E.R. (1999): “Application of Weighted Residual Methods to Dyanmic Economic Models,” *Computational methods for the Study of Dyanmic Economies* R. Marimon and A. Scott (ed.), Oxford University Press.

Miranda, M.J. (1998): “Numerical Solution Strategies for the Nonlinear Rational Expectations Commodity Market Model,” *Computational Economics*, **11** 71–87.

Miranda, M.J. and P.G. Helmberger (1988) “The Effects of Commodity Price Stabilization Programs,” *American Economic Review*, **78** 46-58.

Modigliani, F. (1988): “The Role of Intergenerational Transfers and Life Cycle Saving in the Accumulation of Wealth,” *Journal of Economic Perspectives*, **2-2** 15–40.

Niederreiter, H. (1992): “Random Number Generation and Quasi-Monte Carlo Methods, Vol. 63. Philadelphia: SIAM CBMS-NSF Conference Series in Applied Mathematics.

Papageorgiou, A., and J.F. Traub (1996): “Beating Monte Carlo,” *RISK* **9-6** 63–65.

Papageorgiou, A., and J.F. Traub (1997): “Faster Evaluation of multidimensional integrals,” *Computational Physics*, **11-6** 574–578.

Phelps, E.S. (1962): “The Accumulation of Risky Capital: A Sequential Utility Analysis,” *Eco-nometrica*, **30** 729–743.

Poterba, J.M. (1998): “Population Age Structure and Asset Returns: An Empirical Investigation,” *NBER Working Paper Series,* **6774**.

Press, W. H., S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery (1992): *Numerical Recipes in C*. Second edition. Cambridge: Cambridge University Press.

Puterman, M.L., and S.L. Brumelle (1979): “On the convergence of policy iteration in stationary dynamic programming,” *Mathematics of Operations Research,* **4** 60–69.

Puterman, M.L., and M.C. Shin (1978): “Modified policy iteration algorithms for discounted Markov decision problems,” *Management Science,* **24** 1127–1137.

Rust, J. (1985): “Stationary equilibrium in a Market for durable assets,” *Econometrica,* **53-4** 783–806.

Rust, J. (1986): “When is it optimal to kill off the market for used durable goods?” *Econometrica,* **54-1** 65–86.

Rust, J. (1987): “Optimal Replacement of GMC bus engines: An empirical model of Harold Zurcher,” *Econometrica,* **55-5** 999–1033.

Rust, J. (1996): “Numerical Dynamic Programming in Economics,” in H.M. Amman, D.A. Kendrick and J. Rust (eds.) *Handbook of Computational Economics.* Elsevier, Amsterdam, 619–729.

Rust, J. (1997a): “Using Randomization to Break the Curse of Dimensionality,” *Econometrica,* **65-3** 487–516.

Rust, J. (1997b): “A Comparison of Policy Iteration Methods for Solving Continuous-State, Infinite-Horizon Markovian Decision Problems Using Random, Quasi-random, and Deterministic Discretizations,” manuscript, Yale University.

Rust, J. (1999): “Solving High Dimensional Problems via Decentralization,” Invited Lecture in Computational Economics. Latin American Meetings of the Econometric Society, Cancun, Mexico.

Rust, J. and C. Phelan (1997): “How Social Security and Medicare Affect Retirement Behavior in a World of Incomplete Markets,” *Econometrica,* **65-4** 781–831.

Rust, J. Traub, J.F. and H. Woźniakowski (2000) “No Curse of Dimensionality for Contraction Fixed Points Even in the Worst Case” forthcoming, *Econometrica.*

Santos, M. (1999): “Numerical Solutions of Dynamic Economic Models,” *Handbook of Macroeconomics.* J.B. Taylor and M. Woodford, (eds.).

Sarafidis, Y. (2000): “Strategic Release of Information with Imperfect Memory,” manuscript. Yale University.

Schmalensee, R. (1978): “Entry Deterrence in the Ready-to-Eat Breakfast Cereal Industry,” *Bell Journal of Economics*, **9** 305–327.

Tauchen, G., and R. Hussey (1991): “Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models,” *Econometrica,* **59-2** 371–396.

Taylor, J., and Uhlig, H. (1990): “Solving nonlinear stochastic growth models: a comparison of alternative solution methods,” *Journal of Economics and Business*, **8** 1–17.

Tezuka, S. (1995): *Uniform Random Numbers: Theory and Practice.* Dordrecht, Netherlands: Kluwer.

Traub, J.F., G.W. Wasilkowski, and H. Woźniakowski (1988): *Information Based Complexity,* New York: Academic Press.

Traub, J.F., A.G. Werschulz (1998): *Complexity and Information*, Cambridge University Press.

Wilhelm, M.O. (1996): "Bequest Behavior and the Effect of Heirs' Earnings: Testing the Altruistic Model of Bequests," *American Economic Review*, **86-4** 874–892.

Wright, B. and J. Williams (1982a): "The Economic Role of Commodity Storage," *Economic Journal*, **92** 596-614.

Wright, B. and J. Williams (1982b): "The Roles of Public and Private Storage in Managing Oil Import Diruptions," *Bell Journal of Economics*, **13** 341-353.

Wright, B. and J. Williams (1984): "The Welfare Effects of the Introduction of Storage," *Quarterly Journal of Economics*, **99** 169-182.