



AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search

<http://ageconsearch.umn.edu>

aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

No endorsement of AgEcon Search or its fundraising activities by the author(s) of the following work or their employer(s) is intended or implied.



Queen's Economics Department Working Paper No. 816

Solving Stochastic Dynamic Programming Problems Using Rules Of Thumb

Anthony A. Smith, Jr.

Department of Economics
Queen's University
94 University Avenue
Kingston, Ontario, Canada
K7L 3N6

5-1991

**SOLVING STOCHASTIC DYNAMIC PROGRAMMING
PROBLEMS USING RULES OF THUMB**

Anthony A. Smith, Jr.

Discussion Paper No. 816

May 1991

phone: (613) 545-6659
email: smithaa@qucdn.bitnet

ABSTRACT

This paper develops a new method for constructing approximate solutions to discrete time, infinite horizon, discounted stochastic dynamic programming problems with convex choice sets. The key idea is to restrict the decision rule to belong to a parametric class of functions. The agent then chooses the best decision rule from within this class. Monte Carlo simulations are used to calculate arbitrarily precise estimates of the optimal decision rule parameters. The solution method is used to solve a version of the Brock-Mirman (1972) stochastic optimal growth model. For this model, relatively simple rules of thumb provide very good approximations to optimal behavior.

Keywords: rule of thumb, Monte Carlo simulation, numerical optimization,
stochastic optimal growth

JEL Classification Numbers: 210, 023

1. INTRODUCTION[†]

This paper develops a new method for constructing approximate solutions to discrete time, infinite horizon, discounted stochastic dynamic programming problems with convex choice sets. This class of optimization problems underlies a wide variety of modern work in dynamic, stochastic economic modelling. The solution to this class of optimization problems consists of a time-invariant function that relates the current state of the system to the optimal choices for next period's endogenous state variables. This function, or decision rule, possesses a closed form expression only in the simplest economic environments.¹ This paper develops a new numerical method for calculating a reasonable approximation to the unknown decision rule.

The key idea is to restrict the decision rule to belong to a class of functions characterized by a set of decision rule parameters (e.g., the class of polynomial functions of the state variables). The agent then chooses the best decision rule from within this class; that is, the agent maximizes the specified objective function subject to the functional form restriction on the decision rule by choosing the best set of decision rule parameters (e.g., coefficients on the polynomials).²

This approach to constructing approximate solutions to recursive stochastic dynamic programming problems can be given a behavioral interpretation: agents are boundedly rational in choosing the functional form of the decision rule, but are fully rational in selecting the optimal set of decision rule parameters. This approach therefore places a tight structure

[†]This paper is drawn from Chapter 1 of my Ph.D. dissertation at Duke University. I would like to thank my thesis supervisor, John Geweke, for invaluable guidance and support. I would also like to thank Dan Bernhardt and Gregor Smith for helpful comments and discussions. All errors are mine.

on the way in which agents behave suboptimally: although the choice of functional form is restricted, the choice of the optimal decision rule parameters is not. Thus the "rules of thumb" employed by agents, though not fully optimal, nonetheless solve a more limited optimization problem.

Monte Carlo methods can be used to calculate arbitrarily precise estimates of the optimal settings for the decision rule parameters, given a functional form for the decision rule. The central idea of the numerical algorithm is to simulate the dynamic behavior of the state variables for different values of the decision rule parameters. The agent then chooses the set of decision rule parameters that, on average, produces the best results from the viewpoint of the agent's objective function.

More precisely, it is shown that as the number of independent simulations increases to infinity, the numerical estimates of the optimal decision rule parameters converge in probability to the (true) optimal decision rule parameters. Asymptotically, the estimates are normally distributed with a covariance matrix that can itself be estimated consistently. This distributional result can be used to assess the numerical error associated with a finite number of simulations. A specific numerical application shows that very accurate estimates of the optimal decision rule parameters (given a functional form for the decision rule) can be calculated using a relatively small number of simulations.

By enlarging the class of functions from which the best decision rule is chosen, one can in principle approximate arbitrarily well the fully optimal but unknown decision rule. The practical usefulness of this approach, however, hinges on the ability of relatively low order approximations to provide good approximations to the unknown decision rule. This paper does not address the difficult problem of characterizing in general the

relationship between the accuracy and the order (or "complexity") of approximate decision rules. A numerical example with a known analytical solution, however, suggests that relatively simple rules of thumb can provide very good approximations to optimal behavior.

The paper is organized as follows. Section 2 formally defines the class of optimization problems addressed by this paper and explains the numerical algorithm for constructing approximate solutions to these problems. Section 3 develops the details of the asymptotic theory underlying the numerical algorithm and discusses some issues of implementation. Section 4 extends the numerical algorithm to incorporate antithetic acceleration, a powerful variance reduction technique that provides large increases in accuracy for a given number of simulations.³ Section 5 implements the solution method for the stochastic optimal growth problem studied by Brock and Mirman (1972). Section 6 concludes and suggests some avenues for future research. Proofs of the asymptotic results in Sections 3 and 4 are gathered in an Appendix.

2. THE SOLUTION METHOD

2.1 The Problem

The class of optimization problems addressed by this paper can be defined formally as follows (the notation and assumptions are drawn largely from Chapter 9 of Stokey, Lucas, and Prescott (1989)):

$$\begin{aligned} \max_{\{x_t\}_{t=1}^{\infty}} \quad & E_0 \sum_{t=0}^{\infty} \beta^t r(x_t, x_{t+1}, z_t) \quad \text{given } x_0, z_0 \quad (1) \\ \text{subject to:} \quad & x_{t+1} \in \Gamma(x_t, z_t) \quad \text{for all } t, \end{aligned}$$

where $\{z_t\}_{t=1}^{\infty}$ is a Markov process with stationary transition density

$v(z_{t+1}|z_t)$. z_t is the $l \times 1$ period t exogenous state vector; $z_t \in Z$ for all t , where Z is a compact, convex subset of \mathbb{R}^l . x_t is the $p \times 1$ endogenous state vector; $x_t \in X$ for all t , where X is a closed, convex subset of \mathbb{R}^p .

$S = X \times Z$ is the state space; at any time t , $s_t = (x_t, z_t) \in S$ characterizes the current state of the system. The agent observes s_t prior to making period t decisions (i.e. prior to choosing x_{t+1}). The correspondence $\Gamma: S \rightarrow X$ determines the set of feasible choices for next period's endogenous state vector, x_{t+1} , given the current state s_t . The discount parameter $\beta \in (0,1)$. E_0 is the mathematical expectation operator, conditional on the period 0 information set (i.e. s_0).

The return function $r: A \rightarrow \mathbb{R}$, where $A = \{(x_t, x_{t+1}, z_t): x_{t+1} \in \Gamma(x_t, z_t)\}$, is assumed to be continuous in (x_t, x_{t+1}, z_t) and concave in (x_t, x_{t+1}) for all $z_t \in Z$. The correspondence Γ is assumed to be nonempty, compact- and convex- valued, and continuous. The convexity of $\Gamma(s_t)$ for all $s_t \in S$ rules out problems with discrete choice sets. The assumption that Z is a compact, convex set can be replaced with the assumption that Z is a countable set.

These assumptions, together with the boundedness of $r(\cdot, \cdot, \cdot)$ ⁴, imply the existence of a unique, time-invariant decision rule $w: S \rightarrow X$ which expresses the optimal choice for next period's endogenous state vector (x_{t+1}) as a function of the current state (s_t). Moreover, $w(x_t, z_t)$ is continuous in x_t for all $z_t \in Z$.

2.2 The Approach

The optimal decision rule $w(\cdot, \cdot)$ is, in general, unknown. The approach pursued in this paper replaces $w(\cdot, \cdot)$ with a "rule of thumb" characterized by a $k \times 1$ vector of decision rule parameters $\Psi: x_{t+1} = h(x_t, z_t; \Psi)$. The decision rule $h(x_t, z_t; \cdot)$ closes the model. Fix an initial condition $s_0 = (x_0, z_0)$ and

choose a set of decision rule parameters Ψ . Let $z = \{z_t\}_{t=1}^T$ be a sequence of exogenous state vectors. Then, given s_0 , z , and Ψ , recursive iterations on the decision rule define a sequence of choices $x = \{x_t\}_{t=1}^{T+1}$ for the endogenous state vectors. These iterations proceed in the obvious way:

$$\begin{aligned} x_1 &= h(x_0, z_0; \Psi) \\ x_2 &= h(x_1, z_1; \Psi) \\ &\vdots \\ x_T &= h(x_{T-1}, z_{T-1}; \Psi) \\ x_{T+1} &= h(x_T, z_T; \Psi). \end{aligned}$$

Let $x = q(z; \Psi, s_0)$ represent the mapping from z to x determined by this set of iterations.⁵

Define $b(x, z; s_0) \equiv \sum_{t=0}^T \beta^t r(x_t, x_{t+1}, z_t)$. $b(x, z; s_0)$ is the return (or utility) delivered by a sequence x of endogenous state vectors and a sequence z of exogenous state vectors, given initial condition s_0 .⁶ Then, given z , $g(z, \Psi) \equiv b(q(z; \Psi, s_0), z; s_0)$ is the return realized by the agent when the agent uses decision rule $h(x_t, z_t; \cdot)$ with decision rule parameters Ψ to choose (recursively) the endogenous state variables. (The dependence of g on the initial condition s_0 has been suppressed to maintain notational simplicity.) Given the functional form restriction on the decision rule, the agent chooses the best set of decision rule parameters Ψ_0 . Formally,

$$\Psi_0 = \underset{\Psi}{\operatorname{argmax}} E g(z, \Psi), \quad (2)$$

where, in this context, the operator E means to evaluate the expectation of $g(z, \Psi)$ with respect to the joint density of z , conditional on z_0 (i.e.

integrate with respect to $\prod_{t=1}^T v(z_t|z_{t-1}) \equiv f(z|z_0)$, where z can be written as a $T \times 1$ vector).

Although Problem (2) is a simplification of Problem (1), Problem (2) nonetheless cannot, in general, be solved analytically. Since $g(z, \Psi)$ is typically nonlinear in z , the chief difficulty lies in evaluating the integral associated with the expectation operator in Problem (2). This paper therefore uses Monte Carlo methods to calculate a consistent and asymptotically normal estimate of Ψ_0 . The key idea is to simulate the behavior of the endogenous state variables for different values of the decision rule parameters Ψ . These simulations can be used to approximate the value of the expectation in Problem (2).

2.3 The Algorithm

Proceed as follows. Using a (pseudo)random number generator, generate n i.i.d. sequences $z^{(i)}$, $i = 1, \dots, n$, where each $z^{(i)} = \{z_t^{(i)}\}_{t=1}^T$ has joint density $f(\cdot|z_0)$. For each of the sequences $z^{(i)}$, evaluate $g(z^{(i)}, \Psi)$. Define $Q_n(\Theta, \Psi) \equiv \sum_{i=1}^n g(z^{(i)}, \Psi)$, where $\Theta \equiv \{z^{(i)}\}_{i=1}^n$ is the collection of the n sequences $z^{(i)}$ of exogenous state vectors. $n^{-1} Q_n(\Theta, \Psi)$ is an estimate of the expected return $E g(z, \Psi)$ associated with decision rule parameters Ψ . More precisely, by the strong law of large numbers, $n^{-1} Q_n(\Theta, \Psi)$ converges almost surely to $E g(z, \Psi)$ as the number of independent simulations n increases to infinity.

To compute a consistent estimate $\hat{\Psi}_n$ of Ψ_0 , hold Θ fixed and vary Ψ so as to maximize the estimated expected return $n^{-1} Q_n(\Theta, \Psi)$. Formally,

$$\hat{\Psi}_n = \underset{\Psi}{\operatorname{argmax}} n^{-1} Q_n(\Theta, \Psi). \quad (3)$$

Since Θ is held fixed, Problem (3) is a well-defined deterministic optimization problem. Standard hillclimbing methods can be employed to compute $\hat{\Psi}_n$ numerically. As shown in Section 3, the numerical error associated with the estimate $\hat{\Psi}_n$ can be made as small as desired by choosing n appropriately.

3. ASYMPTOTIC RESULTS

This section states a set of assumptions on the real-valued function $g(z, \Psi)$ (defined in Section 2) under which $\hat{\Psi}_n$ is consistent for Ψ_0 as the number of independent simulations n increases to infinity. This section also develops the asymptotic distribution of $n^{1/2} (\hat{\Psi}_n - \Psi_0)$. A related set of asymptotic results permits the consistent estimation of $E g(z, \Psi_0)$, the optimal value of the problem given a functional form for the decision rule.

Let the $k \times 1$ vector of decision rule parameters Ψ belong to C , a compact subset of \mathbb{R}^k . As in Section 2, let $\Theta = \{z^{(i)}\}_{i=1}^n$ be an i.i.d. sequence of $T \times 1$ random vectors $z^{(i)}$, each with joint density $f(\cdot | z_0)$.

Definition 1 identifies a class of real-valued functions $d(z^{(i)}, \Psi)$ to which uniform laws of large numbers can be applied (this is Definition 1 of Tauchen (1985)).

Definition 1 $d(z^{(i)}, \Psi)$ is said to be *regular* if:

(a) $d(z^{(i)}, \Psi)$ is measurable in $z^{(i)}$ for all $\Psi \in C$.

(b) d is separable (see Huber (1967)).⁷

(c) $d(z^{(i)}, \Psi)$ is dominated (i.e. there exists a real-valued function $b(z^{(i)})$ such that $\int b(z^{(i)}) f(z^{(i)} | z_0) dz^{(i)} < \infty$ and $|d(z^{(i)}, \Psi)| \leq b(z^{(i)})$ for all $\Psi \in C$).

(d) $d(z^{(i)}, \Psi)$ is continuous in Ψ for all $z^{(i)}$.

If $d(\cdot, \cdot)$ is regular, then $E d(z^{(i)}, \Psi) = \int d(z^{(i)}, \Psi) f(z^{(i)} | z_0) dz^{(i)}$ exists and is continuous in Ψ . Moreover, $n^{-1} \sum_{i=1}^n d(z^{(i)}, \Psi)$ converges (in n) almost surely uniformly in $\Psi \in C$ to $E d(z^{(i)}, \Psi)$. (See Tauchen (1985) for proofs of these results.) These results underly the proofs of the asymptotic results presented below.

Assumptions 1-4 place the required structure on $g(z^{(i)}, \Psi)$:

Assumption 1 $g(z^{(i)}, \Psi)$ is twice continuously differentiable in Ψ for all $z^{(i)}$.

Assumption 2

(a) $g(z^{(i)}, \Psi)$ and its first and second partial derivatives (with respect to Ψ) are regular functions.

(b) The functions $\frac{\partial g(z^{(i)}, \cdot)}{\partial \Psi_j} \frac{\partial g(z^{(i)}, \cdot)}{\partial \Psi_\ell}$, $j, \ell = 1, \dots, k$, are regular.

(c) The function $(g(z^{(i)}, \Psi))^2$ is regular.

Assumption 3 $E g(z^{(i)}, \Psi)$ (whose existence is guaranteed by Assumption 2(a)) is uniquely maximized at Ψ_0 , an interior point of C .

Assumption 4 $E \left[\frac{\partial^2 g(z^{(i)}, \Psi_0)}{\partial \Psi \partial \Psi'} \right]$ (whose existence is guaranteed by

Assumption 2(a)) is invertible.

Assumption 1 requires that the return function $r(x_t, x_{t+1}, z_t)$ be twice continuously differentiable in its first two arguments (this is a stronger assumption than the one made in Section 2), and that the decision rule $h(x_t, z_t; \Psi)$ be twice continuously differentiable in x_t and Ψ . Consider the first partial derivative of $g(z, \Psi)$ with respect to Ψ :

$$\frac{\partial g(z, \cdot)}{\partial \Psi} = \sum_{t=0}^T \beta^t \left[\frac{\partial x_t}{\partial \Psi} \frac{\partial r(\cdot, x_{t+1}, z_t)}{\partial x_t} + \frac{\partial x_{t+1}}{\partial \Psi} \frac{\partial r(x_t, \cdot, z_t)}{\partial x_{t+1}} \right], \quad (4)$$

where the $p \times k$ Jacobian matrices $\frac{\partial x_t}{\partial \Psi'}$, $t = 0, \dots, T+1$, can be defined

recursively using the decision rule $x_{t+1} = h(x_t, z_t; \Psi)$. For $t = 0, \dots, T$,

$$\frac{\partial x_{t+1}}{\partial \Psi'} = \frac{\partial h(x_t, z_t; \cdot)}{\partial \Psi'} + \frac{\partial x_t}{\partial \Psi'} \frac{\partial h(\cdot, x_{t+1}; \Psi)}{\partial x_t}, \quad (5)$$

where $\frac{\partial x_0}{\partial \Psi'} = 0$.

The recursive equation (5) captures two effects of varying the decision rule parameters Ψ . The first term on the right hand side of (5) states that, given x_t , a change in Ψ will lead to a direct change in x_{t+1} . The second term states that a change in Ψ will also lead to a different choice for today's endogenous state x_t (since x_t also depends on Ψ), thereby indirectly affecting tomorrow's endogenous state vector x_{t+1} .

Equations (9) and (10) show that the existence of $\frac{\partial g(z, \cdot)}{\partial \Psi}$ hinges on the

existence of $\frac{\partial r(\cdot, x_{t+1}, z_t)}{\partial x_t}$, $\frac{\partial r(x_t, \cdot, z_t)}{\partial x_{t+1}}$, $\frac{\partial h(x_t, z_t; \cdot)}{\partial \Psi'}$, and $\frac{\partial h(\cdot, z_t; \Psi)}{\partial x_t}$.

Expressions analogous to (4) and (5) can be derived for the second partial derivatives of $g(z, \Psi)$ with respect to Ψ .

In the context of Problem (1), the conditions under which the optimal decision rule $w(x_t, z_t)$ is twice continuously differentiable with respect to x_t are unknown. The asymptotic theory underlying the proposed solution algorithm, however, requires only that the rule of thumb $h(x_t, z_t; \Psi)$ used to approximate w be twice continuously differentiable with respect to x_t (as well as Ψ). Assumptions 1-4 place no restrictions on w itself.

As in Section 2, let $Q_n(\Theta, \Psi) = \sum_{i=1}^n g(z^{(i)}, \Psi)$ and define

$$\hat{\Psi}_n = \underset{\Psi}{\operatorname{argmax}} n^{-1} Q_n(\Theta, \Psi).$$

Proposition 1 Under Assumptions 1-4, the following results hold:

$$\operatorname{plim}_{n \rightarrow \infty} \hat{\Psi}_n = \Psi_0 \tag{6}$$

$$n^{1/2} (\hat{\Psi}_n - \Psi_0) \rightarrow N(0, V(\Psi_0)), \tag{7}$$

where $V(\Psi_0) = A(\Psi_0)^{-1} B(\Psi_0) A(\Psi_0)^{-1}$, $A(\Psi_0) = E \left[\frac{\partial^2 g(z^{(i)}, \Psi_0)}{\partial \Psi \partial \Psi'} \right]$, and

$$B(\Psi_0) = E \left[\frac{\partial g(z^{(i)}, \Psi_0)}{\partial \Psi} \frac{\partial g(z^{(i)}, \Psi_0)}{\partial \Psi'} \right].$$

Proof: See the Appendix.

Proposition 1 states that $\hat{\Psi}_n$ is a consistent estimate of the optimal decision rule parameters Ψ_0 . Proposition 1 also states that $n^{1/2} (\hat{\Psi}_n - \Psi_0)$

is asymptotically normally distributed, with an asymptotic covariance matrix that is defined in terms of the first and second partial derivatives of $g(z^{(i)}, \Psi)$ evaluated at Ψ_0 . The distributional result (7) can be used to calculate standard errors for the elements of $\hat{\Psi}_n$ since, for n large, $\hat{\Psi}_n$ is approximately normally distributed with mean Ψ_0 and covariance matrix $n^{-1} V(\Psi_0)$. The estimate $\hat{\Psi}_n$ can therefore be made arbitrarily precise by a suitable increase in n .

Proposition 2 Under Assumptions 1-4, the following results hold:

$$\text{plim}_{n \rightarrow \infty} n^{-1} Q_n(\Theta, \hat{\Psi}_n) = E g(z^{(i)}, \Psi_0) \quad (8)$$

$$n^{1/2} (n^{-1} Q_n(\Theta, \hat{\Psi}_n) - E g(z^{(i)}, \Psi_0)) \rightarrow N(0, \text{var } g(z^{(i)}, \Psi_0)) \quad (9)$$

Proof: See the Appendix.

Proposition 2 states that $n^{-1} Q_n(\Theta, \hat{\Psi}_n)$ is a consistent estimate of $E g(z^{(i)}, \Psi_0)$, the optimal value of the problem given a functional form for the decision rule.

Given the number n of independent simulations, the distributional result (9) can be used to construct confidence intervals for $E g(z^{(i)}, \Psi_0)$. These confidence intervals provide a way to assess the accuracy of $\hat{\Psi}_n$ as an estimate of Ψ_0 , thereby providing some guidance in the choice of n . In particular, for n large, the following statement holds:

$$\Pr (E g(z^{(i)}, \Psi_0) > n^{-1} Q_n(\Theta, \hat{\Psi}_n) + s_\nu (n^{-1/2} \hat{\sigma}_n)) = \nu, \quad (10)$$

where $\hat{\sigma}_n^2$ is a consistent estimate of $\text{var } g(z^{(i)}, \Psi_0)$ and $\Pr (w > s_\nu) = \nu$ for $w \sim N(0, 1)$.

In words, (10) states that the agent is $100(1-\nu)\%$ confident that the welfare loss from estimating $E g(z^{(i)}, \Psi_0)$ using a finite number of simulations n is no greater than $s_\nu (n^{-1/2} \hat{\sigma}_n)$. The upper bound $s_\nu (n^{-1/2} \hat{\sigma}_n)$ can be made arbitrarily small simply by increasing n . Given a significance level ν , the number of independent simulations n is therefore sufficiently "large" if $s_\nu (n^{-1/2} \hat{\sigma}_n)$ is sufficiently "small". The meaning of "small" depends, of course, on the nature of the optimization problem. Section 5 suggests an operational meaning of "small" for a numerical example drawn from the stochastic optimal growth literature.

Proposition 3 shows how to construct consistent estimates of the asymptotic covariance matrix in (7) and the asymptotic variance in (9).

Proposition 3 Define $\hat{\sigma}_n^2 \equiv n^{-1} \sum_{i=1}^n (g(z^{(i)}, \hat{\Psi}_n))^2 - (n^{-1} Q_n(\Theta, \hat{\Psi}_n))^2$,

$\hat{A}_n(\hat{\Psi}_n) \equiv n^{-1} \frac{\partial^2 Q_n(\Theta, \hat{\Psi}_n)}{\partial \Psi \partial \Psi'}$, and $\hat{B}_n(\hat{\Psi}_n) \equiv n^{-1} \sum_{i=1}^n \frac{\partial g(z^{(i)}, \hat{\Psi}_n)}{\partial \Psi} \frac{\partial g(z^{(i)}, \hat{\Psi}_n)}{\partial \Psi'}$. Then,

under Assumptions 1-4, $\text{plim}_{n \rightarrow \infty} \hat{A}_n(\hat{\Psi}_n)^{-1} \hat{B}_n(\hat{\Psi}_n) \hat{A}_n(\hat{\Psi}_n)^{-1} = V(\Psi_0)$ and

$\text{plim}_{n \rightarrow \infty} \hat{\sigma}_n^2 = \text{var } g(z^{(i)}, \Psi_0)$.

Proof: See the Appendix.

Calculating $\hat{\Psi}_n$ numerically (using gradient hillclimbing methods) requires at the very least the calculation of the gradient $\frac{\partial g(z, \cdot)}{\partial \Psi}$ for different values of Ψ . Equations (4) and (5) can be used recursively to compute this gradient exactly. Alternatively, $\frac{\partial g(z, \cdot)}{\partial \Psi}$ can be calculated

using a numerical approximation. Similarly, the second partial derivatives necessary for the calculation of $\hat{A}_n(\hat{\Psi}_n)$ can be computed either numerically or analytically by developing the counterparts to equations (4) and (5).

4. ANTITHETIC ACCELERATION

This section modifies the solution algorithm described in Sections 2 and 3 to incorporate antithetic variates. This modification leads to significant increases in accuracy for a given number of simulations n .

Rather than generate n i.i.d. sequences $z^{(i)}$ (each with joint density $f(\cdot|z_0)$), generate instead $n/2$ i.i.d. antithetic pairs $(z^{(i)}, \bar{z}^{(i)})$, where $\bar{z}^{(i)} = F^{-1}(1 - F(z^{(i)}))$ and $F(\cdot)$ is the cumulative distribution function of $z^{(i)}$ (conditional on z_0). It is easy to show that $E G(z^{(i)}) = E G(\bar{z}^{(i)})$ for any function $G(\cdot)$, provided the latter expectation exists. If $f(\cdot|z_0)$ is symmetric about the mean vector $\mu_z = \int z f(z|z_0) dz$, then $\bar{z}^{(i)}$ simply equals $2\mu_z - z^{(i)}$.

Redefine $Q_n(\Theta, \Psi) \equiv \sum_{i=1}^{n/2} [g(z^{(i)}, \Psi) + g(\bar{z}^{(i)}, \Psi)]$, where

$\Theta = \{(z^{(i)}, \bar{z}^{(i)})\}_{i=1}^{n/2}$ in this case. As before, let $\hat{\Psi}_n = \underset{\Psi}{\operatorname{argmax}} n^{-1} Q_n(\Theta, \Psi)$.

Note that, although there are only $n/2$ antithetic pairs $(z^{(i)}, \bar{z}^{(i)})$, the evaluation of $Q_n(\Theta, \Psi)$ still requires n evaluations of the function $g(\cdot, \Psi)$.

By the strong law of large numbers, for any $\Psi \in C$,

$$\operatorname{plim}_{n \rightarrow \infty} n^{-1} Q_n(\Theta, \Psi) = \frac{1}{2} E [g(z^{(i)}, \Psi) + g(\bar{z}^{(i)}, \Psi)] = E g(z^{(i)}, \Psi). \quad (11)$$

By virtue of (11), when $Q_n(\Theta, \Psi)$ is evaluated using antithetic variates, the consistency results (6) and (8) continue to hold under Assumptions 1-4. The distributional results (7) and (9), however, must be modified. Proposition 4

summarizes these results.

Proposition 4 Under Assumptions 1-4, when $Q_n(\Theta, \Psi)$ is evaluated using antithetic variates, $\text{plim}_{n \rightarrow \infty} \hat{\Psi}_n = \Psi_0$ and $\text{plim}_{n \rightarrow \infty} n^{-1} Q_n(\Theta, \hat{\Psi}_n) = E g(z^{(i)}, \Psi_0)$.

Moreover,

$$n^{1/2} (\hat{\Psi}_n - \Psi_0) \rightarrow N \left(0, A(\Psi_0)^{-1} \left[B(\Psi_0) + \frac{1}{2} (C(\Psi_0) + C(\Psi_0)') \right] A(\Psi_0)^{-1} \right) \quad (12)$$

and

$$n^{1/2} (n^{-1} Q_n(\Theta, \hat{\Psi}_n) - E g(z^{(i)}, \Psi_0)) \rightarrow N \left(0, \text{var } g(z^{(i)}, \Psi_0) + \text{cov} [g(z^{(i)}, \Psi_0), g(\bar{z}^{(i)}, \Psi_0)] \right), \quad (13)$$

where $C(\Psi_0) = E \left[\frac{\partial g(z^{(i)}, \Psi_0)}{\partial \Psi} \frac{\partial g(\bar{z}^{(i)}, \Psi_0)}{\partial \Psi'} \right]$. ($A(\Psi_0)$ and $B(\Psi_0)$ are defined in

Proposition 1.)

Proof: See the Appendix.

Comparing (13) to (9), antithetic acceleration reduces the asymptotic standard error of $n^{1/2} (n^{-1} Q_n(\Theta, \hat{\Psi}_n) - E g(z^{(i)}, \Psi_0))$ to the extent that $g(z^{(i)}, \Psi_0)$ and $g(\bar{z}^{(i)}, \Psi_0)$ are negatively correlated. This correlation in turn is negative to the extent that $g(z^{(i)}, \Psi_0)$ is linear in $z^{(i)}$ and $f(\cdot | z_0)$ is symmetric about μ_z .⁸ In many applications, both of these conditions are met: the transition density $v(z_t | z_{t-1})$ specified by the economic model is symmetric about $E(z_t | z_{t-1})$ and the return function $g(z^{(i)}, \cdot)$ is nearly linear in $z^{(i)}$.

Similarly, comparing (12) to (7), antithetic acceleration leads to a smaller (in the matrix sense) asymptotic covariance matrix for $n^{1/2} (\hat{\Psi}_n - \Psi_0)$ to the extent that $(C(\Psi_0) + C(\Psi_0)')$ is negative definite.

Using antithetic variates, consistent estimates of the asymptotic covariance matrix in (12) and the asymptotic variance in (13) can be constructed as follows. Let $g_i(\Psi) = g(z^{(i)}, \Psi)$ and let $\bar{g}_i(\Psi) = g(\bar{z}^{(i)}, \Psi)$. Then, letting ' \xrightarrow{P} ' denote convergence in probability:

$$n^{-1} \sum_{i=1}^{n/2} \left[(g_i(\hat{\Psi}_n))^2 + (\bar{g}_i(\hat{\Psi}_n))^2 \right] + (n/2)^{-1} \sum_{i=1}^{n/2} g_i(\hat{\Psi}_n) \bar{g}_i(\hat{\Psi}_n) - 2 \left(n^{-1} Q_n(\theta, \hat{\Psi}_n) \right)^2 \xrightarrow{P} \text{var } g_i(\Psi_0) + \text{cov} [g_i(\Psi_0), \bar{g}_i(\Psi_0)]$$

$$n^{-1} \frac{\partial^2 Q_n(\theta, \hat{\Psi}_n)}{\partial \Psi \partial \Psi'} \xrightarrow{P} A(\Psi_0)$$

$$n^{-1} \sum_{i=1}^{n/2} \left[\frac{\partial g_i(\hat{\Psi}_n)}{\partial \Psi} \frac{\partial g_i(\hat{\Psi}_n)}{\partial \Psi'} + \frac{\partial \bar{g}_i(\hat{\Psi}_n)}{\partial \Psi} \frac{\partial \bar{g}_i(\hat{\Psi}_n)}{\partial \Psi'} \right] \xrightarrow{P} B(\Psi_0)$$

$$(n/2)^{-1} \sum_{i=1}^{n/2} \frac{\partial g_i(\hat{\Psi}_n)}{\partial \Psi} \frac{\partial \bar{g}_i(\hat{\Psi}_n)}{\partial \Psi'} \xrightarrow{P} C(\Psi_0) .$$

The proof of Proposition 3 contains examples of the arguments required to prove the above results.

5. A NUMERICAL EXAMPLE

This section implements the solution algorithm developed in Sections 2-4 for the classic stochastic dynamic optimal growth problem. In particular, this section uses the proposed solution method to construct approximate solutions to a variant of a problem studied by Brock and Mirman (1972).

A social planner maximizes a representative agent's expected lifetime utility of consumption subject to technology constraints:

$$\max_{\{x_t\}_{t=1}^{\infty}} E_0 \sum_{t=0}^{\infty} \beta^t \log(y_t - x_{t+1}) \quad \text{given } x_0, z_0 \quad (14)$$

subject to (for all t):

$$y_t = A x_t^{\alpha} \exp(z_t)$$

$$x_{t+1} \in (0, y_t)$$

$$z_{t+1} = \rho z_t + \varepsilon_{t+1}$$

$$\varepsilon_{t+1} \sim \text{IIDN}(0, \sigma_{\varepsilon}^2)$$

Note that the transition density $v(z_t | z_{t-1})$ is normal with mean ρz_{t-1} and variance σ_{ε}^2 . The parameters α , β , A , and ρ satisfy: $A > 0$, $\alpha, \beta \in (0, 1)$, and $|\rho| < 1$. x_t is the capital stock at the beginning of period t , $\exp(z_t)$ is the period t technology shock, and y_t is period t output. The capital stock depreciates fully in every period, so period t consumption is given by $y_t - x_{t+1}$. The state of the system in period t is characterized by the pair $s_t = (x_t, z_t)$. At every point in time, the social planner optimally chooses next period's capital stock, x_{t+1} , as a function of the current state s_t .

The optimal decision rule for this problem possesses a known closed form solution (see Sargent (1987)): $x_{t+1} = A \alpha \beta x_t^{\alpha} \exp(z_t)$. For this problem, therefore, optimal behavior can be compared to rule-of-thumb behavior. Three rules of thumb are considered:

Rule 1 is *quadratic* in the state vector:

$x_{t+1} = b_0 + b_1 x_t + b_2 z_t + b_3 x_t^2 + b_4 z_t^2 + b_5 x_t z_t$, where the decision rule parameters b_0, \dots, b_5 are chosen optimally.

Rule 2 is *linear* in the state vector: $x_{t+1} = a_0 + a_1 x_t + a_2 z_t$, where

the decision rule parameters a_0 , a_1 , and a_2 are chosen optimally.

Rule 3 is a "partial adjustment of the capital stock" rule:

$x_{t+1} = (1-\lambda) x_t + \lambda x_t^*$, where the decision rule parameter $\lambda \in (0,1)$ is chosen optimally and $x_t^* = (A \alpha \beta \exp(z_t))^{1/(1-\alpha)}$; i.e. x_t^* is the optimal steady state capital stock if the shock z_t were to remain indefinitely at its current level. Under Rule 3, next period's capital stock x_{t+1} is a weighted average of today's capital stock x_t and a (myopically calculated) "desired" level of the capital stock x_t^* .

For the numerical results reported below, the model parameters are set at the following values:⁹

A	β	α	ρ	$\frac{\sigma_\varepsilon}{\varepsilon}$
0.25	0.98	0.33	0.95	0.04

The infinite horizon is truncated at $T = 800$. The initial condition x_0 is set at the deterministic steady state value of x_t (i.e. $(A \alpha \beta)^{1/(1-\alpha)}$); the initial condition z_0 is set at the unconditional mean of z_t (i.e. 0).

For each of the rules of thumb, the objective function $Q_n(\theta, \cdot)$ is evaluated using the same set $\Theta = \{(z^{(i)}, \bar{z}^{(i)})\}_{i=1}^{100}$ of one hundred antithetic pairs, where $z^{(i)} = \{z_t^{(i)}\}_{t=1}^T$ and $\bar{z}^{(i)} = \{\bar{z}_t^{(i)}\}_{t=1}^T$ (hence $n = 200$ in this case). To generate $z^{(i)}$, generate a (pseudo)random sequence $\{\varepsilon_t^{(i)}\}_{t=1}^T$, where $\varepsilon_t^{(i)} \sim \text{IIDN}(0, \sigma_\varepsilon^2)$ for all i and t , and iterate on the recursive equation $z_{t+1}^{(i)} = \rho z_t^{(i)} + \varepsilon_{t+1}^{(i)}$, beginning with $t = 0$. To generate $\bar{z}^{(i)}$, simply reverse the signs on the $\varepsilon_t^{(i)}$'s and iterate on the same equation:

$$\bar{z}_{t+1}^{(i)} = \rho \bar{z}_t^{(i)} - \varepsilon_{t+1}^{(i)}.$$

Consistent estimates of the optimal decision rule parameters for Rules

1-3 (with consistently estimated standard errors in parentheses) are tabulated in Table 1.¹⁰

Rules 1-3 can be compared to the optimal rule along several dimensions. Table 2 compares the various decision rules according to a welfare metric. That is, Table 2 contains consistent estimates of the optimal value of the problem (i.e. $n^{-1} Q_n(\theta, \hat{\psi}_n)$) under Rules 1-3 and under the optimal rule.¹¹ Consistently estimated standard errors are in parentheses.

Table 2 shows that, from a welfare perspective, the three rules of thumb are nearly indistinguishable from the optimal rule. For example, the welfare loss from using the optimal linear rule rather than the truly optimal rule is equivalent to losing only 0.0102% of per period consumption (uniformly across all periods of the planning horizon).¹² The corresponding losses for the optimal quadratic rule and the "partial adjustment of the capital stock" rule are, respectively, $8.3 \times 10^{-5}\%$ and $2.0 \times 10^{-5}\%$.

The figures in Table 2 can be used to assess the amount of approximation error associated with the use of 100 antithetic pairs of sequences to evaluate $Q_n(\cdot, \cdot)$. For example, using (10), the agent (social planner) is 95% confident that the welfare loss (in expected utility terms) from using a finite number of simulations to estimate the optimal set of linear decision rule parameters is no greater than $1.645 \times 5.803 \times 10^{-6} = 9.546 \times 10^{-6}$. This welfare loss is equivalent to losing only 0.0029% of per period consumption. From an economic point of view, approximation error is therefore very small. Indeed, according to the metric underlying (10), a value of n smaller than 200 (say $n = 50$, or 25 antithetic pairs) would suffice to obtain very accurate estimates of the optimal decision rule parameters.

The usefulness of antithetic acceleration in reducing standard errors for a given number of simulations n can be evaluated by calculating the

sample correlation between $g(z^{(i)}, \hat{\Psi}_n)$ and $g(\bar{z}^{(i)}, \hat{\Psi}_n)$. For all four decision rules, this correlation is very close to -1, suggesting that, despite the apparent nonlinearity in Problem (14), $g(z^{(i)}, \hat{\Psi}_n)$ is very nearly linear in $z^{(i)}$, regardless of the choice of decision rule. It is precisely in this case that antithetic acceleration proves most valuable.¹³

Table 3 compares the four decision rules from the perspective of the time series behavior of output, consumption, and the capital stock. Using each of the decision rules, simulated time series with 25,000 observations were generated for each of output, consumption, and the capital stock. For all four decision rules, these time series were computed using the same set of shocks $\{\varepsilon_t\}_{t=1}^{25000}$. Table 3 contains a set of summary statistics for the simulated time series. The figures in Table 3 show that, with the exception of the capital stock series under the linear rule, the four decision rules yield time series with very similar unconditional first and second moments.

A final way to compare the four decision rules is to compare their choices for x_{t+1} given different states of the system $s_t = (x_t, z_t)$. Given optimal behavior, some states are more likely to occur than others. The set of states at which to compare the different decision rules is therefore determined by simulating a time series $\{s_t\}_{t=1}^S$ using the optimal decision rule, where $S = 100,000$. Let $x_{t+1}^j(s_t)$ be the capital stock chosen by Rule j given s_t and let $\tilde{x}_{t+1}(s_t)$ be the optimal choice for the capital stock given

s_t . Then $d_j(s_t) = 100 \left| \frac{x_{t+1}^j(s_t) - \tilde{x}_{t+1}(s_t)}{\tilde{x}_{t+1}(s_t)} \right|$ is the percentage deviation

between the optimal decision and the decision made by Rule j , given s_t . For each of the rules of thumb ($j = 1, 2, 3$), Table 4 computes the sample mean of the series $\{d_j(s_t)\}_{t=1}^S$ and the percentage of states in the sequence $\{s_t\}_{t=1}^S$

for which $d_j(s_t)$ is less than, respectively, 10, 5, 1, 0.5, 0.1 and 0.01.

The ranking of the three rules of thumb using the behavioral metric underlying the figures in Table 4 accords with the welfare ranking in Table 2: the "partial capital stock adjustment" rule tracks optimal behavior most closely, followed by, in order, the optimal quadratic rule and the optimal linear rule. The "partial capital stock adjustment" rule and the optimal quadratic rule yield capital stock decisions that deviate only rarely by more than 0.5% from optimal behavior.

The results of this section show that, in the context of Problem (14), parsimoniously parameterized rules of thumb can mimic optimal behavior very closely, according to a variety of metrics.¹⁴ Indeed, a surprising finding is that a one-parameter family of decision rules (the "partial capital stock adjustment" rule) outperforms a six-parameter family (the optimal quadratic rule) along all dimensions considered. The success of the "partial capital stock adjustment" rule shows that rules of thumb which incorporate some of the economic structure underlying the optimization problem can perform better than "brute force" polynomial expansions.

6. CONCLUSION

This paper develops a new method for solving discrete time, discounted stochastic dynamic programming problems with infinite horizons and convex choice sets. The key idea of the solution method is to restrict the decision rule to belong to a parametric class of functions. The agent then chooses from within this class the best decision rule; that is, the agent maximizes the objective function subject to the functional form restriction on the decision rule by choosing the optimal set of decision rule parameters.

This paper shows how Monte Carlo methods can be used to calculate

arbitrarily precise estimates of the optimal decision rule parameters, given a functional form for the decision rule. The central idea of the Monte Carlo methods is to simulate the dynamic behavior of the state variables for different values of the decision rule parameters. These simulations can be used to approximate the value of the expectation that appears in the agent's objective function. The incorporation of a simulation technique known as antithetic acceleration increases greatly the practical usefulness of the solution method by providing large increases in accuracy for a given number of simulations.

Section 5 of this paper uses the solution method to solve a stochastic dynamic programming problem with a known analytical solution: the stochastic optimal growth model with logarithmic preferences and full depreciation of the capital stock. This paper finds that relatively simple rules of thumb can provide very good approximations to optimal behavior, according to both behavioral and welfare metrics. Indeed, a carefully chosen one-parameter family of decision rules outperforms a six-parameter family of second-order polynomial expansions. The rules of thumb considered in Section 5 perform well despite the fact that the standard deviation of the shock is four times the value usually specified in business cycle models.

Future research will explore the performance of rules of thumb in a variety of more complicated environments. This line of research seeks to identify economic environments in which relatively simple rules of thumb do or do not provide good approximations to optimal behavior. A related goal of this line of research is to isolate features of the economic environment that have a significant influence on the ability of rules of thumb to track optimal behavior.

For example, Smith (1990) uses the solution method to solve a version of

Problem (14) with more general preferences and with a depreciation rate less than 1. Although the optimal decision rule for this problem does not possess a closed form expression, it appears that in this environment rules of thumb based on second-order polynomial expansions perform quite well. However, the performance of simple rules of thumb deteriorates as the shock becomes more volatile and as the coefficient of relative risk aversion increases.

Future research will apply the solution method to more complicated versions of the stochastic growth model, including models with both exogenous and endogenous growth. Since the solution method relies on independent simulations of the model's behavior, nonstationarities in any individual simulation do not pose a problem.

The solution method as developed here focuses on infinite horizon problems. Another item for future research is to extend the method to handle problems with finite horizons, perhaps by explicitly incorporating time as an additional exogenous state variable. This extension appears to be straightforward.

A final item for future research is to adapt the solution method to handle problems with nonconvex choice sets, in particular, discrete choice problems. For these problems, the smoothness conditions specified by Assumptions 1 and 2 in Section 3 will not be satisfied. Moreover, calculating an estimate of the optimal decision rule parameters will require the numerical optimization of a nonsmooth function, so that conventional gradient methods cannot be applied. Tackling discrete choice problems will therefore entail the development of a more general asymptotic theory than that contained in Section 3, as well as the use of more sophisticated numerical hillclimbing methods. The problems posed by these extensions, though not insurmountable, are nonetheless nontrivial.

FOOTNOTES

1. For problems with quadratic objective functions and linear constraints, exact linear decision rules can be computed using the principle of certainty equivalence. Some non-linear-quadratic problems can be solved analytically, for instance, a simple version of the stochastic growth model with logarithmic utility and full depreciation of the capital stock in every period. See Chapter 1 of Sargent (1987) and Section 5 of this paper.

2. After completing this work, I discovered that Peris (1982) uses a similar approach to solve a dynamic model of the cyclical behavior of employment with unemployment insurance. Peris' dissertation, however, does not develop the asymptotic theory that underlies and justifies the approach (while at the same time providing a way to assess simulation error). Also, Peris' implementation of the approach does not incorporate antithetic acceleration, a variance reduction technique that greatly increases the usefulness of the approach by providing significant increases in accuracy.

3. This section is motivated by Geweke (1988), in which Geweke coins the phrase "antithetic acceleration".

4. The boundedness assumption can be relaxed by making an appropriate set of alternative assumptions. The proof of the existence and uniqueness of a stationary decision rule also requires a technical condition on the transition density $v(z_{t+1}|z_t)$. See Chapter 9 of Stokey, Lucas, and Prescott (1989) for further details.

5. Care must be taken that the decision rule $x_{t+1} = h(x_t, z_t; \Psi)$ does not violate the constraint $x_{t+1} \in \Gamma(x_t, z_t)$ for any $s_t = (x_t, z_t) \in S$. For some problems (such as the one studied in Section 5), this requirement is of little practical import since the probability of attaining a state $s_t \in A$, where $A = \{s_t \in S : h(x_t, z_t; \Psi) \notin \Gamma(x_t, z_t)\}$, is negligible when Ψ is near the optimal decision rule parameters Ψ_0 (as defined by Problem (2)). For other problems (such as those with occasionally binding inequality constraints), the decision rule must embed the constraint $x_{t+1} \in \Gamma(x_t, z_t)$.

6. Note that the infinite horizon in Problem (1) has been truncated at some large T . The solution to the infinite horizon problem can be approximated arbitrarily closely by increasing T suitably.

7. As discussed in Tauchen (1985), conditions (a) and (b) are "weak and essentially non-restrictive side conditions" (p. 422).

8. As an extreme example, let $w \sim N(\mu, \sigma^2)$ and suppose we want to estimate $E G(w)$ using Monte Carlo methods, where $G(w) = a + b w$. Let $(w^{(1)}, \bar{w}^{(1)})$, where $\bar{w}^{(1)} = 2\mu - w^{(1)}$, be an antithetic pair drawn from the density of w . Then $\frac{1}{2} [G(w^{(1)}) + G(\bar{w}^{(1)})] = a + b \mu = E G(w)$, i.e. antithetic

acceleration eliminates all error in the estimate of $E G(w)$ with only one draw of an antithetic pair.

9. In the real business cycle literature, a more conventional setting for σ_ε is 0.01. The larger value of σ_ε used here is intended to draw out the nonlinearities in Problem (14), thereby providing a more stringent test of the ability of rules of thumb to track optimal behavior.

10. The optimal decision rule parameters for the quadratic rule are not well identified. In particular, several other sets of parameter choices for the quadratic rule were found, all of which yielded the same value of the objective function to 8 decimal places. The time series implied by these alternative sets of parameter choices are essentially indistinguishable. The failure of identification is therefore not a weakness of the solution method, but reflects instead the (very) small amount of nonlinearity in the problem being studied. Because of the lack of identification, the estimated standard errors for the quadratic rule parameters should be treated with caution.

11. As for Rules 1-3, the optimal value of the problem under the optimal rule is estimated using Monte Carlo methods. In this case, the optimal decision rule parameters Ψ_0 are known with certainty: $x_{t+1} = c_0 x_t^{c_1} (\exp(z_t))^{c_2}$, with $c_0 = A\alpha\beta$, $c_1 = \alpha$, and $c_2 = 1$. The optimal value of the problem under the optimal rule could also be computed analytically. Given initial conditions x_0 and z_0 , this value equals $a + b \log(x_0) + c z_0$, where $a =$

$$\frac{1}{1-\beta} \left[\frac{\log(A)}{1-\alpha\beta} + \log(1-\alpha\beta) + \frac{\alpha\beta}{1-\alpha\beta} \log(\alpha\beta) \right], \quad b = \frac{\alpha}{1-\alpha\beta}, \quad \text{and} \quad c = \frac{1}{(1-\alpha\beta)(1-\beta\rho)}.$$

Using the initial conditions specified in the text, the optimal value of the problem under the optimal rule is therefore -1.0052597939. (As noted in Table 2, expected utility estimates are scaled by a factor of 150.) The discrepancy between this number and the number in Table 2 for the optimal rule is due to the use of a finite horizon of length $T = 800$. If T is increased to 1600, then the "true" optimal value of the problem under the optimal rule and its estimate in Table 2 will agree to (at least) 10 decimal places.

12. Differences in expected utility can be expressed in terms of consumption by answering the following question: How much more consumption, in terms of a uniform percentage increase across all periods of the planning horizon, would an agent using Rule i need in order to achieve the level of expected utility associated with Rule j ? Let U_k ($k = i, j$) be expected utility under Rule k . Suppose $\Delta = U_j - U_i > 0$. Solve for λ in the following equation (the scaling factor q is set at 150):

$$E_0 \sum_{t=0}^{\infty} \beta^t q^{-1} \ln(c_{it}^*) + \Delta = E_0 \sum_{t=0}^{\infty} \beta^t q^{-1} \ln[(1+\lambda) c_{it}^*], \quad \text{where } \{c_{it}^*\}_{t=0}^{\infty} \text{ is}$$

the optimal consumption sequence under Rule i . The solution is

$$\lambda = \exp[q (1-\beta) \Delta] - 1.$$

13. Comparing (13) to (9), it can be seen that antithetic acceleration reduces the asymptotic standard error of the estimate of the optimal value of the problem (given a decision rule) by a factor of

$(1 + \text{cor} [g(z^{(i)}, \Psi_0), g(\bar{z}^{(i)}, \Psi_0)])^{-1/2}$. For the linear rule, this factor equals approximately 377. For the other rules, the correlation of $g(z^{(i)}, \Psi_0)$ and $g(\bar{z}^{(i)}, \Psi_0)$ is so close to -1 that this factor cannot be computed accurately on a digital computer.

14. These results echo those of Christiano (1990), in the context of a more general version of Problem (14), who finds that linear and log-linear decision rules provide very accurate approximations to optimal behavior, both from a behavioral and a welfare perspective.

APPENDIX

Proof of Proposition 1:

First to show consistency of $\hat{\Psi}_n$. I will verify that Assumptions (A)-(C) of Theorem 4.1.1 in Amemiya (1985) hold. The parameter space C is compact by assumption. $Q_n(\Theta, \Psi)$ is continuous in Ψ for all Θ and is a measurable function of Θ for all Ψ by Assumption 2. Using the result stated in the text after Definition 1, Assumption 2 implies that $n^{-1} Q_n(\Theta, \Psi)$ converges almost surely uniformly in Ψ (and hence in probability uniformly in Ψ) to $E g(z^{(i)}, \Psi)$. Moreover, the nonstochastic function $E g(z^{(i)}, \Psi)$ is continuous in Ψ . Finally, by Assumption 3, $E g(z^{(i)}, \Psi)$ is uniquely maximized at Ψ_0 . Thus, by Theorem 4.1.1 in Amemiya (1985), $\text{plim}_{n \rightarrow \infty} \hat{\Psi}_n = \Psi_0$.

Now to show the asymptotic normality of $n^{1/2} (\hat{\Psi}_n - \Psi_0)$. By definition, $\hat{\Psi}_n$ satisfies $\frac{\partial Q_n(\Theta, \hat{\Psi}_n)}{\partial \Psi} = 0$. Expanding the left hand side of this equation in a first-order Taylor series about Ψ_0 yields:

$$\frac{\partial Q_n(\Theta, \Psi_0)}{\partial \Psi} + \frac{\partial^2 Q_n(\Theta, \Psi_n^*)}{\partial \Psi \partial \Psi'} (\hat{\Psi}_n - \Psi_0) = 0, \quad (A1)$$

where Ψ_n^* lies on the line joining Ψ_0 and $\hat{\Psi}_n$. (To be precise, Ψ_n^* should vary

from row to row of $\frac{\partial Q_n(\Theta, \cdot)}{\partial \Psi \partial \Psi'}$, but this subtlety makes no difference

asymptotically.) Rewriting (A1) gives:

$$n^{1/2} (\hat{\Psi}_n - \Psi_0) = - \left[n^{-1} \sum_{i=1}^n \frac{\partial^2 g(z^{(i)}, \Psi_n^*)}{\partial \Psi \partial \Psi'} \right]^{-1} n^{-1/2} \sum_{i=1}^n \frac{\partial g(z^{(i)}, \Psi_0)}{\partial \Psi}. \quad (A2)$$

By Assumption 2 (using the result stated in the text after Definition

1), $n^{-1} \sum_{i=1}^n \frac{\partial^2 g(z^{(i)}, \Psi)}{\partial \Psi \partial \Psi'}$ converges in probability uniformly in Ψ to

$$E \left[\frac{\partial^2 g(z^{(i)}, \Psi)}{\partial \Psi \partial \Psi'} \right], \text{ and, moreover } E \left[\frac{\partial^2 g(z^{(i)}, \Psi)}{\partial \Psi \partial \Psi'} \right] \text{ is continuous in } \Psi \text{ (and,}$$

in particular, at Ψ_0). Since $\text{plim } \hat{\Psi}_n = \Psi_0$ and Ψ_n^* lies on the line joining Ψ_0 and $\hat{\Psi}_n$, $\text{plim } \Psi_n^* = \Psi_0$. Thus, by Theorem 4.1.5 in Amemiya (1985),

$$\text{plim } n^{-1} \sum_{i=1}^n \frac{\partial^2 g(z^{(i)}, \Psi_n^*)}{\partial \Psi \partial \Psi'} = E \left[\frac{\partial^2 g(z^{(i)}, \Psi_0)}{\partial \Psi \partial \Psi'} \right] = A(\Psi_0), \text{ which is invertible}$$

by Assumption 4. (This result implies that for n sufficiently large,

$$n^{-1} \sum_{i=1}^n \frac{\partial^2 g(z^{(i)}, \Psi_n^*)}{\partial \Psi \partial \Psi'} \text{ is invertible almost surely.)}$$

Now to work out the asymptotic distribution of $n^{-1/2} \sum_{i=1}^n \frac{\partial g(z^{(i)}, \Psi_0)}{\partial \Psi}$.

$$\text{By the strong law of large numbers, } \text{plim } n^{-1} \sum_{i=1}^n \frac{\partial g(z^{(i)}, \Psi_0)}{\partial \Psi} = E \frac{\partial g(z^{(i)}, \Psi_0)}{\partial \Psi}$$

(the existence of this expectation is guaranteed by Assumption 2). By Corollary 5.9 of Bartle (1966), Assumption 2 implies that the integration and differentiation operators in the forgoing expression can be interchanged, so

$$\text{that } E \frac{\partial g(z^{(i)}, \Psi_0)}{\partial \Psi} = \frac{\partial}{\partial \Psi} E g(z^{(i)}, \Psi_0) = 0 \text{ since, by Assumption 3, } \Psi_0$$

maximizes $E g(z^{(i)}, \Psi)$ and is an interior point of C . By Assumption 2(b),

$$E \left[\frac{\partial g(z^{(i)}, \Psi_0)}{\partial \Psi} \frac{\partial g(z^{(i)}, \Psi_0)}{\partial \Psi'} \right] = B(\Psi_0) \text{ exists and is finite. Then, by the}$$

$$\text{Lindbery-Lévy central limit theorem, } n^{-1/2} \sum_{i=1}^n \frac{\partial g(z^{(i)}, \Psi_0)}{\partial \Psi} =$$

$$n^{1/2} \left[n^{-1} \sum_{i=1}^n \frac{\partial g(z^{(i)}, \Psi_0)}{\partial \Psi} - E \frac{\partial g(z^{(i)}, \Psi_0)}{\partial \Psi} \right] \rightarrow N(0, B(\Psi_0)).$$

Applying

Slutsky's theorem to (A2) yields the desired result:

$$n^{1/2} (\hat{\Psi}_n - \Psi_0) \rightarrow N(0, A(\Psi_0)^{-1} B(\Psi_0) A(\Psi_0)^{-1}).$$

Proof of Proposition 2:

We want to show the consistency and asymptotic normality of the estimate $n^{-1} Q_n(\Theta, \hat{\Psi}_n)$ of the optimized value of $E g(z^{(i)}, \Psi)$ (i.e. $E g(z^{(i)}, \Psi_0)$).

Again using the result stated in the text after Definition 1,

$\text{plim } n^{-1} Q_n(\Theta, \Psi) = E g(z^{(i)}, \Psi)$ uniformly in Ψ and, moreover, $E g(z^{(i)}, \Psi)$ is continuous at Ψ_0 . Since $\text{plim } \hat{\Psi}_n = \Psi_0$, we have by Theorem 4.1.5 in Amemiya

(1985) that $\text{plim } n^{-1} Q_n(\Theta, \hat{\Psi}_n) = E g(z^{(i)}, \Psi_0)$.

Expand $n^{-1} Q_n(\Theta, \hat{\Psi}_n)$ in a first-order Taylor series about Ψ_0 :

$$n^{-1} Q_n(\Theta, \hat{\Psi}_n) = n^{-1} Q_n(\Theta, \Psi_0) + n^{-1} \frac{\partial Q_n(\Theta, \Psi_n^*)}{\partial \Psi} (\hat{\Psi}_n - \Psi_0), \quad (\text{A3})$$

where Ψ_n^* lies on the line joining Ψ_0 and $\hat{\Psi}_n$. Rewrite (A3) as:

$$n^{1/2} [n^{-1} Q_n(\Theta, \hat{\Psi}_n) - E g(z^{(i)}, \Psi_0)] = n^{1/2} [n^{-1} Q_n(\Theta, \Psi_0) - E g(z^{(i)}, \Psi_0)] + \left[n^{-1} \sum_{i=1}^n \frac{\partial g(z^{(i)}, \Psi_n^*)}{\partial \Psi} \right] n^{1/2} (\hat{\Psi}_n - \Psi_0). \quad (\text{A4})$$

Assumption 2 implies that $n^{-1} \sum_{i=1}^n \frac{\partial g(z^{(i)}, \Psi)}{\partial \Psi}$ converges in probability

uniformly in Ψ to $E \frac{\partial g(z^{(i)}, \Psi)}{\partial \Psi}$, and, moreover, $E \frac{\partial g(z^{(i)}, \Psi)}{\partial \Psi}$ is continuous at

Ψ_0 . Since $\text{plim } \hat{\Psi}_n = \Psi_0$, $\text{plim } \hat{\Psi}_n^* = \Psi_0$. Thus $\text{plim } n^{-1} \sum_{i=1}^n \frac{\partial g(z^{(i)}, \hat{\Psi}_n)}{\partial \Psi} =$

$E \frac{\partial g(z^{(i)}, \Psi_0)}{\partial \Psi} = 0$. $n^{1/2} (\hat{\Psi}_n - \Psi_0)$ converges in distribution, so by Slutsky's

theorem the second term on the right hand side of (A4) converges in probability to 0. The asymptotic distribution of

$n^{1/2} [n^{-1} Q_n(\Theta, \hat{\Psi}_n) - E g(z^{(i)}, \Psi_0)]$ is therefore identical to that of

$n^{1/2} [n^{-1} Q_n(\Theta, \Psi_0) - E g(z^{(i)}, \Psi_0)]$. Assumption 2(d) guarantees that

$E [g(z^{(i)}, \Psi_0)]^2$ exists and is finite, implying in turn that $\text{var } g(z^{(i)}, \Psi_0)$

exists and is finite. By the Lindberg-Lévy central limit theorem,

$n^{1/2} [n^{-1} Q_n(\Theta, \Psi_0) - E g(z^{(i)}, \Psi_0)] \rightarrow N(0, \text{var } g(z^{(i)}, \Psi_0))$.

Proof of Proposition 3:

By Assumption 2(c), $(g(z^{(i)}, \Psi))^2$ is a regular function.

$n^{-1} \sum_{i=1}^n (g(z^{(i)}, \Psi))^2$ therefore converges in probability uniformly in Ψ to

$E (g(z^{(i)}, \Psi))^2$, a continuous function of Ψ . Since $\text{plim } \hat{\Psi}_n = \Psi_0$,

$\text{plim } n^{-1} \sum_{i=1}^n (g(z^{(i)}, \hat{\Psi}_n))^2 = E (g(z^{(i)}, \Psi_0))^2$. Using arguments given

above, $\text{plim } (n^{-1} Q_n(\Theta, \hat{\Psi}_n))^2 = (E g(z^{(i)}, \Psi_0))^2$. Hence $\text{plim } \hat{\sigma}_n^2 =$

$E (g(z^{(i)}, \Psi_0))^2 - (E g(z^{(i)}, \Psi_0))^2 = \text{var } g(z^{(i)}, \Psi_0)$.

Using arguments given above, it is straightforward to show that

$\text{plim } \hat{A}_n(\hat{\Psi}_n) = A(\Psi_0)$. Finally, by Assumption 2(b),

$n^{-1} \sum_{i=1}^n \frac{\partial g(z^{(i)}, \Psi)}{\partial \Psi} \frac{\partial g(z^{(i)}, \Psi)}{\partial \Psi'}$ converges in probability uniformly in Ψ to

$E \left[\frac{g(z^{(i)}, \Psi)}{\partial \Psi} \frac{\partial g(z^{(i)}, \Psi)}{\partial \Psi'} \right]$, a continuous function of Ψ . Since $\text{plim } \hat{\Psi}_n = \Psi_0$,

$\hat{B}_n(\hat{\Psi}_n) = n^{-1} \sum_{i=1}^n \frac{\partial g(z^{(i)}, \hat{\Psi}_n)}{\partial \Psi} \frac{\partial g(z^{(i)}, \hat{\Psi}_n)}{\partial \Psi'}$ therefore converges in probability

to $B(\Psi_0)$ (by Theorem 4.1.5 in Amemiya (1985)). Hence

$$\text{plim } \hat{A}_n(\hat{\Psi}_n)^{-1} \hat{B}_n(\hat{\Psi}_n) \hat{A}_n(\hat{\Psi}_n)^{-1} = A(\Psi_0)^{-1} B(\Psi_0) A(\Psi_0)^{-1} = V(\Psi_0).$$

Proof of Proposition 4:

Under antithetic acceleration, $n^{-1} Q_n(\theta, \Psi) = (n/2)^{-1} \sum_{i=1}^{n/2} h(z^{(i)}, \Psi)$,

where $h(z^{(i)}, \Psi) = \frac{1}{2} [g(z^{(i)}, \Psi) + g(\bar{z}^{(i)}, \Psi)]$ and $(z^{(i)}, \bar{z}^{(i)})$ is an antithetic pair of sequences. I will show that if $g(z^{(i)}, \Psi)$ satisfies Assumptions 1-4, then $h(z^{(i)}, \Psi)$ does too. Hence, the asymptotic results (6)-(9) continue to hold, with $h(z^{(i)}, \Psi)$ playing the role of $g(z^{(i)}, \Psi)$, and $n/2$ playing the role of n . (Under antithetic acceleration, convergence is therefore in the number of antithetic pairs $n/2$.)

Clearly, $h(z^{(i)}, \Psi)$ satisfies Assumption 1 so long as $g(z^{(i)}, \Psi)$ does. Suppose $b(z^{(i)})$ dominates $g(z^{(i)}, \Psi)$. Then, by definition, $|g(\bar{z}^{(i)}, \Psi)| \leq b(\bar{z}^{(i)})$. Since $E b(z^{(i)}) = E b(\bar{z}^{(i)})$, the existence of $E b(z^{(i)}, \Psi)$ implies the existence of $E b(\bar{z}^{(i)})$. Hence $b(z^{(i)}) + b(\bar{z}^{(i)})$ dominates $h(z^{(i)}, \Psi)$.

Similar arguments apply to the remaining functions specified by Assumption 2. Thus $h(z^{(i)}, \Psi)$ satisfies Assumption 2 if $g(z^{(i)}, \Psi)$ does. Next, note that $E h(z^{(i)}, \Psi) = E g(z^{(i)}, \Psi)$ for all $\Psi \in C$, which implies that $\Psi_0 =$

$$\underset{\Psi}{\text{argmax}} E g(z^{(i)}, \Psi) = \underset{\Psi}{\text{argmax}} E h(z^{(i)}, \Psi). \text{ Finally, since } E \left[\frac{\partial^2 h(z^{(i)}, \Psi_0)}{\partial \Psi \partial \Psi'} \right] =$$

$$\frac{1}{2} E \left[\frac{\partial^2 g(z^{(i)}, \Psi_0)}{\partial \Psi \partial \Psi'} \right] + \frac{1}{2} E \left[\frac{\partial^2 g(\bar{z}^{(i)}, \Psi_0)}{\partial \Psi \partial \Psi'} \right] = E \left[\frac{\partial^2 g(z^{(i)}, \Psi_0)}{\partial \Psi \partial \Psi'} \right],$$

$h(z^{(i)}, \Psi)$ satisfies Assumption 4 if $g(z^{(i)}, \Psi)$ does.

Let $\tilde{A}(\Psi_0)$ and $\tilde{B}(\Psi_0)$ be the counterparts to $A(\Psi_0)$ and $B(\Psi_0)$ under antithetic acceleration. To simplify notation, let $g_i(\Psi) = g(z^{(i)}, \Psi)$ and

$$\bar{g}_i(\Psi) = g(\bar{z}^{(i)}, \Psi). \text{ As shown above, } \tilde{A}(\Psi_0) = E \left[\frac{\partial^2 h(z^{(i)}, \Psi_0)}{\partial \Psi \partial \Psi'} \right] = A(\Psi_0). \text{ On}$$

$$\text{the other hand, } \tilde{B}(\Psi_0) = E \left[\frac{\partial h(z^{(i)}, \Psi_0)}{\partial \Psi} \frac{\partial h(z^{(i)}, \Psi_0)}{\partial \Psi'} \right] =$$

$$E \left[\frac{1}{2} \left[\frac{\partial g_i(\Psi_0)}{\partial \Psi} + \frac{\partial \bar{g}_i(\Psi_0)}{\partial \Psi} \right] \frac{1}{2} \left[\frac{\partial g_i(\Psi_0)}{\partial \Psi'} + \frac{\partial \bar{g}_i(\Psi_0)}{\partial \Psi'} \right] \right] = \frac{1}{2} B(\Psi_0) +$$

$\frac{1}{4} [C(\Psi_0) + C(\Psi_0)']$, where $C(\Psi_0)$ is defined in the text. The asymptotic

covariance matrix of $(n/2)^{1/2} (\hat{\Psi}_n - \Psi_0)$ under antithetic acceleration is

$\tilde{A}(\Psi_0)^{-1} \tilde{B}(\Psi_0) \tilde{A}(\Psi_0)^{-1}$. The asymptotic covariance matrix of $n^{1/2} (\hat{\Psi}_n - \Psi_0)$ is

therefore $2 \tilde{A}(\Psi_0)^{-1} \tilde{B}(\Psi_0) \tilde{A}(\Psi_0)^{-1}$, which yields the expression in equation

(12) of the text.

Finally, to work out the asymptotic variance of

$(n/2)^{1/2} [n^{-1} Q_n(\Theta, \hat{\Psi}_n) - E g(z^{(i)}, \Psi_0)]$ under antithetic acceleration. To

wit, $\text{var } h(z^{(i)}, \Psi_0) = \frac{1}{4} (\text{var } g_i(\Psi_0) + \text{var } \bar{g}_i(\Psi_0)$

$+ 2 \text{cov}[g_i(\Psi_0), \bar{g}_i(\Psi_0)]) = \frac{1}{2} \text{var } g_i(\Psi_0) + \frac{1}{2} \text{cov}[g_i(\Psi_0), \bar{g}_i(\Psi_0)]$.

Doubling this expression yields the expression in equation (13) of the text.

REFERENCES

- Amemiya, Takeshi, 1985, Advanced econometrics (Harvard University Press, Cambridge, MA).
- Bartle, Robert G., 1966, The elements of integration (Wiley, New York).
- Brock, William A. and Leonard J. Mirman, 1972, Optimal economic growth and uncertainty: the discounted case, *Journal of Economic Theory* 4, 497-513.
- Christiano, Lawrence J., 1990, Linear-quadratic approximation and value-function iteration, *Journal of Business and Economic Statistics* 8, 99-113.
- Geweke, John, 1988, Antithetic acceleration of Monte Carlo integration in Bayesian inference, *Journal of Econometrics* 38, 73-90.
- Huber, Peter J., 1967, The behavior of maximum likelihood estimates under nonstandard conditions, in: Lucien M. LeCam and Jerzy Neyman, eds., *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, Vol. 1 (University of California Press, Berkeley, CA), 221-234.
- Peris, Jose Antonio, 1982, A dynamic model of the cyclical behavior of employment with unemployment insurance, Ph.D. dissertation (University of Minnesota).
- Sargent, Thomas J., 1987, *Dynamic macroeconomic theory* (Harvard University Press, Cambridge, MA).
- Smith, Jr., Anthony A., 1990, Three essays on the solution and estimation of dynamic macroeconomic models, Ph.D. dissertation (Duke University).
- Stokey, Nancy L. and Robert E. Lucas, Jr. (with Edward C. Prescott), 1989, *Recursive methods in economic dynamics*, (Harvard University Press, Cambridge, MA).
- Tauchen, George, 1985, Diagnostic testing and evaluation of maximum likelihood models, *Journal of Econometrics* 30, 415-443.

TABLE 1

Estimated Optimal Decision Rule Parameters

(Estimated standard errors in parentheses)

	b_0	b_1	b_2	b_3	b_4	b_5
Rule 1 [†]	0.01566 (0.00001)	0.3325 (0.0003)	0.02056 (0.00032)	-0.03874 (0.00084)	0.01266 (0.00050)	0.1315 (0.0132)
		a_0	a_1	a_2		
Rule 2		0.01607 (0.00007)	0.3237 (0.0026)	0.02227 (0.00010)		
			λ			
Rule 3			0.67070 (0.00002)			

[†]See footnote 10 concerning identification of the optimal parameters for the quadratic rule.

TABLE 2

Welfare Under Different Decision Rules

(Estimated standard errors in parentheses)

	Estimated Expected Lifetime Utility	Welfare Loss [‡] Relative to Optimal Rule
Optimal Rule	-1.005259698 [†] (*)	0.0%
Rule 3 (Partial capital stock adjustment)	-1.005259766 (0.000000005)	$2.0 \times 10^{-5}\%$
Rule 1 (Optimal quadratic)	-1.005259976 (0.000000055)	$8.3 \times 10^{-5}\%$
Rule 2 (Optimal linear)	-1.005293800 (0.000005803)	0.0102%

[†]Expected utility estimates are scaled by a factor of 150.

[‡]Measured as a percentage of per period consumption (see footnote 12).

*The estimated standard error of this estimate is less than 0.5×10^{-10} .

TABLE 3

Time Series Statistics Using Different Decision Rules
 (based on simulated time series with 25,000 observations)

	<u>Sample Mean of Output</u>	<u>Sample Standard Deviation of Output</u>	<u>Coefficient of Variation (Mean/Std. Dev.)</u>
Optimal Rule	0.073229	0.013484	0.18414
Rule 1	0.073230	0.013478	0.18405
Rule 2	0.073172	0.013140	0.17958
Rule 3	0.073245	0.013488	0.18414

	<u>Sample Mean of Consumption</u>	<u>Sample Standard Deviation of Consumption</u>	<u>Coefficient of Variation (Mean/Std. Dev.)</u>
Optimal Rule	0.049547	0.0091233	0.18414
Rule 1	0.049546	0.0091227	0.18412
Rule 2	0.049542	0.0091286	0.18426
Rule 3	0.049547	0.0091236	0.18414

	<u>Sample Mean of Capital Stock</u>	<u>Sample Standard Deviation of Capital Stock</u>	<u>Coefficient of Variation (Mean/Std. Dev.)</u>
Optimal Rule	0.023682	0.0043607	0.18414
Rule 1	0.023683	0.0043555	0.18391
Rule 2	0.023630	0.0040363	0.17081
Rule 3	0.023698	0.0043640	0.18415

TABLE 4

Deviations Between Optimal Behavior and Rule-of-Thumb Behavior
 (based on simulated time series with 100,000 observations)

	Optimal Quadratic Rule (Rule 1)	Optimal Linear Rule (Rule 2)	Partial Capital Stock Adjust- ment Rule (Rule 3)
Mean deviation (in percent) between Rule j and the optimal rule	0.1006%	1.176%	0.0443%
Percentage of deviations smaller than:			
10.00%	100.0%	99.8%	100.0%
5.00%	100.0% [†]	97.8%	100.0%
1.00%	99.7%	51.9%	100.0%
0.50%	98.6%	21.8%	99.9%
0.10%	64.8%	4.1%	86.7%
0.01%	7.1%	0.4%	36.8%

[†]All but two (2) of the deviations were smaller than 5%.