Federal Reserve Bank of Minneapolis
Research Department

Forthcoming in Journal of
Business and Economic Statistics
Volume 8, Number 1, January 1990

LINEAR QUADRATIC APPROXIMATION AND VALUE FUNCTION ITERATION: A COMPARISON

Lawrence J. Christiano*

Working Paper 415

Revised September 1989

ABSTRACT

This article studies the accuracy of two versions of Kydland and Prescott's (1980, 1982) procedure for approximating optimal decision rules in problems in which the objective fails to be quadratic and the constraints fail to be linear. The analysis is carried out using a version of the Brock–Mirman (1972) model of optimal economic growth. Although the model is not linear quadratic, its solution can nevertheless be computed with arbitrary accuracy using a variant of existing value–function iteration procedures. I find the Kydland–Prescott approximate decision rules are very similar to those implied by value–function iteration.

KEYWORDS: Decision rule; Growth model; Markov chain; Optimization; Production function; State space.

*Federal Reserve Bank of Minneapolis. Prepared for presentation to the National Bureau of Economic Research's Nonlinear Rational Expectations Modeling Study Group, Stanford University, February 1988. I have benefited from the advice and suggestions of Terry Fitzgerald and two anonymous referees. I am also grateful to Inga Velde for her editorial assistance.

1. INTRODUCTION

The linear–quadratic (LQ) method proposed by Kydland and Prescott (1980,1982) for approximating the solution to non–LQ optimization problems has been applied in numerous studies (Altug 1986; Christiano 1987c, 1988; Christiano and Eichenbaum 1988; Cooley and Hansen 1988; Hansen 1985; Hansen and Sargent 1988; and King, Plosser, and Rebelo 1988). In general, little is known about the accuracy of this method. This article provides some evidence of its accuracy within the context of a particular example. This is done by comparing the LQ approximate solutions with the solutions obtained by discretizing the underlying state space and applying a variant of the value–function iteration methods described by Bertsekas (1987). Since the grid for the endogenous variables in the state space is very fine, I expect that the solution obtained by value–function iteration approximates very closely the solution in the version of the problem in which the endogenous state variables take on a continuum of values.

The example used here is a version of the Brock–Mirman (1972) one–sector stochastic growth model. A solution to the model is a set of two decision rules. These relate end–of–period capital and current consumption to the current–period state variables. There are two state variables: beginning–of–period capital and the current period's technology shock, which is a realization from a stationary stochastic process. Two versions of the LQ approximation method are studied: The first relates end–of–period capital linearly to the state variables and is called the linear LQ approximation. This is the original method used by Kydland and Prescott (1980, 1982). The second approximation used is log linear in end–of–period capital and the state variables and is therefore called the log–linear

LQ approximation. It is applied by Christiano (1987b,c; 1988) and King, Plosser, and Rebelo (1988). This method has the virtue that in the special case in which the model does admit an analytic solution, then the log–linear decision rules and the exact decision rules coincide (see Remark 1). This case is the model studied by Long and Plosser (1983) in which the depreciation rate on capital is 100%, the production function is Cobb–Douglas, and utility is logarithmic in consumption.

The accuracy of the approximate decision rules is evaluated on four dimensions. First, I compare the LQ decisions with those of the (approximately) exact solution at selected points in the state space. Second, I compare the LQ decision rules' implications for several first and second moments with those of the exact solution. This comparison is particularly relevant, since first and second moments play an important role at the parameter selection and model evaluation stages, respectively, for many who use LQ approximations. Third, I report the amount, expressed as a fraction of the initial stock, that a planner who only knows the LQ decision rule would be willing to pay to learn the exact decision rule. This is a measure of how close to optimal the LQ decision rules are. Finally, graphs of the steady–state distribution of consumption and capital, as implied by all three solutions, are presented.

The results suggest that the LQ approximation is remarkably accurate for the example at hand. In addition, they show that the log–linear and linear LQ approximations are roughly equally accurate. This latter finding illustrates that the relative accuracy of the two decision rules is context specific. For example, in Christiano (1987a,b) and Christiano (1988, n. 18), I show that in a model similar to the one here but with a productivity shock which is a logarithmic random walk, the two approximations are dramatically different. Christiano (1987b) shows that the difference reflects the accuracy of the log–linear approximation and the very poor

accuracy of the linear approximation in that context.

The plan of the paper is as follows. In Section 2, I present the growth model that is studied and its LQ approximate solution. Section 3 formulates the optimization problem as a dynamic programming problem and discusses its solution by value–function iteration. Section 4 describes the model parameters used in the experiments. Section 5 presents the comparison of the LQ and value–function iteration solutions. Section 6 presents my conclusions.

2. THE PROBLEM AND ITS SOLUTION BY LQ APPROXIMATION

The problem I consider is the one good growth model in which the planner maximizes

$$E_0 \Sigma_{t=0}^{\infty} (1-\tau)^{-1} C_t^{(1-\tau)} \beta^t \tag{2.1}$$

subject to

$$C_t + K_t - (1-\delta)K_{t-1} = \exp(x_t)K_{t-1}^{\alpha}. \tag{2.2}$$

Here, C_t and K_t are date–t consumption and the end–of–period–t capital stock. Also, δ, α, and τ are the rate of depreciation on capital, the share of income due to capital, and risk aversion. Throughout, I assume that x_t is a realization from an n_x state, first–order Markov chain with

$$x_t \in \mathcal{X} = \{\mathcal{X}_1, \ldots, \mathcal{X}_{n_x}\} \tag{2.3}$$
$$\Pr\{x_{t+1} = \mathcal{X}_j \mid x_t = \mathcal{X}_i\} = \pi_{ij}$$
$$Ex_t = 0.$$

It is sometimes convenient to refer to the model for x_t in terms of its Wold representation, which has a first–order autoregressive form:

$$x_t = \rho x_{t-1} + \epsilon_t, \tag{2.4}$$

where ϵ_t is mean zero with variance σ_ϵ^2 and is uncorrelated with x_{t-1}. Further details about n_x, $\pi = [\pi_{ij}]$ and $\mathcal{X}$ are given below.

A solution to this problem is a function relating the date–t decision, K_t, to the date–t information variables, x_t and K_{t-1}. The exact solution is known only for the case $\tau = \delta = 1$ (see Long and Plosser 1983). I now describe two variants of the LQ approximation method proposed by Kydland and Prescott (1980, 1982). The first of these, the linear LQ method, approximates the decision rule for K_t with one that is linear in x_t and K_{t-1}. The second, the log–linear LQ method, approximates it with one in which the log of K_t is related linearly to x_t and the log of K_{t-1}. The decision rule delivered by the log–linear LQ method has the virtue of coinciding with the exact decision rule when $\delta = \tau = 1$.

2.1 The Linear LQ Approximation

Express the problem as a standard calculus–of–variations problem by substituting (2.2) into (2.1):

$$\text{maximize } E_0 \, \Sigma_{t=0}^{\infty} \beta^t u(K_{t-1}, K_t, x_t) \tag{2.5}$$

subject to x_0 and K_{-1} given. Here,

$$u(K, K', x) = \frac{1}{1-\tau}[\exp(x)K^{\alpha} + (1-\delta)K - K']^{(1-\tau)}. \tag{2.6}$$

The linear LQ method approximates the solution to (2.5) by the solution to the following linear–quadratic optimization problem:

$$\text{maximize } E_0 \, \Sigma_{t=0}^{\infty} \beta^t U(K_{t-1}, K_t, x_t), \tag{2.7}$$

where U is the second–order Taylor series expansion of u about $K_{t-1} = K_t = K^*$ and $x_t = x^*$. Here, K^* and x^* are the steady–state values of K_t and x_t of the nonstochastic version of (2.5) obtained by setting $\epsilon_t = 0$ for all t. Trivially, $x^* = 0$. Also, it is easy to verify that

$$K^* = \left\{\frac{\beta\alpha\exp(x^*)}{1-(1-\delta)\beta}\right\}^{[1/(1-\alpha)]}. \tag{2.8}$$

It is convenient to define $\tilde{K}_t = K_t - K^*$. At date t, the first–order necessary condition for $\tilde{K}_t$ to solve (2.7) is

$$E_t\tilde{K}_{t+1} - \phi\,\tilde{K}_t + \frac{1}{\beta}\tilde{K}_{t-1} = -(q/\beta)x_t. \tag{2.9}$$

Here,

$$\phi = -\frac{u_{22} + \beta u_{11}}{\beta u_{12}} = 1 + \beta^{-1} + \left[\frac{1-\alpha}{\tau}\right][1 - (1-\delta)\beta](C^*/K^*), \tag{2.10}$$

where C^*/K^* is the steady–state consumption–to–capital ratio, given by

$$C^*/K^* = \frac{\beta^{-1} - 1 + \delta(1-\alpha)}{\alpha}. \tag{2.11}$$

In (2.10), u_{ij} is cross–derivative of u with respect to its i^{th} and j^{th} arguments, evaluated at steady state. It can also be shown that

$$\begin{aligned} q &= \frac{u_{23} + \rho\beta u_{13}}{u_{12}} \\ &= \beta\left\{(1-\rho)\left[\frac{C^*}{K^*} + \delta\right] + \frac{\rho\beta}{\tau}\left[\beta^{-1} - 1 + \delta\right]\frac{C^*}{K^*}\right\} K^*. \end{aligned} \tag{2.12}$$

Let λ be the unique number that satisfies $|\lambda| \leq 1$ and $\lambda^2 - \phi\lambda + (1/\beta) = 0$. Then, the unique solution to (2.7) is the rule

$$\tilde{K}_t = \lambda \tilde{K}_{t-1} + q\,\frac{\lambda}{1-\beta\rho\lambda}\,x_t,$$

or,

$$\begin{aligned} K_t &= (1-\lambda)K^* + \lambda K_{t-1} + q\,\frac{\lambda}{1-\beta\rho\lambda}\,x_t \\ &= f_{linLQ}(K_{t-1}, x_t), \end{aligned} \tag{2.13}$$

say, where $f_{linLQ}(K, x) \equiv (1-\lambda)K^{*} + \lambda K + [q\lambda/(1-\beta\rho\lambda)]x$. Substituting this decision rule into (2.2), we get the linear LQ approximate decision rule for consumption, g_{linLQ}:

$$C_t = g_{linLQ}(K_{t-1}, x_t)$$

$$\equiv \exp(x_t)K_{t-1}^{\alpha} + (1-\delta)K_{t-1} - f_{linLQ}(K_{t-1}, x_t). \tag{2.14}$$

This decision rule gets its name from the fact that the decision rule for K_t is linear in its arguments. Clearly, g_{linLQ} is not itself linear.

2.2 The Log–Linear LQ Approximation

Let $k_t \equiv \log(K_t)$ and define

$$r(k, k', x) \equiv u(\exp(k), \exp(k'), x), \tag{2.15}$$

where u is given in (2.6). Then, an equivalent way to write (2.5) is

$$\text{maximize } E_0 \, \Sigma_{t=0}^{\infty} \beta^t r(k_{t-1}, k_t, x_t), \tag{2.16}$$

with respect to decision rules for k_t. The log–linear LQ method approximates this decision rule by the one that solves

$$\text{maximize } E_0 \, \Sigma_{t=0}^{\infty} \beta^t R(k_{t-1}, k_t, x_t), \tag{2.17}$$

where R is the second–order Taylor series expansion of r about $k_{t-1} = k_t = \log(K^*)$ and $x_t = x^*$. Let $\tilde{k}_t \equiv k_t - k^*$. The first–order necessary condition for $\tilde{k}_t$ to solve (2.17) is

$$E_t\tilde{k}_{t+1} - \phi\tilde{k}_t + (1/\beta)\tilde{k}_{t-1} = -[q/(\beta K^*)]x_t, \tag{2.18}$$

where ϕ and q are as defined in (2.10) and (2.12), respectively. Equations (2.10) and (2.12) are relevant here despite the fact that they involve u_{ij}, since $(r_{22} + \beta r_{11})/(\beta r_{12}) = (u_{22} + \beta u_{11})/(\beta u_{12})$ and $(r_{23} + \beta\rho r_{13})/r_{12} = (u_{23} + \beta\rho u_{13})/(u_{12}K^*)$. Here, r_{ij} is the cross–derivative of r with respect to its i^{th} and j^{th} arguments, evaluated at $k_{t-1} = k_t = \log(K^*)$.

The solution to (2.17), then, is

$$\tilde{k}_t = \lambda\tilde{k}_{t-1} + (q/K^*)\frac{\lambda}{1-\beta\rho\lambda}x_t,$$

or,

$$\begin{aligned} K_t &= (K^*)^{(1-\lambda)}\exp\left[\frac{q}{K^*}\frac{\lambda}{1-\beta\rho\lambda}x_t\right]K_{t-1}^{\lambda} \\ &\equiv f_{\log LQ}(K_{t-1}, x_t). \end{aligned} \tag{2.19}$$

When $\tau = \delta = 1$, then $\lambda = \alpha$, $q/K^* = (1-\beta\rho\alpha)/\alpha$, and $(K^*)^{(1-\lambda)} = \alpha\beta$, so that (2.19) reduces to

$$K_t = \alpha\beta\exp(x_t)K_{t-1}^{\alpha}, \tag{2.20}$$

which is the exact solution to (2.5). (See Long and Plosser 1983.) This result is of sufficient interest to deserve emphasis.

Remark 1. If $\tau = \delta = 1$, then the log–linear approximation is exact.

An analogous remark applies in the context of the linear LQ decision rule. In particular, when $\tau = \delta = 1$, then (2.13) reduces to $K_t = (1-\alpha)K^* + \alpha K_{t-1} + K^* x_t$. But this is the first–order Taylor series expansion of the right side of (2.20) about $x_t = 0$ and $K_{t-1} = K^*$. Thus, we have

Remark 2. If $\tau = \delta = 1$, then the linear LQ decision rule is the first–order Taylor series expansion of the exact decision rule about the steady–state values of x_t and K_{t-1}.

Denote the log–linear LQ decision rule by f_{logLQ}:

$$K_t = f_{logLQ}(K_{t-1}, x_t) = (K^*)^{(1-\lambda)} \exp\{x_t q\lambda/[K^*(1-\beta\rho\lambda)]\}\, K_{t-1}^{\lambda}. \quad (2.21)$$

The implied log–linear decision rule for consumption is g_{logLQ}:

$$\begin{aligned} C_t &= g_{logLQ}(K_{t-1}, x_t) \\ &= \exp(x_t) K_{t-1}^{\alpha} + (1-\delta)K_{t-1} - f_{logLQ}(K_{t-1}, x_t). \end{aligned} \quad (2.22)$$

Unless $\tau = \delta = 1$, g_{logLQ} is not itself log linear.

3. THE SOLUTION BY VALUE–FUNCTION ITERATION

In the problem posed in Section 2, only the exogenous shock x_t was assumed to lie on a discrete grid $\mathcal{X}$. In particular, the capital stock was implicitly assumed to be able to take on a continuum of values. Value–function iteration methods require that the capital stock lie on a discrete grid and therefore do not, strictly speaking, apply to that problem. However, one expects that by choosing a sufficiently fine grid for K_t, denoted $\mathcal{K}$, an arbitrarily accurate approximation to the underlying continuous problem can be obtained. In the calculations here, $\mathcal{K}$ is in fact extremely fine.

Problem (2.5) is expressed as a dynamic programming problem:

$$v(K_{t-1}, x_t) = \max_{K_t \in A(K_{t-1}, x_t)} \{u(K_{t-1}, K_t, x_t) + \beta E_t v(K_t, x_{t+1})\}, \qquad (3.1)$$

subject to $\{x_t\}$ being generated by the n_x state, first–order Markov chain given in (2.3). In (3.1), **A** denotes the feasible set of possible choices of K_t. Feasibility is determined by the nonnegativity constraint on C_t and by the requirement that $K_t \in \mathcal{K}$. Formally,

$$\mathbf{A}(K, x) = \{K' \in \mathcal{K}: \exp(x)K^{\alpha} + (1-\delta)K - K' \geq 0\}. \qquad (3.2)$$

The set $\mathcal{K}$, described below, is a discrete interval of ℓ positive numbers. I now discuss the method used to find v in (3.1) and explain how v is used to solve (2.1)–(2.3), the problem of interest. The following discussion is not solely relevant

for the example here of one endogenous state variable, one control variable, and one exogenous shock. In particular, for the rest of this section, K_t and x_t can be thought of as vectors. Also, $u(\cdot,\cdot,\cdot)$ can be thought of as an indirect utility function, after control variables (such as hours worked) that directly affect current utility (but not next period's state) have been maximized out.

The date–t state variables (K_{t-1}, x_t) can only take on the $m = n_x \ell$ possible values in $(\mathcal{K} \times \mathcal{X})$, which I denote by $s_1, \ldots, s_m$. Let $v \equiv [v(s_1), \ldots, v(s_m)]' \in \mathbb{R}^m$. We can think of the expression to the right of the equality in (3.1) as defining an ordinary function T mapping points in $\mathbb{R}^m$ into $\mathbb{R}^m$. We can also think of v as the fixed point of that function, i.e., $v = T(v)$. In addition, the "max" in (3.1) refers to m maximization problems, one for each of the m values of $(K_{t-1}, x_t) \in (\mathcal{K} \times \mathcal{X})$.

There exist several iterative methods for finding the fixed point v. One of these, *standard value–function iteration*, starts with some initial $v_0 \in \mathbb{R}^m$ and then evaluates $v_{j+1} = T(v_j)$ for $j = 0, 1, 2, \ldots$. The limit of this sequence is v. (This method is also called *successive approximation* in Bertsekas 1987 and *contraction iteration* in Rust 1987, 1988a,b.) Because of this algorithm's reputation for being slow, I used an alternative, *hybrid value–function iteration*. This method is based on the function T_p, which maps $\mathbb{R}^m$ into $\mathbb{R}^m$, and where p is a strictly positive integer. For a given $v_j \in \mathbb{R}^m$, $T_p(v_j)$ is defined as follows: Let $f^j_{DP}(s) \in A(s)$ be the argmax indicated in (3.1) when v is replaced by v_j and $s \in \{s_1, \ldots, s_m\}$. Let G_j denote the m×m state transition matrix induced by π and f^j_{DP}. Specifically, the μ, ν^{th} element of G_j is the probability that $(K_t, x_{t+1}) = s_\nu$ given that $(K_{t-1}, x_t) = s_\mu$. Let $u_j = [u(s_1, f^j_{DP}(s_1)), \ldots, u(s_m, f^j_{DP}(s_m))] \in \mathbb{R}^m$. In this notation, it is easy to verify that $T(v_j) = u_j + \beta G_j v_j$. Then, $T_p(v_j) \equiv u_j + \beta G_j u_j + \ldots + (\beta G_j)^{p-1} u_j + (\beta G_j)^p v_j$. In words, $T_p(v_j)$ is the value of following the decision rule f^j_{DP} for p periods given that

the p+1 period stock of capital is valued according to $\beta^p v_j$. Hybrid value–function iteration involves evaluating $v_{j+1} = T_p(v_j)$ for $j = 0, 1, 2, \ldots$, starting from some initial $v_0 \in \mathbb{R}^m$. The limit of this sequence is tested to determine whether it is the fixed point of T that is sought. Obviously, standard value–function iteration is a limiting version of hybrid value–function iteration, with $p = 1$. The motivation for using hybrid value–function iteration is based on considering the other extreme, as $p \to \infty$: $T_\infty(v_j) = [I - \beta G_j]^{-1} u_j$. It is easy to verify that $[I - \beta G_j]^{-1} u_j = v_j + [I - T'(v_j)]^{-1}[T(v_j) - v_j]$, where $T'(v_j) = \beta G_j$ is the derivative of T with respect to v_j. Using these relations, it is easy to verify that $T_\infty(v_j)$ is the fixed point of the linear approximation of T about v_j. This is why it seems reasonable to call this method *Newton value–function iteration.* [This method appears under different names in the literature. Rust (1987, 1988a,b) calls it *Newton–Kantorovich iteration*; Bertsekas 1987 calls it *policy iteration*, which is not to be confused with a very different method having the same name used by Danthine, Donaldson, and Smith (1987). Sargent (1987, p.47) calls it *Howard policy improvement.*] Under Newton value–function iteration, one evaluates $v_{j+1} = T_\infty(v_j)$ for $j = 0, 1, 2, \ldots$, starting from some initial $v_0 \in \mathbb{R}^m$. A well–known virtue of Newton value–function iteration is that this sequence converges in a finite number of steps to v (Bertsekas 1987.) The computational difficulty with Newton value–function iteration is that it requires inverting the m×m matrix $[I - T'(v_j)]$. When $n_x = 1$ in the problem addressed here, then this matrix can be inverted very rapidly by a recursive algorithm, and Newton value–function iteration is feasible. (For another class of problems whose structure permits inverting $[I - T'(v_j)]$ rapidly by a recursive algorithm, see Rust 1987, 1988a,b.) However, I have not found a way to recursively invert this matrix when $n_x > 1$. The motivation for using hybrid value–function iteration to do the calculations here is that, roughly speaking, it approximates

Newton value–function iteration by replacing $[I - T'(v_j)]^{-1}$ by a truncated geometric sum. In several computational experiments, I found that hybrid value–function iterations with p = 10 generated a 25% reduction in central processor unit (CPU) time over the p = 1 case. This reduction is relatively modest, partly reflecting the fact that the maximization in (3.1) is relatively simple. Christiano and Fitzgerald (1989) consider a problem in which the maximization step is much more time consuming. They report experiments in which hybrid value–function iteration led to more than a tenfold reduction in CPU time over standard value–function iteration.

The function v is used to compute a decision rule for capital f_{DP}:

$$f_{DP}(K_{t-1}, x_t) = \underset{K_t \in A(K_{t-1}, x_t)}{\operatorname{argmax}} \{u(K_{t-1}, K_t, x_t) + \beta E_t v(K_t, x_{t+1})\}. \tag{3.3}$$

A decision rule for C_t can be obtained from f_{DP} by

$$g_{DP}(K_{t-1}, x_t) = \exp(x_t)K_{t-1}^{\alpha} + (1-\delta)K_{t-1} - f_{DP}(K_{t-1}, x_t). \tag{3.4}$$

The decision rules (3.3)–(3.4) solve a version of the maximization problem posed in (2.1)–(2.3), in which the constraint $K_t \in \mathcal{K}$ is imposed (Harris 1987).

I now discuss the choice of capital grid $\mathcal{K}$. Let $\underline{x}$ and $\bar{x}$ denote the smallest and largest possible values of x_t. Let $\underline{K}$ be the limit of the sequence $\{K_i\}$, where K_0 is the smallest element in $\mathcal{K}$ and $K_i = f_{DP}(K_{i-1}, \underline{x})$, i = 1, 2, ... Similarly, $\bar{K}$ is the limit of the sequence K_j defined by the condition that K_0 is the largest element in $\mathcal{K}$ and $K_j = f_{DP}(K_{j-1}, \bar{x})$, j = 1, 2, ... Then, since $f_{DP}(\cdot, x)$ is increasing in x for all examples studied, it follows that $\mathcal{K}_{DP} \equiv \{K \in \mathcal{K}: \underline{K} \leq K \leq \bar{K}\}$

forms an ergodic set relative to the DP decision rule. That is, if $K_t \in \mathcal{K}$ should begin outside this set, it then moves inside and stays there. The steady–state probability of K_t lying outside this set is zero. In a similar way, the linear and log–linear LQ decision rules also imply ergodic sets, which I label $\mathcal{E}_{\text{linLQ}}$ and $\mathcal{E}_{\text{logLQ}}$. Using (2.13), it is easily confirmed that $\mathcal{E}_{\text{linLQ}} = \{K^* + \underline{x}q\lambda/[(1-\beta\rho\lambda)(1-\lambda)], K^* + \bar{x}q\lambda/[(1-\beta\rho\lambda)(1-\lambda)]\}$. A similar calculation can be used to compute $\mathcal{E}_{\text{logLQ}}$ using (2.19).

One way to choose the grid $\mathcal{K}$ is to make its smallest (or largest) point slightly less (or greater) than the lowest (or highest) value in $\mathcal{E}_{\text{DP}}$. Since knowledge of $\mathcal{E}_{\text{DP}}$ requires the decision rules, one could proceed by first getting a rough guess of $\mathcal{E}_{\text{DP}}$ according to decision rules obtained by using a coarse grid and based on $\mathcal{E}_{\text{linLQ}}$ or $\mathcal{E}_{\text{logLQ}}$. The second–stage calculations can then be based on a very fine grid that contains few points outside $\mathcal{E}_{\text{DP}}$. This is the strategy I followed for the calculations here.

4. MODEL PARAMETERIZATION

This section reports the model parameterizations used in the experiments. I begin by describing the Markov–chain models used to model the exogenous shock.

4.1 Markov–Chain Models for $\{x_t\}$

In the experiments, I used two Markov–chain models for $\{x_t\}$: a two–state model ($n_x = 2$) and a three–state model ($n_x = 3$). The Markov chain is completely described by the state space of x_t, $\mathcal{X}$, and the transition probability matrix π.

The Two–State Markov–Chain Model for $\{x_t\}$. In this case,

$$\pi = \begin{bmatrix} \phi & 1-\phi \\ 1-\phi & \phi \end{bmatrix}, \quad \mathcal{X} = \{-\sigma, \sigma\}. \tag{4.1}$$

The Wold representation (2.4) associated with this Markov chain is

$$\rho = 2\phi - 1, \quad E[\epsilon_t | x_{t-1}] = 0, \quad E[\epsilon_t^2 | x_{t-1}] = \sigma_\epsilon^2 = \sigma^2(1-\rho^2). \tag{4.2}$$

Also, the steady–state probabilities of $x_t = -\sigma$ and $x_t = \sigma$ are each 1/2. Evidently, values for ρ and σ_ϵ^2 completely determine the parameters of the two–state Markov chain. In the experiments, I set $\sigma_\epsilon^2 = (.01)^2$ and $(.10)^2$, and $\rho = .95$. This implies σ = .032 and .32, and ϕ = .975. Prescott (1986, p. 15) argues that a value of σ_ϵ which is a little under .01 is empirically plausible. The large value of σ_ϵ (.10) was also used to see how large the shocks must be before the LQ approximation deteriorates significantly. The large value of ρ corresponds well with Prescott's (1986) empirical finding that technology shocks are highly serially correlated.

The Three–State Markov–Chain Model for $\{x_t\}$. In this case,

$$\pi = \begin{bmatrix} \phi & \gamma & 1-\phi-\gamma \\ \psi & 1-2\psi & \psi \\ 1-\phi-\gamma & \gamma & \phi \end{bmatrix}, \quad \mathcal{X} = \begin{bmatrix} -x \\ 0 \\ x \end{bmatrix}. \tag{4.3}$$

The Wold representation corresponding to this Markov chain is also (2.4) and

$$\rho = 2\phi + \gamma - 1, \ \mathrm{E}[\epsilon_t | x_{t-1}] = 0, \ \kappa = 1 + .5\gamma/\psi, \tag{4.4}$$

$$\mathrm{Var}(x_t) = x^2/\kappa, \ \sigma_\epsilon^2 = \mathrm{Var}(x_t)(1-\rho^2),$$

where $\kappa \equiv \mathrm{E}(x_t^4)/[\mathrm{E}(x_t^2)]^2$ is kurtosis. Unlike in the two–state model, ϵ_t is conditionally heteroscedastic in the three–state Markov–chain model. The steady–state probabilities for $x_t = -x$, $x_t = 0$, and $x_t = x$, are P, (1–2P), and P, where $P = (2\kappa)^{-1}$. To determine this model, values must be assigned to four parameters: ϕ, γ, ψ, and x. Thus, unlike in the two–state case, values for σ_ϵ^2 and ρ are not sufficient. To determine values for this Markov–chain example, I set σ_ϵ^2, ρ, κ, and γ. In the experiments, I set $\sigma_\epsilon^2 = (.01)^2$ and $(.10)^2$, $\rho = .95$, $\kappa = 3$, and $\gamma = .040$. These parameter values imply $\pi_{11} = \pi_{33} = .955$, $\pi_{31} = \pi_{13} = .005$, $\pi_{12} = \pi_{32} = .040$, $\pi_{21} = \pi_{23} = .010$, and $\pi_{22} = .980$ for each value of σ_ϵ^2. I set $\kappa = 3$ so that the model would resemble the normal distribution, for which $\kappa = 3$.

4.2 Other Model Parameters

I analyzed five parameterizations of the model. The first four of them set $\beta = .98$, $\tau = .5$, $\rho = .95$, $\alpha = .33$, and they incorporate one of the four Markov–chain models for x_t: the low–variance ($\sigma_\epsilon = .01$) and high–variance ($\sigma_\epsilon = .10$) two–state Markov chain and the low– and high–variance three–state Markov chain. By comparing these model results, we can judge their robustness to the number of states in the Markov chain and to the variance of the technology shock innovation. A fifth model was studied to determine how risk aversion affects the results. In that model, risk aversion is very high, with $\tau = 3$. Otherwise, the fifth

model is parameterized in the same way as the model with $\tau = .5$ and a three–state low–variance Markov–chain model for x_t.

The ergodic sets associated with each of the three capital decision rules, corresponding to each of my five models, are reported in Table 1. It is striking how similar $\mathscr{E}_{\text{logLQ}}$ and $\mathscr{E}_{\text{DP}}$ are. In contrast, $\mathscr{E}_{\text{linLQ}}$ is shifted to the left of $\mathscr{E}_{\text{DP}}$ in the high shock variance cases. Table 1 also reports the boundaries of the capital grid $\mathscr{K}$ used when solving each model by value–function iteration. In each case, $\mathscr{K}$ contains 20,000 points. The interval between grid points is reported in column (6) of Table 1. Column (7) of Table 1 reports the number of minutes of CPU time used in solving the model by value–function iteration. The time used to solve the three–state exogenous–shock models exceeds by about 50 percent the time required for the two–state exogenous–shock models. This excess reflects the fact that the number of points in the state space of the three–shock models (60,000) exceeds that in the two–shock models (40,000) by 50 percent. In each case, the value–function iterations were started with $v_0 \equiv 0$ and were considered to have converged when $\sup \{|v_j - v_{j-1}|/|v_{j-1}|\} \times 100$ was less than .000001.

5. COMPARISON OF LQ AND DP DECISION RULES

This section reports comparisons of the LQ and DP decision rules for each of the five models defined in Section 4.2 and column (1) of Table 1. There seems to be no best metric for comparing decision rules, so I use several. The first metric compares the LQ and DP decision rules directly by tabulating their values at alternative points in the state space. The second compares several first– and second–moment properties of the decision rules. The third measures the amount an agent who uses an LQ decision rule would be willing to pay to learn the DP decision

rule. (The amount is expressed as a fraction of initial capital.) This is a direct measure of how close to optimal the LQ decision rules are. The fourth compares the steady–state distribution of C_t and K_t implied by the LQ and DP decision rules.

5.1 Tabulation of Decision Rules

Tables 2 and 3 tabulate the DP and two LQ approximate decision rules at various points in the state space for the five models whose solutions were computed. The (K_{t-1}, x_t) combinations represented in Tables 2 and 3 include all possible x_t's in the relevant Markov chain and five representative K_{t-1}'s. Of these, the middle one is always K^* and the least and greatest ones are the end points of $\mathcal{E}_{DP}$, taken from Table 1. The other two points are halfway between K^* and these end points. To aid in comparing the LQ and DP decision rules, cases where they differ by between 1 and 10 percent are indicated by an asterisk (*), cases where they differ by between 10 and 20 percent are marked by a dagger (†), and cases where they differ by more than 20 percent are marked by a double dagger (††).

First, consider Table 2, which reports results for both two–state Markov–chain models. Panel A in that table shows that the DP rule and the two LQ decision rules are all approximately identical in the low–variance case. In particular, if the capital and consumption decisions are rounded to one digit after the decimal, the decisions are identical. Not surprisingly, the decision rules diverge somewhat for the high–variance two–state Markov model, whose results are reported in Panel B. The divergence is fairly minor in the case of the capital decision, where a difference exceeding 1% occurs just once. The differences are larger in the context of consumption, where the level of consumption implied by the LQ decision rules tends to overstate optimal consumption, in one case by over 20%.

It is also interesting to compare the decision rules according to whether they are increasing or decreasing in the state variables. In all cases in Table 2, the DP decision rules for C_t and K_t are increasing over the reported values of x_t and K_{t-1}. This property is shared by the LQ decision rules for K_t, a fact that can be verified analytically from the appropriate formulas in Section 2. Over the reported values of x_t, K_{t-1}, the LQ decision rules for C_t are also increasing in K_{t-1}. However, only the linear LQ decision rule for C_t is increasing in x_t. In particular, in the high–variance case, the log–linear decision rule for C_t is *decreasing* in x_t for $K_{t-1} = 86.19$ and 108.69.

Digressing slightly, note that the linear LQ approximate decision rule for K_t is not monotone in ϵ_t when $\rho = 1$ and the approximation is taken by first transforming the model so that the planner's choice variables are $c_t = C_t/\exp(x_t)$, $k_t = K_t/\exp(x_t)$. Then, even though k_t is monotone in ϵ_t, K_t is not. The log–linear LQ approximate decision rule for K_t is, by contrast, monotone in this case. (For a fuller explanation and a demonstration of the quantitative significance of these differences, see Christiano 1987a; 1988, n. 18.)

Next consider Table 3, which reports results for the three–state Markov–chain models. Basically, the same picture that emerged from Table 2 emerges here as well. In particular, for the low–variance version of the model with $\tau = .5$ (Panel A), the DP and LQ decision rules are virtually identical. As in Table 2, the differences show up in the high–variance case (Panel B), principally in the consumption decision rule. Panel C shows that the high accuracy of the LQ decision rule when $\tau = .5$ and $\sigma_\epsilon = .01$ also obtains when $\tau = 3$.

All decision rules in Table 3 are monotone increasing over the reported values of the state variables, with the exception of the LQ decision rules for C_t in the high–variance version of the model. Those rules are monotone increasing in

K_{t-1}, but they fail to be monotone increasing in x_t. Since they differ in this respect from the corresponding DP rules, this difference reflects approximation error.

A feature of the three–shock models that the two–shock models lack is that both the low– and high–variance models with $\tau = .5$ share some common points in the state space. One of these, $(K_{t-1}, x_t) = (63.69, 0.0)$, is reported in Table 3. Because their construction imposes certainty equivalence, the LQ decisions for C_t and K_t at this point is the same for both the low– and high–variance models. However, the exact problem does not satisfy certainty equivalence, so there is no reason to expect the DP rules to share this property. In fact, Table 3 indicates that $f_{DP}(63.69, 0.0)$ is 63.69 and 63.68 for the low– and high–shock models. Also, $g_{DP}(63.69, 0.0)$ is 3.94 and 3.95 in these two cases. Thus, while certainty equivalence does not hold exactly, it appears to hold approximately. This may be one of the reasons that the LQ approximations are so accurate.

In sum, the evidence in Tables 2 and 3 suggests that for reasonable shock variance (e.g., the low–variance case), the LQ approximation is very accurate, even with high risk aversion of $\tau = 3$. When the shock variances get very large, then—not surprisingly—the quality of the approximations begins to deteriorate. Based only on the evidence in Tables 2 and 3, it is hard to say which approximation—the linear or the log–linear LQ—performs better in the high–variance case. On the one hand, there is evidence that the log–linear LQ approximation performs poorly at more points in the state space. For example, there are more †'s associated with the logLQ results than with the linLQ results in Tables 2 and 3. On the other hand, there is evidence that the states in which the log–linear LQ decision rules perform worst also have lower probability than the states in which the linear LQ decision rules perform poorly. This possibility is suggested by the fact that the logLQ decision rules perform worst in states with low

K and high x and with high K and low x combinations, whereas the reverse is true for the linLQ decision rules. Given that K_{t-1} and x_t are positively correlated (which they must be, given the high positive autocorrelation of x_t), then—other things being equal—this would cause the log–linear decision rule to dominate the linear one in a weighted, overall sense.

5.2 First– and Second–Moment Implications of LQ and DP Decision Rules

Tables 4–7 report selected first– and second–moment properties of the DP and LQ decision rules, obtained by Monte Carlo simulation. I simulated 100 data sets on C_t, K_t, $Y_t \equiv C_t + K_t - K_{t-1}$, the risk–free rate of interest R_t, and the marginal product of capital $MP_{k,t}$. Each data set has length 10,050, but the first 50 observations were discarded before computing first and second moments. The risk–free rate of interest R_t is defined in the usual way as $R_t = u'(C_t)/[\beta E_t u'(C_{t+1})] - 1$, where the conditional expectation is evaluated relative to the appropriate consumption decision rule and Markov–chain probabilities, and $u(C_t) \equiv C_t^{(1-\tau)}/(1-\tau)$. The marginal product of capital is $MP_{k,t} \equiv \theta Y_{t+1}/K_t$, where $Y_t \equiv \exp(x_t)K_{t-1}^{\alpha}$ is output.

In performing the simulations, I found that the linear LQ decision rule occasionally implies a negative value for C_t. This happened only when the exogenous shocks were drawn from the three–state, high–variance Markov chain. To accommodate this, I redefined f_{linLQ} and g_{linLQ} in a way that whenever they implied a negative C_t, then C_t was set to .01 and K_t was adjusted appropriately. These redefined linear LQ decision rules were also used to compute the risk–free rate of interest. Of the 1 million total values of C_t computed, 92 had to be adjusted in this way. Similarly, negative values of C_t were encountered in computing 0.67% of

the R_t's.

First Moments. Table 4 reports first–moment properties of the DP and LQ decision rules for the $\tau = .5$ models, as indicated in the column headings. The statistic column contains the variable whose mean is reported in the other columns. Those columns contain the average, across 100 simulations, of the mean value of the variable. Numbers in parentheses in the statistic column are steady–state values of the associated variable. Numbers in parentheses in the other columns are the standard deviation across the 100 simulations. The small size of the standard deviations reflects the large number of observations per simulation (i.e., 10,000). By comparing the first–moment properties of the DP rules with the corresponding steady states, we can assess an assumption implicit in many applications that utilize the LQ approximation. This assumption—that steady states and unconditional means roughly coincide—plays a role in two places in applied work. First, there would be little sense in approximating a model about steady state if the model's variables did not fluctuate about this point in the stochastic version of the problem. Second, many empirical researchers who use the LQ approximation select model parameter values by equating nonstochastic steady–state properties of their model with corresponding sample statistics (Kydland and Prescott 1982.) This method of assigning parameter values would be inappropriate if the nonstochastic steady state diverged substantially from the mean of the stochastic version of the model.

First consider capital, consumption, and output. Table 4 indicates that, for the high–variance economy, the mean of these variables is roughly 10% higher than their steady–state values. In addition, the mean of the capital output ratio is about 7% higher than its steady–state value. Presumably, the larger average capital stock in the stochastic economy reflects households' efforts to insure themselves

against the risk associated with the production technology. Recall, however, that the innovation to the technology shock in the high–variance economy is more than 10 times what is plausible empirically. In the more empirically plausible low–variance economy, $EK_t/Y_t \cong K^*/Y^*$, $EK_t \cong K^*$, $EC_t \cong C^*$, and $EY_t \cong Y^*$ (stars indicate steady–state quantities).

The mean value of capital implied by the linear LQ decision rule is roughly equal to K^*, as it must, given that it is linear. The log–linear LQ decision rule implies a larger mean value of K_t in the high–variance economies because of the convexity of the exponential function. Thus, in the high–variance economies, the mean of the logLQ capital stock lies between that of the DP and linLQ decision rules. The same is true for C, Y, and K/Y.

Next, consider R_t and $MP_{k,t}$. In nonstochastic steady state these quantities are both $\beta^{-1} - 1 = .0204$. In the stochastic version of the model, one expects $EMP_{k,t} > ER_t$. This inequality reflects that $MP_{k,t}$ is the return on a riskier investment than is R_t, since the states in which the former pays off the most are those in which consumption is valued least, i.e., $Cov_t(MP_{k,t}, u'(C_{t+1})) < 0$. As it turns out, both ER_t and $EMP_{k,t}$ are approximately $\beta^{-1} - 1$, even in the high–variance model. The fact that the average equity premium $E[MP_{k,t} - R_t]$ is roughly zero in this model is reminiscent of a similar result obtained by Mehra and Prescott (1985) for an endowment economy.

In sum, the evidence in Table 4 for the four $\tau = .5$ models suggests that the steady–state properties of the nonstochastic version of the model approximate closely the corresponding first–moment properties of the stochastic version of the model, as long as the innovation variance to the technology shock is of plausible magnitude (i.e., $\sigma_\epsilon = .01$). Table 7 reports first–moment results for the $\tau = 3$, low–variance model. As in the $\tau = .5$ low–variance models, there is little difference

between steady states and unconditional means in the $\tau = 3$ model.

Second Moments. Tables 5–7 report second–moment properties of the models. There, σ_x denotes the standard deviation of the variable X_t. In addition, $\rho_{x,y}(\tau)$ denotes the correlation between X_t and $Y_{t-\tau}$, and Δc_t signifies $C_t - C_{t-1}$. Numbers not in parentheses are the average of the associated statistic, across the 100 data sets. Numbers in parentheses are the corresponding standard deviation.

First consider Table 5, which reports results for the $\tau = .5$ models. A striking feature of that table is that the results differ so little between decision rules and models. The only quantitatively meaningful differences lie in σ_r/σ_y and $\rho_{r,\Delta c}(0)$. With regard to the former, the standard deviation of σ_r/σ_y is higher for the linLQ decision rule and the high–variance three–shock economy than for the other cases. With regard to $\rho_{r,\Delta c}(0)$, the LQ versions of that statistic are smaller in the high–variance models while the DP versions appear scale independent. Since the LQ and DP versions are equal for the low–variance economies and scale independence seems plausible, I interpret this to reflect approximation error in the LQ approximation. The importance of the very fine grid used in the paper showed up in calculating $\rho_{r,\Delta c}(0)$. For example, when I computed the DP decision rules with a grid of .01 between capital points, then I got the following results for $\rho_{r,\Delta c}(0)$ in the four $\tau = .5$ models: .787 (.006), .365 (.006), .780 (.007), and .406 (.013) in the high– and low–variance two–state models and the high– and low–variance three–state models, respectively. Thus, using a grid coarser than the one underlying the results in Table 5, I found that the DP rules imply some scale dependence.

Next, consider the dynamic correlations reported in Table 6 for the $\tau = .5$

economies. As in the other tables, there are few significant discrepancies between solution methods and models. One discrepancy is that the correlations based on the linLQ solution to the high–variance three–state Markov model are all smaller than the other correlations, presumably reflecting approximation error. Another discrepancy is the scale dependence in the LQ versions of $\rho_{\Delta c,\Delta c}(i)$, for i=1,2.

Table 7 contains the second–moment results for the $\tau = 3$ model. There are virtually no noticeable discrepancies between LQ and DP second moments. One exception is $\rho_{\Delta c,r}(2)$, which is lower for the DP decision rule than for the LQ decision rules.

An interesting feature of the results in Table 7 is that the correlations between consumption changes and lagged variables is close to zero when $\tau = 3$ and much further from zero in the $\tau = .5$ case. In this respect, the $\tau = 3$ results are close to what I found in a version of this model that I have studied elsewhere (Christiano 1987b,c), in which $\rho = \tau = 1$, $\delta = .018$, $\beta = .99$, and in which hours are variable. In that model, consumption changes are also approximately uncorrelated with lagged variables.

To summarize, in the low–variance economies, the first– and second–moment properties implied by the LQ approximations and the value–function iteration solution are roughly identical, even with high risk–aversion. Discrepancies occur for very large shock variances. An interesting feature of the results is that second–moment properties seem relatively insensitive to whether the exogenous shock is drawn from a two– or three–state Markov chain.

5.3 The Value of the DP Rule to an LQ Decision Maker

Table 8 reports the amount, as a fraction of K_{t-1}, that a planner using an

LQ decision rule would be willing to pay to learn the DP rule in the four $\tau = .5$ models. This quantity was computed for the same (K_{t-1}, x_t) combinations used in Tables 2 and 3. How I did this is discussed below.

To place the LQ and DP decision rules on a comparable basis, I redefined the LQ rules slightly. Instead of allowing f_{logLQ} and f_{linLQ} to map (K_{t-1}, x_t) onto the real line, I replaced (for purposes of computation in Table 8) $f_{logLQ}(K_{t-1},x_t)$ and $f_{linLQ}(K_{t-1},x_t)$ by the nearest point in $A(K_{t-1},x_t)$, defined in (3.2). Given the fine grid $\mathcal{K}$, this adjustment presumably has negligible effect.

I computed the v_{logLQ} and v_{linLQ} functions that solve the following functional equations:

$$v_{logLQ}(K_{t-1}, x_t) = u(K_{t-1}, f_{logLQ}(K_{t-1},x_t), x_t) + \beta E_t v_{logLQ}(K_t, x_{t+1}) \tag{5.1}$$

and

$$v_{linLQ}(K_{t-1}, x_t) = u(K_{t-1}, f_{linLQ}(K_{t-1}, x_t), x_t) + \beta E_t v_{linLQ}(K_t, x_{t+1}), \tag{5.2}$$

for u defined in (2.6). In each case, the expectation operator was evaluated relative to the relevant two– or three–state Markov chain with high or low variance. (To avoid complicating the notation, I do not index the v functions by the Markov–chain model.)

Relative to a given model (defined by the Markov chain for the exogenous shock) and specified initial conditions, I computed the loss of using the LQ decision

rules as follows. Let $v^* = v_{logLQ}(K_{t-1}, x_t)$. Then let K' be defined by the property $v_{DP}(K', x_t) = v^*$, i.e., K' is that level of capital ($K' \leq K_{t-1}$) such that a planner starting with K', x_t and knowing the DP rule achieves the same utility as a planner starting with K_{t-1}, x_t who uses the log–linear LQ decision rule. Evidently, the LQ planner would be willing to pay no more than $K_{t-1}-K'$ to acquire knowledge of the DP rule. Table 8 reports this as a percentage of K_{t-1}, i.e., $[(K_{t-1}-K')/K_{t-1}]\times 100$. Similar calculations were carried out for the linear LQ decision rule.

It is not surprising, in view of the preceding results, that for the low–variance shock distributions with $\tau = .5$, the LQ planner would not pay anything to acquire the DP decision rule (see Table 8, Panels A and C.) Thus, for these shock distributions, the LQ rules are roughly optimal. I obtained exactly the same results for the low–variance, $\tau = 3$ model as for the low–variance, $\tau = .5$ models.

With high shock variances, the results in Panels B and D of Table 8 show that the LQ planner would pay a positive, though still very small, amount to acquire the DP rule. Except when very far from steady state, the planner would pay less than 1% of initial capital.

5.4 Steady–State Capital and Consumption Distribution

Figures 1 and 2 plot steady–state capital and consumption, respectively, for the high–variance two–state Markov–chain model for $\tau = .5$. The same is done in Figures 3 and 4 for the high–variance three–state Markov–chain model. The low–variance steady–state distributions are not plotted because they actually coincide.

A distinguishing feature of these figures is the bimodal distribution when the Markov chain has two states and the unimodal distribution for the three–state Markov process. In the two–state case, note the pronounced lack of symmetry in all but one of the distributions. This reflects the nonlinearity of all but one of the decision rules. The exception, f_{linLQ}, produces a roughly symmetric steady–state distribution for capital.

An interesting feature of these charts is the left shift in the linLQ distributions for C_t and K_t relative to the logLQ and DP distributions. In view of this shift, it is not surprising that the nonnegativity constraint on C_t proved occasionally binding when the three–state high–variance Markov–chain version of the model was solved by linear LQ approximation.

6. CONCLUSION

The purpose of this article was to evaluate the accuracy of two methods for approximating the consumption and capital decision rules that solve a version of the Brock and Mirman (1972) optimal growth problem. The strategy I took was to define the growth problem in such a way that numerical dynamic programming methods could be used to obtain arbitrarily accurate approximations to the exact decision rules. This involved positing a discrete distribution for the exogenous shocks of the model and forcing the capital stock to lie on a very fine grid. The accuracy of the LQ approximate decision rules was evaluated by comparing them along four dimensions with the presumed exact solutions obtained by dynamic programming (DP) methods.

I found that the LQ approximation works well in the model economy studied here. In this respect, the conclusions are similar to those reached by others

who looked at different examples (Christiano 1986, 1987b; Danthine, Donaldson, and Mehra 1988; Rebelo and Rouwenhorst 1989).

Blank Page

REFERENCES

Altug, S. (1986), "Time to Build and Aggregate Fluctuations: Some New Evidence," manuscript, Federal Reserve Bank of Minneapolis, Research Dept.

Bertsekas, D. (1987), *Dynamic Programming: Deterministic and Stochastic Models*, Englewood Cliffs: NJ Prentice–Hall.

Brock, W. A., and Mirman, L. J. (1972), "Optimal Economic Growth and Uncertainty: The Discounted Case," *Journal of Economic Theory*, 4, 479–513.

Christiano, L. J. (1986), "On the Accuracy of Linear Quadratic Approximations: An Example," Research Department Working Paper 303, Federal Reserve Bank of Minneapolis.

_____ (1987a), "Dynamic Properties of Two Approximate Solutions to a Particular Growth Model," Research Department Working Paper 338, Federal Reserve Bank of Minneapolis.

_____ (1987b), "Intertemporal Substitution and the Smoothness of Consumption," Research Department Working Paper 427, Federal Reserve Bank of Minneapolis.

_____ (1987c), "Is Consumption Insufficiently Sensitive to Innovations in Income?," *American Economic Review Papers and Proceedings*, 77, 337–341.

_____ (1988), "Why Does Inventory Investment Fluctuate So Much?," *Journal of Monetary Economics*, 21, 247–280.

Christiano, L. J., and Eichenbaum, M. (1988), "Is Theory Really Ahead of Measurement? Current Real Business Cycle Theories and Aggregate Labor Market Fluctuations," Research Department Working Paper 412, Federal Reserve Bank of Minneapolis.

Christiano, L. J., and Fitzgerald, T. J. (1989), "The Magnitude of the Speculative Motive for Holding Inventories in a Real Business Cycle Model," Discussion Paper 10, Institute for Empirical Macroeconomics, Federal Reserve Bank of Minneapolis.

Cooley, T. F., and Hansen, G. D. (1988), "The Inflation Tax in a Real Business Cycle Model," Working Paper 496, UCLA, Dept. of Economics.

Danthine, J.–P., Donaldson, J. B., and Mehra, R. (1988), "On Some Computational Aspects of Real Business Cycle Theory," Working Paper 1984–88, Massachusetts Institute of Technology, Sloan School of Management.

Danthine, J.–P., Donaldson, J. B., and Smith, L. (1987), "On the Superneutrality of Money in a Stochastic Dynamic Macroeconomic Model," *Journal of Monetary Economics*, 20, 475–499.

Hansen, G. D. (1985), "Indivisible Labor and the Business Cycle," *Journal of Monetary Economics*, 16, 309–327.

Hansen, G. D., and Sargent, T. J. (1988), "Straight Time and Overtime in Equilibrium," *Journal of Monetary Economics*, 21, 281–308.

Harris, M. (1987), *Dynamic Economic Analysis*, New York: Oxford University Press.

King, R. G., Plosser, C. I., and Rebelo, S. T. (1988), "Production, Growth and Business Cycles: I. The Basic Neoclassical Model," *Journal of Monetary Economics*, 21, 195–232.

Kydland, F., and Prescott, E. C. (1980), "A Competitive Theory of Fluctuations and the Feasibility and Desirability of Stabilization Policy," (with comments), in *Rational Expectations and Economic Policy*, ed. S. Fischer, Chicago: University of Chicago Press, pp. 169–198.

_____ (1982), "Time to Build and Aggregate Fluctuations," *Econometrica*, 50, 1345–1370.

Long, J. B., and Plosser, C. I. (1983), "Real Business Cycles," *Journal of Political Economy*, 91, 39–69.

Mehra, R., and Prescott, E. C. (1985), "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, 15, 145–161.

Prescott, E. C. (1986), "Theory Ahead of Business Cycle Measurement," *Federal Reserve Bank of Minneapolis Quarterly Review*, 10, 9–22. Also in *Real Business Cycles, Real Exchange Rates and Actual Policies*, ed. K. Brunner and A. H. Meltzer, Carnegie–Rochester Conference Series on Public Policy, 25, 1–44.

Rebelo, S. T., and Rouwenhorst, K. G. (1989), "Linear Quadratic Approximations Versus Discrete State Space Methods: A Numerical Evaluation," manuscript, University of Rochester, Dept. of Economics.

Rust, J. (1987), "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 55, 999–1033.

_____ (1988a), "Maximum Likelihood Estimation of Discrete Control Processes," *SIAM Journal of Control and Optimization*, 26, 1006–1024.

_____ (1988b), "A Dynamic Programming Model of Retirement Behavior," in *The Economics of Aging*, ed. D. Wise, Chicago: University of Chicago Press (for NBER).

Sargent, T. J. (1987), *Dynamic Macroeconomic Theory*, Cambridge, MA: Harvard University Press.

Table 1. Decision Rule Information

(1) Model	(2) $\mathscr{E}_{\text{logLQ}}$	(3) $\mathscr{E}_{\text{linLQ}}$	(4) $\mathscr{E}_{\text{DP}}$	(5) $\mathscr{K}$	(6) Increment	(7) CPU Time
Two–state $\sigma_\epsilon = .10$ $\tau = .5$	{37, 110}	{29, 98}	{37, 109}	{35, 115}	.00400	213.84
Two–state $\sigma_\epsilon = .01$ $\tau = .5$	{60, 67}	{60, 67}	{60, 67}	{55, 70}	.00075	214.02
Three–state $\sigma_\epsilon = .10$ $\tau = .5$	{25, 163}	{4, 123}	{25, 161}	{20, 165}	.00725	310.58
Three–state $\sigma_\epsilon = .01$ $\tau = .5$	{58, 70}	{58, 70}	{58, 70}	{55, 75}	.00100	314.80
Three–state $\sigma_\epsilon = .01$ $\tau = 3.0$	{49, 83}	{47, 80}	{49, 83}	{45, 85}	.00200	342.08

NOTE: Column (1): Number of states and σ_ϵ value indicate the Markov–chain model of the x_t's. The only other parameter that differs between models is τ.
Columns (2)–(4): First and second numbers associated with $\mathscr{E}$ are the upper and lower boundaries, respectively, of the ergodic set associated with the linLQ, logLQ, or DP capital decision rules, as indicated by the subscript.
Column (5): The first and second numbers associated with $\mathscr{K}$ are the boundaries of the capital grid used in the value–function iteration calculations.
Column (6): Increment between adjacent values of capital in $\mathscr{K}$.
Column (7): Time, in central processor unit minutes, used to solve the associated model by value–function iteration on the Federal Reserve Bank of Minneapolis's Amdahl dual 5860 mainframe computer. Details of the solution method are described in Section 3. In the case of the first four models, p = 10 for all j, using notation presented Section 3. In the case of the last model, p = 10 for j = 1, ..., 108 and p = ϖ for j = 109, ..., 118, whereupon convergence occurred. The convergence criterion is reported in Section 4.2.

Table 2. Capital and Consumption Rules: Two–State Markov Process

Panel A: $\sigma_\epsilon = .01$, $\tau = .5$ (Low Variance)

K_{t-1} grid ↓	x_t=−.03	x_t=.03	x_t=−.03	x_t=.03	x_t=−.03	x_t=.03
	—f_{DP}—		—f_{logLQ}—		—f_{linLQ}—	
60.32	60.32	60.53	60.32	60.52	60.31	60.53
62.00	61.95	62.16	61.95	62.16	61.95	62.16
63.69	63.58	63.79	63.58	63.79	63.58	63.79
65.46	65.30	65.51	65.29	65.51	65.30	65.51
67.23	67.01	67.23	67.01	67.23	67.02	67.23
	—g_{DP}—		—g_{logLQ}—		—g_{linLQ}—	
60.32	3.75	3.78	3.75	3.79	3.75	3.78
62.00	3.83	3.87	3.83	3.87	3.84	3.87
63.69	3.92	3.96	3.92	3.96	3.92	3.96
65.46	4.01	4.05	4.02	4.05	4.01	4.05
67.23	4.10	4.14	4.11	4.14	4.10	4.14

Panel B: $\sigma_\epsilon = .10$, $\tau = .5$ (High Variance)

K_{t-1} grid ↓	x_t=−.32	x_t=.32	x_t=−.32	x_t=.32	x_t=−.32	x_t=.32
	—f_{DP}—		—f_{logLQ}—		—f_{linLQ}—	
36.78	36.78	38.63	36.79	38.05*	36.55	38.70
50.24	49.79	51.82	49.76	51.47	49.58	51.73
63.69	62.77	64.95	62.62	64.77	62.61	64.76
86.19	84.45	86.85	83.95	86.83	84.41	86.56
108.69	106.11	108.69	105.11	108.72	106.21	108.36
	—g_{DP}—		—g_{logLQ}—		—g_{linLQ}—	
36.78	2.39	2.68	2.38	3.25††	2.62*	2.61*
50.24	3.09	3.43	3.12	3.78†	3.30*	3.52*
63.69	3.78	4.16	3.93*	4.34*	3.93*	4.35*
86.19	4.90	5.33	5.40†	5.35	4.94	5.62*
108.69	5.99	6.47	6.99†	6.45	5.89*	6.80*

NOTE: The table reports capital and consumption decisions for various points in the state space based on the DP, logLQ, and linLQ decision rules. Rows correspond to values of initial capital and columns correspond to technology shock values, as indicated. Let z denote the ratio of an LQ decision to a DP decision at a given point in the state space. Let z' be $|(z-1)\times 100|$, rounded to the nearest integer, where $|\cdot|$ denotes the absolute value operator. Then * indicates $1 \leq z' < 10$, † indicates $10 \leq z' < 20$, and †† indicates $z' > 20$.

Table 3. Capital and Consumption Rules: Three–State Markov Process

Panel A: $\sigma_\epsilon = .01$, $\tau = .5$ (Low Variance)

$K_{t-1}\downarrow$	x_t=−.06	x_t=.0	x_t=.06	x_t=−.06	x_t=.0	x_t=.06	x_t=−.06	x_t=.0	x_t=.06
	f_{DP}			f_{logLQ}			f_{linLQ}		
57.96	57.96	58.14	58.32	57.96	58.13	58.30	57.95	58.14	58.32
60.82	60.73	60.91	61.10	60.73	60.91	61.09	60.73	60.91	61.10
63.69	63.51	63.69	63.88	63.50	63.69	63.87	63.50	63.69	63.87
66.82	66.54	66.72	66.92	66.53	66.72	66.92	66.54	66.73	66.91
69.96	69.58	69.76	69.96	69.55	69.76	69.96	69.58	69.77	69.95
	g_{DP}			g_{logLQ}			g_{linLQ}		
57.96	3.61	3.64	3.67	3.61	3.65	3.70	3.62	3.64	3.67
60.82	3.76	3.79	3.82	3.76	3.79	3.83	3.77	3.79	3.82
63.69	3.91	3.94	3.97	3.91	3.94	3.98	3.91	3.94	3.98
66.82	4.07	4.10	4.13	4.08	4.10	4.13	4.07	4.10	4.14
69.96	4.23	4.26	4.29	4.25	4.27	4.29	4.23	4.26	4.30

Panel B: $\sigma_\epsilon = .10$, $\tau = .5$ (High Variance)

$K_{t-1}\downarrow$	x_t=−.55	x_t=.0	x_t=.55	x_t=−.55	x_t=.0	x_t=.55	x_t=−.55	x_t=.0	x_t=.55
	f_{DP}			f_{logLQ}			f_{linLQ}		
24.60	24.60	25.65	27.53	24.61	25.34*	26.09*	23.95*	25.81	27.68
44.14	43.45	44.71	46.97	43.36	44.65	45.97*	42.89*	44.75	46.61
63.69	62.27	63.68	66.20	61.85	63.69	65.58	61.82	63.69	65.55
112.26	108.95	110.63	113.64	107.12*	110.30	113.57	108.89	110.75	112.61
160.84	155.60	157.48	160.84	151.76*	156.26	160.90	155.95	157.81	159.67
	g_{DP}			g_{logLQ}			g_{linLQ}		
24.60	1.65	1.83	2.08	1.64	2.14†	3.52††	2.30††	1.66*	1.93*
44.14	2.69	2.92	3.25	2.78*	2.98*	4.25††	3.26††	2.88*	3.61†
63.69	3.68	3.95	4.34	4.10†	3.94	4.97†	4.12†	3.94	5.00†
112.26	6.04	6.38	6.90	7.87††	6.71*	6.96	6.10*	6.26*	7.92†
160.84	8.30	8.70	9.31	12.15††	9.92†	9.25	7.96*	8.37*	10.47†

Panel C: $\sigma_\epsilon = .01$, $\tau = 3.0$ (Low Variance)

$K_{t-1}\downarrow$	x_t=−.06	x_t=.0	x_t=.06	x_t=−.06	x_t=.0	x_t=.06	x_t=−.06	x_t=.0	x_t=.06
	f_{DP}			f_{logLQ}			f_{linLQ}		
48.95	48.95	49.09	49.23	48.95	49.07	49.19	48.93	49.09	49.24
56.32	56.25	56.39	56.54	56.25	56.38	56.52	56.23	56.39	56.54
63.69	63.54	63.69	63.85	63.53	63.69	63.84	63.53	63.69	63.84
73.24	72.99	73.15	73.32	72.96	73.14	73.32	72.99	73.15	73.30
82.78	82.44	82.61	82.78	82.38	82.58	82.78	82.45	82.61	82.76
	g_{DP}			g_{logLQ}			g_{linLQ}		
48.95	3.42	3.48	3.54	3.42	3.49	3.58	3.43	3.47	3.53
56.32	3.65	3.71	3.78	3.65	3.72	3.80	3.66	3.71	3.78
63.69	3.87	3.94	4.00	3.88	3.94	4.01	3.88	3.94	4.01
73.24	4.15	4.21	4.28	4.17	4.22	4.28	4.14	4.21	4.29
82.78	4.41	4.47	4.54	4.46	4.49	4.54	4.39	4.47	4.56

NOTE: See notes to Table 2.

Table 4. First-Moment Properties: $\tau = .5$ Economy

	Two-State $\sigma_\varepsilon = .10$			Two-State $\sigma_\varepsilon = .01$			Three-State $\sigma_\varepsilon = .10$			Three-State $\sigma_\varepsilon = .01$		
Statistic[a]	DP	Log	Lin	DP	Log	Lin	DP	Log	Lin	DP	Log	Lin
C (3.94)	4.34 (.15)	4.27 (.15)	4.19 (.15)	3.94 (.014)	3.94 (.014)	3.94 (.014)	4.31 (.13)	4.24 (.13)	4.16 (.13)	3.94 (.012)	3.94 (.012)	3.94 (.012)
Y (3.94)	4.34 (.15)	4.27 (.15)	4.19 (.15)	3.94 (.014)	3.94 (.014)	3.94 (.014)	4.31 (.13)	4.24 (.13)	4.16 (.13)	3.94 (.012)	3.94 (.012)	3.94 (.012)
K (63.69)	70.85 (2.64)	67.61 (2.65)	64.00 (2.53)	63.78 (.25)	63.75 (.25)	63.72 (.25)	70.44 (2.34)	67.06 (2.25)	63.42 (2.10)	63.73 (.21)	63.70 (.21)	63.66 (.21)
K/Y (16.17)	17.30 (.068)	16.73 (.059)	16.08 (.103)	16.18 (.005)	16.18 (.005)	16.17 (.005)	17.29 (.098)	16.71 (.006)	15.99 (.131)	16.18 (.004)	16.17 (.004)	16.17 (.005)
MP_K (.0204)	.0204 (.65E-4)	.0211 (.78E-4)	.0220 (.13E-3)	.0204 (.66E-5)	.0204 (.67E-5)	.0204 (.67E-5)	.0204 (.55E-4)	.0211 (.86E-4)	.0224 (.35E-3)	.0204 (.56E-5)	.0204 (.57E-5)	.0204 (.60E-5)
R (.0204)	.0204 (.88E-4)	.0203 (.12E-3)	.0204 (.94E-4)	.0204 (.90E-5)	.0204 (.89E-5)	.0204 (.90E-5)	.0204 (.74E-4)	.0203 (.14E-3)	.0200 (.72E-3)	.0204 (.75E-5)	.0204 (.75E-5)	.0204 (.75E-5)

NOTE: Results are based on 100 data sets, each of length 10,000, using the indicated solution (DP, log-linear LQ, or linear LQ) to the version of the growth model that incorporates the indicated probability model for the exogenous shock, x_t (two- or three-state Markov chain, with high or low variances). Initial conditions were randomized across simulations. Standard deviations across simulations appear in parentheses. For parameter values, see Section 3 in the text.

[a] C, Y, K, R, MP_K denote the means of C_t, Y_t, K_t, R_t, and $MP_{K,t}$, respectively. Numbers in parentheses in this column are steady-state values.

Table 5. Second-Moment Properties: $\tau = .5$ Economy

Statistic	Two-State $\sigma_\varepsilon = .10$			Two-State $\sigma_\varepsilon = .01$			Three-State $\sigma_\varepsilon = .10$			Three-State $\sigma_\varepsilon = .01$		
	DP	**Log**	Lin	**DP**	Log	Lin	**DP**	Log	Lin	DP	Log	Lin
σ_c/σ_y	.78 (.02)	.79 (.02)	.78 (.02)	.78 (.02)	.78 (.02)	.78 (.02)	.77 (.02)	.78 (.02)	.80 (.02)	.78 (.02)	.78 (.02)	.78 (.02)
σ_{dk}/σ_y	.56 (.02)	.57 (.02)	.54 (.02)	.55 (.02)	.55 (.02)	.55 (.02)	.56 (.02)	.54 (.02)	.51 (.02)	.55 (.02)	.55 (.02)	.55 (.02)
σ_r/σ_y	.003 (.0001)	.003 (.0001)	.004 (.0002)	.003 (.0001)	.003 (.0001)	.003 (.0001)	.003 (.0001)	.003 (.0002)	.002 (.017)	.003 (.0001)	.003 (.0001)	.003 (.0001)
σ_y	1.62 (.02)	1.61 (.02)	1.59 (.02)	.155 (.002)	.155 (.002)	.155 (.002)	1.77 (.10)	1.74 (.10)	1.69 (.08)	.155 (.007)	.155 (.007)	.155 (.007)
$\rho_{r,c}(0)$	-.05 (.01)	-.08 (.03)	-.06 (.01)	-.05 (.013)	-.05 (.013)	-.05 (.013)	-.05 (.02)	-.12 (.03)	-.05 (.07)	-.05 (.016)	-.05 (.016)	-.05 (.017)
$\rho_{dk,c}(0)$	.11 (.005)	.06 (.010)	.11 (.007)	.10 (.005)	.10 (.005)	.10 (.005)	.11 (.008)	.09 (.015)	.13 (.022)	.10 (.007)	.10 (.007)	.10 (.007)
$\rho_{r,\Delta c}(0)$	.80 (.006)	.49 (.015)	.55 (.014)	.81 (.004)	.82 (.007)	.82 (.005)	.78 (.007)	.32 (.019)	.01 (.073)	.81 (.005)	.81 (.007)	.80 (.008)

NOTE: σ_x denotes the standard deviation of $\{X_t\}$; $\rho_{x,y}(\tau)$ denotes the correlation between X_t and $Y_{t-\tau}$; Δc_t denotes $C_t - C_{t-1}$; and $dk_t \equiv K_t - K_{t-1}$. See also the notes to Table 4.

Table 6. Dynamic Correlations: $\tau = .5$ Economy

Statistic	Two-State $\sigma_\varepsilon = .10$			Two-State $\sigma_\varepsilon = .01$			Three-State $\sigma_\varepsilon = .10$			Three-State $\sigma_\varepsilon = .01$		
	DP	Log	Lin	DP	Log	Lin	DP	Log	Lin	DP	Log	Lin
$\rho_{\Delta c,r}(1)$	.53 (.014)	.51 (.027)	.47 (.022)	.54 (.014)	.54 (.016)	.54 (.014)	.52 (.016)	.39 (.021)	.14 (.065)	.54 (.018)	.54 (.018)	.53 (.020)
$\rho_{\Delta c,r}(2)$	.48 (.016)	.45 (.024)	.43 (.020)	.49 (.017)	.50 (.018)	.50 (.017)	.48 (.018)	.34 (.020)	.13 (.057)	.49 (.019)	.49 (.020)	.49 (.021)
$\rho_{\Delta c,\Delta c}(1)$	.44 (.014)	.23 (.021)	.30 (.022)	.44 (.014)	.44 (.017)	.45 (.014)	.45 (.018)	.13 (.022)	.08 (.013)	.44 (.018)	.44 (.019)	.43 (.021)
$\rho_{\Delta c,\Delta c}(2)$	.40 (.016)	.21 (.018)	.28 (.022)	.41 (.015)	.41 (.017)	.41 (.016)	.41 (.020)	.11 (.022)	.07 (.013)	.40 (.020)	.40 (.021)	.39 (.022)
$\rho_{\Delta c,y}(1)$	.27 (.012)	.24 (.021)	.22 (.017)	.27 (.012)	.28 (.013)	.28 (.012)	.28 (.013)	.16 (.017)	.07 (.014)	.27 (.012)	.27 (.013)	.27 (.014)
$\rho_{\Delta c,y}(2)$	.23 (.011)	.20 (.018)	.18 (.015)	.23 (.011)	.23 (.012)	.23 (.011)	.23 (.012)	.13 (.015)	.05 (.013)	.23 (.012)	.22 (.012)	.22 (.013)
$\rho_{y,r}(1)$	.48 (.028)	.46 (.040)	.42 (.026)	.49 (.028)	.49 (.028)	.49 (.028)	.47 (.034)	.35 (.025)	.12 (.053)	.49 (.033)	.49 (.033)	.49 (.034)
$\rho_{y,r}(2)$	.47 (.027)	.45 (.039)	.41 (.026)	.48 (.027)	.48 (.027)	.48 (.027)	.46 (.033)	.34 (.024)	.12 (.052)	.47 (.032)	.47 (.031)	.47 (.033)
$\rho_{y,r}(3)$	.45 (.027)	.43 (.039)	.40 (.026)	.46 (.027)	.46 (.027)	.46 (.027)	.44 (.032)	.33 (.023)	.12 (.051)	.46 (.030)	.46 (.030)	.46 (.032)

NOTE: See notes to Tables 4 and 5.

Table 7. Stochastic Properties: Three-State, $\sigma_\varepsilon = .01$, $\tau = 3.0$ Economy

Statistic	DP	Log	Lin	Statistics	DP	Log	Lin	Statistic	DP	Log	Lin
C (3.94)	3.94 (.020)	3.94 (.020)	3.94 (.020)	σ_c/σ_y	.76 (.026)	.76 (.026)	.76 (.026)	$\rho_{\Delta c,r}(1)$	.107 (.012)	.090 (.011)	.090 (.012)
Y (3.94)	3.94 (.020)	3.94 (.020)	3.94 (.020)	σ_{dk}/σ_y	.49 (.022)	.49 (.022)	.48 (.022)	$\rho_{\Delta c,r}(2)$	.053 (.011)	.087 (.011)	.087 (.012)
K (63.69)	63.97 (.59)	63.72 (.59)	63.61 (.59)	σ_r/σ_y	.006 (.0001)	.005 (.0001)	.005 (.0001)	$\rho_{\Delta c,\Delta c}(1)$	.026 (.011)	.028 (.011)	.028 (.011)
K/Y (16.17)	16.22 (.059)	16.17 (.069)	16.16 (.069)	σ_y	.170 (.010)	.169 (.010)	.169 (.010)	$\rho_{\Delta c,\Delta c}(2)$	.024 (.013)	.023 (.014)	.024 (.013)
MP_k (.0204)	.0204 (.86E-4)	.0204 (.87E-4)	.0205 (.88E-4)	$\rho_{r,c}(0)$	-.40 (.049)	-.47 (.057)	-.48 (.062)	$\rho_{\Delta c,y}(1)$	-.005 (.008)	-.005 (.009)	-.005 (.009)
R (.0204)	.0204 (.88E-4)	.0204 (.87E-4)	.0204 (.89E-4)	$\rho_{dk,c}(0)$	.26 (.022)	.26 (.022)	.26 (.022)	$\rho_{\Delta c,y}(2)$	-.009 (.008)	-.009 (.008)	-.009 (.009)
				$\rho_{r,\Delta c}(0)$	.25 (.012)	.30 (.015)	.30 (.017)	$\rho_{y,r}(1)$	-.031 (.069)	-.035 (.082)	-.041 (.089)
								$\rho_{y,r}(2)$	-.036 (.068)	-.042 (.081)	-.047 (.088)

NOTE: The model economy underlying the simulations in this table is the same as the three-state low-variance economy in Tables 4-6, with the exception that here, $\tau = 3.0$. See notes to Tables 4 and 5 for further information on notation and details of the simulation experiments.

Table 8. Cost, in Percent of K_{t-1}, of Using LQ Decision Rules (How Much an LQ Planner Would Pay to Learn the DP Rule)

Panel A: σ_ϵ = .01, Two–state Markov chain, τ = .5

K_{t-1}	x_t=–.032		x_t=.032	
	Log	Lin	Log	Lin
60.32	0.0	0.0	0.0	0.0
62.00	0.0	0.0	0.0	0.0
63.69	0.0	0.0	0.0	0.0
65.46	0.0	0.0	0.0	0.0
67.23	0.0	0.0	0.0	0.0

Panel B: σ_ϵ = .10, Two–state Markov chain, τ = .5

K_{t-1}	x_t=–.32		x_t=0.32	
	Log	Lin	Log	Lin
36.78	0.3	0.3	0.6	0.3
50.24	0.2	0.3	0.2	0.2
63.69	0.1	0.2	0.1	0.2
86.19	0.2	0.2	0.1	0.2
108.69	0.3	0.1	0.2	0.2

Panel C: σ_ϵ = .01, Three–state Markov chain, τ = .5

K_{t-1}	x_t=–.06		x_t=0.00		x_t=0.06	
	Log	Lin	Log	Lin	Log	Lin
57.96	0.0	0.0	0.0	0.0	0.0	0.0
60.82	0.0	0.0	0.0	0.0	0.0	0.0
63.69	0.0	0.0	0.0	0.0	0.0	0.0
66.82	0.0	0.0	0.0	0.0	0.0	0.0
69.96	0.0	0.0	0.0	0.0	0.0	0.0

Panel D: σ_ϵ = .10, Three–state Markov chain, τ = .5

K_{t-1}	x_t=–.55		x_t=0.00		x_t=0.55	
	Log	Lin	Log	Lin	Log	Lin
24.60	1.0	2.6	1.1	1.2	5.1	1.5
44.14	0.3	2.0	0.3	0.8	1.1	1.2
63.69	0.3	1.2	0.2	0.6	0.4	1.1
112.26	0.7	0.5	0.2	0.4	0.2	0.8
160.84	1.4	0.3	0.4	0.3	0.3	0.6

NOTE: Let $v^* = v_{LQ}(K_{t-1},x_t)$, for given K_{t-1}, x_t, where LQ refers either to log–linear LQ or linear LQ, as indicated in the column headings. Let K^* be such that $v_{DP}(K^*,x_t) = v^*$. The entries in the table are $[(K_{t-1}–K^*)/K_{t-1}]\times100$.

FIGURE TITLES

Figure 1. Steady-State Distribution of Capital: Two-State Markov Shocks With High Variance.

Figure 2. Steady-State Distribution of Consumption: Two-State Markov Shocks With High Variance.

Figure 3. Steady-State Distribution of Capital: Three-State Markov Shocks With High Variance.

Figure 4. Steady-State Distribution of Consumption: Three-State Markov Shocks With High Variance.

FIGURE 1

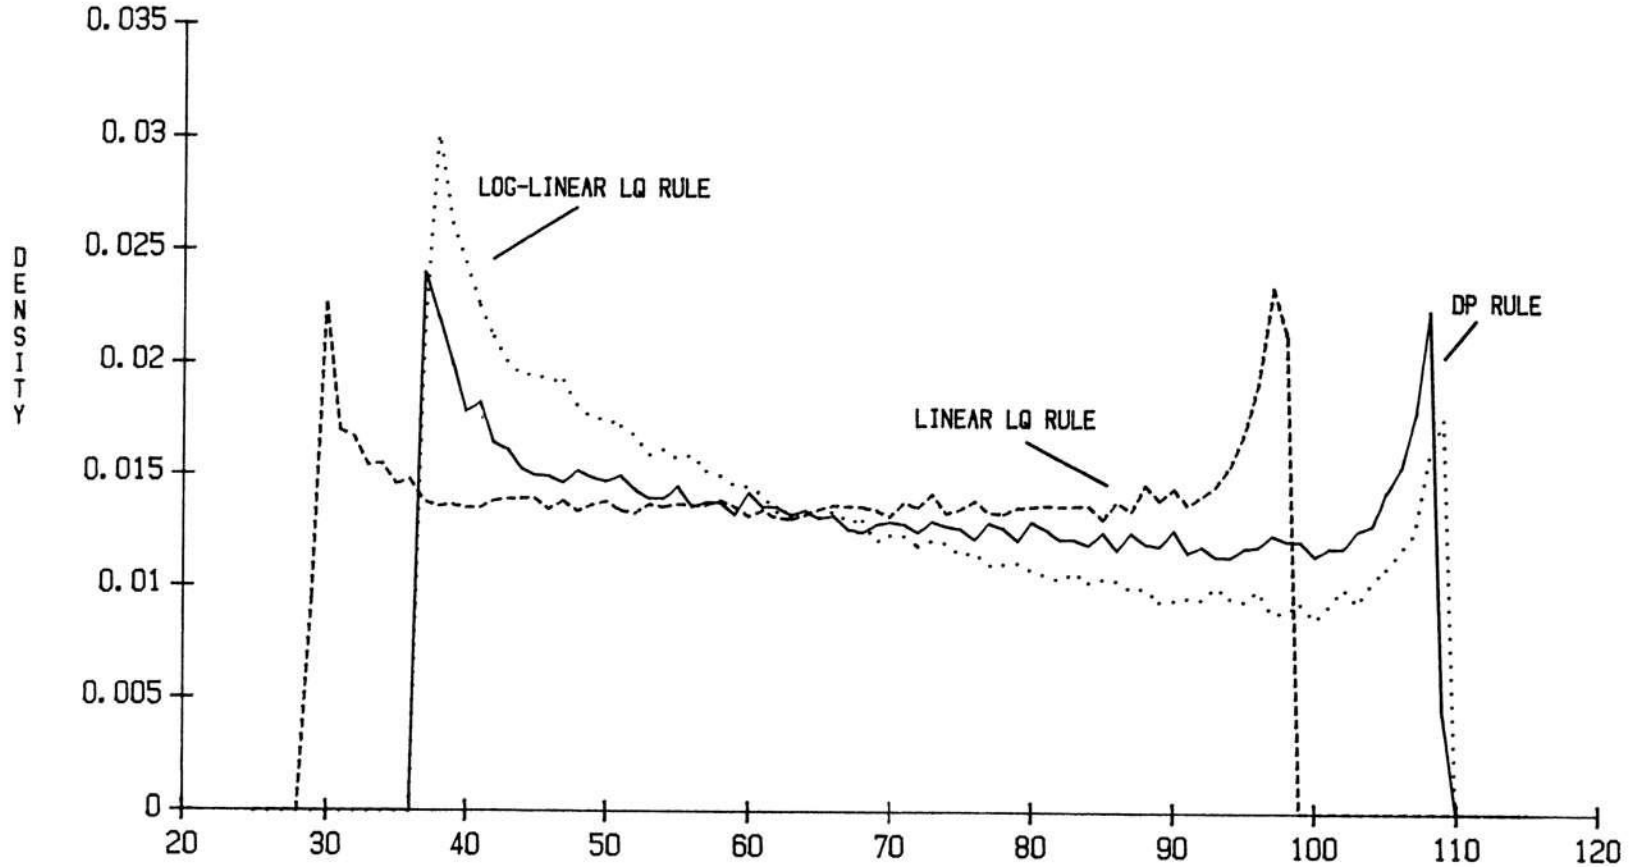
0.035
0.03
0.025
0.02
0.015
0.01
0.005
0
DENSITY
LOG-LINEAR LQ RULE
LINEAR LQ RULE
DP RULE
20
30
40
50
60
70
80
90
100
110
120

FIGURE 2

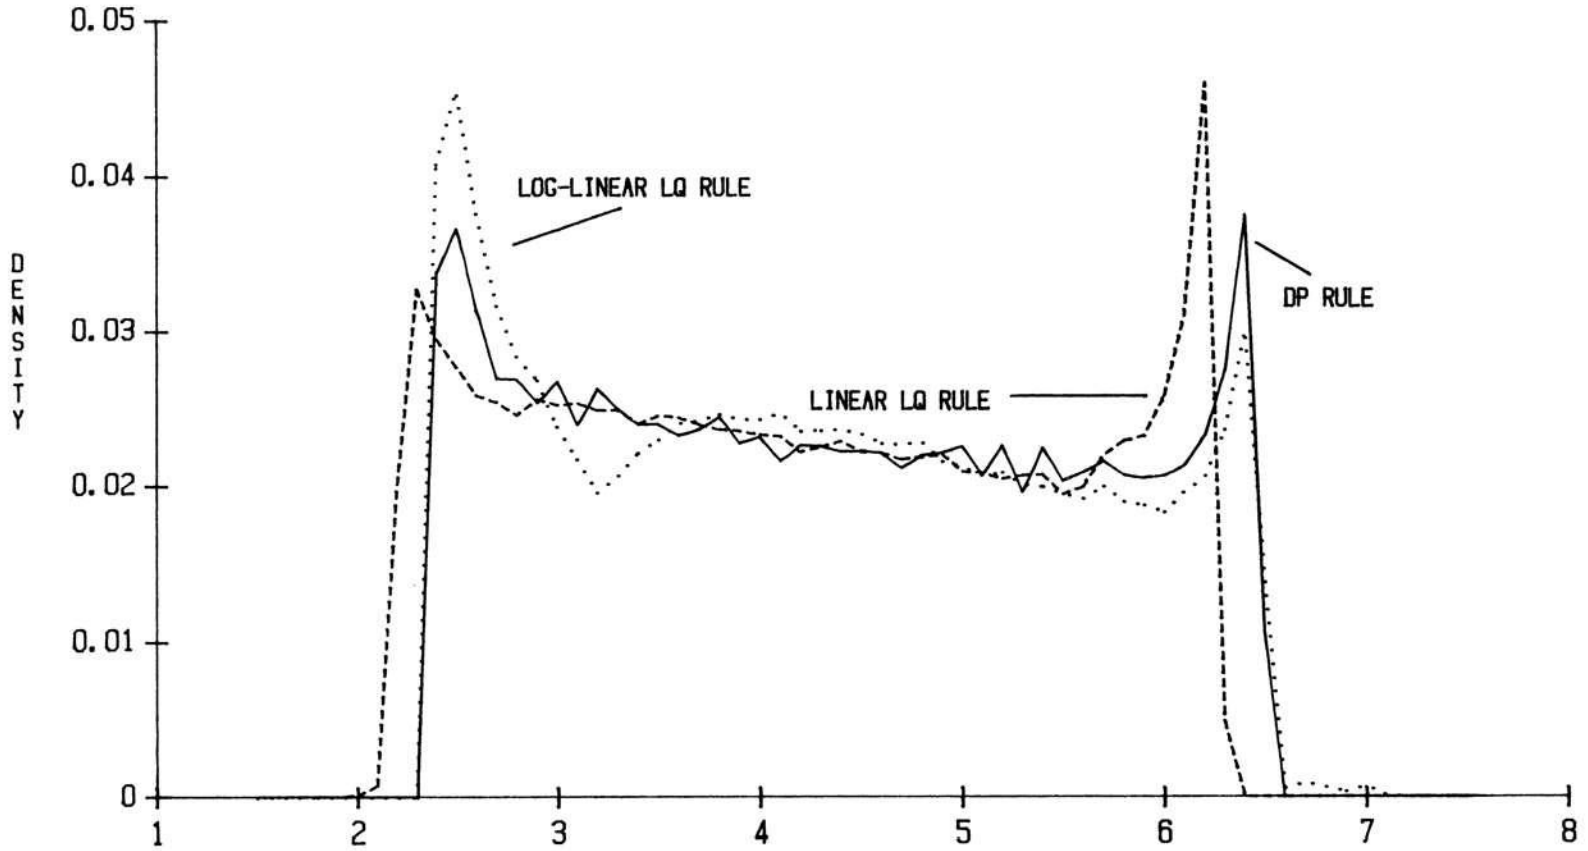
0.05
0.04
0.03
0.02
0.01
0
DENSITY
LOG-LINEAR LQ RULE
LINEAR LQ RULE
DP RULE
1
2
3
4
5
6
7
8

FIGURE 3

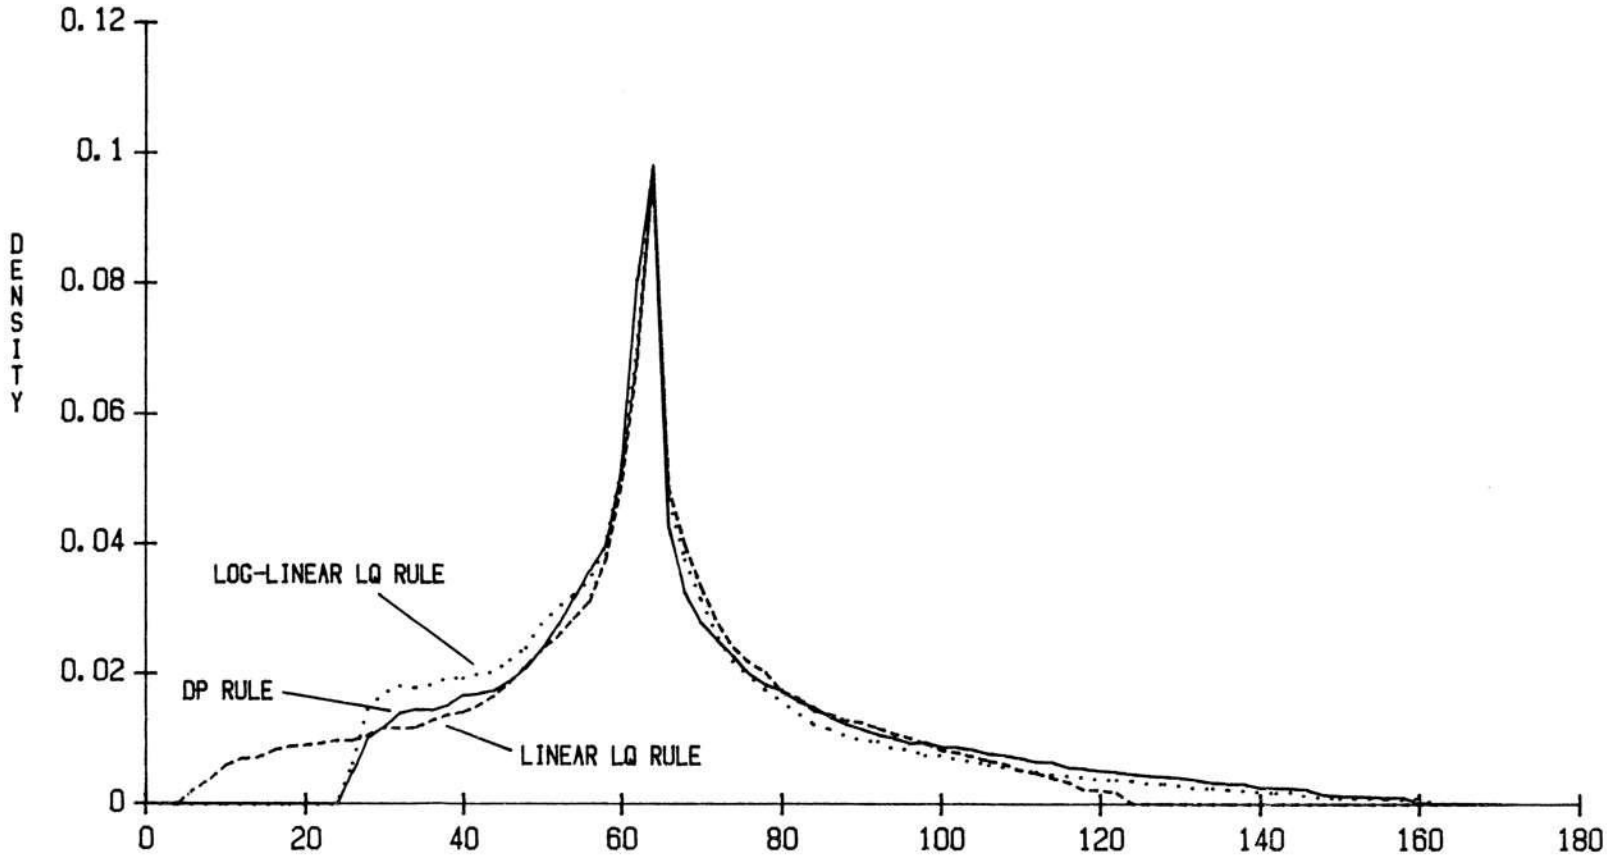
0.12
0.1
0.08
0.06
0.04
0.02
0
DENSITY
LOG-LINEAR LQ RULE
DP RULE
LINEAR LQ RULE
0
20
40
60
80
100
120
140
160
180

FIGURE 4

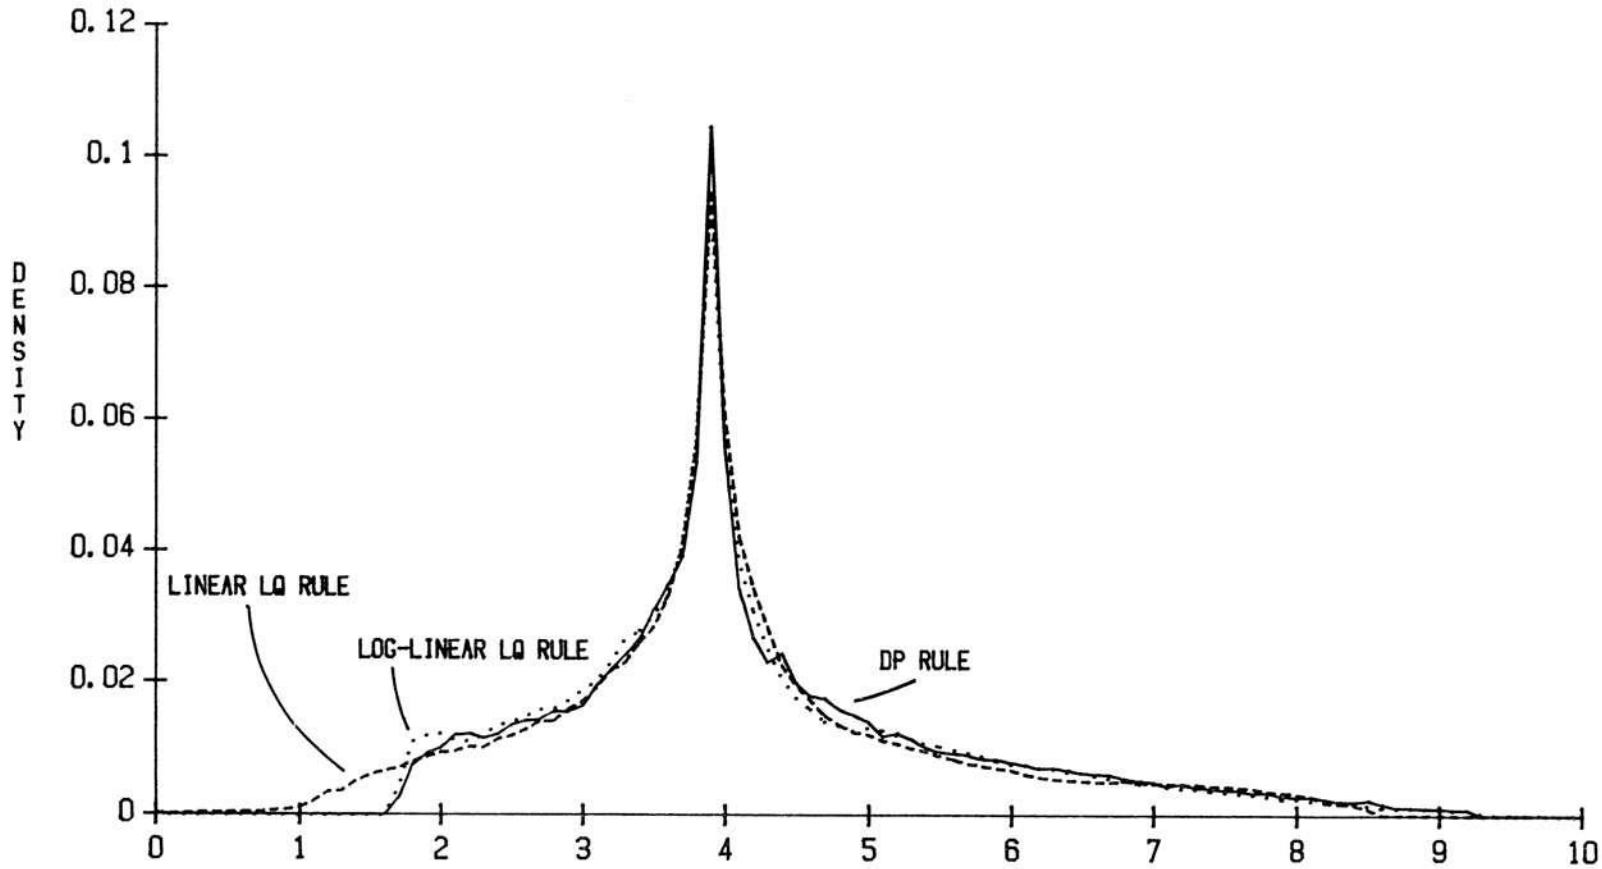
0.12
0.1
0.08
0.06
0.04
0.02
0
DENSITY
LINEAR LQ RULE
LOG-LINEAR LQ RULE
DP RULE
0
1
2
3
4
5
6
7
8
9
10