Federal Reserve Bank of Minneapolis
Research Department Staff Report 104

March 1986

Models of Policy under Stochastic Replanning

William Roberds

Federal Reserve Bank of Minneapolis

Thanks are due to Robert Avery, V.V. Chari, Dennis Epple, Lars Hansen, Rodolfo Manuelli, Thomas Sargent, Christopher Sims, and Charles Whiteman for comments on earlier drafts. Any remaining errors are my own.

Abstract

This paper considers a policy environment in which policy is not set by a single policymaker, but by a sequence of policymaking administrations. Administration turnover is determined by a simple random process. The consequences of administration turnover are traced through for two versions of a linear rational expectations model, and numerical simulations of various policy environments are presented.

Current address of author:

Reseach Department
Federal Reserve Bank of Minneapolis
250 Marquette Avenue
Minneapolis, Minnesota 55480

612-340-7775

1. Introduction

Since the appearance of Kydland and Prescott (1977) , henceforth K-P, macroeconomists have increasingly turned to game theory in constructing both positive and normative models of policy.[1] There seems to be a considerable amount of disagreement, however, as to the role of game theory in "real world" policymaking.

K-P presents two kinds of game theoretic models of policy: time consistent (or discretionary) leader-follower games and time inconsistent (or precommitment) games. In the K-P view, real world policy is likely to correspond to the leader's equilibrium strategy in a time consistent game, in which the policy authority represents the Stackelberg leader and the public represents the follower(s). The outcomes of such games are by construction time consistent, in the sense that the policy authority has no incentive to alter its policy rule as time passes. However, there will almost always exist policy rules that outperform (yield a better value of the policy objective than) any consistent rule.[2] Although these better-performing rules are socially desirable, they will also not be time consistent; as a result, the policy authority would be tempted to alter these rules as time passes. In the absence of explicit constraints on the actions of the policy authority, K-P argues that such policy rules are not credible. Thus, according to this argument, there typically exists room for social gain by adopting some fixed rule for policy, or by imposing constraints on the actions of policymakers. K-P demonstrates the potential for such gain by considering games in in which the policy authority has complete credibility, i.e., precommitment or open loop leader-follower games. In these games, policy rules are set once and for all at the beginning of the game. These rules are never altered, no matter how strong the incentive to renege.

Given K-P's results, a reasonable question to ask is, to what extent do the two policy environments considered (discretion and precommitment) represent the menu of alternatives in real world policy situations? One interesting line of research that addresses this question is given in papers by Barro and Gordon (1983a, 1983b), Backus and Driffill (1985a, 1985b), Canzoneri (1985), and Tabellini (1983). In each of these papers, positive models are constructed in which monetary policy is described as resulting from games where the reputation or credibility of the monetary authority is determined endogenously. In these models, equilibrium outcomes often result in values of the policy objective intermediate between the values corresponding to the discretionary and precommitment outcomes. The theoretical potential for social gain via implementation of fixed rules is therefore less than in the K-P setup.

Sims (1982) has attempted to address this question from an empirical direction. In Sims's view (1982, p. 109), it is unrealistic to characterize the formation of policy as a once-and-for-all decision. Due to the inherently political nature of the policy process, Sims finds it unfruitful to compare discretionary outcomes to outcomes under fixed rules, since implementation of a fixed rule would not be politically feasible. Accepting this limitation on the policy environment raises the interesting possibility that, in many cases, the historical performance of policy may have approximated the best politically feasible performance.[3]

The analysis presented below represents an attempt at interpreting Sims's view of policymaking by broadening the spectrum of game theoretic models of policy available to macroeconomists. This is done by considering a class of models which are in some sense intermediate between the discretionary and precommitment models. This class of models is intended as neither purely positive nor normative, but is instead presented as a first attempt at understanding the effects of political uncertainty in a rational expectations setting. Although "political uncertainty" as defined in this class of models doesn't exactly correspond to Sims's (1982) definition, there are some interesting similarities between reduced forms implied by the models below and those advocated by Sims.

The basic idea of the models is to model policy formation not as a game between the public and a single policymaker, but rather as a game played between the public and a sequence of policymaking administrations. The administrations come into power at random intervals. Upon coming into power, each administration announces a policy plan that will hold as long as that administration stays in power. With the (random) arrival of a new administration, replanning occurs, hence the term "stochastic replanning."

For purposes of tractability, the analysis below simply assumes each administration to be completely credible. Although this concession to tractability might be seen as overly generous, the approach taken here could be viewed as roughly complementary to that taken in the reputational models of Barro and Gordon and others, which abstract from the possibility of administration turnover.[4] A complete theory of policy would probably involve neither of these abstractions.

The rest of the paper is organized as follows. Section 2 illustrates the idea of stochastic replanning in a two period model. Section 3 derives a stochastic replanning equilibrium in a simple rational expectations environment. Section 4 extends the analysis to the case where the public does not know the decision of the current policy authority. Section 5 offers some interpretations of the models presented, as well as some numerical simulations. Section 6 concludes the paper and offers some possible directions for future research. Mathematical proofs are relegated to appendices.

2. An Example

Consider the simple dynamic game depicted in Figure 1. In this game there are two players (P1 and P2) and two time periods (0 and 1). Payoffs occur at the end of period 1 and each player seeks to maximize his payoff. The players must choose either decision 0 or 1 at the beginning of each period. P1 is dominant in the sense

Figure 1: A Two Period Dynamic Game

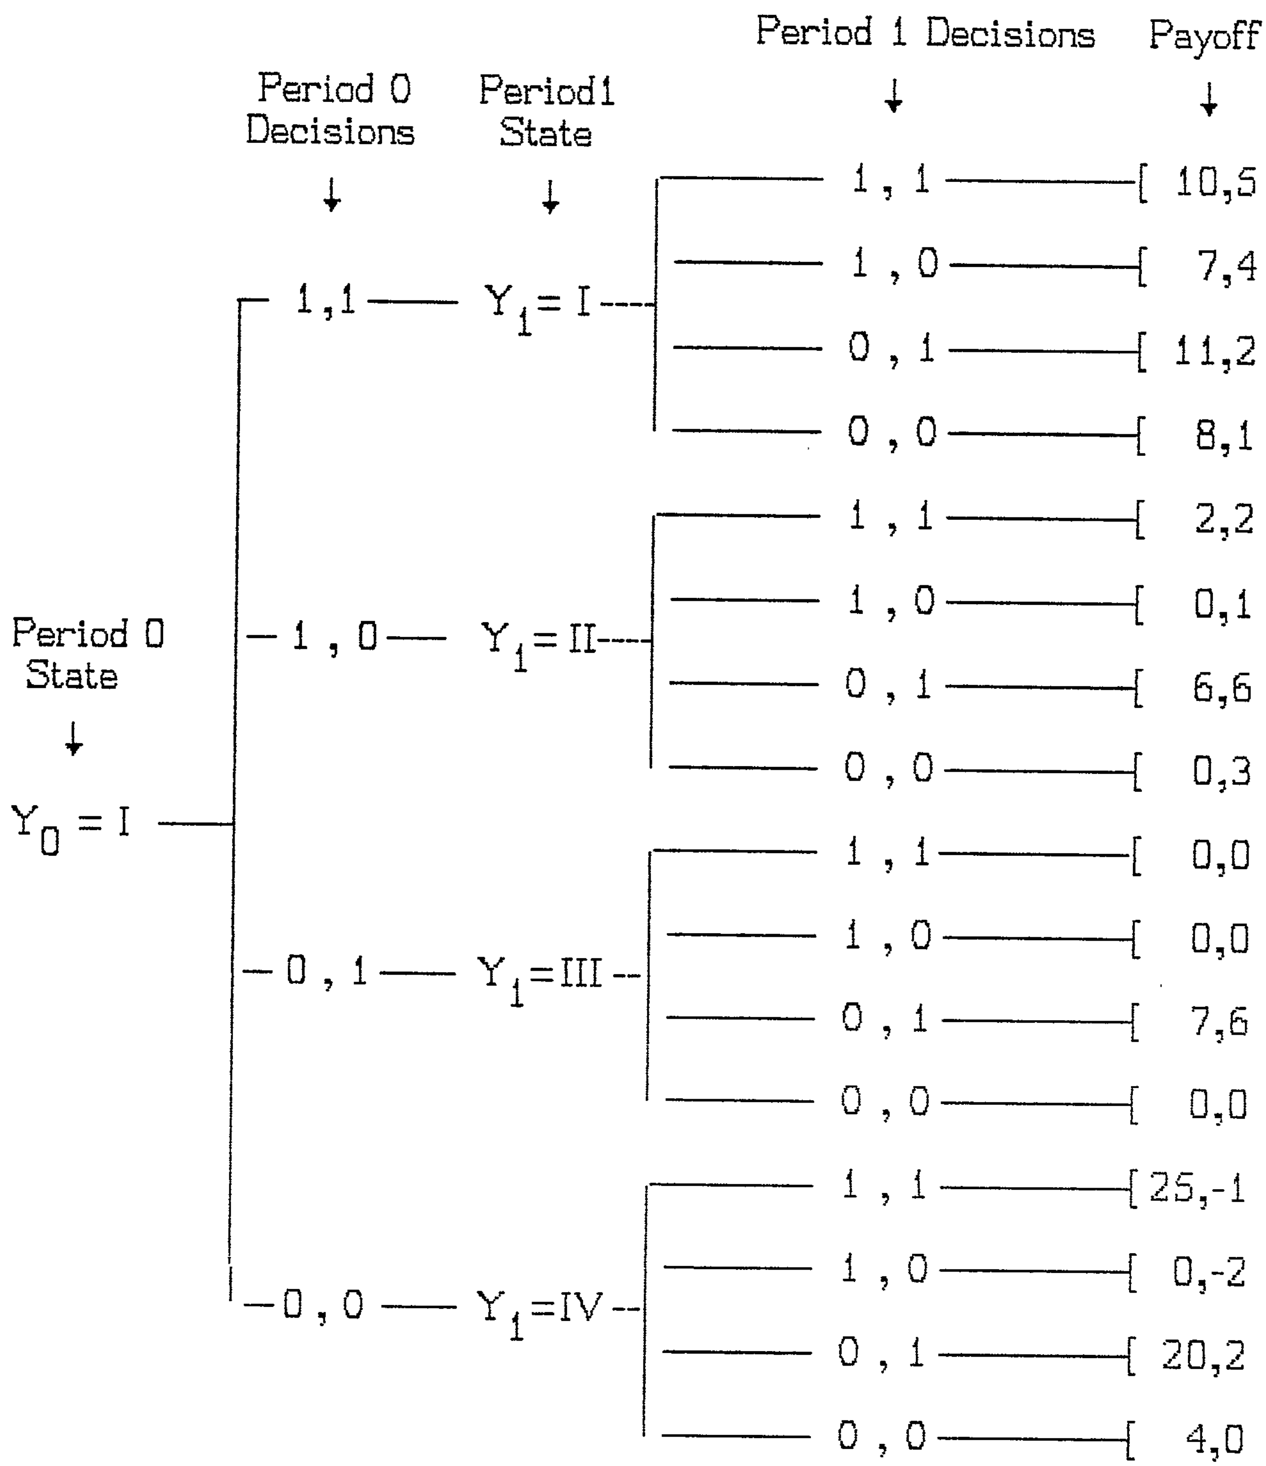

that he is first to announce his strategies, which P2 must take as given. In this sense, P1 "goes first" in each period although nothing is changed if we assume that decisions are taken simultaneously by both players in equilibrium.

The precommitment (or open loop) solution to this game can be found using the payoff matrix of Figure 2. In this solution, P1 announces a sequence of decisions which P2 must take as given in deciding which sequence to play. By enumeration, the best sequence for P1 is (1,1), in which case the optimal strategy for P2 is also (1,1), resulting in equilibrium payoffs (10,5). The time inconsistency of this solution is evident: P1 clearly has an incentive to change his period 1 decision to zero, once period 0 has passed. If this were to happen, P2 would play decision one, resulting in payoffs (11,2).

The game in Figure 1 also admits a time consistent Stackelberg solution.[5] That solution may be found by backward induction, as outlined in Figure 3. We first solve for the equilibria of each of the four possible period 1 subgames, again under the assumption that P1 goes first. Having done this, first period equilibrium decisions are then computed as in a static Stackelberg game, taking the equilibrium payoffs of the period 1 subgames as given. Proceeding in this fashion yields equilibrium strategies (0,0) for P1 and (1,1) for P2, and equilibrium payoffs (7,6).

If we seek to describe policymaking in terms of a dynamic

Figure 2: Precommitment Solution

(Player 1 Is Dominant)

P1 \ P2	0 , 0	0 , 1	1 , 0	1 , 1
0 , 0	4 , 0	20 , 2	0 , 0	7 , 6
0 , 1	0 , -2	25 , -1	0 , 0	0 , 0
1 , 0	0 , 3	6 , 6	8 , 1	11 , 2
1 , 1	0 , 1	2 , 2	7 , 4	* 10 , 5

* = equilibrium outcome

leader-follower game, the consistent solution seems the more realistic of the two solutions outlined above. The intuitive appeal of the consistent solution is increased if it is viewed as the outcome of a noncooperative game with three players: P2, representing the public, and two policy administrations, one acting at time 0 and one acting at time 1 (call them A0 and A1). Each administration seeks to maximize the payoff accruing to the period 1 administration. While the time 0 administration can predict what the time 1 administration will do, it cannot control the future administration's actions, which precludes any precommitment.[6]

Replanning Equilibrium

In the stochastic replanning solution to this game, there are also two policy administrations. However, the period 0 administration can now commit itself to period 1 actions, which will be realized with exogenously specified probability α, where $1 > \alpha > 0$. Viewed in another way, there is a probability α that A0 will not be removed from power in period 1. Conditional on this last event, the probability of A0 sticking to its original plan is one. Under this setup, the objective of the A1 is, as in the consistent game, to maximize its end of period payoff. The objective of A0 is taken as to maximize the expected payoff accruing to either administration at the end of period 1.

The stochastic planning solution can be derived by backward induction. If A1 comes into power, then it chooses its (period 1) strategies as in the consistent solution (Figure 3), irrespective of the value of α. The period 0 administration (A0) must take into account that its period 1 strategies may not be implemented, as must the public (P2). For $\alpha = 1/2$, the resulting period 0 expected payoff matrix is given in Figure 4. This matrix is calculated by averaging the entries of the matrix in Figure 2 with the appropriate subgame outcomes. Inspection of the matrix in Figure 4 reveals that the equilibrium of this new game occurs when A0 plays (0,0), A1 plays 0, and P2 plays (1,1). The equilibrium outcome for this value of α is thus the same as in the consistent solution. For sufficiently large α ($> 4/7$), however, the period 0 administration would find it optimal to play the precommitment strategies (1,1) . The reader is invited to verify this using an appropriate value of α.

3. Stochastic Replanning with an Infinite Horizon

The idea of stochastic replanning will now be extended to a more complex dynamic game. An example will be considered where private agents and policy administrations have potentially infinite planning horizons. The example to be considered will be the "generic" one considered by Whiteman (1986) in deriving expressions for the precommitment and time consistent solutions to policy games with linear rational expectations models. The example is simple enough to be tractable, yet captures the essential features of more complex models.[7] In this example, policymakers are confronted with

Figure 3: Consistent Solution

(Player 1 is dominant)

At t = 1 :

P1 \ P2	0	1
0	4 , 0	20 , 2
1	0 , -2	*25 , -1

$Y_1 = \text{IV}$

P1 \ P2	0	1
0	0 , 0	*7 , 6
1	0 , 0	0 , 0

$Y_1 = \text{III}$

P1 \ P2	0	1
0	0 , 3	*6 , 6
1	0 , 1	2 , 2

$Y_1 = \text{II}$

P1 \ P2	0	1
0	8 , 1	*11 , 2
1	7 , 4	10 , 5

$Y_1 = \text{I}$

At t = 0 :

P1 \ P2	0	1
0	25 , -1	*7 , 6
1	6 , 6	11 , 2

$Y_0 = \text{I}$

* = equilibrium outcome

Figure 4: Replanning Solution

($\alpha = 1/2$)

A0 \ P2	0 , 0	0 , 1	1 , 0	1 , 1
0 , 0	14.5, -.5	22.5, .5	3.5, 3	* 7 , 6
0 , 1	12.5, -1.5	25, -1	3.5, 3	3.5, 3
1 , 0	3, 4.5	6 , 6	9.5, 1.5	11 , 2
1 , 1	3, 3.5	4 , 4	9 , 3	10.5, 3.5

* = equilibrium outcome

$$E_t\, y(t+1) - \rho\, y(t) = x(t) + e(t) \tag{1}$$

where $y(t)$ is an endogenous random variable reflecting decisions made by the public, ρ is a parameter with $|\rho| > 1$, $x(t)$ is a variable controllable by the current policy administration, and $e(t)$ is the current realization of a forcing process. The process $\{e(t)\}$ is assumed to follow the AR(1) law

$$e(t) = \gamma\, e(t-1) + u(t) \tag{2}$$

where $|\gamma| < 1$ and $\{u(t)\}$ is Gaussian white noise. The symbol E_t is the conditional expectation operator, where the conditioning set is explained below.[8]

One macroeconomic model that gives rise to equation (1) is a version of Cagan's portfolio balance schedule:

$$\eta\, [\, E_t\, p(t+1) - p(t)\,] = m(t) - p(t) \tag{3}$$

where $\eta < 0$, $p(t)$ = log of the price level, and $m(t)$ = log of the money stock. Manipulation of equation (3) yields equation (1), via the substitutions $\rho = (\eta-1)/\eta$, $y(t) = p(t)$, and $x(t) + e(t) = m(t)/\eta$. One could interpret $-x(t)$ as proportional to the log of the monetary base, and $-e(t)$ as proportional to the log of the base multiplier.

Policy Dynamics

Turnover of administrations is governed by the sequence of independent Bernoulli random variables $\{S(t)\}$, where

$$S(t) \equiv \begin{cases} 0 & \text{if replanning occurs (a new administration comes into power) at time } t; \\ 1 & \text{otherwise.} \end{cases}$$

There is always probability α that the current administration will continue in power next period, which together with the serial independence assumption implies the convenient properties:

(C1) $\Pr\{\, S(t+n) = j \mid S(t), S(t-1), \cdots \} = \Pr\{\, S(t+n) = j \,\}$

$$= \left\{ \begin{array}{ll} 1 - \alpha & \text{for } j = 0 \\ \alpha & \text{for } j = 1 \end{array} \right\}, \quad \text{for } n = 1, 2, 3, \cdots$$

(C2) The probability of any current administration being in power n periods in the future is α^n.

The public and the administration currently in power are assumed to know the complete histories of $\{e(t)\}$ and $\{S(t)\}$ up to time t. The processes $\{e(t)\}$ and $\{S(t)\}$ are assumed to be independent at all leads and lags.

The objective of the administration coming to power at time t will be taken as to minimize

$$\lim_{J\to\infty} \tfrac{1}{2} E_t \sum_{j=0}^{J} \beta^j \left[y(t+j)^2 + \lambda\, x(t+j)^2 \right] ,$$

where $\lambda > 0$ and $1 > \beta > 0$, by choice of a plan for $\{x(t+j)\}_{j=0}^{\infty}$, taking the plans of future administrations as given, and taking into account the public's decision rule.[9] The rule for $x(t+j)$ may only depend on j and $e(t+j)$, $e(t+j-1)$, $\cdots$.

A stochastic replanning equilibrium for this model is a pair of sequences ($\{x^*(t)\}$, $\{y^*(t)\}$) such that for any realized path of $\{S(t)\}$ and $\{e(t)\}$:

(E1) The public's expectational equation (1) holds;

(E2) For all t, $x^*(t)$ reflects the optimal choice of the administration currently in power.

The cases where $\alpha = 0$ and $\alpha = 1$ correspond respectively to the time consistent and precommitment Stackelberg games analyzed by Whiteman (1986).

Although this example is greatly simplified, it incorporates important features of policy problems under rational expectations. The objective of the policymaker(s) is to minimize the discounted weighted sum of squared fluctuations in the "target variable" $y(t)$ and the "policy instrument" $x(t)$. In considering the effects of policy, each policymaker must take into account the effect of policy choice on the public's anticipations . This can be seen by manipulating equation (1) as in Sargent (1979, p.269) to yield

$$y(t) = -\rho^{-1} \sum_{j=0}^{\infty} \rho^{-j} E_t [x(t+j) + e(t+j)] . \tag{4}$$

The choice of $y(t)$ thus reflects the public's anticipations of all future actions by policymakers.

Derivation of the Stochastic Replanning Solution[10]

To derive the stochastic replanning solution, I now make two additional simplifying assumptions. First, following Whiteman (1986), the discount factor β is taken to equal 1. This assumption does not affect the qualitative properties of the solution and reduces notational complexity. Alternatively, we could think of the results below as applying to the case where each administration has the average cost objective

$$\lim_{J \to \infty} \frac{1}{2J} E_t \sum_{j=0}^{J} [\, y(t+j)^2 + \lambda\, x(t+j)^2 \,] .$$

This interpretation is somewhat problematic, however, because each administration, when $1 > \alpha \geq 0$, is optimizing over a sequence of events that ultimately has zero probability. This causes its optimization problem to be ill-defined.

Second, I assume certainty equivalence and solve the model first for the case that $\mathrm{var}(u(t)) = 0$. This is common practice for linear-quadratic-Gaussian models. Because each administration's objective function turns out to be time nonseparable, however, the usual theorems justifying this practice will not apply.[11] Absent more general results, the analysis below constrains each administration's policy rule to be linear in the shocks $e(t), e(t-1), \cdots$.[12]

To evaluate expectations when $\{e(t)\}$ is nonstochastic, first define the process $w(t)$ as

$$w(t) \equiv \min_{j} \{\, t - j \,\} \quad \text{s.t. } j \leq t,\ S(j) = 0.$$

In other words, $w(t)$ is the age of the administration in power at time t. Denoting the administration coming to power at time t as $A(t)$, the administration currently in power at a given time τ is $A[\tau - w(\tau)]$.

The process $\{w(t)\}$ is a Markov chain with state space equal to $\{0, 1, 2, \cdots\}$, and transition matrix $\prod$ given by[13]

$$\prod \equiv \begin{bmatrix} (1-\alpha) & \alpha & 0 & 0 & \cdots \\ (1-\alpha) & 0 & \alpha & 0 & \cdots \\ (1-\alpha) & 0 & 0 & \alpha & \cdots \\ \vdots & \vdots & \vdots & \vdots & \end{bmatrix}$$

The model is constrained so that an administration coming into power at time t can only maximize along the path $w(t+j) = j$ for $j = 0, 1, 2, \cdots$. Hence $x(t)$ and $y(t)$ will be $w(t)$-measurable, and can be written $x(t,w(t))$, and $y(t,w(t))$. The plan by $A(t)$ for $x(t+j)$ will be denoted $x(t+j, j)$.

Exploiting the Markov nature of $w(t)$, equation (4) can then be rewritten in the form

$$y(t,w(t)) - e(t) / (\gamma - \rho)$$

$$+ \rho^{-1} \sum_{k=0}^{\infty} \rho^{-k} \sum_{m=0}^{w(t)+k} x(t+k, m) \Pr\{w(t+k) = m \mid w(t)\} = 0 \qquad (5)$$

That is, the form of the transition matrix Π guarantees that k periods from now, w(t) must take on a value between 0 and w(t) + k. The term E_t x(t+k) can then be evaluated by weighting the w(t) + k possible values of x(t+k) by their probabilities and summing. Similarly, an administration coming to power at time t will have "objective function"

$$J(t) \equiv \tfrac{1}{2}\sum_{j=0}^{\infty} \left\{ \sum_{n=0}^{j} [\, y(t+j, n)^2 + \lambda x(t+j, n)^2 \,] \right\}$$

$$\times \Pr \{ w(t+j) = n \mid w(t) \}$$

Following the approach used in Hansen, Epple, and Roberds (1985), the problem of administration A(t) can be solved using the Lagrangian

$$\mathcal{I}(t) \equiv J(t) + \sum_{j=0}^{\infty} \left\{ \sum_{n=0}^{j} \theta(t+j, n)\, c(t+j, n) \right\}$$

$$\times \Pr \{ w(t+j) = n \mid w(t) = 0 \}$$

where c(t,w(t)) ≡ LHS of equation (5), and $\theta(\cdot)$ is a random Lagrange multiplier. This is done by differentiating $\mathcal{I}(t)$ with respect to x(t+j,w(t+j)) and y(t+j,w(t+j)) along the path w(t+j) = j. By differentiating $\mathcal{I}(t)$ with respect to y(t+j,j) the administration A(t) is allowed to choose y(t) along one path for w(t). However, *y(t)* must be chosen in a manner consistent with the public's decision rule, as expressed in equation (5).

By carrying out the differentiation described above the following result can be shown (see Appendix A):

Lemma 1. In equilibrium, the relations

$$y(t) = (\lambda\rho)\, x(t), \text{ when } S(t) = 0 \tag{6}$$

or

$$y(t) = (\lambda\rho)\, x(t) - \lambda\, x(t-1)\ , \text{ when } S(t) = 1 \tag{7}$$

must hold. □

Lemma 1 implies that an equilibrium pair of sequences $(x^{\star}, y^{\star})$ must satisfy conditions (5), (6), and (7). Equations (6) and (7) give a representation of the optimal plan for $x(t)$ as a function of $y(t)$ and lagged $x(t)$. The form of this representation is independent of α.

Simultaneously solving equations (5), (6), and (7) then yields the following result (see Appendix A):

Theorem 1. One stochastic replanning equilibrium is defined implicitly by the equations

$$x^{\star}(t) = f_0\, e(t) + f_1\, S(t)\, x^{\star}(t-1) \tag{8}$$

$$y^{\star}(t) = b_0\, e(t) + b_1\, x^{\star}(t) \tag{9}$$

where f_0, f_1, b_0, and b_1 are functions of λ, α, ρ, and γ. This equilibrium is unique in the class of equilibria satisfying the stability condition

$$E_t\, x(t+j)\ ,\ E_t\, y(t+j) \to 0 \quad \text{as } j \to \infty\ .\ \square$$

Note that by manipulating equation (8), one can obtain the representation

$$x^{\star}(t) = f_0 \sum_{j=0}^{w(t)} f_1{}^{j}\, e(t-j) \tag{10}$$

implying that $x^{\star}(t)$ and $y^{\star}(t)$ are, as explained above, $w(t)$-measurable when $\{e(t)\}$ is nonstochastic.

4. The Delayed Information Model

Using the same approach as in Section 3, it should be possible to derive stochastic replanning equilibria for a number of different policy environments. For example, the setup of the previous section can be modified to reflect an assumption that the current policy administration knows more than the public.

Following the setup of Sims (1985), suppose that each time period is split into two subperiods, $t - \frac{1}{2}$ and t. The public chooses y(t) at time $t - \frac{1}{2}$. At time t, w(t) and e(t) are realized, and x(t) is chosen. Each of the last three events is observed by the public. At time $t + \frac{1}{2}$, the public chooses y(t+1), etc. Under this setup, the public does not know the time t values of e(t) nor w(t) when making its decision for time period t, but the current (time t) administration does. For this reason, the model of this section will be referred to as the "delayed information model." The model of the previous section will be referred to as the "contemporaneous information model."

To understand the effects of delayed information, it may be useful to reconsider the two stage game depicted in Figure 1. In this game, the assumption of delayed information amounts to reversing the order of play at each stage, so that P2 goes first at each stage. P1 is still dominant in the sense of being first to announce strategies. Thus, in the case where $\alpha = 1$, reversing the order of play has no impact on the outcome of the game. However, in the replanning case ($1 > \alpha > 0$), this assumption is less innocuous. To see this, consider A1's problem at the second stage of the game (see Figure 5). In this stage, P2 moves first, without knowing whether A1 has come to power. All that A1 can do is then choose an optimal reaction to P2's second stage move. The resulting outcomes are marked with an asterisk in Figure 5.[14] As with the contemporaneous information model, equilibrium strategies for A0 and P2 can then be found by considering a weighted average of the outcomes shown in Figure 5 and the payoff matrix in Figure 2, the weights

Figure 5: Replanning Solution

(Follower Has Delayed Information)

At t = 1 :

A1 \ P2	0	1
0	*4 , 0	20 , 2
1	0 , -2	*25 , -1

Y_1 = IV

A1 \ P2	0	1
0	*0 , 0	*7 , 6
1	0 , 0	0 , 0

Y_1 = III

A1 \ P2	0	1
0	*0 , 3	*6 , 6
1	0 , 1	2 , 2

Y_1 = II

A1 \ P2	0	1
0	*8 , 1	*11,2
1	7 , 4	10, 5

Y_1 = I

* = optimal reaction of A1

being respectively α and $1-\alpha$. Performing this computation for $\alpha = 0.9$ yields the expected payoff matrix shown in Figure 6. Inspection of this matrix reveals that the equilibrium strategies of A0 and P2 are the same as in the precommitment game. The expected payoff of A0, however, is slightly higher, 10.1 versus 10 for the case where $\alpha = 1$.[15] Finally, we can derive the consistent solution for this case by taking limits as $\alpha \downarrow 0$.

Returning to the stabilization example, in the delayed information case the public's decision rule can be represented as

$$y(t) = -\rho^{-1} \sum_{j=0}^{\infty} \rho^{-j} E_{t-1}[\, x(t+j) + e(t+j) \,] \tag{11}$$

The policy administration's objective functions remain as before. In Appendix B, the following result (analagous to Theorem 1) is shown for the delayed information model:

Theorem 2. For the delayed information model, the unique stable stochastic replanning equilibrium is defined implicitly by the equations

$$x^{\star}(t) = g_0\, S(t)\, e(t-1) + g_1\, S(t)\, x^{\star}(t-1) \tag{12}$$

$$y^{\star}(t) = m_0\, e(t-1) + m_1 x^{\star}(t-1) \tag{13}$$

where g_0, g_1, d_0, and d_1 are functions of λ, α, ρ, and γ.□

Figure 6: Replanning Solution

(Follower Has Delayed Information, and $\alpha = 0.9$)

A0 \ P2	0 , 0	0 , 1	1 , 0	1 , 1
0 , 0	4, 0	20.5, 1.7	0, 0	7 , 6
0 , 1	.4, -1.8	25, -1	0, 0	.7, .6
1 , 0	0, 3	6 , 6	8, 1	11 , 2
1 , 1	0, 1.2	2.4 , 2.4	7.1, 3.7	* 10.1, 4.7

* = equilibrium outcome

An interesting feature of the delayed information model is that the optimal value of x(t) is always zero for an administration coming to power at time t. Correspondingly, the optimal time consistent strategy for this model consists of setting $x(t) = 0$ for all t, i.e., of never intervening under any circumstances.

5. Discussion and Numerical Simulations

The models presented in the two previous sections posess equilibria that are, loosely speaking, intermediate between the precommitment and time consistent models. When $1 > \alpha > 0$, the optimal strategy of each administration is time inconsistent. The public knows, however, that the current administration's strategy cannot remain in place forever, and takes that fact into account in formulating its decisions.

A characteristic feature of the stochastic replanning equilibria is the nonlinearity of the equilibrium laws of motion for x(t), as seen in equations (8) and (11). Alternatively, we could think of these equations as being linear with randomly time varying coefficients. In this sense, the models presented above give rise to reduced forms that resemble those suggested by Sims (1982) as useful for predicting the effects of policy. As with Sims's setup, optimal predictors of future x(t) and y(t) can be constructed using the time varying VAR model implied by equations (8) and (11). However, the form of the optimal predictors will be linear (in the x and e processes) and time invariant.[16]

Evaluating Relative Policy Performance

The nonlinear nature of the law of motion for x(t) also renders somewhat cumbersome closed form comparisons of policy performance (i.e., of J(t) for different values of α). For this reason, Monte Carlo simulations were used for these comparisons. The results of 21 such simulations are reported in Table 1. Simulations 1-9 are of the contemporaneous information model, while simulations 10-21 are of the delayed information model. In each of the simulations, the parameter values $\rho = 1.1$, $\gamma = 0.9$, and var(u(t)) = 1.0 were assumed. The parameter λ was allowed to take on the values 1.0, 10, and 0.1. For simulations 1-18, the parameter α was set to 0.0, 0.5, and 1.0, while the value $\alpha = 0.9$ was assumed for simulations 19-21. A random number generator was used to construct artificial time series of length T = 500 for e(t), and S(t) in the cases where $1 > \alpha > 0$.[17] As an approximation to $2T^{-1}J(t)$, the statistic

$$S(\alpha,\lambda) \equiv \text{svar}(y) + \lambda \text{svar}(x)$$

was calculated for each simulation, where "svar" means sample variance. A sample performance index was then calculated as

$$P(\alpha,\lambda) \equiv 100\,[S(\alpha,\lambda)/S(0,\lambda)]$$

Table 1: Evaluation of Policy Performance

Simulation	λ	α	Performance Index P (%)
(Contemporaneous Information Model)			
1	1	0	100.0
2	1	0.5	76.9
3	1	1	58.6
4	10	0	100.0
5	10	0.5	74.6
6	10	1	44.0
7	0.1	0	100.0
8	0.1	0.5	95.8
9	0.1	1	91.2
(Delayed Information Model)			
10	1	0	100.0
11	1	0.5	3.15
12	1	1	2.91
13	10	0	100.0
14	10	0.5	27.8
15	10	1	20.2
16	0.1	0	100.0
17	0.1	0.5	.204
18	0.1	1	.319
19	1	0.9	2.50
20	10	0.9	18.1
21	0.1	0.9	.274

The index P gives the sample performance of policy as a percentage of the performance of the best consistent policy for that model.

For the contemporaneous information model, simulations 1-9 suggest that the performance of policy improves (i.e., the relative index P decreases) monotonically as α increases. For example, when $\lambda = 1$, increasing α from 0 to 0.5 yields a 33% drop in the policy loss function, while increasing α from 0 to 1 yields a 42% decrease (see simulations 1-3). Simulations 1-9 also suggest that the importance of precommitment increases with the weight (λ) attached to fluctuations in the policy instrument x(t). This is evidenced by the values of P obtained when $\alpha = 0.5$ or 1.0 for the various values of λ that were tried. For example, when $\alpha = 1$, improvement over the time consistent case is only 8.8% when $\lambda = 0.1$, but increases to 66% when $\lambda = 10$ (see simulations 6 and 9). In fact, I was unable to reverse these orderings, even after trying many different parameter values. It would therefore seem reasonable to conclude that they will always hold for this particular model.

Simulations 10-21 indicate that different orderings can hold for the delayed information model. While interventionist policies (corresponding to $\alpha > 0$) appear to dominate the passive policy of always setting $x(t) = 0$ (corresponding to $\alpha = 0$), the performance of policy no longer monotonically improves with increasing α. For example, when $\lambda = 0.1$, the performance of policy when $\alpha = 0.5$ is better than the performance of the precommitment ($\alpha = 1.0$) policy (see simulations 17 and 18). For all values of λ considered, the performance of policy when $\alpha = 0.9$ dominates the performance under $\alpha = 1.0$ (compare simulations 19 and 12, 20 and 15, 21 and 18). Intuitively, this can happen because in the delayed information model, policymakers intervene only after the public has committed itself. With a positive cost attached to policy interventions, it is not surprising that a policy environment in which administrations randomly decline to intervene can result in a better performance than an environment of perfect precommitment. Still, it should be noted that these gains are relatively small.

This last result is not really surprising when viewed from a game theoretic perspective. If we aggregate the administrations into a single player, then this result simply seems to imply the dominance of mixed strategies for this player, in this particular game. The sort of randomization considered, however, is not arbitrary, but corresponds roughly to the notion of "political uncertainty" (i.e., uncertainty about who will be setting policy in the future).

The delayed information model also contrasts with the contemporaneous information model in terms of the effects of changes in the λ parameter on relative performance. Performance relative to the time consistent case improves (again, P declines) as λ decreases in simulations 10-21.

Overall, the numerical simulations illustrate the extreme sensitivity of relative policy performance to changes in the model parameters and to changes in assumptions regarding the public's information.

Simulating the Effects of Policy

Some of the interesting time series properties of the stochastic replanning equilibria can be readily illustrated by simulating the contemporaneous information model for an admittedly extreme case. Figures 7-11 depict one set of simulations for this model, for the parameter values $\lambda = 0.5$, $\rho = 1.01$, and $\gamma = 0.99$. These values were chosen so as to highlight the dynamic properties of the model. The value of γ close to 1 causes shocks to e(t) to take on an almost permanent character, while the value of ρ close to 1 guarantees that potential future fluctuations in x(t) and e(t) will not be highly discounted. The relatively small value of λ indicates that only small penalties are associated with fluctuations in the policy instrument x(t).

Using artificial e(t) and S(t) time series, the model was simulated for 1000 time periods, for α equal to 0.0, 0.5, and 1.0. For graphical clarity, the resulting x(t) and y(t) series were sampled at every fourth observation. Figure 7 depicts the x(t) series for $\alpha = 1.0$ and 0.5. Figure 8 shows the x(t) series for $\alpha = 0.5$ and 0.0. A striking feature of these graphs is the similarity of the x(t) series for the various values of α. In fact, the time *paths* for x(t) are virtually identical across policy environments. This similarity is deceptive, however. From the standpoint of the public, the

Figure 7

Sampled x(t) Series for the Contemporaneous Information Model

($\lambda = 0.5, \quad \rho = 1.01, \quad \gamma = 0.99$)

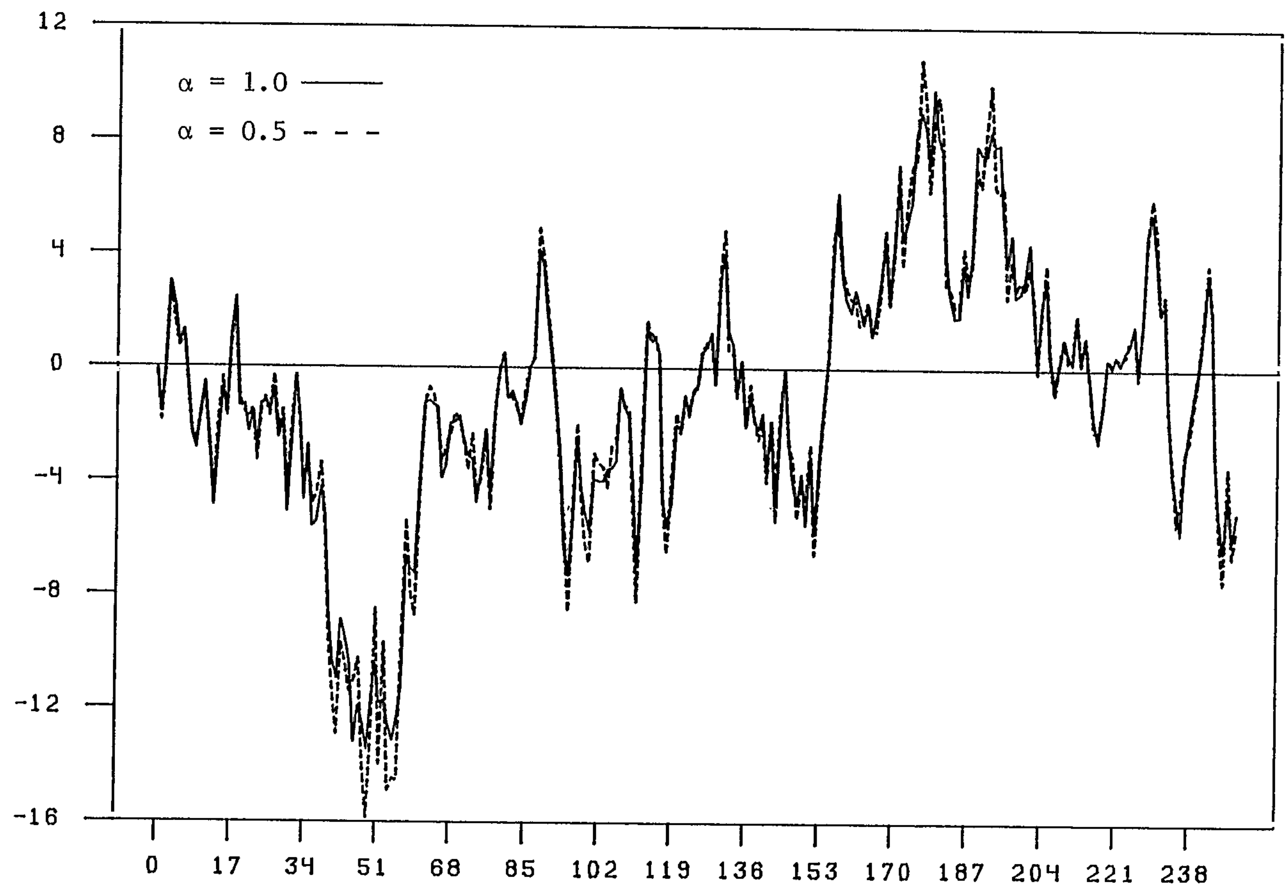

Figure 8

Sampled x(t) Series (continued)

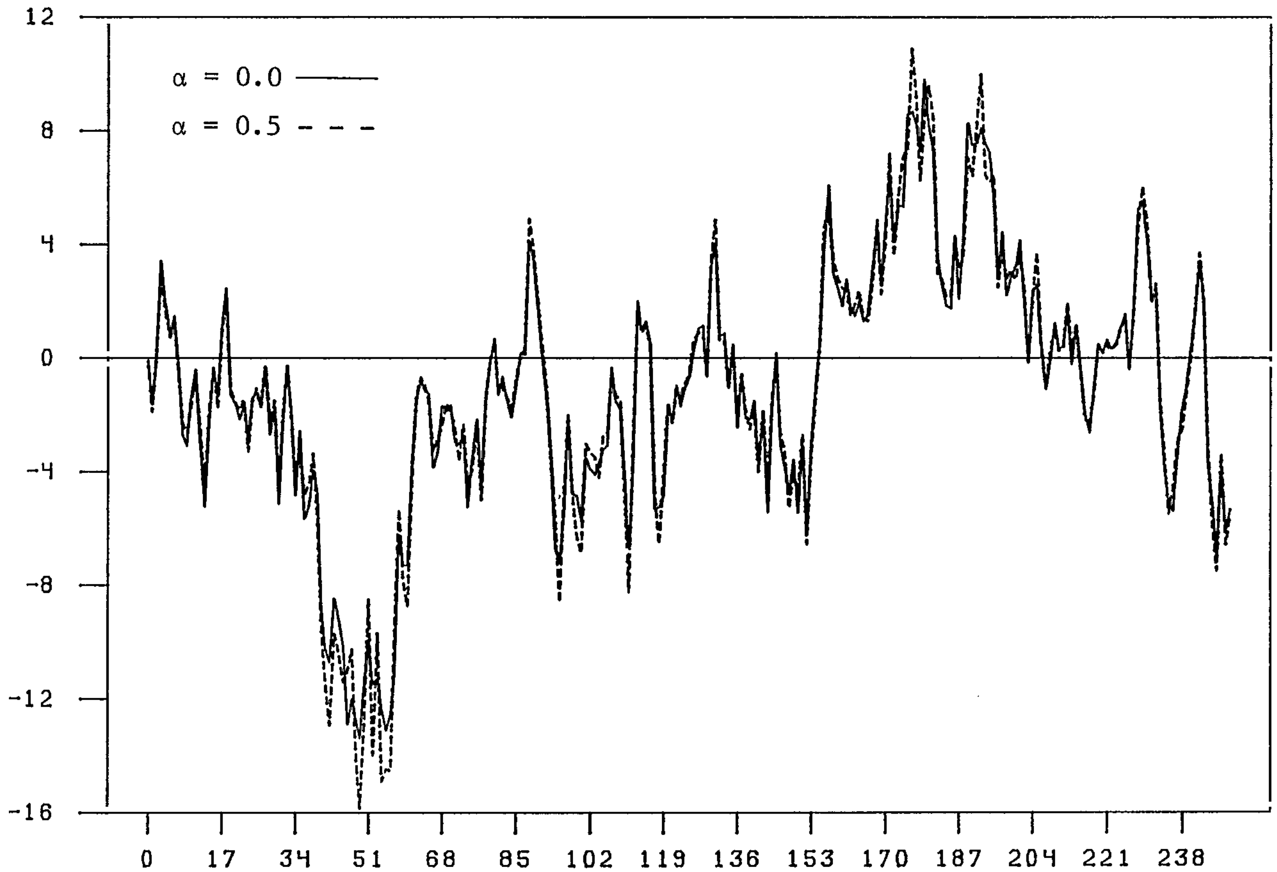

Figure 9

Sampled y(t) Series for the Contemporaneous Information Model

($\lambda = 0.5$, $\rho = 1.01$, $\gamma = 0.99$)

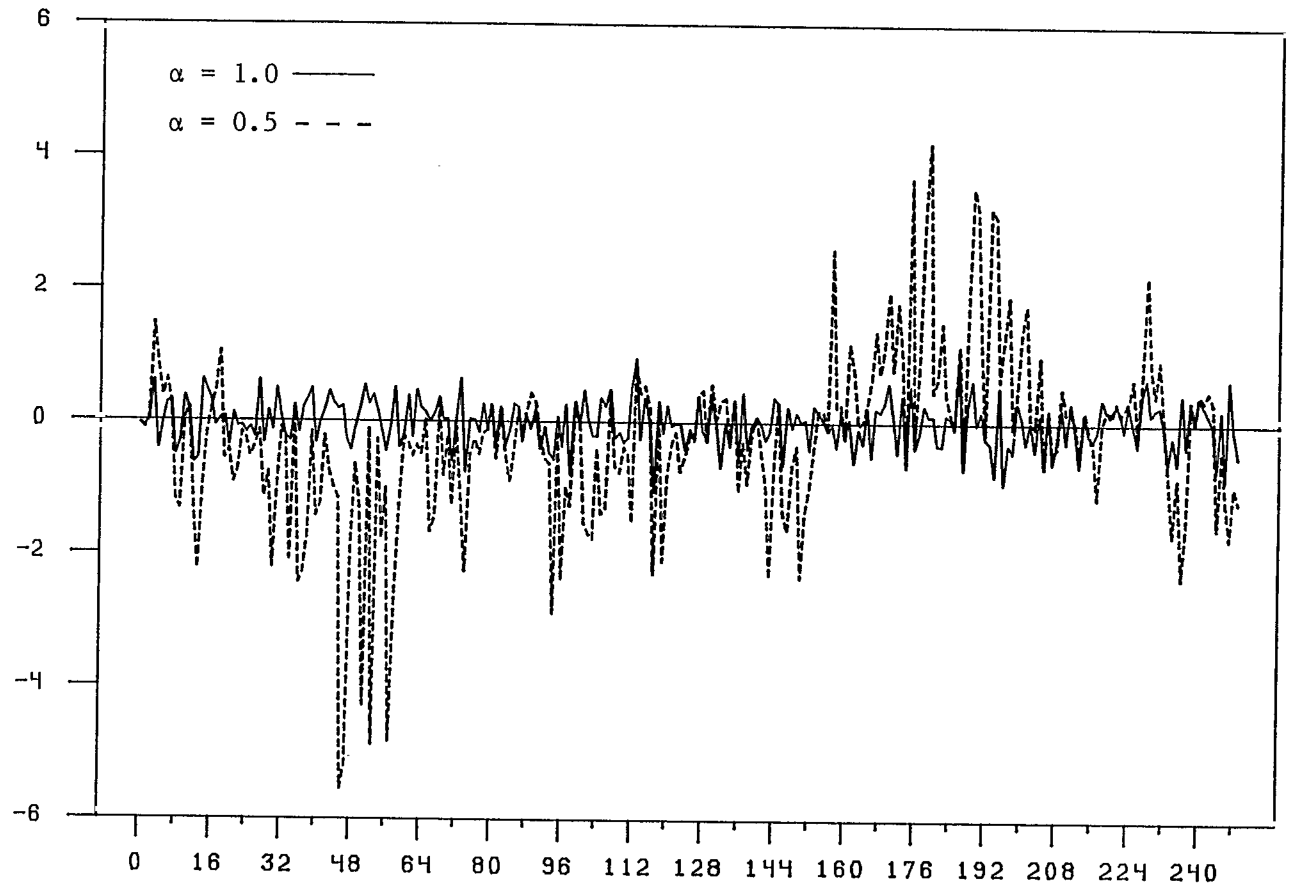

Figure 10

Sampled y(t) Series (continued)

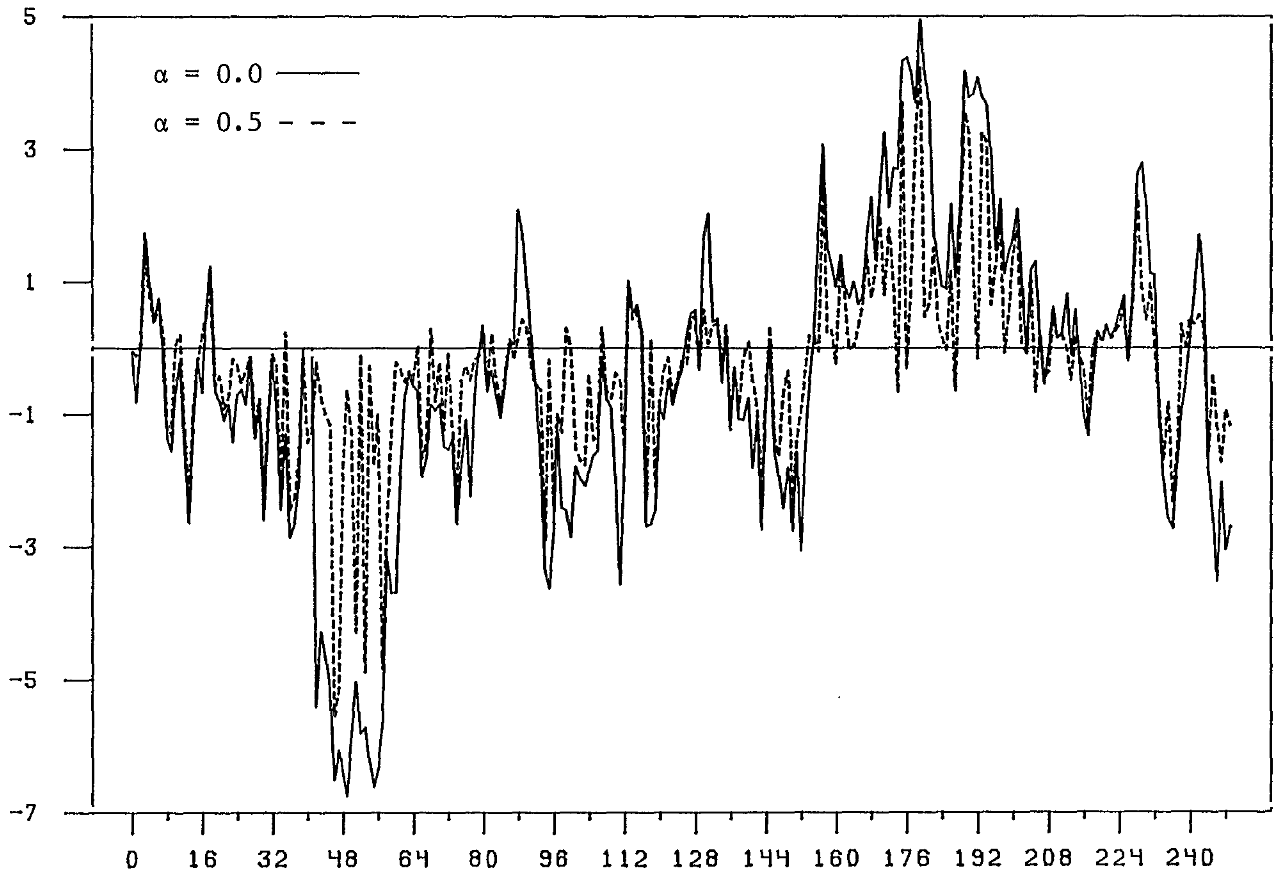

Figure 11

Logs of Estimated Spectra for the y(t) Series

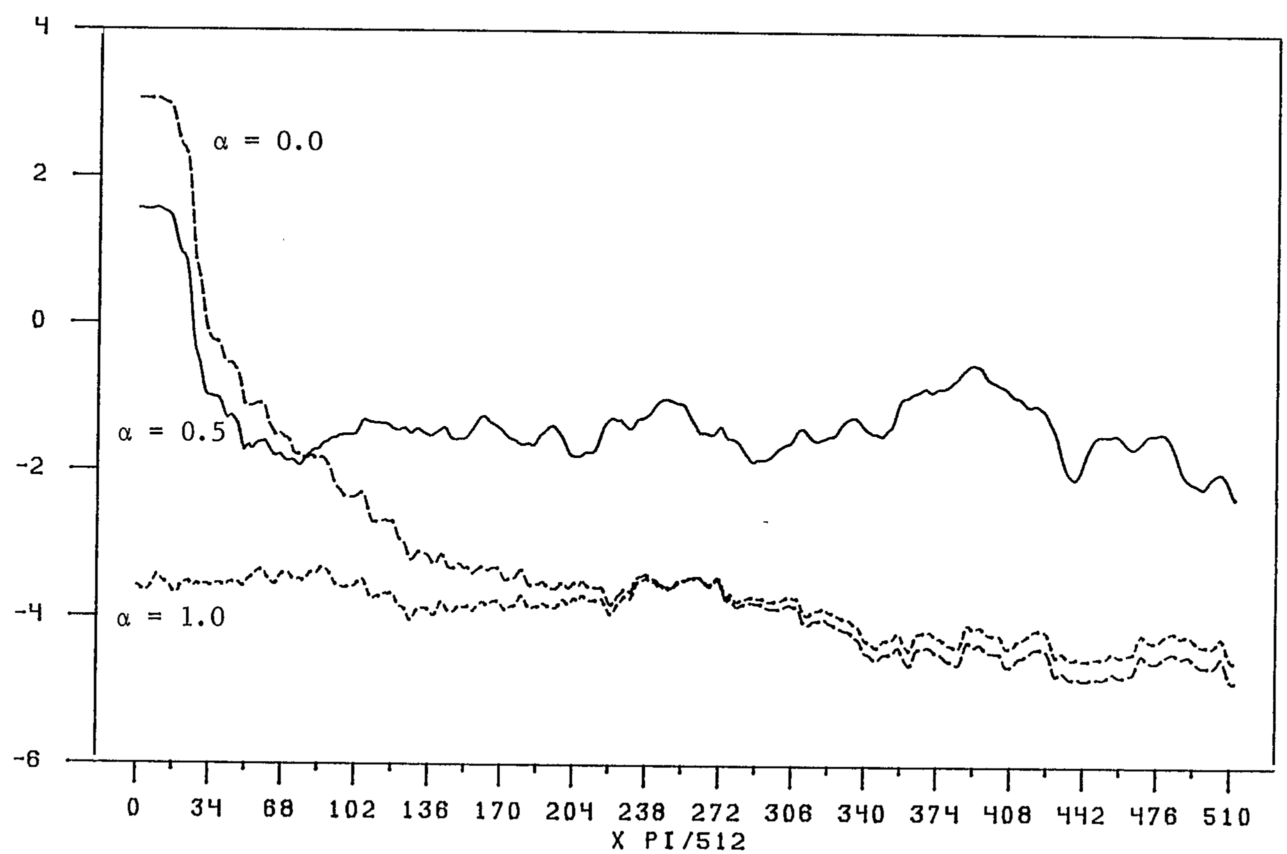

policy process x(t) becomes slightly more predictable as α goes from 0 to 1. This slight increase in predictability causes the variance of the target variable y(t) to decrease as the probability of policy continuation increases.

It is also revealing to note the manner in which this decrease occurs. Figures 9 and 10 depict the sampled y(t) time series for this set of simulations. In contrast to the x(t) series, the y(t) series differ markedly across policy environments . Comparing the cases where $\alpha = 1.0$ and 0.5 (see Figure 9) indicates that the average magnitude ("amplitude") of fluctuations in y(t) increases dramatically as α decreases over this range. Figure 10 shows that as α decreases from 0.5 to 0.0, the magnitude of the fluctuations remains about the same, while the persistence of the fluctuations increases. Figure 11 offers a comparison of the three y(t) series in the frequency domain. For each value of α, Figure 11 shows the log of an estimated spectrum of the (original nonsampled) y(t) series. Although the variance or power of y(t) falls as α increases, this reduction in power is not uniform across frequencies. As α goes from 0.0 to 0.5, the power of y(t) apparently increases at the higher end of the spectrum. The transition from a time consistent to a stochastic replanning policy environment thus causes the overall variance of y(t) to decrease, but also introduces more short term fluctuation into this process. In the case of perfect precommitment, the power of y(t) is uniformly low across all frequencies.

Again, these very different y(t) series were generated from virtually identical $x(t)$ series. The processes that generated these series, however, are quite different from the point of view of the public: the higher is the value of α, the more predictable is the course of future policy. The fact that $x(t)$ becomes more predictable, however, doesn't necessarily result in a smoother $y(t)$ series, at least for values of α close to 0.5. For such values of α, the simulations show that stochastic replanning policies trade off increased short term fluctuations in the target variable $y(t)$ for reduced persistence of these fluctuations.

Finally, it should be noted that time series properties similar to those described above were found in simulations of the model for many different parameter values. However, differences among the $x(t)$ series tended to be greater, and among the $y(t)$ series less, than in the example discussed above. Intuitively, this is what one would expect, as moving the parameters away from the extreme values assumed for the example generally reduces the efficacy of policy interventions. Increasing the value of λ, for example, generally has the effect of shifting the policy instrument $x(t)$ towards 0, and thereby reduces the effect of an optimal policy on the target variable $y(t)$. Driving ρ away from 1 has a similar effect, since this causes the public to discount more heavily future expected values of $x(t)$. Also, reducing the value of γ reduces the need for predictable policy, as the destabilizing shocks $e(t)$ become less predictable.

6. Summary and Conclusion

The implications of exogenous, stochastic regime turnover have been explored for two versions of a simple model of stabilization under rational expectations. The probability of regime turnover has been shown to be a potentially important determinant of policy performance, and of the serial correlation properties of the models' endogenous variables. Below, two potential roles are suggested for these models in interpreting the issues outlined in the Introduction.[18]

One view would be that some dynamic game with stochastic replanning constitutes a reasonable approximation of actual macroeconomic policymaking. This view abstracts from the credibility issue discussed in the Introduction. Under this interpretation, the potential gains to precommitment should be evaluated not by contrasting outcomes under the consistent and precommitment ($\alpha = 0$ and 1) cases, but by contrasting the prevalent ($1 > \alpha > 0$) replanning and precommitment outcomes. The numerical exercises of Section 5 then suggest that such gains, correctly evaluated, could be relatively small if not negative.

Another interpretation would reject stochastic replanning models as positive models of policy in favor of the time consistent model. Because of the political necessity of replanning, however, the performance of policy for some stochastic replanning model would serve as a first best standard against which alternative policies could be compared. The simulations reported in Table 1 suggest that in this case, comparisons of historical policy performance against some replanning standard could differ significantly from comparisons based on a precommitment standard.

Given the preliminary nature of the models considered in this paper, the validity of either interpretation is open to question, and other interpretations are possible. Still, the examples presented above strongly suggest that in considering the effects of policy, environments other that pure discretion and precommitment should be analyzed.

Suggestions for Future Research

The models of this paper represent a first attempt at analyzing the implications of stochastic replanning in a rational expectations policy environment. It would be clearly desirable to extend this analysis to more complicated and presumably more realistic models. One nontrivial extension would be to apply this paper's methods to general linear-quadratic models. This would be a nontrivial extension because these methods exploit the fact that for the models considered, the optimal strategy of each administration will be independent of initial conditions. In more general models, such independence is unlikely to obtain.[19] A second extension would be to allow for a more general probability structure for the process governing administration turnover. A third important extension would be to allow for reputational effects. Finally, administration turnover could be modeled as the endogenous outcome of a process of rational political choice.

Notes

1. See any of the nonmathematical references at the end of this paper, e.g., Sims(1982).

2. The existence of such rules clearly depends on the existence of a policy environment that would make such rules credible.

3. In Sims (1985), an example is given where a "naive" policymaker, using a reduced form for policy analysis, ends up choosing policy in a socially optimal fashion. Also see Cooley, Leroy, and Raymon (1984) for a discussion of this issue.

4. This point is made by Taylor (1983). These papers also abstract from all dynamics other than reputational effects.

5. This is the feedback solution discussed by Kydland (1977) and Basar and Olsder (1982).

6. The administration interpretation of the time consistent solution is due to Lars Hansen. Hansen, Epple, and Roberds (1985) show that this idea can be used to construct time consistent Stackelberg solutions to certain linear-quadratic games.

7. A key simplifying feature of this example is that equation (1) only admits a zero stable homogeneous solution. This means that initial conditions can be ignored in solving the partial dif-

ference equations that correspond to each administration's first order conditions (see Appendix A).

8. Solutions to equation (1) are discussed by Whiteman (1983) and Watson (1985), for the case where $x(t) = 0$ for all t. For reasons of tractability, "bubble" or "sunspot" solutions of the sort considered by Watson are excluded in the analysis that follows.

9. In the terminology of Kydland (1977), the admissible strategies are "open loop." Other authors, e.g., Basar and Olsder (1982) and Buiter (1981), do not allow open loop strategies to depend on random shocks.

10. Readers not interested in the technical details of the solution could skip this section, except for the statement of Theorem 1.

11. A linear-quadratic-Gaussian structure does not automatically guarantee certainty equivalence. For a useful summary of certainty equivalence results see Witsenhausen (1971).

12. Basar and Bagchi (1981) and Levine and Currie (1984) have shown that certainty equivalence holds in related models (essentially continuous time versions of the case $\alpha = 1$ above).

13. See, for example, Chung (1975) or Hoel, Port, and Stone (1972) for a treatment of Markov chains.

14. In Figure 5, possible nonuniqueness of A1's strategy has been arbitrarily resolved.

15. The value of the leader's objective could be driven even higher if the administrations were allowed to use "threat" or "forcing" strategies, as in the closed loop Stackelberg game described by Basar and Olsder (1982). However, such strategies would seem of limited use in macroeconomic settings and are thus not considered here.

16. To see this, note that the law of iterated expectations implies that the optimal one-step-ahead predictor for $x(t)$ will be given by replacing $S(t)$ with α on the RHS of equations (8) and (12). Using the linearity of equations (9) and (13) and an induction argument, it follows that k-step-ahead predictors of $x(t)$ and $y(t)$ are linear and invariant with respect to $S(t)$.

17. Sample values of the performance statistic P were generally robust to the use of different artificial time series.

18. The two interpretations offered below derive from the discussion by Sargent (1984a).

19. See, for example, the proportional taxation model analyzed by Sargent (1984b).

Appendix A. Solution of the Contemporaneous Information Model

Proof of Lemma 1: Differentiating $\mathcal{I}(t)$ with respect to $x(t+j,j)$ and $y(t+j,j)$ yields the first order conditions

$$\frac{\partial \mathcal{I}(t)}{\partial y(t+j,j)} = \left[y(t+j,j) + \theta(t+j,j) \right] \alpha^{j}$$

$$\frac{\partial \mathcal{I}(t)}{\partial x(t+j,j)} = \left[\lambda x(t+j,j) + \rho^{-1} \sum_{k=0}^{j} \rho^{-k}\, \theta(t+j-k,j-k) \right] \alpha^{j}.$$

Manipulation of the above two equations yields

$$y(t+j,j) = \lambda\rho\, x(t+j,j), \text{ when } j=0 \tag{A.1}$$

$$y(t+j,j) = \lambda\rho\, x(t+j,j) - \lambda\, x(t+j-1,j-1), \text{ when } j > 0\,. \tag{A.2}$$

The last two equations are equations (6) and (7) of the text.□

Note that equations (A.1) and (A.2) are in fact *partial* (multi-index) difference equations, as is equation (5) of the text. This causes the solution of the replanning models to be somewhat more complicated than the cases where $\alpha = 0$ or 1.

Proof of Theorem 1: First, it will be shown that

$$E_t\, y(t+j) = \lambda\rho\, E_t\, x(t+j) - \alpha\lambda\, E_t\, x(t+j-1) \tag{A.3}$$

for $j > 1$. To show equation (A.3), note that Lemma 1 implies

$$E\left[y(t+j,w(t+j)) \mid w(t+j) = 0,\ w(t) = k\right]$$

$$= (\lambda\rho)\, E\left[x(t+j,w(t+j)) \mid w(t+j) = 0,\ w(t) = k\right] \tag{A.4}$$

and

$$E\left[y(t+j,w(t+j)) \mid w(t+j) > 0,\ w(t) = k\right]$$

$$= (\lambda\rho)\, E\left[x(t+j,w(t+j)) \mid w(t+j) > 0,\ w(t) = k\right]$$

$$- \lambda\, E\left[x(t+j-1,w(t+j-1)) \mid w(t+j) > 0,\ w(t) = k\right]. \tag{A.5}$$

Now use the following facts:

$$E\left[y(t+j,w(t+j)) \mid w(t) = k\right]$$

$$= E\left[y(t+j,w(t+j)) \mid w(t+j) > 0, w(t)=k\right] \Pr\left\{w(t+j) > 0 \mid w(t)=k\right\}$$

$$+ E\left[y(t+j,w(t+j)) \mid w(t+j) = 0, w(t)=k\right] \Pr\left\{w(t+j) = 0 \mid w(t)=k\right\} \tag{A.6}$$

and similarly for $x(t+j)$. This is just the law of iterated expectations.

$$\Pr\left\{w(t+j)=0 \mid w(t) = k\right\}=1-\alpha \text{ for all } k \geq 0, \text{ and}$$

$$\Pr\left\{w(t+j)=1 \mid w(t) = k\right\} = \alpha . \tag{A.7}$$

This is a restatement of property (C1) of the text.

$$E\left[x(t+j-1,w(t+j-1)) \mid w(t+j) > 0,\ w(t) = k\right]$$

$$= E\left[x(t+j-1,w(t+j-1)) \mid w(t) = k\right] \tag{A.8}$$

That is, knowing $w(t+j) > 0$ does not provide any information about $w(t+j-1)$. This follows from a simple application of Bayes' Theorem. Equation (A.3) then obtains by evaluating (A.6) as $(1-\alpha)$ (A.4) + α (A.5) , then using (A.8).

Now, using equation (A.3), the law of iterated expectations, and equation (1) of the text, one can eliminate the $y(t)$ terms from equation (1) of the text to obtain

$$\lambda\rho\, E_t\, x(t+j+1) - (1+\lambda\alpha+\lambda\rho^2)\, E_t x(t+j)$$
$$+ \alpha\lambda\rho\, E_t\, x(t+j-1) = e(t+j) \tag{A.9}$$

for $j = 1, 2, 3, \cdots$. For a fixed value of t, equation (A.9) is a nonstochastic (recall that $e(t+j)$ is assumed to be known) difference equation in $E_t x(t+j)$ that can be solved by the methods described in

Sargent (1979, chapter 9). The characteristic polynomial of (A.9) is given by

$$c(z) = \lambda\rho z^{-1} - (1+\lambda\alpha+\lambda\rho^2) + \alpha\lambda\rho z.$$

It is easy to show that for $\alpha \in (0,1]$, $c(z)$ always admits a factorization of the form

$$c(z) = c_0 (1-c_1 z^{-1}) (1-c_2 z)$$

where $c_1, c_2 \in (0,1)$. (For the case $\alpha = 0$, $c_2 = 0$.) Imposing the stability condition of the theorem yields the following first order difference equation in $E_t x(t+j)$:

$$E_t\, x(t+j) = c_2\, E_t\, x(t+j-1) + c_0^{-1} \gamma^j\, e(t) / (1-c_1\gamma)$$

for $j \geq 1$, which has unique solution

$$E_t\, x(t+j) = c_2^j\, [\, x(t) - \tilde{c}\, e(t)\,] + \gamma^j\, \tilde{c}\, e(t) \qquad \text{(A.10)}$$

where $\tilde{c} = [\, c_0\, (1-c_1\gamma)\, (1-c_2\gamma^{-1})\,]^{-1}$.

Equation (A.10) may now be used to evaluate the expectations term of equation (5) of the text. This has the effect of eliminating all terms involving expectations of future values of $x(t+j)$. Performing this operation yields

$$y(t) = b_0\, e(t) + b_1\, x(t) \tag{A.11}$$

where

$$b_0 = -\rho^{-1} a_0 + (\gamma-\rho)^{-1}$$

$$b_1 = -\rho^{-1}\, a_1$$

$$a_0 = \tilde{c}\,[\, (1-\gamma\rho^{-1})^{-1} - (1-c_2\rho^{-1})^{-1}]$$

$$a_1 = (1-c_2\rho^{-1})^{-1}.$$

Finally, when $S(t)=1$, we can equate the RHS of (A.11) and that of equation (7) to obtain

$$x(t) = f_0\, e(t) + f_1\, x(t-1) \tag{A.12}$$

where

$$f_0 = [1-(\lambda\rho)^{-1} b_1\,]^{-1}\, (\lambda\rho)^{-1}\, b_0$$

$$f_1 = \rho^{-1}\, [1-(\lambda\rho)^{-1} b_1\,]^{-1}.$$

Similarly, when $S(t)=0$, equating the LHS of equations (6) and (A.11) yields

$$x(t) = f_0\, e(t)\ . \tag{A.13}$$

Equations (A.11) - (A.13) describe one solution to equations (5), (6), and (7). Uniqueness in the class of stable solutions follows from the fact that (A.10) is the unique stable solution of (A.9).□

Appendix B. Solution of the Delayed Information Model

Proof of Theorem 2: Again making use of certainty equivalence, write $x(t)$ as $x(t,w(t))$ and $y(t)$ as $y(t,w(t-1))$. Employing techniques directly analagous to those employed in the contemporaneous information model, we can show that the first order conditions of the time t administration's problem are equivalent to

$$x(t,0) = 0 \tag{B.1}$$

$$y(t+j,j-1) = (\rho\lambda\alpha^{-1})\, x(t+j,j) - (\lambda\alpha^{-1})\, x(t+j-1,j-1) \tag{B.2}$$

for $j \geq 1$. Equations (B.1) and (B.2) correspond to equations (A.1) and (A.2) of Appendix A.

Mimicking the proof of Theorem 1, it can then be shown that

$$E_{t-1} y(t+j, w(t+j-1)) = (\lambda\rho\alpha^{-1})\ E_{t-1}\, x(t+j,\ w(t+j))$$
$$- (\lambda\alpha^{-1})\, E_{t-1}\, x(t+j,\ w(t+j-1)) \tag{B.3}$$

for $j \geq 1$. The analog of equation (A.9) will be

$$(\lambda\rho\alpha^{-1})\, E_{t-1}\, x(t+j+1) - (1+\lambda\alpha^{-1}+\lambda\rho^2\alpha^{-1})\, E_{t-1}\, x(t+j)$$
$$+ (\lambda\rho\alpha^{-1})\, E_{t-1} x(t+j-1) = e(t) \tag{B.4}$$

for $j \geq 0$, which will have unique stable solution given by

$$E_{t-1}\, x(t+j) = d_2^{j+1}\, x(t-1) + \tilde{d}\, [\gamma^{j+1} - d_2^{j+1}]\, e(t-1) \qquad \text{(B.5)}$$

where d(z) is the characteristic polynomial of (B.4), with canonical factorization $d(z) = d_0\,(1 - d_1 z^{-1})\;(1 - d_2\, z)$, and $\tilde{d} = [\, d_0\,(1-d_1\gamma)\;(1-d_2\gamma^{-1})\,]$. Solving for y(t) yields

$$y(t) = m_0\, e(t-1) + m_1\, x(t-1) \qquad \text{(B.6)}$$

where

$$m_0 = [\, -\rho^{-1} h_0 + (\gamma-\rho)^{-1}\gamma^{-1}]$$

$$m_1 = -\rho^{-1} h_1$$

$$h_0 = d_2\;(1-d_2\,\rho^{-1})$$

$$h_1 = \tilde{d}\;[\,\gamma\;(1-\gamma\rho^{-1}\,)^{-1} - d_2\;(1-d_2\rho^{-1})^{-1}\,]^{-1}.$$

By equating the RHS of (B.6) and (B.2), it follows that when S(t)=1,

$$x(t) = g_0\, e(t-1) + g_1\, x(t-1) \qquad \text{(B.7)}$$

where

$$g_0 = \alpha\gamma m_0 \;/\; (\lambda\rho)$$

and

$$g_1 = [\ \rho^{-1} + \alpha\ m_1 \;/\; (\lambda\rho)\ ]\ . \qquad \square$$

In the case where e(t) is stochastic, it might seem that x(t) should depend on e(t) rather than e(t-1), since the administration in power at time t can observe e(t). However, "updating" equation (B.7) by replacing γe(t-1) with e(t) would increase the variance of x(t) without changing y(t), since y(t) is decided before e(t) is realized. Hence such updating would be suboptimal.

References

Backus, David and Driffill, John. 1985a. Rational Expectations and Policy Credibility Following a Change in Regime. *Review of Economic Studies* 52 (169): 211-221.

Backus, David and Driffill, John. 1985b. Inflation and Reputation. *American Economic Review* 75 (3): 530-538.

Barro, Robert J. and Gordon, David B. 1983a. A Positive Theory of Monetary Policy in a Natural Rate Model. *Journal of Political Economy* 91 (4): 589-610.

Barro, Robert J. and Gordon, David B. 1983b. Rules, Discretion and Reputation in a Model of Monetary Policy. *Journal of Monetary Economics* 12 (1): 101-121.

Basar, T. and Bagchi, A. 1981. Stackelberg Strategies in Linear-Quadratic Stochastic Differential Games. *Journal of Optimization Theory and Applications* 35 (3): 443-463.

Basar, T. and Olsder, G. J. 1982. *Dynamic Noncooperative Game Theory*. New York: Academic Press.

Buiter, Willem H. 1981. The Superiority of Contingent Rules over Fixed Rules in Models with Rational Expectations. *The Economic Journal* 91 (363): 647-670.

Canzoneri, Matthew B. 1985. Monetary Policy Games and the Role of Private Information. *American Economic Review* 75 (5): 1056-1070.

Cooley, Thomas F.; LeRoy, Stephen F.; and Raymon, Neil. 1984. Econometric Policy Evaluation: Note. *American Economic Review* 74 (3): 467-470.

Chung, Kai L. 1975. *Elementary Probability Theory with Stochastic Processes*. New York: Springer-Verlag.

Hansen, Lars P.; Epple, Dennis; and Roberds, William. 1985. Linear-Quadratic Duopoly Models of Resource Depletion. In *Energy, Foresight, and Strategy*, ed. Thomas J. Sargent, 101-142. Washington, D.C.: Resources for the Future.

Hoel, Paul G.; Port, Sydney C.; and Stone, Charles J. 1972. *Introduction to Stochastic Processes*. Boston: Houghton Mifflin.

Kydland, Finn. 1977. Equilibrium Solutions in Dynamic Dominant-Player Models. *Journal of Economic Theory* 15 (2): 307-324.

Kydland, Finn E. and Prescott, Edward C. 1977. Rules Rather than Discretion: The Inconsistency of Optimal Plans. *Journal of Political Economy* 85 (3): 473-491.

Levine, Paul and Currie, David. 1984. The Design of Feedback Rules in Linear Stochastic Rational Expectations Models. University of London. Working Paper.

Sargent, Thomas J. 1979. *Macroeconomic Theory.* New York: Academic Press.

Sargent, Thomas J. 1984a. Autoregressions, Expectations, and Advice. *American Economic Review* 74 (2): 408-415.

Sargent, Thomas J. 1984b. Dynamic Optimal Taxation. Unpublished manuscript.

Sims, Christopher A. 1982. Policy Analysis with Econometric Models. *Brookings Papers on Economic Activity* 1: 107-152.

Sims, Christopher A. 1985. A Rational Expectations Framework for Short Run Policy Analysis. University of Minnesota. Working Paper.

Tabellini, Guido. 1983. Accommodative Monetary Policy and Central Bank Reputation. University of California, Los Angeles. Working Paper.

Taylor, John B. 1983. Comments on 'Rules, Discretion and Reputation in a Model of Monetary Policy,' by Robert J. Barro and David B. Gordon. *Journal of Monetary Economics* 12 (1): 123-125.

Watson, Mark W. 1985. Recursive Solution Methods for Dynamic Linear Rational Expectations Models. Harvard University. Working Paper.

Whiteman, Charles H. 1983. *Linear Rational Expectations Models: A User's Guide*. Minneapolis: University of Minnesota Press.

Whiteman, Charles H. 1986. Analytical Policy Design Under Rational Expectations. *Journal of Economic Dynamics and Control*, forthcoming.

Witsenhausen, Hans. 1971. Separation of Estimation and Control for Discrete Time Systems. *Proceedings of the IEEE* 59 (11): 1557-1566.