Economic Review

Federal Reserve Bank of San Francisco

1993 Number 1

Adaptive Forecasts, Hysteresis, and Endogenous Fluctuations

George W. Evans and Seppo Honkapohja

Professors of Economics, University of Edinburgh and University of Helsinki, respectively. We are indebted to the Economics Research Department of the FRBSF for numerous helpful comments. The first draft of this paper was written while Evans was a Visiting Scholar at the Federal Reserve Bank of San Francisco. The research was also partially funded by the SPES program of the EC. The first author acknowledges support from the National Science Foundation.

This paper considers fluctuations and policy in an economic model with multiple steady states due to a production externality. In the absence of policy changes, the driving forces generating fluctuations are exogenous random productivity shocks. However, because there are multiple steady states, large productivity shocks can shift the economy between high- and low-level equilibria, providing an additional endogenous source of fluctuations. The scope for macroeconomic policy is large since changes in policy can also shift the economy between equilibria. In this setting macroeconomic policy exhibits hysteresis (irreversibilities) and threshold effects and can be used to eliminate endogenous fluctuations.

A relatively recent major focus of macroeconomic theories has been on nonlinear models with multiple self-fulfilling equilibria and the potential for "endogenous fluctuations." This category can be interpreted quite broadly to encompass models with solutions following regular periodic cycles, "sunspot" solutions depending on extraneous random variables, or multiple steady states arising from coordination failures.[1]

These models may be contrasted with the recent "real business cycle" theory and more generally with linear models with exogenous random shocks generating business cycles around a unique equilibrium. Such models do not generate as wide a range of dynamic time series patterns as is possible in nonlinear models which can generate asymmetries endogenously and manifest additional types of persistence. Models with multiple equilibria are also a potential explanation for the empirical results that the economy appears to exhibit regime-switching.[2]

In the model below we combine aspects of both approaches. In the absence of policy changes, the driving forces generating fluctuations are exogenous random productivity shocks. Without policy changes or productivity shocks, the economy would settle down to a nonstochastic steady state. Clearly productivity shocks will generate fluctuations around a steady state even if it is unique. However, for some cases of the model below, the economy has multiple steady states and large shocks can shift the economy between them. This provides an additional source of fluctuations which is endogenous in the sense that it arises from the structure of the model.

---

1. A very limited selection of well-known papers includes Azariadis (1981), Diamond (1982), Grandmont (1985), Cooper and John (1988), and Woodford (1990). See also the survey paper by Guesnerie and Woodford (1991). The endogenous fluctuations in our model are closest to those described in Howitt and McAfee (1992). Their "animal spirits" cycles are generated by an extraneous sunspot variable, whereas the fluctuations considered here are generated by intrinsic productivity shocks. They do not consider the impact of policy, which is a major focus of this paper.

2. Such results are documented by Hamilton (1989), Boldin (1990) and Potter (1992). They show that real GNP appears to switch probabilistically between high and low growth regimes.

The scope for policy in models with multiple equilibria is potentially large, since the levels of policy variables can affect the likelihood of the economy being in alternative equilibria and since changes in policy can shift the economy between equilibria. This is a fundamental issue for policy. In linear models there is a continuous map from control variables to the expected values of targets. This is no longer so in these nonlinear models. Although policy changes within a certain range have a continuous response, beyond some threshold the economy may be displaced from one equilibrium to another.

In this paper our objective is to consider these issues in a specific model which has multiple stochastic steady state equilibria arising from an aggregate production externality. In the model we develop below, the level of economic activity in the current period is positively related to the level expected in the following period. Figure 1 shows the relationship between current and future aggregate employment, $n_t = F(n_{t+1})$ for different values of the policy parameter γ (γ measures government purchases financed by seignorage, that is, by printing money). The shape of the curve arises from the positive production externality: Each agent's productivity is higher when aggregate output is higher. At low n, diminishing returns yield the usual concave shape to F. Above some threshold the production externality generates increasing social returns and a steeper slope to F. This leads to the nonconcavity shown, though F again becomes concave at sufficiently high n since the magnitude of the external effects is bounded.

FIGURE 1

Effect of Policy on Steady States

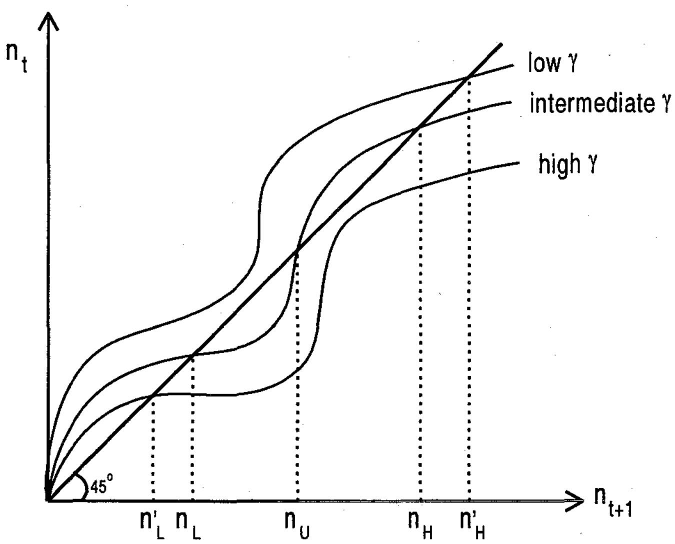

$n_t = F(n_{t+1})$ for nonstochastic model

There is thus the possibility of multiple interior steady states, shown in Figure 1 for the intermediate γ case. Because the model is forward-looking, the evolution of the economy is determined in part by the forecast rules employed by agents, and following recent literature we assume that forecasting is based on adaptive learning algorithms. The steady states labeled n_L and n_H are stable under adaptive learning, while n_U is not. These steady states can be stochastic in the sense of there being fluctuations around them due to productivity shocks. Furthermore, endogenous fluctuations can arise, under certain learning rules, when productivity shocks are large enough (and of the right sign) to move the economy between n_L and n_H. These endogenous fluctuations could be eliminated by using policy to shift up F until only the high-level steady state n_H remained.

Adaptive learning also makes the economy path dependent in the sense that the current equilibrium is determined by initial expectations, the forecast rules, and the shocks. In such environments economic policy can exhibit hysteresis (irreversibilities). Suppose the economy starts at n_H (with intermediate γ) as shown in Figure 1. As γ is increased n_H will decrease and this will be tracked by actual employment under adaptive learning. When γ is large enough, n_H and n_U disappear and employment falls to n_L. If γ is decreased to its original (intermediate) value, employment will only increase to n_L instead of returning to n_H. Thus, since a change in policy can move the economy to a different equilibrium, reversing the policy need not restore the equilibrium that prevailed before the change.

This paper develops in detail the results sketched above. Section I specifies the model and in Section II we describe the forecast rules and how they adapt to forecast errors under learning. Sections III and IV present simulations which illustrate hysteresis effects and endogenous fluctuations, respectively.

I. The Model

The Basic Overlapping Generations Model

We use a generalization of the overlapping generations model incorporating a production externality, developed in Evans and Honkapohja (1991), to derive the model outlined above.[3] The externality leads to social increasing returns over a certain range so that, for some structural parameter values, there is the possibility of multiple steady states. In

3. For brevity we will subsequently denote the references to our own work by EH91, EH92a, EH92b, and EH92c.

the version employed here we will also allow for random productivity shocks and for fiscal–monetary policy.

Before turning to the detailed specification, we emphasize that we will be considering a highly stylized model. No attempt has been made in this paper to make either the structural parameters or the time series properties of the solution paths empirically realistic, though we think that the time series properties for output are "suggestive." Our main objectives are to show the potential of such models to exhibit endogenous fluctuations generated by intrinsic random shocks and to illustrate the effects of policy in this setting.

In the basic overlapping generations (OG) model representative agents (who are producers—consumers) live two periods. An agent born at time t maximizes utility $W=U(c_{t+1})+Z(g_t,g_{t+1})-V(n_t)$, where c_{t+1} is private consumption when old, n_t is labor supply when young, and g_t is public consumption at t. g_t and g_{t+1} are taken as exogenously determined by the government. The budget constraints for the agent are $p_{t+1}c_{t+1}=M_t$ and $p_tq_t=M_t$, where M_t is the money stock, q_t is the quantity of output produced, and p_t is the price of output. In the standard formulation of the OG model, one unit of labor produces one unit of output, that is, $q_t=n_t$, but we will modify this assumption below. The household thus works and produces output when young, and exchanges the goods produced for money (held by the old) at price p_t. This money is then carried forward to the following period when it is exchanged at the possibly different price p_{t+1} for goods to consume when old. In the standard version of the OG model, the stock of (fiat) money is held constant and there are no government purchases.

We choose to analyze a version of the OG model because it is one of the simplest fully specified dynamic general equilibrium models in which expectations matter. The young agent must decide how hard to work, or equivalently, since all income is saved as money, how much to save when young. Since the rate of return on money (the only permitted vehicle for saving in the model) is p_t/p_{t+1}, the expected price in the following period, or more accurately the probability distribution of p_{t+1}, is crucial to the agent's optimal decision.

It should be pointed out that the standard OG model has the disadvantage that the time unit serves several distinct purposes: the length of the working life, the length of the retirement, and the frequency at which economic data are generated. Clearly, we adopt such a model only for tractability and ease of exposition. In principle there is no difficulty constructing analogous models with distinct horizons for these different time periods, as is done in some empirical models. We anticipate that all the phenomena illustrated in this paper will arise in such more realistic models.

The Model with Increasing Social Returns

We now introduce two modifications to the standard OG model. First we allow for government consumption. The government is assumed to purchase the proportion γ_t of output at t, that is, $g_t=\gamma_t q_t$. For convenience we assume that there are no explicit taxes, so that these purchases are entirely financed by seignorage.[4] Thus the government budget constraint is

$$M_{t+1} = M_t + p_{t+1}g_{t+1}.$$

Using $p_tq_t=M_t$ it follows that

$$M_{t+1}/M_t = (1-\gamma_{t+1})$$

and

$$p_t/p_{t+1} = (1-\gamma_{t+1})\, q_{t+1}/q_t.$$

The other modification concerns the production function. A positive externality is introduced into production. Moreover, we allow for random productivity shocks. Thus the production function is assumed to have the form

$$q_t = f(n_t, N_t)\nu_t, \tag{1}$$

where $N_t=Kn_t$ is aggregate employment, K is the number of households in each generation, and ν_t is a positive identically and independently distributed random productivity shock.

The N_t term represents a positive production externality, and we adopt the particular form developed in EH91:

$$f(n, N) = An^{\alpha}\{\max(I^*, \lambda N(1 + a\lambda N)^{-1})\}^{\beta}.$$

This form arises as follows. Individual output is assumed to depend on "ideas" as well as labor effort. Ideas are generated (and "broadcast" to other agents) at a rate λ proportional to labor effort, and the "complementary" ideas obtained from other agents, beyond some threshold I^*, exert a positive external effect on productivity. This effect generates increasing returns over a range at the aggregate level. However, because there is a fixed time cost, a, for accessing a suitable complementary idea, the range of increasing returns is bounded.

Because $0 < \alpha < 1$, the individual faces diminishing marginal returns to individual labor effort (taking N as exogenous) and we adopt a competitive model. The parameterization of the model is completed by assuming the isoelastic forms for the utility functions $U(c)=c^{1-\sigma}/(1-\sigma)$

4. Only a minor modification would be required to allow for using lump-sum taxes to raise part of the revenue.

and $V(n) = n^{1+\epsilon}/(1+\epsilon)$. It can be shown that the law of motion for the economy satisfies[5]

$$E_t((1-\gamma_{t+1})q_{t+1})^{1-\sigma} = n_t^{1+\epsilon}/\alpha. \tag{2}$$

Here E_t denotes the expectations held by agents at time t. In a rational expectations equilibrium this will be equal to the true conditional expectation at time t.

Since by (1) output next period is given by $q_{t+1} = f(n_{t+1}, Kn_{t+1})\nu_{t+1}$ it can be seen that (2) determines current employment as a function of the expected state of the economy next period (together with the current productivity shock this also determines current output and the price level). That is, the reduced form of the model can be written as

$$n_t = H(E_t X(n_{t+1}, \nu_{t+1}, \gamma_{t+1})), \tag{3}$$

where H and X depend on the utility function and production technology parameters according to:

$$X(n, \nu, \gamma) = ((1-\gamma) f(n, Kn)\, \nu)^{1-\sigma}, \tag{4}$$

and

$$H(X) = \alpha X^{1/(1+\epsilon)}. \tag{5}$$

Multiple Steady States and Coordination Failures

The economics of the model can be most easily understood by examining the nonstochastic case in which (1) reduces to $q_t = f(n_t, Kn_t)$. Under perfect foresight, (3) relates n_t to a function of n_{t+1}, that is, $F(n_{t+1}) = n_t$, where we have also incorporated a constant policy parameter $\gamma_{t+1} = \gamma$ into F. Provided $\sigma < 1$ the substitution effect dominates the income effect and function F is upward sloping. The S-shape shown in Figure 1 arises because of the production externality: Below the "kink" point (corresponding to the threshold of "free" ideas I^*) F is concave because of diminishing returns. Above this point social increasing returns set in, generating a nonconcavity in F. However the region of increasing social returns is bounded, and F eventually again becomes concave.

If the externality is sufficiently strong relative to other parameters there can be multiple steady states. Depending on the various parameters, the model can have 0, 1, 2, or 3 interior perfect foresight steady states and the number can be affected by the policy parameter γ. Similar results arise in the stochastic case with random productivity shocks. Provided the range of the shocks is not too "large,"[6] each of the perfect foresight steady states will, in the stochastic case, correspond to a rational (stochastic) steady state $n_t = \bar{n}$[7] and $q_t = f(\bar{n}, K\bar{n})\nu_t$.[8]

When there are three steady states, the welfare in the three steady states depends on the value of government consumption. If government consumption yields no (or sufficiently low) utility, then steady states with high n Pareto dominate those with lower n and the n_L and n_U steady states represent coordination failures. The interpretation of the multiple steady states is straightforward: When other agents work hard, this raises the marginal product of individual effort and induces a higher work effort. At n_L and n_U agents work less hard than at n_H only because *other* agents are working less hard than at n_H. It would be more efficient if agents could coordinate on the high effort level n_H.

The Effect of Policy on Steady States

An increase in γ rotates the F function down around the fixed origin. Figure 1 illustrates one possibility, which we will focus on in this paper: For γ sufficiently small, there is only one interior steady state, the high-level equilibrium n_H. As γ is increased we at some point enter a regime with three interior steady states, n_H, n_U, and n_L. Finally, when γ is sufficiently large, only the interior steady state n_L remains.

Recalling that the policy parameter γ represents a mixed fiscal–monetary policy, it will be noted that, interpreted as fiscal policy, the effects are anti-Keynesian in the following sense. As we discuss below, only the steady states n_L and n_H are stable under learning. If the economy is at a stable steady state, for example, n_L, an increase in γ lowers the level of steady state employment (and output). This result is due to a supply side effect. The higher money growth required to finance increased government consumption leads to a higher level of inflation and therefore a

5. The first order condition for utility maximization is

$$E_t U'(p_t q_t/p_{t+1})(p_t/p_{t+1}) f_1(n_t, N_t)\, \nu_t = V'(n_t),$$

where $f_1(\cdot)$ is the partial derivative with respect to n. Using the market clearing condition $p_t/p_{t+1} = (1-\gamma_{t+1})\, q_{t+1}/q_t$ we obtain

$$E_t U'((1-\gamma_{t+1})\, q_{t+1})(1-\gamma_{t+1})\, q_{t+1} = V'(n_t) f(n_t, N_t)/f_1(n_t, N_t).$$

Substituting the assumed parametric forms for U, V and $f(\cdot,\cdot)$ we obtain (2).

6. Technically suppose that ν_t has bounded support. Then not "large" essentially means that the length of the support is sufficiently small. See EH92a.

7. If the productivity shock were not proportional, then in a rational steady state n_t itself would be a function of ν_t.

8. The model can also have other rational expectations solutions, solutions which depend on exogenous sunspot processes and nonstationary solutions. Consideration of these solutions is not needed for the analysis in this paper.

lower rate of return on work and saving. Because we assume the substitution effect dominates the income effect, this leads to less work effort at higher γ.

II. Forecast Rules and Learning

Expectations and Learning Rules

The possibility of multiple rational expectations solutions (for example, multiple steady states) appears awkward for the pure rational expectations approach. A now widely used approach which overcomes the "multiple equilibria" problem is to replace the assumption of rational expectations with the specification of a learning rule for expectation formation.[9] This may in any case be a more realistic view of expectation formation. The model is thus written as

$$(6) \qquad n_t = H(X(n_{t+1}, \nu_{t+1}, \gamma_{t+1})^e),$$

where the superscript e denotes the expectations of X held by the agents at time t on the basis of a forecasting rule which has been estimated using observed data.[10]

This way of looking at the economy converts a rational expectations model with multiple equilibria into a model with path dependence in which the actual evolution of the economy depends on:

(i) the adaptive forecast rules used by the agents,
(ii) the initial parameter estimates and forecasts held by the agents, and
(iii) the sequence of stochastic shocks and structural shifts.

It may be noted that typically not all equilibria are stable outcomes of adaptive learning processes. Requirement of convergence provides a stability condition which may be used to select equilibria of interest.[11]

We thus depart from strict rational expectations, though for appropriate adaptive forecast rules, expectations may converge to rational expectations over time. Consider forecast rules in which agents treat the law of motion as a stochastic steady state with an unknown mean. Suppose agents estimate the unknown population mean using the sample mean. Such forecast rules are adaptive in the sense that key parameters are altered in response to forecast errors.

A convenient way to write this forecast rule is:

$$(7) \qquad X^e_{t+1} = X^e_t + \delta_t (X_{t-1} - X^e_t),$$

where $\delta_t = 1/t$ and $X^e_0 = X_0$. The formula (7) is the same as the conventional adaptive expectations formula,[12] except that the coefficient specifying the size of the revision to the forecast error, δ_t, goes to 0 at rate $1/t$. This reflects the fact that each new data point provides proportionately less information compared to the history of data.[13] δ_t will be referred to as the "gain" parameter at time t.

Will this learning rule converge to a rational steady state? The problem is not straightforward to answer, since the system is "self-referential" in the sense used by Marcet and Sargent (1989): Agents change their expectations in response to the evolution of the system, and the evolution of the system depends in turn on the expectation rules the agents use.

For the case at hand it is possible to characterize the possible asymptotic outcomes when there are three steady states (see EH92a). The adaptive rule which forecasts X by means of the average of its past values can lead n_t to converge to either n_H or n_L, depending on initial conditions and the sequence of random shocks. That is, either of these two rational steady states can be the outcome of an adaptive learning rule. In contrast, the middle steady state n_U is not stable under learning. When there is only one steady state it will be stable under learning.[14]

Structural Change and Constant Gain Estimators

In the learning rule just described (estimation using the sample mean), the gain parameter δ_t decreases at rate t. The choice $\delta_t \rightarrow 0$, often referred to as "decreasing gain," is appropriate if agents confidently believe that they are in an economy in which X_t, the variable being forecasted, has a constant mean over time. While this would be reasonable if agents believe that the structure of the economy never changes, such an assumption does not seem realistic in practice.

How should the learning rule be modified if agents

9. See EH92c for a recent review of the literature.

10. Note that in considering learning rules we are straining the overlapping generations interpretation of the model. Implicitly we are assuming that agents inherit forecast rules from their "parents," which they then update. Alternatively, it may be possible to reinterpret the model in terms of infinitely lived agents facing finance constraints, as in Woodford (1988).

11. However, more than one equilibrium may be stable under learning; see below for an example.

12. We have introduced a one-period lag into the expectation formula (7) in order to avoid simultaneity between (6) and (7).

13. For an early discussion of the adaptive expectations formula with a possibly nonconstant gain, see Turnovsky (1969).

14. More exotic equilibria, for example, periodic solutions and "sunspot" solutions depending on extraneous variables, can be stable under learning for certain parameter values and appropriate choices of learning rules. See EH92a,b for details.

believe that the structure of the economy may be subject to change? In the context of recursive algorithms for parameter estimation, this is a general problem which has been considered in the statistical and engineering literature; see Benveniste, et al. (1990, ch. 1 and 4, part I). There are two approaches. The first is for agents to build a model, with hyperparameters, of how the system is evolving over time, and to estimate simultaneously both parameters and hyperparameters. This approach requires knowledge of the form of structural change.

The alternative approach, which appears more robust and which we will adopt in this paper, is to replace the assumption $\delta_t \to 0$ by the assumption that δ_t is equal to (or approaches) some fixed value $\delta > 0$. This procedure, known as "constant gain," involves a trade-off between bias and variance when used to adapt to an exogenous time-varying process. A larger value of δ will allow changes in structure to be tracked more rapidly, but will also produce more noisy forecasts.

Although the choice of the gain parameter δ is subject to this trade-off, and the optimal choice of δ will depend on the size and frequency of structural change, the use of a constant gain learning rule, in preference to a decreasing gain rule, is clearly indicated when the structure is subject to change.

An Example with a Time-Varying Policy

To illustrate the importance of using constant gain estimators when structural change is present, consider the behavior of the economy if the monetary–fiscal policy parameter, γ_t, varies systematically over time. In particular, suppose that the share of government purchases is made to vary according to:

$$\gamma_t = \alpha_0 + \alpha_1 \cos(\omega t). \tag{8}$$

Here α_0 specifies the mean level of γ_t, $2\alpha_1$ is the range over which γ_t varies, and ω is the frequency.

Agents are assumed not to know the path (8), but to allow for the structural change by using a constant gain estimator in (7). Of course, a regular sinusoidal pattern in γ_t should be easy to detect, but our point would apply just as well if the pattern for γ_t were highly irregular and difficult to predict.[15] In the simulations of this section we choose α_0 and α_1 so that, given the other parameters of the model, there is a unique steady state n_L.[16]

The other crucial part of the specification is the distribution of the iid proportional productivity shock ν. We choose

$$\nu_t = 1 + \tau(0.5 - u_t),$$

where u_t is iid uniform over the unit interval (τ is restricted to $0 \leq \tau \leq 2$). We set $\tau = 0.20$.

Figure 2 shows a simulation over 1,000 periods when the policy parameters are $\alpha_0 = 0.07$, $\alpha_1 = 0.02$, and $\omega = 0.04$ and when agents use the gain parameter $\delta = 0.15$. The path of employment over time reflects the combined effects of the time variation of policy, random productivity shocks, and the adaption of expectations through the learning rule.

To see the importance of using a constant gain learning rule rather than a decreasing gain rule (such as averaging, that is, $\delta_t = 1/t$) we can compare the quality of the forecasts. For convenience we adopt the mean square forecast error criterion $MSE = T^{-1}\Sigma_{t=1}^{T}(X_t - X_t^e)^2$ and we choose $T = 10{,}000$ periods. Suppose first that agents use a constant gain estimator with $\delta = 0.15$. Then simulations indicate that an individual agent would obtain a much higher MSE with a decreasing gain estimator (0.0206 vs. 0.0149). Even if all other agents were using a decreasing gain estimator, a single agent could somewhat lower his MSE by using an appropriate constant gain estimator (for example $\delta = 0.05$ yields 0.0148 vs 0.0150 with decreasing gain). Thus with time-varying structure there is a forecasting advantage in using a constant gain estimator. We discuss the choice of $\delta = 0.15$ in the next subsection.

Equilibria in Learning Rules

The point just developed merits some further discussion. Is a gain parameter $\delta = 0.15$ a good choice from a statistical point of view? On the basis of the data shown in Figure 2, agents could consider whether another choice of the gain parameter δ would have been better in terms of the mean square forecast error,[17]

$$MSE(\delta) = T^{-1}\Sigma_{t=1}^{T}(X_t(\delta_0) - X_t^e(\delta))^2. \tag{9}$$

Equation (9) is interpreted as follows: The data are generated by the model (4)-(8) with agents using the gain

15. We are not allowing agents to condition their forecasts of X_t on γ_t. This assumption is justifiable if the data on γ_t are infrequent (compared to data on X_t) and of poor quality. Recent data on X_t would then provide most of the information relevant for forecasts.

16. Throughout the paper we use the following structural parameters: $\epsilon = 0.25$, $\sigma = 0.1$, $A = 0.0805$, $a = 0.025$, $\alpha = 0.9$, $\lambda = 0.5$, $K = 40$, $I^* = 19.5$, $\beta = 1.007$.

17. Other possible criteria could be devised based on utility losses.

FIGURE 2

EFFECT OF TIME-VARYING POLICY

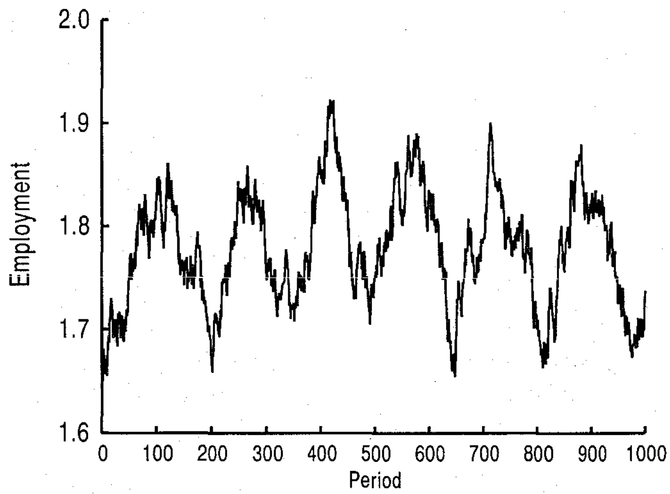

parameter $\delta = \delta_0$ (in our case $\delta_0 = 0.15$). At the end of T periods, agents consider whether they have made a good choice of the gain parameter, *given* the data (that is, given the choice of the gain parameter $\delta = \delta_0$ made by other agents).

Indeed, we can take this line of thought one step further. For each δ_0 we can look for the value of δ which minimizes (9). If the minimum of (9), given δ_0, is attained at δ_0 itself, then we have an equilibrium learning rule with parameter δ_0, in the usual sense. No agent would want to alter his gain parameter δ (based on the MSE criterion), given the choice made by others.

We make no attempt to establish the formal existence of such an equilibrium, but present the evidence for the case at hand, using sample estimates of the MSE for $T = 4{,}000$. Table 1 provides the simulation results. Table 1 shows MSE(δ_0) for various values of δ_0 and also the value of δ, and corresponding MSE, which minimizes (9) for each choice of δ_0. It can be seen that there does appear to be an equilibrium learning rule with a gain parameter of approximately $\delta_0 = 0.15$. This is no accident: We chose our value of δ on the basis of Table 1. It is also worth noting that a wide range of δ_0 would be "reasonable" choices in the sense that the MSE loss of using the wrong δ would be small.

We close this section with one final point: The "equilibrium" δ will depend on the policy parameters. For example, a higher frequency of change ω can be expected to lead to a higher equilibrium value of δ.

TABLE 1

MSE WITH TIME-VARYING POLICY

ACTUAL δ_0	MSE (δ_0)	ARGMIN MSE (δ)	MIN MSE (δ)	DIFFERENCE IN MSE, %
0.01	0.0169	0.07	0.0155	8.89
0.05	0.0149	0.09	0.0146	2.01
0.10	0.0147	0.13	0.0146	0.48
0.15	0.0146	0.15	0.0146	0.00
0.20	0.0147	0.18	0.0147	0.15
0.25	0.0147	0.20	0.0147	0.49
0.30	0.0149	0.23	0.0147	0.85
0.35	0.0150	0.26	0.0148	1.17
0.40	0.0151	0.31	0.0149	1.40
0.45	0.0152	0.36	0.0150	1.53
0.50	0.0154	0.41	0.0151	1.62
0.55	0.0155	0.46	0.0153	1.69
0.60	0.0157	0.51	0.0154	1.74
0.65	0.0158	0.56	0.0156	1.78
0.70	0.0160	0.61	0.0157	1.81
0.75	0.0162	0.66	0.0159	1.84
0.80	0.0163	0.71	0.0160	1.86
0.85	0.0165	0.76	0.0162	1.87
0.90	0.0166	0.81	0.0163	1.89
0.95	0.0168	0.86	0.0165	1.89

NOTE: Table shows MSE for gain δ when data generated in a model with time-varying policy and agents use actual gain δ_0. Simulations are over 4,000 periods.

III. HYSTERESIS EFFECTS

In the preceding section the variation in γ_t was restricted to a range over which the system had a unique steady state n_L. Over this range, neglecting random shocks and transitional learning dynamics, there is a continuous relationship between policy and employment and the policy is "reversible." However, an important feature of our model is that certain variations in the policy parameter γ_t will induce discontinuous responses. To illustrate this aspect of policy we again investigate the effects of a time-varying policy of the form (8), but now set $\alpha_0 = 0.04$, $\alpha_1 = 0.02$, and $\omega = 0.01$. γ_t thus varies continuously over the range 0.02 to 0.06. We use a gain parameter of $\delta = 0.35$, which is approximately the equilibrium value in the sense of the preceding section.

The values of $\gamma = 0.02, 0.04$, and 0.06 correspond to the "low," "intermediate," and "high" γ cases shown in Figure 1. Consider the effects as policy moves from a low value of γ to a high one. Starting from the low γ case of

Figure 1, estimators will continuously track the mean of $X(n_H,\nu,\gamma)$ as we move through the low and intermediate γ cases. However, when γ becomes sufficiently high, n_H and n_U coalesce and then disappear. The system bifurcates to the high γ case, inducing a discontinuous change in the attracting steady state employment level to n_L (a "catastrophe" phenomenon).

From a policy point of view, some of the more interesting features are the hysteresis effects illustrated in Figure 3. Here we show the relationship between γ_t and n_t over one complete cycle of γ_t (from 0.06 to 0.02 to 0.06). Over most of this range there are two distinct branches to the policy relationship, with the lower branch corresponding to the n_L steady state and the upper branch corresponding to n_H.

The branch on which the system lies at some point in time is determined by history. On a given branch, over a range of γ, policy is reversible in the sense that an increase in γ followed by an equal decrease in γ will return the system to its original position (if an allowance is made for random productivity shocks and for transitional learning dynamics).

However, changes in γ_t beyond a certain point induce policy irreversibilities when the system is forced onto the other branch. Starting with $\gamma_t = 0.06$ and $n_t \approx 1.9$, the system moves (clockwise) along the lower branch until at low values of γ_t employment becomes forced onto the upper branch (when the low steady state disappears). When γ_t begins to increase from its minimum of 0.02, it remains on the upper branch until γ_t is sufficiently high.

The message for policy is this: If the system is trapped into a low-level steady state, the policy variable can have a strongly nonlinear response. Decreases in γ may have initially small effects on employment, while beyond some threshold value the induced response can be much larger as the economy is pushed from the low-level to the high-level steady state.

FIGURE 3

HYSTERESIS EFFECTS

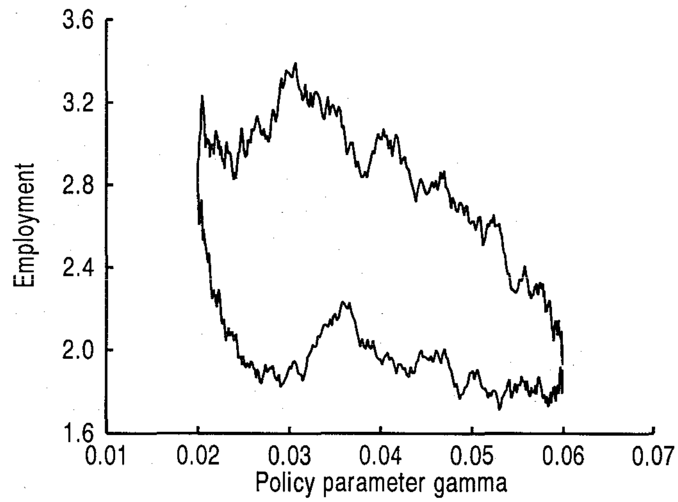

IV. ENDOGENOUS CYCLES

Endogenous Shifts in Expectations

In the foregoing we have considered a model in which a key driving force is variation in the policy variable γ_t. However, the qualitative results obtained suggest the following additional possibility. Suppose that agents use a constant gain forecast rule. Could random productivity shocks lead to shifts between high-level and low-level steady states, via induced changes in forecasts, even in the absence of structural or policy shifts? As we will see, the answer can be yes.

Thus consider the system (4)-(7) with γ_t fixed at $\gamma =$ 0.04. With this value of the policy parameter there are two stable rational steady states. If agents use (7) with decreasing gain, for example, with $\delta_t = c/t$, for some constant c, then, as pointed out in Section II, the system will converge to one of the two rational steady states corresponding to n_L or n_H (and expectations will converge to the corresponding fixed rational forecast).

However, if agents use a constant gain $\delta_t = \delta > 0$ this will not be so—agents' forecasts X^e_{t+1} will retain some randomness even in the limit because of their sensitivity (through $\delta_t X_{t-1}$) to random productivity shocks. Furthermore, there is now the possibility that, say, with the system starting near the low-level steady state, a large favorable productivity shock leads to a sufficiently large revision in X^e_{t+1} so that in subsequent periods the system is drawn, for an extended period of time, to the high-level steady state.

This phenomenon is illustrated in Figure 4, which presents the results of a simulation with $\delta = 0.15$, $\gamma = 0.04$ and $\tau = 0.20$ (other parameters are unchanged). The system appears to alternate between two noisy steady states, centered near n_H and n_L. Occasionally productivity shocks are sufficiently large and in the right direction to move actual and thus subsequent expected aggregate economic activity from one region of attraction to the other.

Why should agents use a constant gain expectations rule in this situation? There are two reasons. First, although in the simulation presented the policy parameter was held fixed, agents may use a constant gain forecast rule because they are concerned about the *possibility* of structural or policy shifts. It seems plausible that agents would want to make some allowance for this by maintaining δ above some minimum positive level.

FIGURE 4

ENDOGENOUS CYCLES WITH CONSTANT GAIN ESTIMATOR

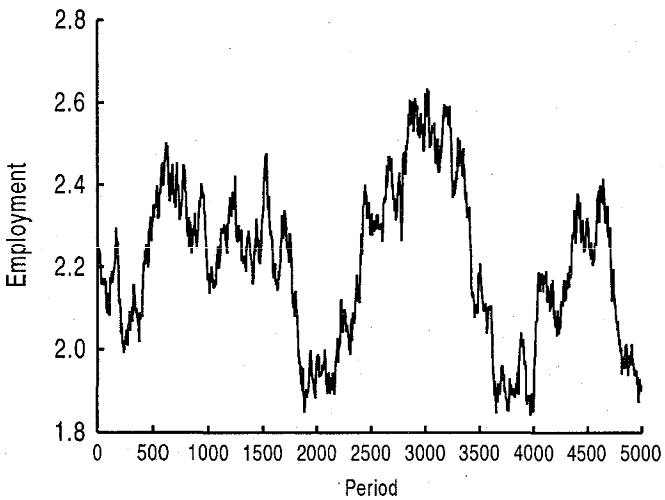

TABLE 2

MSE WITH ENDOGENOUS CYCLES

δ	MSE
0.05	0.0268
0.10	0.0248
0.15	0.0246
0.20	0.0248
0.25	0.0253
0.30	0.0258
0.40	0.0272
0.50	0.0288
0.60	0.0307
0.70	0.0329
0.90	0.0386

NOTE: The MSE for a model with endogenous cycles was generated by a fixed policy parameter $\gamma=0.04$ and constant gain $\delta_0=0.15$. Simulations are over 4,000 periods.

Second, and more fundamental, the choice of a constant gain learning rule may be an equilibrium learning rule in the sense defined in Section II. That is, the choice δ minimizes the forecast mean square error for each agent, given that other agents use that δ. In this case we would have a self-fulfilling prophecy in learning rules, with expectations adapting to fluctuations in economic activity, and the changes in expectations in turn inducing fluctuations in the economy. Even if δ is not strictly an equilibrium value, it may still yield an MSE which is nearly optimal. Table 2 shows that this is indeed true for $\delta=0.15$.[18]

Although the model is highly stylized, the results of this section are attractive as a model of economic fluctuations in the following sense: Unlike "sunspot" equilibria, in which the solution depends on extraneous variables, the precipitating variables here, as in real business cycle (RBC) models, are productivity shocks. The difference from RBC models is that a sequence of large shocks can induce a self-fulfilling overreaction in which the economy moves between its two stable steady states.

The policy implications are again straightforward. Faced with an economy undergoing endogenous fluctuations, the policy parameter γ can be shifted to a level at which there is a unique steady state.

Regime-Switching Models

But would agents stick with a constant gain estimator if they observe the process shown in Figure 4? They might, since the existence of two regimes may not be apparent in the presence of the random shocks, and since the use of a constant gain estimator is designed to allow for and adapt to unspecified structural shifts. However, it is of interest to know whether endogenous cycles would continue to exist if agents *did* infer the existence of two regimes (corresponding to n_L and n_H) and estimated a regime-switching model in an attempt to improve their forecasts.

Thus suppose that agents believe that the conditional mean level of X follows a two-state Markov process. We will assume that agents believe the regime is triggered by the recent average level of economic activity, rather than by some hidden variable, and so use the "self-exciting" framework of Potter (1992). Based on Figure 4, we choose a regime switching parameter X^* corresponding to $n=2.1$ ($X^* = 2.1^{(1+\epsilon)/\alpha}$). Agents assume that the economy is in state 1 if $\bar{X}_{t-1}$, the average of X of the recent past, is less than X^* and in state 2 if it exceeds X^*. We again use a recursive estimation procedure, but now assume that agents estimate both the conditional mean value of X in each state and the conditional probability of being in each state. Let p_i for $i=1,2$ be the probability, given that we are in state i at s, of staying in state i at $s+1$. p_i at time t is essentially estimated by the corresponding actual proportion $p_{i,t}$ through time t.[19] The conditional means in the two

18. The choice of $\delta_0=0.15$ is only approximately an equilibrium, because the δ that minimizes MSE for this δ_0 lies between 0.14 and 0.15.

19. The estimation is actually done using the associated recursive formula and initial estimates $p_1=p_2=1$.

states are estimated by

$$\left.\begin{aligned} X1_{t+1} &= X1_t + \delta(X_t - X1_t) \\ X2_{t+1} &= X2_t \end{aligned}\right\} \quad \text{if } \bar{X}_t \leq X^*$$

$$\left.\begin{aligned} X2_{t+1} &= X2_t + \delta(X_t - X2_t) \\ X1_{t+1} &= X1_t. \end{aligned}\right\} \quad \text{if } \bar{X}_t > X^*$$

For simplicity, the switch point X^* is fixed exogenously and not estimated. In the simulations we assume that $\bar{X}_t$ is computed using the average over the last three periods. We continue to use a constant gain estimator,[20] on the assumption that agents still want to allow for the possibility of structural/policy shifts.

Agents then forecast X_{t+1} at time t according to

$$X^e_{t+1} = p_{1,t} X1_t + (1 - p_{1,t}) X2_t \quad \text{if } \bar{X}_{t-1} \leq X^*,$$

and

$$X^e_{t+1} = (1 - p_{1,t}) X1_t + p_{2,t} X2_t \quad \text{if } \bar{X}_{t-1} > X^*.$$

Figure 5 shows the results of a simulation (over 2,000 periods) with unchanged structural parameters and with the same values for γ and δ. It is apparent that the broad pattern of endogenous cycles remains. The main effect of the "more sophisticated" forecast procedure is to speed up the transition between regimes.

There are numerous other ways in which agents might attempt to capture the dynamics of the system, but the broad point seems clear. If agents allow for the possibility of changes in regime when making their forecasts, this reinforces the potential of productivity shocks and other sources of intrinsic noise to induce, periodically, large self-fulfilling changes in the level of economic activity.

V. Conclusions

We have considered a macroeconomic model, incorporating production externalities and random productivity shocks, which has the potential to generate two stable stochastic steady states. For the structural parameter values chosen, the number of stable steady states depends on the monetary—fiscal policy parameter γ. When γ is low (high), only the high- (low-) level steady state exists. Both steady states coexist for intermediate values of γ.

If γ_t varies sufficiently over time, and if agents use adaptive forecasts with an appropriate constant gain parameter, aggregate economic activity will periodically shift between the high-level and low-level regimes. The economy will exhibit hysteresis effects over the "business cycle" in the response of output to policy. If instead γ_t is held fixed at an intermediate value, the economy can still exhibit fluctuations between the two regimes, driven now by the productivity shocks themselves.

FIGURE 5

Endogenous Cycles with Regime-Switching Estimators

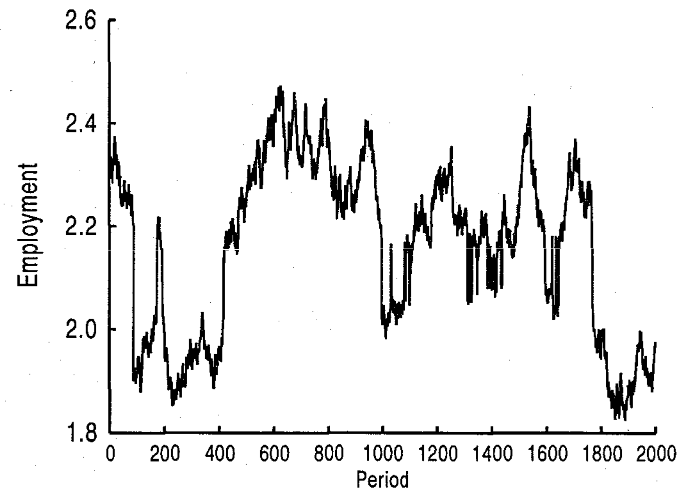

There is thus the potential for models with multiple steady states to explain the empirical regime-switching results documented by Hamilton (1989), Boldin (1990) and Potter (1992).[21]

In our model the "switches" are determined by fundamentals, either by intrinsic productivity or taste shocks or by policy changes. We emphasize, however, that no attempt has been made in this paper to use empirically realistic parameters or to fit macroeconomic data.[22] A large gap currently exists between theoretical models of endogenous fluctuations and their empirical implementation, and it may be desirable to attempt calibration to observed fluctuations in future research.

We have emphasized the possibility that policy may have a highly nonlinear response in these models, since it can sometimes shift the economy between high-level and low-level steady states if the policy variable exceeds some threshold. This is clearly an important phenomenon and indicates that it would be worthwhile to examine the performance of monetary feedback rules in such models.

20. If agents use a decreasing gain estimator, the system might converge to a self-fulfilling solution in which the regime is a function of the productivity shock.

21. Howitt and McAfee (1992) and Boldin (1990) have also noted this potential connection.

22. A wider range of policy variables also could be incorporated.

REFERENCES

Azariadis, C. 1981. "Self-fulfilling Prophecies." *Journal of Economic Theory* 25, pp. 380-396.

Benveniste, A., M. Metivier, and P. Priouret. 1990. *Adaptive Algorithms and Stochastic Approximations*. New York: Springer Verlag.

Boldin, M. D. 1990. "Characterizing Business Cycles with a Markov Switching Model: Evidence of Multiple Equilibria." Research Paper No. 9037. Federal Reserve Bank of New York.

Cooper, R., and A. John. 1988. "Coordinating Coordination Failures in Keynesian Models." *Quarterly Journal of Economics* 113, pp. 441-464.

Diamond, P.A. 1982. "Aggregate Demand Management in Search Equilibria." *Journal of Political Economy* 90, pp. 881-894.

Evans, G., and S. Honkapohja. 1991. "Increasing Social Returns, Learning and Bifurcation Phenomena." Discussion Paper. London School of Economics. (Forthcoming in *Learning and Rationality in Economics*, ed. A. Kirman and M. Salmon. Oxford: Basil Blackwell.)

__________, and __________. 1992a. "Local Convergence of Recursive Learning to Steady States and Cycles in Stochastic Nonlinear Models." Mimeo.

__________, and __________. 1992b. "On the Local Stability of Sunspot Equilibria under Adaptive Learning Rules." Mimeo.

__________, and __________. 1992c. "Adaptive Learning and Expectational Stability: An Introduction." Discussion Paper. London School of Economics. (Forthcoming in *Learning and Rationality in Economics*, ed. A. Kirman and M. Salmon. Oxford: Basil Blackwell.)

Grandmont, J-M. 1985. "On Endogenous Competitive Business Cycles." *Econometrica* 53, pp. 995-1045.

Guesnerie, R., and M. Woodford. 1991. "Endogenous Fluctuations." Mimeo. Paris: DELTA.

Hamilton, J. 1989. "A New Approach to Economic Analysis of Economic Time Series." *Econometrica* 57, pp. 357-384.

Howitt, P., and P. McAfee. 1992. "Animal Spirits." *American Economic Review* 82, pp. 493-507.

Marcet, A., and T. Sargent. 1989. "Convergence of Least Squares Learning Mechanisms in Self-referential Stochastic Models." *Journal of Economic Theory* 48, pp. 337-368.

Potter, S. 1992. "A Nonlinear Approach to U.S. GNP." Mimeo. UCLA.

Turnovsky, S. 1969. "A Bayesian Approach to the Theory of Expectations." *Journal of Economic Theory* 1, pp. 220-227.

Woodford, M. 1988. "Expectations, Finance Constraints and Aggregate Instability." In *Finance Constraints, Expectations and Macroeconomics*, eds. M. Kohn and S.C. Tsiang, pp. 230-261. New York: Oxford University Press.

__________, 1990. "Learning to Believe in Sunspots." *Econometrica* 58, pp. 277-307.