This PDF is a selection from a published volume from the National Bureau of Economic Research

Volume Title: NBER Macroeconomics Annual 2001, Volume 16

Volume Author/Editor: Ben S. Bernanke and Kenneth Rogoff, editors

Volume Publisher: MIT Press

Volume ISBN: 0-262-02520-5

Volume URL: http://www.nber.org/books/bern02-1

Conference Date: April 20-21, 2001

Publication Date: January 2002

Title: Evolving Post-World War II U.S. Inflation Dynamics

Author: Timothy Cogley, Thomas J. Sargent

URL: http://www.nber.org/chapters/c11068

Timothy Cogley and Thomas J. Sargent
ARIZONA STATE UNIVERSITY; AND
STANFORD UNIVERSITY AND HOOVER INSTITUTION

Evolving Post–World War II U.S. Inflation Dynamics

1. Introduction

This paper uses a nonlinear stochastic model to describe inflation–unemployment dynamics in the United States after World War II. The model is a vector autoregression with coefficients that are random walks with reflecting barriers that keep the VAR stable. The innovations in the coefficients are arbitrarily correlated with each other and with innovations to the observables. The model enables us to detect features that have been emphasized in theoretical analyses of inflation–unemployment dynamics. Those analyses involve coefficient drift in essential ways.

Thus, DeLong (1997), Taylor (1997, 1998), and Sargent (1999) interpreted the broad movements of the inflation rate in terms of the monetary authority's changing views about the Phillips curve. According to them, the runup in inflation in the late 1960s and 1970s occurred because the monetary authority believed that there was an exploitable trade-off between inflation and unemployment. Its beliefs induced the monetary authority to accept the temptation to inflate more and more until eventually it had attained Kydland–Prescott (1977) time-consistent inflation rates. But the observations of the 1970s taught Volcker and

Greenspan the natural-rate hypothesis, which they eventually acted upon to reduce inflation.

Another mechanism was posited by Parkin (1993) and Ireland (1999), who argued that the inflation–unemployment dynamics are driven by exogenous drift in the natural rate of unemployment, for example due to demographic changes. Because the time-consistent inflation rate varies directly with the natural rate of unemployment, Parkin and Ireland attributed the drift in the inflation rate to drift in the natural rate of unemployment.

The DeLong–Taylor–Sargent story makes contact with various elements in Lucas's (1976) critique. It makes the drift in inflation–unemployment dynamics a consequence of the monetary authority's evolving views about the economy. The story attributes alterations in the law of motion for inflation and unemployment to the changing behavior of the monetary authority, which emerges in turn from its changing beliefs. This story is consistent with one way that Lucas (1976) has been read, namely, as an invitation to impute observed drift in coefficients of econometric models to time-series variation in government policy functions.

Sargent's (1999) version of the story focuses on how the coefficient drift over time affected the results of time-series tests of the natural-rate hypothesis. In the late 1960s, Robert Solow and James Tobin proposed a test of the natural-rate hypothesis. Using data through the late 1960s, that test rejected the natural-rate hypothesis in favor of a permanent trade-off between inflation and unemployment. Lucas (1972) and Sargent (1971) criticized that test for not properly stating the implications of the natural-rate hypothesis under rational expectations. In particular, the Solow–Tobin test was correct only if inflation exhibited a unit root. Before the 1970s, postwar U.S. inflation data did not exhibit a unit root, rendering invalid (in the opinion of Lucas and Sargent) Solow's and Tobin's interpretation of their test. However, in the 1970s, just when U.S. inflation seems to have acquired a unit root, the Solow–Tobin test began accepting the natural-rate hypothesis. Building on Sims (1988) and Chung (1990), Sargent (1999) constructs an adaptive model of the government's learning and policymaking that centers on the process by which the government learns an imperfect version of the natural-rate hypothesis, cast in terms of Solow and Tobin's representation.

Parts of Sargent's adaptive story acquire credibility when it is noted how the Solow–Tobin characterization of the natural-rate hypothesis has endured, despite the criticism of Lucas and Sargent. As Hall (1999) and Taylor (1998) lament, that faulty characterization continues to be widely used. For example, see Rudebusch and Svensson (1999) for a widely

cited model that represents the natural-rate hypothesis in the Solow–Tobin form. Fisher and Seater (1993), King and Watson (1994, 1997), Fair (1996), Eisner (1997), and Ahmed and Rogers (1998) construct tests of long-run neutrality that are predicated on the assumption of a unit root in inflation.[1] Estrella and Mishkin (1999) use the Solow–Tobin characterization to estimate the natural rate of unemployment. In the discussion following the paper by Estrella and Mishkin, John Williams confesses that the Federal Reserve Board's large-scale macroeconometric model also incorporates this characterization. Hall questions its validity for U.S. data after 1979 and sharply criticizes its continued use.

Taylor (1998) warns that adherence to the erroneous econometric characterization of the natural-rate hypothesis will eventually cause policy to go astray. Because of the diminished serial correlation that he sees in recent inflation data, Taylor is concerned that the disappearance of a unit root in inflation means that the faulty test may soon signal an exploitable trade-off that will once again tempt the monetary authority. The theme of both Hall and Taylor is that failure to remember the theoretical and econometric lessons of the 1970s is likely to resuscitate pressure to inflate emanating from the empirical Phillips curve. In the same symposium, Friedman (1998) and Solow (1998) made a number of assertions that may have contributed to Taylor's worries. Friedman asserted that the real effects of monetary policy are so long-lasting that "for all practical purposes they might just as well be permanent." Solow (1998) expressed skepticism about the natural-rate hypothesis and suggested that the supporting evidence is specific to the U.S. economy since 1970. He argued that monetary policy can affect the natural rate of unemployment and that the experience of the United States in the 1960s suggests that persistent high unemployment would yield to a revival of aggregate demand. Taylor's concern is that low inflation would be hard to sustain if belief in a long-run trade-off were again to become influential.

The object of this paper is to develop empirical evidence that is relevant to this discussion.[2] Section 2 describes a Bayesian model that we use to summarize the evolution of inflation dynamics. Section 3 reports

1. Many of these authors pretest for a unit root and apply the Solow–Tobin test only if they fail to reject the null hypothesis. But pretesting could result in a more subtle version of the Lucas–Sargent trap. Unit-root tests have low power and may fail to detect circumstances in which the Solow–Tobin test is inappropriate.
2. Albanesi, Chari, and Christiano (2000) model the inception and termination of inflation in the 1970s with a sunspot variable that shifts expectations between two regimes. Their equilibrium excludes the concerns about model misspecification that are the focus of the present discussion. It is possible that a regime-switching model like theirs can confront the observations about comovements between inflation persistence and mean inflation that we document below.

stylized facts about this evolution, and Section 4 discusses test statistics for the Solow–Tobin version of the natural-rate hypothesis. Section 5 considers Taylor's warning about recidivism on the natural-rate hypothesis. The paper concludes with a summary.

2. A RANDOM-COEFFICIENTS REPRESENTATION

We use a Bayesian vector autoregression with time-varying parameters to describe the evolution of the law of motion for inflation. We are interested in a random-coefficients representation for some of the reasons expressed in the initial sections of Lucas (1976). The Bayesian framework treats coefficients as random variables, making it attractive for modeling data from economies in which important decision makers, including the monetary authority, are learning.[3]

2.1. NOTATION AND STATE-SPACE REPRESENTATION

The model has a nonlinear state-space representation. The measurement equation is

$$y_t = X_t'\theta_t + \varepsilon_t, \tag{2.1}$$

where y_t is an $N \times 1$ vector of endogenous variables, θ_t is a $K \times 1$ vector of coefficients, X_t' is an $N \times K$ matrix of predetermined and/or exogenous variables, and ϵ_t is an $N \times 1$ vector of prediction errors. The vector y_t includes inflation and variables useful for predicting inflation. In this paper, we use (2.1) to represent a vector autoregression, so that the right-hand variables are lags of y_t. In an unrestricted vector autoregression, each equation contains the same right-hand variables, $X_t' = (I_N \otimes x_t')$.

We treat the coefficients of the VAR as a hidden state vector. The state vector θ_t evolves according to

$$p(\theta_{t+1}|\theta_t, V) \propto I(\theta_{t+1})f(\theta_{t+1}|\theta_t, V), \tag{2.2}$$

where $I(\theta_t) = 0$ if the roots of the associated VAR polynomial are inside the unit circle and 1 otherwise; V is a covariance matrix defined below; and

$$f(\theta_{t+1}|\theta_t, V) \sim N(\theta_t, Q). \tag{2.3}$$

3. Our focus in this paper is on the evolution of reduced-form relationships. Structural models involve nonlinear cross-equation restrictions on the evolving parameters, and they require nonlinear filtering methods. We are currently studying nonlinear filters.

Thus, $f(\theta_{t+1}|\theta_t, V)$ can be represented as the driftless random walk

$$\theta_t = \theta_{t-1} + v_t, \tag{2.4}$$

where v_t is an i.i.d. Gaussian process with mean 0 and covariance Q. The economy changes over time when news arrives, making θ_t vary in an unpredictable way. Throughout this paper, we use $f(\cdot)$ to denote a normal density, and $p(\cdot)$ to denote a more general density.

We assume that the innovations, $(\varepsilon_t', v_t')'$, are identically and independently distributed normal random variables with mean zero and covariance matrix

$$E_t \begin{bmatrix} \varepsilon_t \\ v_t \end{bmatrix} \begin{bmatrix} \epsilon_t' & v_t' \end{bmatrix} = V = \begin{pmatrix} R & C' \\ C & Q \end{pmatrix}, \tag{2.5}$$

where R is the $N \times N$ covariance matrix for measurement innovations, Q is the $K \times K$ covariance matrix for state innovations, and C is a $K \times N$ cross-covariance matrix. Following the Bayesian literature, we call the θ's *parameters* and the elements of R, Q, and C *hyperparameters*.

We assume that the hyperparameters and initial state θ_0 are independent, that the initial state is a truncated Gaussian random variable, and that the hyperparameters come from an inverse-Wishart distribution. We adopted these parts of the prior mostly because of their convenience in being natural conjugates for our Gaussian virtual prior f.

Let $f(\theta_0) = N(\bar{\theta}, \bar{P})$ represent a normal prior with mean $\bar{\theta}$ and variance $\bar{P}$. The prior for the initial state is

$$p(\theta_0) \propto I(\theta)N(\bar{\theta}, \bar{P}). \tag{2.6}$$

Our prior for the hyperparameters is

$$p(V) = \text{IW}(\bar{V}^{-1}, T_0), \tag{2.7}$$

where $\text{IW}(S, \text{df})$ represents the inverse-Wishart distribution with scale matrix S and degrees of freedom df. This is a convenient form because it yields an inverse-Wishart posterior when combined with a Gaussian likelihood. Collecting the pieces, the joint prior for θ_0, V can be represented as

$$p(\theta_0\ V) \propto I(\theta)N(\bar{\theta}, \bar{P})\ \text{IW}\ (\bar{V}^{-1}, T_0). \tag{2.8}$$

Both pieces are informative, but in the empirical section we set $\bar{\theta}$, $\bar{P}$, $\bar{V}$, and T_0 so that they are only weakly informative.

We use the following notation to denote partial histories of the variables Y_t and θ_t. The vectors

$$Y^T = [y'_1, \ldots, y'_T]' \tag{2.9}$$

and

$$\theta^T = [\theta'_1, \ldots, \theta'_T]' \tag{2.10}$$

represent the history of data and states up to date T, and

$$Y^{T+1,T+H} = [y'_{T+1}, \ldots, y'_{T+H}]' \tag{2.11}$$

and

$$\theta^{T+1,T+H} = [\theta'_{T+1}, \ldots, \theta'_{T+H}]' \tag{2.12}$$

represent potential future trajectories from date T onward.

We can use (2.2) to assemble the joint density

$$p(\theta^T|V) \propto I(\theta^T)f(\theta^T|V), \tag{2.13}$$

where

$$f(\theta^T|V) = f(\theta_0|V) \prod_{t=0}^{T-1} f(\theta_{t+1}|\theta_t, V) \tag{2.14}$$

and

$$I(\theta^T) = \prod_{t=0}^{T} I(\theta_t). \tag{2.15}$$

We call f our *virtual prior*, and p the prior. The virtual prior f makes θ a driftless random walk. Multiplying $f(\theta^T|V)$ by $I(\theta^T)$ puts zero probability on sample paths of $\{\theta_t\}$ for which θ_t for any $t \geq 0$ corresponds to unstable VAR coefficients.[4]

---

4. An appendix shows that the model formed by (2.3), (2.13), (2.14), and (2.15) implies the nonlinear transition equation (2.2).

In (2.2), the truncation of $f(\theta_t|\theta_{t-1}, V)$ through multiplication by $I(\theta_t)$ reflects our opinion that explosive representations are implausible for the United States. An unrestricted normal density $f(\theta^T|V) = f(\theta_0) \prod_{t=0}^T f(\theta_{t+1}|\theta_t, V)$ for the history of states θ^T implies a positive probability of explosive autoregressive roots, but an explosive representation implies an infinite variance for inflation, which cannot be optimal for a central bank that minimizes a loss function involving the variance of inflation.[5] We restrict the prior to put zero probability on explosive states.

This representation resembles some of the models in Doan, Litterman, and Sims (1984), but with a different prior. Doan *et al.* were primarily interested in forecasting and recommended a "random walk in variables" prior for the sake of parsimony. We are less interested in forecasting and more interested in summarizing the data in a relatively unconstrained fashion, so we chose the prior described above.

2.2 A LIMITATION OF OUR MODEL: NO STOCHASTIC VOLATILITY

For macroeconomic variables and a period similar to ours, Bernanke and Mihov (1998a, 1998b) and Sims (1999) presented evidence that favors a vector autoregression with time-invariant autoregressive coefficients but a covariance matrix of innovations that fluctuates over time. In contrast, our specification allows the coefficients to vary and assumes a time-invariant but unknown innovation covariance matrix V. While our prior fixes V, our statistical methods nevertheless allow the data to speak up for volatility or drift in V, albeit in a restricted and adaptive way. Our estimates of V conditioned on time t data fluctuate over time in ways that we shall discuss.

We chose our specification partly because we want to focus attention on the coefficient-drift issues raised by Lucas (1976). Our model is rigged to let us detect drifts in the systematic parts of government and private behavior rules that show up in the systematic parts of vector autoregressions. Our prior embodies a prejudice that monetary policy changed systematically during the years that we study. In contradistinction, the interpretation of the evidence favored by Bernanke and Mihov (1998a, 1998b) and Sims (1999) is consistent with a view that while distributions of shocks have evolved, agents' responses to them have been stable.[6]

---

5. Alternatively, explosive representations cannot result if the monetary policy rule ensures that inflation is bounded. We do not claim that an integrated representation for inflation is implausible on statistical grounds, only that drift in inflation is hard to reconcile with purposeful central-bank behavior.
6. See Sims (1982) and Sargent (1983) for theoretical settings that, by assuming that the historical sample was produced by optimizing government behavior and stable private-sector responses to it, can explain such a pattern.

2.3 POSTERIOR PREDICTIVE DENSITY

As Bayesians, our goal is to summarize the posterior density for the objects of interest. We are mostly interested in a forward-looking perspective in inflation, so we want posterior *predictive* densities.

In this model, there are four sources of uncertainty about the future. The terminal state θ_T and the hyperparameters V are unknown and must be estimated. In addition, as time goes forward, the state vector will drift away from θ_T, and the measurement equation will be hit by random shocks. Conditional on prior beliefs and data through date T, beliefs about the future can be expressed by the joint posterior distribution,

$$p(Y^{T+1,T+H}, \theta^{T+1,T+H}, \theta^T, V \mid Y^T). \tag{2.16}$$

Our objective is to characterize (2.16). This is a complicated object, but it can be decomposed into more tractable components. We begin by factoring (2.16) into the product of a conditional and a marginal density,

$$p(Y^{T+1,T+H}, \theta^{T+1,T+H}, \theta^T, V \mid Y^T)$$

$$= p(\theta^T, V \mid Y^T)\, p(Y^{T+1,T+H}, \theta^{T+1,T+H} \mid \theta^T, V, Y^T). \tag{2.17}$$

This expression splits the joint density into a factor that represents beliefs about the past and present and another that represents beliefs about the future. The first factor is the joint posterior density for hyperparameters and the history of states. It summarizes current knowledge about system dynamics, based on data and prior beliefs. The second factor reflects the uncertainty about the future that would be present even if the current state and hyperparameters were known with certainty. This factor reflects the influence of future innovations to the state and measurement questions.

Analytical expressions for each piece are unavailable, even for simple cases. Instead, we use Monte Carlo methods to simulate them. The algorithm is split into two parts, corresponding to the components of (2.17). The first part uses the Gibbs sampler to simulate a draw of θ^T and V from the marginal density, $p(\theta^T, V \mid Y^T)$. The second step plugs that draw into the conditional density $p(Y^{T+1,T+H}, \theta^{T+1,T+H} \mid \theta^T, V, Y^T)$ and generates a trajectory for future data and states.

2.4 BELIEFS ABOUT THE PAST AND PRESENT

The posterior density for states and hyperparameters can be expressed as

$$p(\theta^T, V \mid Y^T) \propto p(Y^T \mid \theta^T, V)p(\theta^T, V),$$

$$\propto f(Y^T \mid \theta^T, V)p(\theta^T|V)p(V),$$

$$\propto I(\theta^T)\,[f(Y^T \mid \theta^T, V)f(\theta^T|V)p(V)]. \tag{2.18}$$

The first line follows from Bayes's theorem: $p(\theta^T, V)$ represents a joint prior for hyperparameters and states and $p(Y^T \mid \theta^T, V)$ is a conditional likelihood. Conditional on states and hyperparameters, the measurement equation is linear in observables and has normal innovations. Thus, the conditional likelihood is Gaussian, $p(Y^T|\theta^T, V) = f(Y^T \mid \theta^T, V)$, as shown in the second line. The joint prior for hyperparameters and states can be factored into a marginal prior for V and a conditional prior for θ^T, and substituting $I(\theta^T)f(\theta^T|V)$ for $p(\theta^T|V)$ delivers the expression on the third line.

Notice that the expression in brackets on the last line is the joint posterior kernel that would result if the restriction on unstable roots were not imposed. If not for this restriction, the model would have a linear Gaussian state-space representation, with transition equation $f(\theta^T|V)$. The posterior kernel associated with this linear transition law is

$$p_L(\theta^T, V \mid Y^T) \propto f(Y^T \mid \theta^T, V)f(\theta^T|V)p(V). \tag{2.19}$$

Substituting this relation into the last equation, the posterior density for the nonlinear model can be expressed as a truncation of the posterior for the unrestricted linear model,

$$p(\theta^T, V \mid Y^T) \propto I(\theta^T)p_L(\theta^T, V \mid Y^T). \tag{2.20}$$

Among other things, this means that $p(\theta^T, V \mid Y^T)$ can be represented and simulated in two steps. First, we derive the posterior associated with linear transition equation, $p_L(\theta^T, V \mid Y^T)$, and then we multiply by $I(\theta^T)$ to rule out explosive outcomes. In the Monte Carlo simulation, this is implemented by simulating the unrestricted posterior and rejecting draws that violate the stability condition. The next subsection describes our method for simulating $p_L(\theta^T, V \mid Y^T)$, and the one after that confirms the validity of our rejection sampling procedure.

2.5 SIMULATING THE UNRESTRICTED POSTERIOR

Following Kim and Nelson (1999), we use the Gibbs sampler to simulate draws from $p_L(\theta^T, V \mid Y^T)$. The Gibbs sampler iterates on two operations. First, conditional on the data and hyperparameters, we draw a history of

states from $p_L(\theta^T | Y^T, V)$. Then, conditional on the data and states, we draw hyperparameters from $p_L(V | Y^T, \theta^T)$. Subject to regularity conditions (see Roberts and Smith 1992), the sequence of draws converges to a draw from the joint distribution, $p_L(\theta^T, V | Y^T)$.

2.5.1. Gibbs Step 1: States Given Hyperparameters Conditional on data and hyperparameters, the unrestricted transition law is linear and has normal innovations. Thus, the virtual states are Gaussian,

$$p_L(\theta^T | Y^T, V) = f(\theta^T | Y^T, V). \tag{2.21}$$

This density can be factored as[7]

$$f(\theta^T | Y^T, V) = f(\theta_T | Y^T, V) \prod_{t=1}^{T-1} f(\theta_t | \theta_{t+1}, Y^t, V). \tag{2.22}$$

The leading factor is the marginal posterior for the terminal state, and the other factors are conditional densities for the preceding time periods. Since the conditional densities on the right-hand side are Gaussian, it is enough to update their conditional means and variances. This can be done via the Kalman filter.

Deriving forward and backward recursions for $f(\theta^T | Y^T, V)$ is straightforward. Going forward in time, let

$$\theta_{t|t} \equiv E(\theta_t | Y^t, V),$$

$$P_{t|t-1} \equiv \text{Var}(\theta_t | Y^{t-1}, V), \tag{2.23}$$

$$P_{t|t} \equiv \text{Var}(\theta_t | Y^t, V).$$

represent conditional means and variances. These are computed recursively, starting from $\bar{\theta}$ and $\bar{P}$, by iterating on

$$K_t = (P_{t|t-1} X_t + C)(X_t' P_{t|t-1} X_t + R + X_t' C + C' X_t)^{-1},$$

$$\theta_{t|t} = \theta_{t-1|t-1} + K_t(y_t - X_t' \theta_{t-1|t-1}),$$

$$P_{t|t-1} = P_{t-1|t-1} + Q, \tag{2.24}$$

$$P_{t|t} = P_{t|t-1} - K_t(X_t' P_{t|t-1} + C').$$

7. See Kim and Nelson (1999, Chapter 8).

The matrix K_t is the Kalman gain.[8] At the end of the sample, these iterations yield the conditional mean and variance for the terminal state,

$$f(\theta_T \mid Y^T, V) = N(\theta_{T|T}, P_{T|T}). \tag{2.25}$$

This pins down the first factor in (2.22).

The remaining factors in (2.22) are derived by working backward through the sample, updating means and variances to reflect the additional information about θ_t contained in θ_{t+1}.[9] Let

$$\begin{aligned} \theta_{t|t+1} &\equiv E(\theta_t \mid \theta_{t+1}, Y^t, V), \\ P_{t|t+1} &\equiv \mathrm{Var}(\theta_t \mid \theta_{t+1}, Y^t, V), \end{aligned} \tag{2.26}$$

represent backward estimates of the mean and variance, respectively. Because the states are conditionally normal, these can be expressed as

$$\begin{aligned} \theta_{t|t+1} &= \theta_{t|t} + P_{t|t} \; P_{t+1|t}^{-1} \; (\theta_{t+1} - \theta_{t|t}), \\ P_{t|t+1} &= P_{t|t} - P_{t|t} P_{t+1|t}^{-1} P_{t|t}. \end{aligned} \tag{2.27}$$

Therefore the remaining elements in the (2.22) are

$$f(\theta_t \mid \theta_{t+1}, Y^T, V) = N(\theta_{t|t+1}, P_{t|t+1}). \tag{2.28}$$

Notice that the smoothed covariances depend only on the output of the Kalman filter, but the smoothed conditional means depend on realizations of θ_{t+1}. Accordingly, a random trajectory for states may be drawn from a backward recursion. First, draw θ_T from (2.25), using (2.24) to compute the mean and variance. Next, conditional on its realization, draw θ_{T-1} from (2.28), using (2.27) to compute the mean and variance. Then draw θ_{T-2} conditional on the realization of θ_{T-1}, and so on back to the beginning of the sample.

2.5.2 Gibbs Step 2: Hyperparameters Given States Conditional on Y^T and θ^T, the innovations are observable. Under the unrestricted linear transition law, these are identically and independently distributed normal random variables, and their conditional likelihood is Gaussian. When an

8. The formula for K_t differs from that given in Anderson and Moore (1979) for the case of correlated innovations because of a difference in assumptions about the timing of innovations.
9. Notice that the backward recursions are *not* determined by the Kalman smoother. We want the mean and variance for $f(\theta_t|\theta_{t+1}, Y^t, V) = f(\theta_t|\theta_{t+1}, Y^T, V)$. The Kalman smoother computes the mean and variance for $f(\theta_t|Y^T, V)$.

inverse-Wishart prior is combined with a Gaussian likelihood, the posterior is also an inverse-Wishart density,

$$p(V|Y^T, \theta^T) = IW(V_1^{-1}, T_1),$$ (2.29)

where

$$\begin{aligned} T_1 &= T_0 + T, \\ V_1 &= \bar{V} + \bar{V}_T, \end{aligned}$$ (2.30)

and $\bar{V}_T$ is proportional to the usual covariance estimator,

$$\frac{1}{T}\bar{V}_T = \frac{1}{T}\sum_{t=1}^{T} \begin{pmatrix} \varepsilon_t \\ v_t \end{pmatrix} (\varepsilon_t' \; v_t').$$ (2.31)

The posterior degree-of-freedom parameter is the sum of the prior degrees of freedom, T_0, plus the degrees of freedom in the sample, T. The posterior scale matrix is the sum of the prior and sample sum-of-squares matrices.[10]

To sample from an inverse-Wishart distribution, we exploit two facts. First, if a matrix V is distributed as IW(S, df), then V^{-1} is a Wishart matrix with scale matrix S and degrees of freedom df. Second, to simulate a draw from the Wishart distribution, we take df independent draws of a random vector η_i from a $N(0, S)$ density and form the random matrix $V^{-1} = \sum_{i=1}^{df} \eta_i\eta_i'$. Since V^{-1} is a draw from a Wishart density, V is a draw from an inverse-Wishart density.

2.5.3 Summary of the Gibbs Sampler To summarize, the Gibbs sampler iterates on two simulations, drawing states conditional on hyperparamaters and then hyperparameters conditional on states. After a transitional or "burn-in" period, the sequence of draws approximates a sample from the virtual posterior, $p_L(\theta^T, V | Y^T)$.

2.6 REJECTION SAMPLING

The final step is to impose the stability condition, which is done by checking the autoregressive roots at each date and rejecting draws with roots inside the unit circle. The rejection step ensures that the posterior density puts zero probability on explosive outcomes.

---

10. See Gelman *et al.* (1995).

To confirm the validity of this procedure, we check the conditions associated with rejection sampling.[11] The normalized target density is

$$p(\theta^T, V \mid Y^T) = \frac{I(\theta^T)p_L(\theta^T, V \mid Y^T)}{\iint I(\theta^T)p_L(\theta^T, V \mid Y^T)\, d\theta^T\, dV}. \qquad (2.32)$$

To perform rejection sampling, we need a candidate density, $g(\theta^T, V)$, that satisfies three properties. The candidate density must be non-negative and well defined for all (θ^T, V) for which $p(\theta^T, V \mid Y^T) > 0$, it must have a finite integral, and the importance ratio $R(\theta^T, V)$ must have a known upper bound M:

$$R(\theta^T, V) = \frac{p(\theta^T, V \mid Y^T)}{g(\theta^T, V)} \le M < \infty. \qquad (2.33)$$

A natural candidate density is the virtual posterior, $p_L(\theta^T, V \mid Y^T)$. Because this is a probability density, it is non-negative and integrates to 1. Since it is an unrestricted analogue of the target density, it is also well defined for all (θ^T, V) which occur with positive probability. Finally, the importance ratio is bounded by the reciprocal of the probability of obtaining a stable draw from the virtual posterior,

$$R(\theta^T, V) = \frac{I(\theta^T)}{\iint I(\theta^T)p_L(\theta^T, V \mid Y^T)\, d\theta^T\, dV}$$

$$\qquad (2.34)$$

$$\le \frac{1}{\iint I(\theta^T)p_L(\theta^T, V \mid Y^T)\, d\theta^T\, dV} = M.$$

The denominator is the expected value of $I(\theta^T)$ under the virtual posterior, or the probability of a stable draw from the unrestricted density. M is finite as long as this probability is nonzero.

Rejection sampling proceeds in two steps: draw a trial (θ_i^T, V_i) from the virtual posterior, and then accept the draw with probability $R(\theta_i^T, V_i)/M$. Since $R(\theta_i^T, V_i)/M = I(\theta_i^T)$, the second step is equivalent to accepting the trial draw whenever it satisfies the stability condition, and rejecting it when it does not.

2.7 BELIEFS ABOUT THE FUTURE

Having processed data through date T, the next step is to simulate future data and states. Conditional on hyperparameters and the current

---

11. See, e.g., Gelman *et al.* (1995, pp. 303–305).

state of the system, the posterior density for future data and states is quite tractable. This density can be factored into the product of a marginal distribution for future states and a conditional distribution for future data,

$$p(Y^{T+1,T+H}, \theta^{T+1,T+H} | \theta_T, V, Y^T)$$

$$= p(\theta^{T+1,T+H} | \theta_T, V, Y^T) \, p(Y^{T+1,T+H} | \theta^{T+1,T+H}, \theta_T, V, Y^T). \qquad (2.35)$$

Because the states are Markov, the first factor can be factored in turn into

$$p(\theta^{T+1,T+H} | \theta_T, V, Y^T) = \prod_{i=1}^{H} p(\theta_{T+i} | \theta_{T+i-1}, V, Y^T). \qquad (2.36)$$

Apart from the restriction on explosive autoregressive roots, θ_{T+1} is conditionally normal with mean θ_T and variance Q. Similarly, conditional on θ_{T+1}, V, and Y^T, θ_{T+2} is normally distributed with mean θ_{T+1} and variance Q, and so on. Therefore, to sample from the virtual posterior for future states, we take H random draws of v_i from the $N(0, Q)$ density and iterate on the state equation,

$$\theta_{T+i} = \theta_{T+i-1} + v_i. \qquad (2.37)$$

The stability restriction is implemented in the same way as in the Gibbs sampler, by checking the autoregressive roots associated with each draw and rejecting explosive draws.

Given a trajectory for future states, all that remains is to simulate future data. The second factor in (2.35) can be factored in turn into

$$p(Y^{T+1,T+H} | \theta^{T+1,T+H}, \theta_T, V, Y^T) = \prod_{i=1}^{H} p(y_{T+i} | Y^{T+1,i-1}, \theta^{T+1,T+H}, \theta_T, V, Y^T). \qquad (2.38)$$

Conditional on θ_T, V, Y^T, and a trajectory for future states, the measurement innovation ε_{T+1} is normally distributed with mean $C'Q^{-1}v_{T+1}$ and variance $R - C'Q^{-1}C$. Hence y_{T+1} is conditionally normal with mean $X'_{T+1}\theta_{T+1} + C'Q^{-1}v_{T+1}$ and variance $R - C'Q^{-1}C$. Similarly, ε_{T+2} is conditionally normal with mean $C'Q^{-1}v_{T+2}$ and variance $R - C'Q^{-1}C$, and so on. Therefore, to sample from (2.38), we take H random draws of ε_i from a $N(C'Q^{-1}v_{T+i}, R - C'Q^{-1}C)$ density and iterate on the measurement equation,

$$y_{T+i} = X'_{T+i}\theta_{T+i} + \varepsilon_i, \qquad i = 1, \dots, H, \tag{2.39}$$

using lags of y_{T+i} to compute X_{T+i}.

2.8 COLLECTING THE PIECES

Combining the results of the previous sections, (2.16) can be expressed as

$$p(Y^{T+1,T+H}, \theta^{T+1,T+H}, \theta^T, V \mid Y^T) = p(\theta^T, V \mid Y^T)$$

$$\times \prod_{i=1}^{H} p(\theta_{T+i} \mid \theta_{T+i-1}, V, Y^T)$$

$$\times \prod_{i=1}^{H} p(y_{T+i} \mid Y^{T+1,i-1}, \theta^{T+1,T+H}, \theta_T, V, Y^T). \tag{2.40}$$

To sample from this distribution, we use the Gibbs sampler to simulate a draw from $p(\theta^T, V \mid Y^T)$. Then, conditional on that draw, we simulate a trajectory for future states, and conditional on both of those we simulate a trajectory for future data. This provides the raw material for our analysis.

3. *Stylized Facts about the Evolving Law of Motion*

We study data on inflation, unemployment, and a short-term nominal interest rate. Inflation is measured using the CPI for all urban consumers, unemployment is the civilian unemployment rate, and the nominal interest rate is the yield on 3-month Treasury bills. The inflation and unemployment data are quarterly and seasonally adjusted, and the Treasury-bill data are the average of daily rates in the first month of each quarter. The sample runs from 1948.1 to 2000.4. We work with a VAR(2) specification for inflation, the logit of unemployment, and the ex post real interest rate.[12]

To calibrate the prior, we estimate a time-invariant vector autoregression using data for 1948.1–1958.4. The mean of the virtual prior, $\bar{\theta}$, is the point estimate; $\bar{P}$ is its asymptotic covariance matrix; and $\bar{R}$ is the innovation covariance matrix. To initialize the other hyperparameters, we assume that $\bar{C} = 0$ and that $\bar{Q}$ is proportional to $\bar{P}$. To begin conservatively, we start with a minor perturbation from a time-invariant representation, setting $\bar{Q} = (0.01)^2 \bar{P}$. In other words, our prior is that time variation ac-

---

12. The unemployment rate is bounded between 0 and 1, and the logit transformation maps this into $(-\infty, \infty)$, which is more consonant with our Gaussian approximating model. To ensure that posterior draws for unemployment lie between 0 and 1, we simulate logit(u_t) and use the inverse logit transformation. The non-negativity bound on nominal interest rates is implemented by rejection sampling.

counts for only 1% of the standard deviation of each parameter.[13] The prior degrees of freedom, T_0, are equal to those in the preliminary sample.

This is an informative prior, but only weakly so. Because the preliminary sample contains only 4.5 data points per VAR parameter, the prior mean is just a ballpark number and the prior variance allows for a substantial range of outcomes. As time passes, the prior becomes progressively less influential and the likelihood comes to dominate the posterior.

The simulation strategy follows the algorithm described above. Starting in 1965.4, we compute posterior densities for each year through 2000, for a total of 36 years. At each date, we perform 10,000 iterations of the Gibbs sampler, discarding the first 2000 to let the Markov chain converge to its ergodic distribution.[14] Then, conditional on those outcomes, we generate 8000 trajectories of future data and states. Each posterior trajectory is 120 quarters long and contains information about both short- and long-run features of the data.

3.1 OBJECTS OF INTEREST

We initially focus on three features of the data: long-horizon forecasts of inflation and unemployment, the spectrum for inflation, and selected parameters of a version of the Taylor rule for monetary policy. The long-horizon forecasts approximate core inflation and the natural rate of unemployment, the spectrum encodes information about the variance, persistence, and predictability of inflation, and the Taylor-rule parameters summarize the changes in monetary policy that underlie the changing nature of inflation.

We are interested in these features because they play a role in theories about the rise and fall of U.S. inflation. For example, Parkin (1993) and Ireland (1999) point out that the magnitude of inflationary bias in the Kydland–Prescott (1977) and Barro–Gordon (1983) model depends positively on the natural rate of unemployment. Taylor (1997, 1998) and Sargent (1999) argue that core inflation depends on the monetary authority's beliefs about the natural-rate hypothesis, which in turn depend on the degree of inflation persistence. In particular, the model presented in Sargent (1999) imposes a definite restriction on the joint evolution of core inflation and the degree of persistence, which we discuss below. Changes

---

13. The Gibbs sampler quickly adds more time variation to the system.
14. Recursive mean graphs suggest rough convergence, though some wiggling persists beyond the burn-in period. We checked our results by performing a much longer simulation based on data through 2000.4. The longer simulation involved 106,000 draws from the Gibbs sampler, the first 18,000 of which were discarded to allow for convergence. Smoothed estimates based on this simulation were qualitatively similar to the filtered estimates reported in the text. Indeed, we also performed calculations based on a burn-in period of 98,000 and found that the results were much the same.

in beliefs about the natural-rate hypothesis should also be reflected in Taylor-rule parameters.

3.2 CORE INFLATION AND THE NATURAL RATE OF UNEMPLOYMENT

Beveridge and Nelson (1981) define a stochastic trend in terms of long-horizon forecasts. For a driftless random variable like inflation or unemployment, the Beveridge–Nelson trend is defined as the value to which the series is expected to converge once the transients die out,

$$\tau_t = \lim_{h \to \infty} E_t x_{t+h}. \tag{3.1}$$

Assuming that expectations of inflation and unemployment converge to the core and natural rates as the forecast horizon lengthens, the latter can be approximated using this measure.[15] Because the posterior distributions are skewed and have fat tails, we modify the Beveridge–Nelson definition by substituting the posterior median for the mean. We approximate core inflation and the natural rate of unemployment by setting $h = 120$ quarters and finding the median of the posterior predictive density,

$$\pi_{ct} = \text{med}_t(\pi_{t+120}),$$
$$u_{nt} = \text{med}_t(u_{t+120}). \tag{3.2}$$

Estimates of core inflation and the natural rate are shown in Figure 1. The circles represent inflation, and the crosses unemployment. According to this measure, core inflation was between 1.75% and 4% in the late 1960s. It rose throughout the 1970s and peaked at roughly 8% in 1979–1980. Thereafter it fell quickly, and it has fluctuated between 2.25% and 3.25% since the mid-1980s. Core inflation was just shy of 3% at the end of 2000.

The natural rate of unemployment also rose throughout the 1970s, reaching a peak of 6.6% in 1980. It declined gradually in the early 1980s and fluctuated between 5.5% and 6% from the mid-1980s to the mid-1990s. The natural rate again began to fall after 1994 and was a bit less than 5% at the end of 2000.

A scatterplot, shown in Figure 2, provides a better visual image of the association between the two. The simple correlation is 0.63, which is rather remarkable given the difficulty of measuring these components.

15. Hall (1999) recommends an unconditional mean of unemployment as an estimator of the natural rate of unemployment.

Figure 1 CORE INFLATION AND THE NATURAL RATE OF
UNEMPLOYMENT

Figure 2 CORE INFLATION AND THE NATURAL RATE OF
UNEMPLOYMENT

The two series rise and fall together, in accordance with Parkin and Ireland's theory.

As a reality check for the model, Figures 3 and 4 report the cyclical components of inflation and unemployment, measured by subtracting the median Beveridge–Nelson trend estimates from the actual values. We include these plots to confirm that the model captures important features of the data. The first figure shows that the estimated peaks and troughs occur at the right times and are of plausible magnitude. For example, unemployment was well above the natural rate following the recessions of 1975 and 1982. Using Okun's law as a rule of thumb, these estimates correspond to "output gaps" of roughly 6.75% and 12.5% respectively. The model also correctly predicts that the high inflation of 1974–1975 and 1980–1981 would be partially reversed.

Figure 4 shows a scatterplot of the cyclical components and illustrates two other characteristics of the data. The first is that the components are asymmetric, with large positive deviations occurring more often than large negative ones. Second, from 1967 until 1983, there were large counterclockwise loops in inflation and unemployment, with increases in inflation leading increases in unemployment. After 1986, the loops were smaller but still mostly counterclockwise. The direction of the

Figure 3 CYCLICAL COMPONENTS OF INFLATION AND
UNEMPLOYMENT

Figure 4 CYCLICAL COMPONENTS OF INFLATION AND
UNEMPLOYMENT

loops is consistent with other evidence on the cyclical relation between inflation and economic activity, e.g. as summarized by Taylor (1999).

Beveridge–Nelson measures often suggest that all the variation is in the trend, a feature to which many economists object. Our model does not have this feature.

3.3 THE PERSISTENCE, VARIANCE, AND PREDICTABILITY OF INFLATION

Next we consider the evolution of the second moments of inflation. This information is encoded in the spectrum, and its evolution is illustrated in Figures 5 through 7.

Figure 5 shows the median posterior spectrum for each year in the sample. This figure was generated as follows. For each year, we estimated a spectrum for each inflation trajectory in the posterior predictive density. Then we computed a median spectrum by taking the median of the estimates on a frequency-by-frequency basis.[16] This yields a single

16. The ordinates are asymptotically independent across frequencies.

Figure 5 MEDIAN POSTERIOR SPECTRUM FOR INFLATION
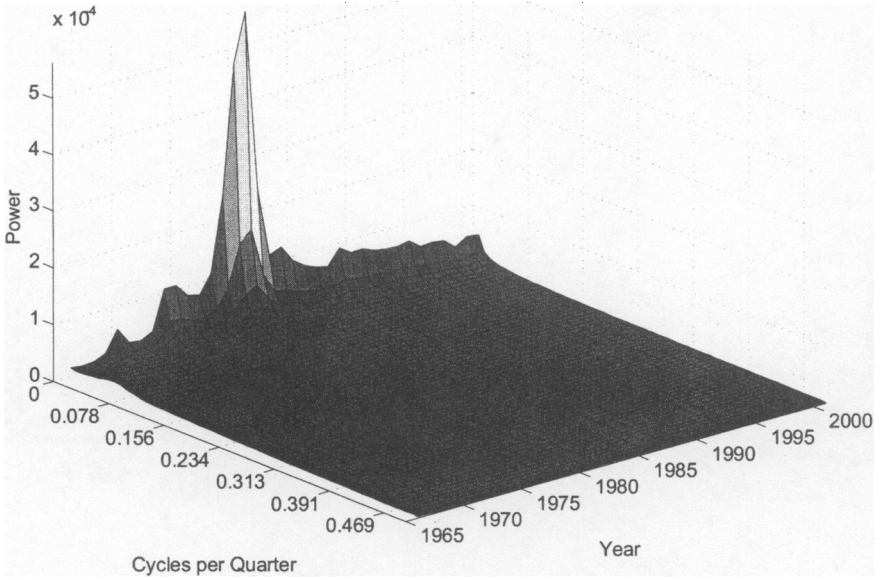
Power is measured in basis points, the units of measurement for the variance of inflation.

Figure 6 MEDIAN POSTERIOR SPECTRUM FOR INFLATION IN SELECTED YEARS

Figure 7 LOG OF THE MEDIAN POSTERIOR SPECTRUM FOR INFLATION
IN SELECTED YEARS

slice of the figure, relating power to frequency for a given year. By repeating this for each year, we produced the three-dimensional surface shown in the figure. We emphasize that these are predictive measures, which represent expected variation going forward in time. That is, the slice associated with a given year represents a prediction about how inflation is likely to vary in the future, conditional on data up to the current date.[17]

The most significant feature of this graph is the variation over time in the magnitude of low-frequency power. Since the spectral densities have Granger's (1966) typical shape, we can interpret low-frequency power as a measure of inflation persistence. According to this measure, inflation was weakly persistent in the 1960s and 1990s, when there was little low-frequency power, but strongly persistent in the late 1970s, when there was a lot. Indeed, the degree of persistence peaked in 1979–1980, at the same time as the peak in core inflation.

Figures 6 and 7 report results for selected years. Here, circles represent 1965, crosses 1979, and asterisks 2000. Figure 6 plots the spectrum,

17. We also calculated an alternative local linear approximation using the VAR representation and the mean posterior state at each date. The results were similar to those shown in the figure.

and Figure 7 plots its logarithm. To interpret the figures, recall that the total variance is the integral of the spectrum,

$$\sigma_\pi^2 = \frac{1}{2\pi} \int_{-\pi}^{\pi} f_{\pi\pi}(\omega)\, d\omega, \tag{3.3}$$

and that the log of the univariate innovation variance can be expressed as the integral of the log of the spectrum,

$$\ln \sigma_\epsilon^2 = \frac{1}{2\pi} \int_{-\pi}^{\pi} \ln f_{\pi\pi}(\omega)\, d\omega. \tag{3.4}$$

The function $f_{\pi\pi}(\omega)$ is the spectrum at frequency ω, σ_π^2 is the variance, and σ_ϵ^2 is the error variance for one-step-ahead univariate forecasts of inflation. The former measures long-run uncertainty about inflation; the latter, short-run uncertainty.

Looking first at Figure 6, we can say something about how the total variance has changed over time. Between 1965 and 1979, inflation became smoother but more persistent. That is, there was less variation at high and medium frequencies, especially those associated with business cycles (say 4 to 20 quarters per cycle), but more variation at low frequencies, especially those corresponding to cycles lasting 5 years or more. The increase in low-frequency power was greater in magnitude than the decrease in high-frequency power, so the total variance was greater. Thus, the increase in variance during the late 1960s and 1970s reflected an increase in inflation persistence.

Between 1979 and 2000, the spectrum for inflation fell at all frequencies, and therefore so did the total variance. But the decline in power was greatest at low frequencies, especially at those greater than 20 quarters per cycle. In other words, the diminished degree of inflation persistence accounted for most of the decline in variance in this period. Thus the evolution of the variance has been closely associated with that of inflation persistence. Inflation became more persistent and more variable in the 1970s, and less persistent and less variable in the 1980s and 1990s.

Figure 7 is relevant for short-term forecasting and tells a somewhat different story. The increase in the log of low-frequency power between 1965 and 1979 was smaller in magnitude than the decrease in the log of high-frequency power. Thus, although inflation became more persistent and more variable during the 1970s, it also became easier to predict one quarter ahead. In other words, although there was more long-term uncertainty in 1979, there was actually less short-term uncertainty. Between

1979 and 2000, the log spectrum fell at all frequencies, and inflation became even easier to forecast one quarter ahead. By 2000 there was less uncertainty at both long and short horizons.

The next two figures provide more information about prediction errors. Figure 8 is a multivariate analogue of Figure 7 and is related to the total prediction variance for the system. To interpret this figure, recall that the total prediction variance, $|V_{\epsilon\epsilon}|$, for a vector time series y_t can be expressed in terms of the log of the determinant of the spectral density,

$$\ln|V_{\epsilon\epsilon}| = \frac{1}{2\pi} \int_{-\pi}^{\pi} \ln|F_{yy}(\omega)| \, d\omega, \tag{3.5}$$

where $V_{\epsilon\epsilon}$ is the covariance matrix for innovations based on the history of y_t, and $F_{yy}(\omega)$ is the spectral density matrix. Whittle (1953) interprets $|V_{\epsilon\epsilon}|$ as a measure of the total random variation entering the system at each date.

Unlike the univariate measure, the total prediction variance increased between 1965 and 1979. For the system as a whole, there was only a slight decrease in variation at business-cycle frequencies, and this was

Figure 8 LOG DETERMINANT OF THE MEAN POSTERIOR SPECTRAL
   DENSITY MATRIX IN SELECTED YEARS

Figure 9 STANDARD DEVIATION OF ONE-STEP-AHEAD VAR PREDICTION
ERRORS

more than offset by a substantial increase in variation at low and high
frequencies. Between 1979 and 2000, the system became more predict-
able, with $\ln|F_{yy}(\omega)|$ falling at all frequencies. This more than reversed the
increase in the earlier period. By the end of 2000, the total prediction
variance was 40% smaller than in 1979 and 30% smaller than in 1965.
Thus, for the system as a whole, the degree of short-term uncertainty
has fallen substantially.

Figure 9 reports the variance of VAR forecast errors over the period
1965–2000 and provides more detail about the evolution of short-run
uncertainty. At each date, the posterior prediction error variance was
computed by averaging across realizations of the posterior predictive
density, one quarter ahead. For inflation and ex post real interest rates,
there has been a downward trend in short-term uncertainty since 1965,
punctuated by an increase in 1974 and again in 1978–1982. According to
this measure, the VAR innovation variance for inflation fell by 21%
between 1979 and 2000 and by 42% for the period as a whole. In contrast,
the forecast error variance for unemployment fluctuated until the early
1980s, rising and falling with the business cycle. Since then it has fallen
steadily to less than one-third its peak level. Changes in short-run uncer-

Figure 10 CORE INFLATION AND INFLATION PERSISTENCE

tainty about unemployment account for much of the rise and fall of the total prediction variance.[18]

Finally, in Figures 10 and 11, we relate changes in core inflation to the evolution of the variance and degree of persistence of inflation. Figure 10 plots core inflation and the spectrum at frequency zero, which summarizes the degree of persistence. The two are very closely related. Both rose in the 1960s and 1970s, and both fell during and after the Volcker disinflation. The simple correlation is 0.915.

Because persistence contributes to variance, core inflation also covaries positively with the long-horizon standard deviation of inflation, as shown in Figure 11.[19] Again, both measures rose during the 1970s and fell during the 1980s and 1990s. The correlation between the mean and standard deviation is 0.783. This is a bit lower than the previous correlation because the variance includes changes in both low- and high-frequency power, and the latter are less highly correlated with changes in core inflation. Thus the well-known positive correlation between the

18. Although our model assumes that V is constant, the figures illustrate that filtered estimates do shift little by little over time, thus introducing a limited degree of variation in shock variances. This variation may reflect a transient adaptation to the kind of shifts emphasized by our discussants.

19. We focus on the long-horizon variance, $\text{var}_t(\pi_{t+120})$, in order to let the transients die out.

Figure 11 CORE INFLATION AND THE STANDARD DEVIATION OF
INFLATION, 30 YEARS AHEAD

mean and variance of inflation reflects an even stronger correlation be-
tween the mean and degree of persistence.

3.4 TAYLOR-RULE PARAMETERS

At the end of the day, we hope to interpret the evolution of inflation
dynamics in terms of the changing behavior of central bankers. Accord-
ingly, we also investigate the evolution of the parameters of a Taylor rule.

A simple form of the Taylor rule posits that the central bank's nominal
interest target, i_t^*, varies positively with inflation and negatively with
unemployment,

$$i_t^* = (r^* + \pi^*) + \beta(\pi_{t-1} - \pi^*) + \gamma(u_{t-1} - u^*), \tag{3.6}$$

where π^*, u^*, and r^* represent target values for inflation, unemploy-
ment, and the real interest rate, respectively. The lags in the relationship
reflect the fact that current observations on inflation and unemployment
are often unavailable to policymakers, especially early in the quarter.[20]

20. This is relevant in our case because the interest rate is sampled in the first month of the
    quarter.

Therefore decisions are based on lagged values of inflation and unemployment. The basic Taylor rule is usually augmented with a policy shock η_t and a partial adjustment formula to allow for interest-rate smoothing,

$$\Delta i_t = \rho(L)(i_t^* - i_{t-1}) + \eta_t. \tag{3.7}$$

Cast in this form, the Taylor rule can be represented as the interest-rate equation in a vector autoregression for inflation, unemployment, and nominal interest rates.

In an alternative form of the Taylor rule, decisions about the ex ante real interest rate depend on lags of inflation, unemployment, and ex post real rates,

$$i_t - E_{t-1}\pi_t = \mu + \beta(L)\pi_{t-1} + \gamma(L)u_{t-1} + \rho(L)(i_{t-1} - \pi_{t-1}) + \eta_t. \tag{3.8}$$

By substituting $\pi_t = E_{t-1}\pi_t + \varepsilon_{\pi t}$, this form can be cast as the real-interest equation in a vector autoregression for inflation, unemployment, and ex post real rates, with a composite innovation consisting of policy shocks and inflation prediction errors,

$$i_t - \pi_t = \mu + \beta(L)\pi_{t-1} + \gamma(L)u_{t-1} + \rho(L)(i_{t-1} - \pi_{t-1}) + (\eta_t - \varepsilon_{\pi t}). \tag{3.9}$$

This is the form of the Taylor rule that we shall study.[21] In response to our discussants, we concede that it is controversial to interpret the systematic part of the monetary policy rule as the projection of real interest rates only on *past* information. By orthogonalizing an innovation covariance matrix in a particular order, many studies attribute part of the contemporaneous covariance among innovations to the monetary rule (i.e., the rule for setting interest rates responds to contemporary information). We also recognize that the shapes of impulse response functions for the response of macroeconomic aggregates to the monetary policy shock can depend sensitively on how much of the contemporaneous innovation volatility is swept into the monetary shock. In defense of our choice, we note that among others McCallum and Nelson (1999) doubt that monetary authorities have timely and reliable enough reports to let them respond to what the vector autoregression measures as contemporaneous information.[22]

---

21. Actually, we substitute the logit of unemployment for unemployment.
22. It would have been possible for us to condition on contemporaneous information by using the time t estimate of the R component of V to orthogonalize R as desired, though we have not done so in this paper.

The literature on monetary policy rules emphasizes several aspects of central-bank behavior. We focus on two elements that are especially relevant to the evolution of the law of motion for inflation. One concerns the evolution of target inflation, π^*, and the other concerns the evolution of the degree of activism.

The value of target inflation cannot be identified from the interest-rate equation alone. But assuming that the central bank adjusts interest rates so that inflation eventually converges to its target, this value can be estimated by computing long-horizon forecasts using the entire vector autoregression. Under this assumption, target and core inflation are synonymous. Evidence on this feature of the policy rule is reported above, in Figure 1.

Another important issue concerns whether a rule is *activist* or *passivist*, a distinction that bears on the determinacy of equilibrium (e.g., see Clarida, Gali, and Gertler, 2000). A policy rule is activist if, other things equal, the central bank increases the nominal interest rate more than one-for-one in response to an increase in inflation, so that the real interest rate increases. A passivist central bank adjusts the nominal interest rate one-for-one or less, so that the real interest remains constant or falls as inflation rises. In the real-interest version of the Taylor rule, the degree of activism can be measured by

$$A = \frac{\beta(1)}{1 - \rho(1)}. \tag{3.10}$$

A policy rule is activist if $A > 0$.

Because our version of the Taylor rule is the real-interest equation in the vector autoregression, the posterior density for the activism coefficient can be computed directly from the posterior density for the states. The output of the Gibbs sampler at date t includes the terminal state, θ_t, and for each draw of the terminal state we calculate the implied value for A. Conditional on data up to date t, this measures the degree of activism that would be forecast going forward from date t.

Posterior beliefs about A are illustrated in Figure 12. Because of outliers in the posterior density, the figure graphs the posterior median and interquartile range.[23] The figure has two salient features. First, as reported by Clarida, Gali, and Gertler, there have been important changes in the degree of activism over time. Judging by the posterior median, which is marked by circles, the degree of activism declined in

---

23. The outliers result from division by $1 - \rho(1)$, which sometimes takes on values close to zero.

Figure 12 POSTERIOR MEDIAN AND INTERQUARTILE RANGE FOR THE
ACTIVISM COEFFICIENT

the late 1960s and was approximately neutral in the early 1970s. For the
remainder of the 1970s, the rule was decidedly passive, allowing real
interest rates to fall as inflation rose. Monetary policy started becoming
activist in 1981 and continued to grow more activist until the end of
Volcker's term. During the first half of Greenspan's term, policy drifted
toward a less active stance, perhaps reflecting the "opportunistic" ap-
proach to disinflation. But policy has again grown more activist since
1993, surpassing the peak achieved at the end of the Volcker years.

The second notable feature concerns the dispersion of beliefs about the
degree of activism. Judging by the interquartile range, beliefs were tightly
concentrated only in the 1970s, when monetary policy was passive. At
that time, there seemed to be little doubt, for better or worse, about how
the Fed was doing business. The periods before and after both involve
more uncertainty about the degree of activism. In the 1960s, the lower end
of the interquartile range straddled the boundary of the activist region. In
the Volcker–Greenspan years, the interquartile range was wider but
safely within the activist region.

Figure 13 shows how the activism parameter has covaried with core
inflation and the degree of inflation persistence.[24] The latter both in-

24. The variables are measured in standard units in order to put them on a common basis.

Figure 13 CORE INFLATION, INFLATION PERSISTENCE, AND POLICY
ACTIVISM

creased during the 1970s experiment with a passivist monetary rule, and
they both fell in the 1980s and 1990s as policy became more activist. The
correlation between the degree of activism and core inflation is −0.69
over the full sample and −0.87 in the Volcker–Greenspan era. Similarly,
the correlation between the activism and persistence measures is −0.46
over the full sample and −0.76 in the Volcker–Greenspan years. Thus,
as one might expect, there is an inverse relation between the degree of
activism on the one hand and core inflation and inflation persistence on
the other.

4. Testing the Natural-Rate Hypothesis

Figures 14 through 16 summarize the consequences of implementing
econometric tests of the natural-rate hypothesis along the lines of Solow
(1968), Tobin (1968), Gordon (1970), and many others. They tested the
natural-rate hypothesis by regressing inflation on its own lags along
with current and lagged values of unemployment,

$$\pi_t = \beta_0 + \beta_1(L)\pi_{t-1} + \beta_2(L)u_t + \varepsilon_t. \tag{4.1}$$

Figure 14 RECURSIVE TESTS OF THE NATURAL-RATE HYPOTHESIS

Circles represent recursive least-squares estimates with decreasing gain, and diamonds represent constant-gain estimates.

They interpreted the condition $\beta_1(1) = 1$ as evidence in favor of the natural-rate hypothesis, and $\beta_1(1) < 1$ as evidence in favor of a long-run trade-off.[25]

The outcomes of recursive natural-rate tests are shown in Figure 14. The initial estimates are based on data from 1948 through 1964, allowing for lags at the beginning of the sample. On the right-hand side of equation (4.1), we include two lags of inflation along with the current value and two lags of unemployment. Starting in 1965.1, new data are added one quarter at a time, and $\hat{\beta}_1(1)$ and its t-ratio are updated using the

25. The thought experiment in play imagines the consequences of a permanent increase in expected inflation, which is proxied by the lagged inflation terms on the right-hand side. In order for this to be neutral in the long run, it must be the case that this has a one-for-one effect on actual inflation, so that $\beta_1(1) = 1$. Assuming that current unemployment is predetermined with respect to current inflation, this regression can be estimated by least squares. King and Watson (1997) point out that the last assumption follows from the structure of vintage 1960s Keynesian models, in which unemployment and inflation were determined in a block recursive fashion. Unemployment was determined by aggregate demand and Okun's law. Taking unemployment as given, inflation was determined by a Phillips-curve relation for wages and a markup equation for prices.

## Figure 15 INFLATION PERSISTENCE AND NRH TEST STATISTICS

Figure 16 CORE INFLATION AND NRH TEST STATISTICS

Kalman filter. The figure plots the resulting sequence of t-statistics for $\hat{\beta}_1(1) - 1$. Points marked with a circle represent OLS estimates, and those marked with a diamond represent discounted least-squares (DLS) estimates. For the latter, the gain parameter was $g_t = \max(1/t, \frac{1}{120})$.[26] The horizontal line marks the 1% critical value for a one-sided test.

Sargent (1971) pointed out that this approach is valid only if the sample used to estimate $\beta_1(1)$ contains permanent shifts in inflation. Otherwise the data are uninformative for the thought experiment, and $\beta_1(1)$ could be less than 1 even if there were no long-run trade-off. Thus, as the degree of inflation persistence in the sample varies over time, so too will outcomes of the test.

Early versions of the test, based on samples in which there was little inflation persistence, found estimates of $\beta_1(1) < 1$ and were interpreted as evidence in favor of a long-run trade-off. As shown in the figure, the natural-rate hypothesis was strongly rejected through 1973. Later versions were based on samples containing more inflation persistence, and they fail to reject long-run neutrality. Indeed, from the mid-1970s until the mid-1980s there was very little evidence against long-run neutrality. Since then, as the degree of inflation persistence has fallen, evidence against the natural-rate hypothesis has grown.

Figure 15 illustrates the relation between inflation persistence and outcomes of the test.[27] The figure confirms that the test statistic is positively related to the degree of persistence, though the relation is nonlinear. Once there was enough persistence to identify the long run trade-off parameter, the test began to accept long-run neutrality, and further increases in persistence no longer increased the t-ratio. Figure 16 shows that the test statistic is also positively related with core inflation. Without alterations, the model of Sims (1988), Chung (1990), and Sargent (1999) cannot explain that pattern. In that model, persistence rises and the natural-rate hypothesis is learned as inflation falls, so the model predicts an inverse relation between core inflation and the outcome of the test. The pattern shown in Figure 16 is more consistent with an alternative story, in which the upward drift in inflation taught the government to accept the natural-rate hypothesis via the Solow–Tobin test.

Thus, though the Solow–Tobin procedure provided a valid test of the natural-rate hypothesis only when inflation had become sufficiently persistent, by the mid-1970s inflation had become persistent enough to let the test detect the natural rate. Therefore the Solow–Tobin econometric

---

26. There are only minor differences between the two estimators within the sample, because until recently $1/t > \frac{1}{120}$. The distinction between constant and decreasing gain estimators matters more when we consider the likely outcomes of future tests.

27. These figures refer to discounted least-squares estimates, but the results for OLS estimates are essentially the same.

procedures gave policymakers information that should have caused them to stabilize inflation if they had the preferences attributed to them, for example, by Kydland and Prescott (1977). For when a policymaker solves the problem of minimizing an expected discounted sum of a quadratic loss function in inflation and unemployment subject to a Phillips curve like (4.1), and when the policymaker accepts the natural-rate hypothesis in the form in which Solow and Tobin cast it, then for discount factors large enough, the policymaker will soon push average inflation to zero.[28] When Volcker took control, the advice to push inflation quickly toward zero came even from those models and optimal-control exercises that took inadequate account of the Lucas critique, because they rested on the Solow–Tobin test.

However, the strong inflation persistence that induced the Solow–Tobin test to detect the natural rate in the mid-1970s depended on the monetary authority's having recently allowed inflation to drift upward, perhaps in response to its earlier erroneous views about an exploitable trade-off. If the government's success in lowering inflation created lower persistence in inflation, the Solow–Tobin test could one day again point to an exploitable trade-off that would tempt later monetary authorities to use inflation to fight unemployment. That possibility has worried John Taylor and others, an issue to which we now turn.

5. Taylor's Warning about Recidivism

Recently, John Taylor (1998) has warned about recidivism on the natural-rate hypothesis. Taylor notes that inflation is lower and more stable in the current monetary regime, and he points out that as such data accumulate, erroneous econometric tests of long-run neutrality may again begin to suggest the existence of a trade-off. To the extent that the tests undermine confidence in the natural-rate hypothesis, they could also undermine support for a low-inflation policy. In this section, we offer quantitative evidence to back up Taylor's warning. The evidence is based on the posterior predictive density conditioned on data through the end of 2000. We use this to make predictions about the probability of rejecting the natural-rate hypothesis going forward in time.

Figure 14 suggests that Taylor's concern has some merit, because by the end of the sample conventional tests were close to rejecting $\beta_1(1) = 1$

---

28. This is a version of the control problem described by Phelps (1967) and Sargent (1999). Long ago, Albert Ando pointed out that good macroeconometric models had confirmed the absence of a long-run inflation–unemployment trade-off by the early or mid-1970s.

against $\beta_1(1) < 1$ at the 5% level. The in-sample evidence is marginal,[29] however, and it is an open question whether stronger evidence will emerge as data from a low-inflation regime accumulate. To address this question, we compute the posterior predictive density of natural rate t-ratios going forward in time from 2000.4. Then we calculate the probability, conditioned on what we know now, of rejecting the natural-rate hypothesis at various dates in the future. In this way, we can quantify the risk of backsliding.

Let $\hat{\tau}^{T+1,T+H}$ represent a potential future sequence of recursive t-statistics for $\beta_1(1) - 1$,

$$\hat{\tau}^{T+1,T+H} = [\hat{\tau}_{T+1}, \ldots, \hat{\tau}_{T+H}]'.$$

We want to make statements about how these sequences are likely to evolve. From a Bayesian perspective, the natural way to proceed is to compute the posterior predictive density for these sequences,

$$p(\hat{\tau}^{T+1,T+H} | Y^T). \tag{5.1}$$

To sample from this density, we start with the posterior predictive density for inflation and unemployment and then exploit the fact that t-statistics are deterministic functions of the data.[30] Hence we can write

$$p(\hat{\tau}^{T+1,T+H} | Y^T) = p(g(Y^{T+1,T+H}, Y^T)|Y^T), \tag{5.2}$$

where the function $g(\cdot)$ is nothing more than the output of the recursive least-squares algorithm initialized with estimates through date T. To draw a realization from (5.2), we first draw a trajectory for future inflation and unemployment from their posterior predictive density and then apply the Kalman filter to compute the associated sequence of test statistics. The probability that the test will reject at some future date h is

$$\int_{-\infty}^{c(\alpha)} p(\hat{\tau}_h|Y^T)\, d\hat{\tau}, \tag{5.3}$$

where $c(\alpha)$ is the normal critical value corresponding to a one-sided test of size α. In terms of our sampling strategy, this is the fraction of simu-

---

29. In our opinion, strong rejections will be needed to reverse the consensus in favor of the natural-rate hypothesis.
30. Remember, from a Bayesian perspective $\beta(1)$ is a random and $\hat{\beta}(1)$ is deterministic.

lated trajectories in which $\hat{\beta}_1(1)$ is significantly less than 1 at date h, where significance is determined by the usual classical criterion. Thus, we are offering a Bayesian interpretation of judgments based on a classical procedure.

Figure 17 reports results for a constant-gain estimator. The results for a recursive OLS estimator are similar. We focus on the constant-gain estimator because this holds the effective sample size constant as data accumulate. Thus the increased probability of rejection does not follow simply from an increase in the number of observations.

As the figure shows, the probability of rejection remains small in the first two years of the forecast. But then it increases quickly, reaches 50% within 9 years, and approaches 85% in 20 years. The increasing probability of rejection reflects the changing nature of inflation–unemployment dynamics along with the fact that data from new and old regimes are being mixed in different proportions. As time moves forward, data from the old high-inflation, strong-persistence regime are discounted more heavily, and data from the new low-inflation, weak-persistence

Figure 17 PROBABILITY OF REJECTING THE NATURAL-RATE HYPOTHESIS, CALCULATED FROM THE POSTERIOR PREDICTIVE DISTRIBUTION FOR THE CONSTANT-GAIN ESTIMATOR

regime increasingly dominate the sample. The identifying information from the 1970s is lost little by little, and the properties of the Volcker–Greenspan era come more and more into play. This confirms an element of Taylor's warning, that the Solow–Tobin test may once again begin to suggest the existence of a trade-off.

6. Concluding Remarks

This paper has used a vector autoregression with random coefficients to measure parameter drift in U.S. inflation–unemployment–interest-rate dynamics. We construct our model to focus on parameter drift because we are sympathetic to the theoretical views expressed in Lucas (1976) and Sargent (1999), which leads us to suspect that evolution in the monetary policy authority's view of the world will make the systematic part of a vector autoregression drift. We have taken seriously our model's description of four sources of uncertainty about the future,[31] and have used computer-intensive Bayesian methods to take those uncertainties into account. We use the model to develop a number of stylized facts about the evolution of postwar U.S. inflation and relate them to important issues about learning to detect the natural-rate hypothesis using imperfect tests, and how the evolving results from those tests were associated with evolution in a description of a monetary policy rule (a Taylor rule). Among other things, we find that the mean and persistence of inflation are strongly positively correlated; that the persistence of inflation is positively associated with statistics that have been used to test for accepting the natural-rate hypothesis; that evolving measures of policy activism in fighting inflation broadly point to more activism with a lag somewhat after test statistics began accepting the natural-rate hypothesis; and that recently the degree of persistence in inflation has been drifting downward as inflation has come under control.

We also study John Taylor's warning about recidivism toward an exploitable trade-off between inflation and unemployment. Unfortunately, our statistical model confirms Taylor's concerns. Our model predicts that as observations of lower, more stable inflation accumulate, econometric evidence against the natural-rate hypothesis is likely to develop.[32] Against

---

31. These are: (1) the unknown current location of the VAR coefficients, (2) the unknown covariance matrix of innovations to VAR coefficients and equations, (3) the future evolution of the VAR coefficients, and (4) the stream of future shocks to the VAR equations.

32. Prospects for a gradual backsliding away from the zero-inflation Ramsey outcome toward the higher Nash inflation rate also permeate the "mean dynamics" in the model of Sargent (1999) and Cho, Williams, and Sargent (2001).

this evidence, we hope that policymakers do not succumb again to the temptation to exploit the Phillips curve.

Appendix. A Nonlinear Transition Equation

Our numerical procedures construct a sample using $p(\theta^T|V)$ defined by (2.13). This appendix verifies that these procedures are consistent with the nonlinear transition function defined in the text. In particular, we verify the nonlinear transition equation, $p(\theta_{t+1}|\theta_t, V) \propto I(\theta_{t+1}) f(\theta_{t+1}|\theta_t, V)$ from equations (2.3), (2.13), (2.14), and (2.15). First consider the transition equation for terminal state,

$$p(\theta_T|\theta_{T-1}, V) = \frac{p(\theta_T, \theta_{T-1}|V)}{p(\theta_{T-1}|V)}. \tag{A.1}$$

The joint density in the numerator can be expressed as

$$p(\theta_T, \theta_{T-1}|V) = \int p(\theta^T|V)\, d\theta^{T-2} \tag{A.2}$$

$$\propto I(\theta_T) f(\theta_T|\theta_{T-1}, V) \int \prod_{t=0}^{T-2} I(\theta_{t+1}) f(\theta_{t+1}|\theta_t, V) d\theta^{T-2}.$$

The marginal density in the denominator of (A.1) can be expressed as

$$p(\theta_{T-1}|V) = \int p(\theta_T, \theta_{T-1}|V)\, d\theta_T \tag{A.3}$$

$$\propto \int I(\theta_T) f(\theta_T|\theta_{T-1}, V)\, d\theta_T \int \prod_{t=0}^{T-2} I(\theta_{t+1}) f(\theta_{t+1}|\theta_t, V) d\theta^{T-2}.$$

The ratio between the two is

$$p(\theta_T|\theta_{T-1}, V) \propto I(\theta_T) f(\theta_T|\theta_{T-1}, V). \tag{A.4}$$

Next consider the transition equation for the penultimate state,

$$p(\theta_{T-1}|\theta_{T-2}, V) = \frac{p(\theta_{T-1}, \theta_{T-2}|V)}{p(\theta_{T-2}|V)}. \tag{A.5}$$

The joint density in the numerator of (A.5) can be expressed as

$$p(\theta_{T-1}, \theta_{T-2}|V) = \int\int p(\theta^{T-1}|V)p(\theta_T|\theta_{T-1}, V)\, d\theta^{T-3}\, d\theta_T$$

$$= \int p(\theta^{T-1}|V)\, d\theta^{T-3} \int p(\theta_T|\theta_{T-1}, V)\, d\theta_T$$

$$= \int p(\theta^{T-1}|V)\, d\theta^{T-3}, \tag{A.6}$$

where the last equality follows from the fact that $p(\theta_T|\theta_{T-1}, V)$ integrates to one. Using the same argument as above, this can be expressed as

$$p(\theta_{T-1}, \theta_{T-2}|V) \propto I(\theta_{T-1})f(\theta_{T-1}|\theta_{T-2}, V) \int \prod_{t=0}^{T-3} I(\theta_{t+1})f(\theta_{t+1}|\theta_t, V)\, d\theta^{T-3}.$$

The marginal density for θ_{T-2} is

$$p(\theta_{T-2}|V) = \int p(\theta_{T-1}, \theta_{T-2}|V)\cdot d\theta_{T-1}$$

$$\propto \int I(\theta_{T-1})f(\theta_{T-1}|\theta_{T-2}, V)d\theta_{T-1} \int \prod_{t=0}^{T-3} I(\theta_{t+1})f(\theta_{t+1}|\theta_t, V)\, d\theta^{T-3}. \tag{A.7}$$

The ratio between the two is

$$p(\theta_{T-1}|\theta_{T-2}, V) \propto I(\theta_{T-1})f(\theta_{T-1}|\theta_{T-2}, V). \tag{A.8}$$

Continuing a backward recursion implies

$$p(\theta_t|\theta_{t-1}, V) \propto I(\theta_t)f(\theta_t|\theta_{t-1}, V). \tag{A.9}$$

Hence, the nonlinear transition equation can indeed be expressed in terms of the truncated linear transition equation.

REFERENCES

Ahmed, S., and J. H. Rogers. (1998). Inflation and the great ratios: Long-term evidence from the U.S. Washington: Board of Governors of the Federal Reserve System. International Finance Discussion Paper 628.

Albanesi, S., V. V. Chari, and L. J. Christiano. (2000). Expectations traps and monetary policy. Northwestern University and University of Minnesota. Mimeo.

Anderson, B. D. O., and J. B. Moore. (1979). *Optimal Filtering*. Englewood Cliffs, NJ: Prentice-Hall.

Barro, R. J., and D. B. Gordon. (1983). A positive theory of monetary policy in a natural rate model. *Journal of Political Economy* 91:589–610.

Bernanke, B. S., and I. Mihov. (1998a). The liquidity effect and long-run neutral-

ity. In *Carnegie-Rochester Conference Series on Public Policy*, Vol. 49, B. T. Mc-
Callum and C. I. Plosser (eds.). Amsterdam: North Holland, pp. 149–194.
————, and ————. (1998b). Measuring monetary policy. *Quarterly Journal of Economics* 113(August):869–902.

Beveridge, S., and C. R. Nelson. (1981). A new approach to decomposition of economic time series into permanent and transitory components with particular attention to measurement of the "Business Cycle." *Journal of Monetary Economics* 7:151–174.

Cho, I. K., N. Williams, and T. J. Sargent. (2001). Escaping Nash inflation. *Review of Economic Studies*, forthcoming.

Chung, H. (1990). Did policy makers really believe in the Phillips curve? An econometric text. University of Minnesota. PhD Dissertation.

Clarida, R., J. Gali, and M. Gertler. (2000). Monetary policy rules and macroeconomic stability: Evidence and some theory. *Quarterly Journal of Economics* 115(1):147–180.

DeLong, J. B. (1997). "America's only peacetime inflation: The 1970s. In *Reducing Inflation: Motivation and Strategy*, C. D. Romer and D. Romer (eds.). NBER Studies in Business Cycles, Vol. 30 Chicago: University of Chicago Press.

Doan, T., R. Litterman, and C. Sims. (1984). Forecasting and conditional projections using realistic prior distributions. *Econometric Reviews* 3(1):1–100.

Eisner, R. (1997). A new view of the NAIRU. In *Improving the Global Economy: Keynesianism and Growth in Output and Employment*, P. Davidson and J. Kregel (eds.). Cheltenham: Edward Elgar.

Estrella, A., and F. S. Mishkin. (1999). Rethinking the role of the NAIRU in monetary policy: Implications of model formulation and uncertainty. In *Monetary Policy Rules*, J. B. Taylor (ed.). NBER Conference Report. Chicago: University of Chicago Press.

Fair, R. (1996). Testing the standard view of the long-run unemployment–inflation relation. Cowles Commission. Discussion Paper 1121.

Fisher, M. E., and J. J. Seater. (1993). Long-run neutrality and superneutrality in an ARIMA framework. *American Economic Review* 83(3):402–415.

Friedman, B. (1998). Introduction. In *Inflation, Unemployment, and Monetary Policy*, R. M. Solow and J. B. Taylor (eds.). Cambridge, MA: MIT Press.

Gelman, A., J. B. Carlin, H. S. Stern, and D. B. Rubin. (1995). *Bayesian Data Analysis* London: Chapman and Hall.

Gordon, R. J. (1970). The recent acceleration of inflation and its lessons for the future. *Brookings Papers on Economic Activity* 1:8–41.

Granger, C. W. J. (1966). The typical spectral shape of an economic variable. *Econometrica* 34(1):150–161.

Hall, R. (1999). Comment on rethinking the role of the NAIRU in monetary policy: Implications of model formulation and uncertainty. In *Monetary Policy Rules*, J. B. Taylor (ed.). NBER Conference Report. Chicago: University of Chicago Press.

Ireland, P. (1999). Does the time-consistency problem explain the behavior of inflation in the United States? *Journal of Monetary Economics* 44(2):279–292.

Kim, C.-J., and C. R. Nelson. (1999). *State-Space Models with Regime Switching*. Cambridge, MA: MIT Press.

King, R. G., and M. W. Watson. (1994). The Post-war U.S. Phillips curve: A revisionist econometric history. *Carnegie-Rochester Conference Series on Public Policy* 41:157–219.

———, and ———. (1997). Testing long-run neutrality, *Federal Reserve Bank of Richmond Economic Quarterly* 83(3):69–101.

Kydland, F., and E. C. Prescott. (1977). Rules rather than discretion: The inconsistency of optimal plans. *Journal of Political Economy* 85(3):473–491.

Lucas, R. E., Jr. (1972). Econometric testing of the natural rate hypothesis. In *The Econometrics of Price Determination*, O. Eckstein (ed.). Washington: Board of Governors of the Federal Reserve System.

———. (1976). Econometric policy evaluation: A critique. In *The Phillips Curve and Labor Markets*, K. Brunner and A. Meltzer (eds.). Carnegie-Rochester Series on Public Policy, Vol. 1.

McCallum, B. T., and E. Nelson. (1999). Performance of operational policy rules in an estimated semiclassical structural model. In *Monetary Policy Rules*, J. B. Taylor (ed.). Chicago: University of Chicago Press, pp. 15–45.

Parkin, M. (1993). Inflation in North America. In *Price Stabilization in the 1990s*, K. Shigehara (ed.).

Phelps, E. S. (1967). Phillips curves, expectations of inflation, and optimal unemployment over time. *Economica* 2(3):22–44.

Roberts, G. O., and A. F. M. Smith. (1992). Simple conditions for the convergence of the Gibbs sampler and Metropolis–Hastings algorithms. *Stochastic Processes and Their Applications* 49:207–216.

Rudebusch, G. D., and L. E. O. Svensson. (1999). Policy rules for inflation targeting. In *Monetary Policy Rules*, J. B. Taylor (ed.). NBER Conference Report. Chicago: University of Chicago Press.

Sargent, T. J. (1971). A note on the accelerationist controversy. *Journal of Money Credit and Banking* 8(3):721–725.

———. (1983). Autoregressions, expectations, and advice. *American Economic Review, Papers and Proceedings* 74(2):408–415.

———. (1999). *The Conquest of American Inflation*. Princeton, NJ: Princeton University Press.

Sims, C. A. (1982). Policy analysis with econometric models. *Brookings Papers on Economic Activity* 1:107–152.

———. (1988). Projecting policy effects with statistical models. *Revista de Analysis Economico* 3:3–20.

———. (1999). Drifts and breaks in monetary policy. Princeton University. Mimeo.

Solow, R. M. (1968). Recent controversy on the theory of inflation: An eclectic view. In *Proceedings of a Symposium on Inflation: Its Causes, Consequences, and Control*, S. Rousseaus (ed.). New York: New York University.

———. (1998). How cautious must the Fed be? In *Inflation, Unemployment, and Monetary Policy*, R. M. Solow and J. B. Taylor (eds.). Cambridge, MA: The MIT Press.

Taylor, J. B. (1997). Comment on America's only peacetime inflation: The 1970s. In *Reducing Inflation: Motivation and Strategy*, C. D. Romer and D. Romer (eds.). NBER Studies in Business Cycles, Vol. 30. Chicago: University of Chicago Press.

———. (1998). Monetary policy guidelines for unemployment and inflation stability. In *Inflation, Unemployment, and Monetary Policy*, R. M. Solow and J. B. Taylor (eds.). Cambridge, MA: The MIT Press.

———. (1999). Staggered price and wage setting in macroeconomics. In *Handbook of Macroeconomics*, Vol. 1B, J. B. Taylor and M. Woodford (eds.). Amsterdam: North Holland.

Tobin, J. (1968). Discussion. In *Proceedings of a Symposium on Inflation: Its Causes, Consequences, and Control*, S. Rousseaus (ed.). New York: New York University.
Whittle, P. (1953). The analysis of multiple stationary time series. *Journal of the Royal Statistical Society, Series B* 15:125–139.

Comment

CHRISTOPHER A. SIMS
Princeton University

1. Introduction

My comments fall under three main headings:

(i) The later, Taylor-rule part of the paper is a structural VAR analysis. It uses nonstandard, and questionable, identifying assumptions without giving us a discussion of why it differs from most of the literature or what motivates the nonstandard specification. It also fails to check its specification as thoroughly as is standard in the structural VAR literature.

(ii) The evidence that monetary policy behavior has changed sharply between early and late postwar periods, or even between interwar and postwar periods, is less strong than might appear from this paper.

(iii) The paper sets a new, and high, standard for descriptive analysis of macroeconomic data. I hope it will be widely copied, and therefore want to be sure to register objections to certain aspects of its technical procedures before it's too late. Some of the questionable aspects of its procedures may have affected its conclusions.

2. Identification

There are several related facts about policy rules and their relation to the data that reflect the identification problem that must be confronted in evaluating claims to estimate a rule.

It is easy to generate "policy shocks" that produce strong price puzzles, particularly in pre-1979 data, as we see from Barth and Ramey's paper in this volume. Identification schemes that produce price puzzles tend also to imply large real effects of monetary policy shocks and small responses of interest rates to lagged inflation—low activism.

No matter what the actual policy rule, it will be possible to estimate a regression of interest rate on "fundamentals" (i.e. not P, M, or other nominal variables: intrinsic state variables) that can play the role of a statistical "interest-rate equation." Yet, in most equilibrium models, if this regression were in fact the policy rule and fiscal policy took the conventionally assumed form, the model's equilibrium would be indeterminate.

Observations from a gold-standard or price-level-targeting policy regime will spuriously imply a nonactivist policy rule unless quite sophisticated simultaneity is recognized in the estimation. This follows because in such regimes high inflation predicts low future inflation, which through the Fisher equation then implies low current nominal interest rates. Such a regime can be generated by a policy reaction function that makes r respond very strongly to the price level or inflation, but the policy reaction function is not recovered by OLS regression.

In other words, there is *always* an identification problem in determining whether policy is active. The identification problem can be solved, but only by bringing in identifying assumptions that are not testable.

One of the identifying assumptions in this paper is that the residual in a VAR ex post real-interest-rate equation with unemployment and CPI on the right is the policy shock, which amounts to a recursive VAR identification scheme. While much of the identified VAR literature relies on this assumption, it can lead to problematic interpretations of the data. Most prominently, price puzzles (inflationary response to monetary tightening) are a common outcome (as e.g. in Barth and Ramey's paper in this volume) when purely recursive identification schemes are applied to pre-1980 U.S. data. As Leeper and Zha (2001) show, policy rules are estimated as stable and without price puzzles when the fact that policy behavior (at least before 1980) involved responses to the money stock is allowed for and the resulting simultaneity is recognized.

The paper also presents its policy reaction function as a "real-interest-rate rule." The unusual timing of the paper's data (r is not a quarterly average, but rather a monthly average from the first month in the quarter, while the other data are quarterly averages) makes this assertion difficult to interpret. In a continuous-time, or cleanly discrete-time, model, when prices are flexible and money is neutral, the monetary authority simply cannot set the real interest rate. A policy equation with the real rate on the left, even if it has lagged inflation on the right, contradicts the mapping from the economy's real state to its real interest rate. With non-neutralities in the model, nonexistence will no longer be a logical necessity, but there will be a range of models, with weak non-neutralities, for which such policy rules raise existence problems. It

seems unwise to impose a policy rule of this form on the data as an *a priori* restriction.

To understand this problem, consider the simple model

$$r_t = E_{t-1}\pi_t + \bar{r}_t \qquad \text{(Fisher relation)},$$
$$r_t = E_{t-1}\pi_t + \alpha\pi_{t-1} + \gamma u_{t-1} + \varepsilon_t \quad \text{(policy rule)}.$$

It is easy to understand that this pair of equations leads to nonexistence of a stable rational-expectations equilibrium, because taking the difference of the two equations would force innovations in the real rate to be exact functions of innovations in the policy equation. If we replaced $E_{t-1}\pi_t$ in the first equation with $E_t\pi_{t+1}$, as would be appropriate if the model's data had conventional timing, the system would be well behaved. But of course, if the data had conventional timing, this specification would no longer represent policy setting the real rate. Replacing $E_{t-1}\pi_t$ in the second equation with π_t itself is no help, however, as the resulting system still has no solution. It would have been better for the paper to stick with a nominal-rate rule, as does the rest of the structural VAR literature. As it is, the interpretation of all the parts of the paper that depend on this identification is problematic.

I agree with the authors that it is reasonable to assert as an identifying assumption that policy responds only to lagged information. This view could have been incorporated into their structure simply by omitting current π_t from the reaction function.

Papers in the structural VAR literature almost universally check identification by examining impulse responses, trying to ensure that the estimated system does not have unreasonable properties. It is easy for apparently reasonable identifying restrictions to lead to estimated systems that are implausible, so this type of check is important. This paper does no such checking. Thus we do not know whether the periods of implied low activism also are accompanied by a price puzzle, whether the implied responses of monetary authorities to private shocks are reasonable, or whether the responses of the economy to the policy shocks are reasonable.

Probably the majority view among macroeconomists (and especially within the Fed system?) is that monetary policy has changed drastically for the better over the last 30 or 40 years—Alan Greenspan is completely different from Arthur Burns. But the most careful statistical assessments of this idea are at best inconclusive, and for the most part suggest on the contrary that changes in the systematic component of policy in this period are modest. Examples of work that comes to this conclusion, using widely different methodologies, are papers by Orphanides (2001), Leeper and Zha (2001), Hanson (2001), and me (Sims, 1999). My own paper

argues that the most important changes between periods can be accounted for as shifts in the variances of the structural disturbances. Time-varying variances are hard to distinguish from parameter variation. Attempts to show shifts in policy behavior should recognize this, in order to come into contact with the literature supporting the opposite view.

3. Time-Varying Descriptive Statistics

The paper implements a novel strategy to summarize the variation in the economy's characteristics over time. It uses descriptive statistics computed from simulated future time paths drawn from the posterior predictive density at each date, displaying how they change over time. The results are thought-provoking and deserve further study. I found particularly interesting the concentration of the posterior on the activism coefficient during the 1970s, followed by widening uncertainty thereafter. Even though the paper's interpretation of its activism coefficient may be dubious, this pattern of increased, then decreased, certainty about important components of inflation dynamics is suggestive. Phenomena like this might have played a role in the inertia of policy at the time and in the subsequent popularity of Monday-morning quarterbacking about it.

The paper sticks entirely to forward-looking data summaries. For many purposes this is appropriate, but such *filtered*, as opposed to *smoothed*, estimates of the stochastic properties of the model contain a component of variation that is learning, rather than actual time variation in the behavior of the economy. Commonly graphs like, say, Figure 11 or 12 show quite different time paths when computed on the basis of smoothed estimates. The difference lets us distinguish between best ex post estimates of what was actually happening and best current estimates at the time of what was happening. It would be interesting to see the work extended in that direction.

4. The "Learning the NRH" Story

The paper's Figure 14 confirms a point that Albert Ando has made for a long time: It is hard to blame the inflation of the 1970s on econometric modelers serving up a long-run inflation trade-off. It is an important result of both Chung's thesis—which this paper cites—and Sargent's book that the story that naive econometric Phillips-curve estimation led to the inflation of the 1970s cannot be sustained.

This paper proposes a new, incompletely articulated theory. It seems to me more a narrative theory than a time-invariant one that could be tested. The theory used in Chung's thesis, in Sargent's book, and in my

(1988) paper specifies both the (incorrect) model the policymakers use and the correct (natural-rate) model relating unemployment and inflation. It works out the consequences of these assumptions. My paper and Chung's thesis show that such a setup can easily lead to very long (at least millennia), possibly permanent periods of near-Ramsey behavior, with interest rates and inflation low on average. Sargent's book and Chung's thesis show that this setup does poorly at explaining U.S. postwar inflation and unemployment data, because it implies that policy authorities quickly realized the Phillips curve is nearly vertical.

It is hard to understand why the paper gives such a prominent role to the t-test for the hypothesis $\beta_1(1) = 1$. Figure 14 shows that the test strongly rejected the null starting in 1973. Not until more than 6 years later, in late 1979, did the "Volcker regime" begin. If the t-test showing neutrality was crucial to producing the Volcker policies, the connection was certainly not a simple one. It seems likely that the connection of this t-test to future changes in policy will be at least as tenuous.

My own view, which agrees in many respects with that of Orphanides (2001), is that unemployment rose and inflation rose because of real disturbances that lowered growth. Faced with the simultaneous rise in these two variables, and believing that unemployment affected inflation with a lag, policymakers had to decide whether the rise in unemployment that had already occurred was enough to exert adequate deflationary pressure. Since such "stagflation" had not occurred before on such a scale, they faced a difficult inference problem, which it took them some years to unravel. Note that in this story it is not $\beta_1(1)$ that is crucial, but the relation between β_0 and $\beta_2(1)$, i.e. the Phillips-curve "natural rate." I think it likely that careful statistical work using the Phillips curve would have demonstrated much earlier than 1979 that the current levels of unemployment were not exerting much downward pressure on inflation. But policy models at the time were estimating "gap" variables by focusing entirely on real factors—production functions and trend rates of growth. Policymakers realized their mistake only slowly because of excessive reliance on a theory that claimed the "gap" was a function of the level of output and the current level of technology. If they had paid more attention to a wider range of data, they would have seen their mistake earlier.

The notion that monetary policy acts on the price level by first affecting unemployment, or a "gap," which then via a Phillips curve affects inflation, is in my view mistaken. But if it had been the basis of a flexibly parameterized dynamic econometric model analyzing inflation, interest rates, and real growth jointly, it probably would not have led to such an acceleration of inflation as actually occurred.

5. Priors

The paper uses a prior that makes no attempt to push the parameter estimates toward the unit-root boundary, centers the prior at an OLS estimate (which will tend to be more stationary than the truth when the truth is near the unit-root boundary), and truncates the parameter space to rule out even mildly unstable roots. This is in the name of being "less informative" than, e.g., Doan, Litterman, and Sims. It is always true that there is no unique way to produce an "uninformative" prior, and this is especially true in VARs. A prior like that proposed here, in a model that conditions on initial observations, implies a lot of weight on stationary models, which in turn generally imply that a great deal of sample history is explained by large initial transients. How this happens is elaborated in some earlier work of mine (Sims, 2000). Such a prior is not uninformative, and may easily lead to strange results.

In the latter part of the paper simulations are used to give us an idea of how long it is likely to be before t-tests of $\beta_1(1) = 1$ are likely again to accept the null hypothesis. But the prior's concentration on stable models, and the time-variation model's insistence on making the model bounce away from the nonstationary boundary, could be strongly influencing the results of these simulations.

6. Conclusion

This paper breaks new ground in interpreting data with a structural VAR and time-varying parameters. Many of the methodological ideas in it are new and worth pursuing. Its choices of prior and identifying assumptions, however, are deviations from standard practice in the structural VAR literature that should not, in my view, be imitated. These aspects of the modeling and interpretation are crucial enough to the paper's substantive conclusions that those conclusions remain doubtful.

REFERENCES

Hanson, M. (2001). Varying monetary policy regimes: A vector autoregressive investigation. Wesleyan University. Discussion Paper.
Leeper, E., and T. Zha. (2001). Modest policy interventions. Indiana University and Federal Reserve Bank of Atlanta. Discussion Paper. http://php.indiana.edu/~eleeper/Papers/lz0101Rev.pdf.
Orphanides, A. (2001). Monetary policy rules, macroeconomic stability, and inflation: A view from the trenches. Board of Governors of the Federal Reserve System. Discussion Paper.
Sims, C. A. (1988). Projecting policy effects with statistical models. Revista de Analysis Economico, pp. 3–20. www.princeton.edu/~sims.

——. (1999). Drift and breaks in monetary policy, Princeton University. Discussion Paper. http://www.princeton.edu/~sims/. Presented at plenary session of the July 1999 meetings of the Econometric Society, Australasian region.
——. (2000). Using a likelihood perspective to sharpen econometric discourse: Three examples. *Journal of Econometrics* 95(2):443–462. http://www.princeton.edu/~sims/.

Comment

JAMES H. STOCK
Kennedy School of Government, Harvard University; and NBER

1. Introduction

Cogley and Sargent have provided a provocative and innovative contribution on an important problem, understanding the history of inflation in the United States and the evolving role of monetary policy in that history. They make many points in their rich paper, some empirical and some methodological.

In this discussion, I focus on four of their most salient empirical findings:

1. The persistence of the postwar inflation process has evolved over the past four decades. In the 1960s, inflation was mean-reverting; in the 1970s and early 1980s, it was highly persistent; and in the past ten to fifteen years it has been mean-reverting, as it was in the 1960s. This view is widely shared—for example, it has also been made by Taylor (1999) and by Brainard and Perry (2000)—and it seems to reflect conventional wisdom across a wide spectrum of views of monetary policy.
2. There is a positive correlation between the level of inflation, as measured by its low-frequency component, and its persistence. This is essentially an implication of the first point, because inflation was low in the 1960s, high in the 1970s and early 1980s, and low again during the 1990s.
3. The inflation process has been unstable, not just as measured by its persistence, but also over its entire spectrum or, equivalently, all its autocorrelations.
4. The reduced-form backward-looking Phillips curve relating inflation to lagged inflation and a measure of real economic activity (in Cogley and Sargent, the unemployment rate) has been unstable over the past four decades.

Cogley and Sargent draw several conclusions from these and related empirical findings. The most immediately relevant for policy bears on Taylor's (1999) warning that the decline in the persistence of inflation might induce revisionism by policymakers, who might return to the belief that there is an exploitable long-run trade-off between unemployment and inflation. The meat of Taylor's warning is that this revisionism—perhaps a better term is recidivism—would lead to the same mistakes and the same bad outcomes that it did in the 1960s and early 1970s. In this, Cogley and Sargent's message is the same as in Sargent's (1999) monograph on the history of U.S. inflation as elaborated on by Cho, Williams, and Sargent (2001).

Most of this discussion is devoted to presenting various pieces of empirical evidence that suggest that the foregoing four empirical findings are less clear-cut than Cogley and Sargent make them out to be. Specifically, I shall present evidence, based on hypothesis tests, confidence intervals, and median-unbiased estimates, that:

1. Inflation persistence has been roughly constant, and high, over the past 40 years in the United States.
2. Therefore, there is no correlation between the level of inflation and its persistence.
3. The autocorrelations of inflation are stable—at least, one cannot reject this hypothesis.
4. The reduced-form Phillips curve is stable, once one allows for a time-varying NAIRU, or if one interprets it not just as a relation between the unemployment rate and the rate of inflation, but more broadly as a relation between real economic activity and inflation.

These conclusions are quite at odds with Cogley and Sargent's, and this raises an interesting econometric question as to why my evidence is so different than theirs. The answer, not surprisingly, lies in differences between Cogley and Sargent's Bayesian methods and my frequentist methods.

2. Evaluating Cogley and Sargent's Empirical Results

Cogley and Sargent use a sophisticated nonlinear multivariate procedure to characterize inflation dynamics. The methods used here are simpler and univariate, but get at the same issues. The inflation data I consider are for the GDP deflator, quarterly from 1959:I to 2000:IV, although the results are robust to using other inflation measures.

2.1 PERSISTENCE OF INFLATION

There are a variety of ways to measure persistence, none perfect. The measure I consider is the largest root of an autoregressive representation of inflation. Cogley and Sargent's emphasis is on measurement, not testing, so to make this analysis parallel I consider median-unbiased estimates of the largest autoregressive root of inflation, constructed by inverting the augmented Dickey–Fuller statistic using the procedure developed in Stock (1991). This procedure produces confidence intervals for the largest root as well.

Recursive median-unbiased estimates of the largest AR root and 90% confidence intervals for this root are plotted in Figure 1 [these estimates are based on AR(4) models estimated recursively using all the data from 1959:I through the date indicated on the horizontal axis]. The striking feature of this plot is the stability of the estimates. Because the number of observations increases with the terminal date, the confidence intervals are tighter towards the end of the sample than at the beginning. At all dates since 1976, these intervals include one (the 90% confidence interval is briefly *above* one in 1975), and the recursive median-unbiased estimate is typically just less than one.

The recursive estimates in Figure 1 use all the historical data through the terminal date, and this might miss changes in persistence towards the end of the sample. Figure 2 therefore plots rolling median-unbiased estimates of the largest AR root and the associated 90% confidence interval for AR(4) models estimated using 12 years of data terminating at the date on the horizontal axis. The median-unbiased point estimates and confidence intervals evidently are quite noisy—not surprisingly, because each estimate is based on just 48 observations, quite few for performing inference about large autoregressive roots. Still, the evidence is striking (and is robust to changing the inflation series, the window length, and the number of lags). With one brief exception for the samples ending near 1994, the 90% confidence intervals contain a unit root, and the median-unbiased estimate, while variable, exceeds one almost as often as it is less than one. Notably, the median-unbiased estimate exceeds one early in the sample, for 12-year periods ending in 1972 through 1976, and late in the sample, for 12-year periods ending in 1997 through 2000.

2.2 RELATION BETWEEN PERSISTENCE AND THE LEVEL OF INFLATION

The results in Figure 2 suggest that there will be no particular relation between the level of inflation and its persistence as measured by the rolling median-unbiased AR root, because this root is estimated to be

Figure 1 RECURSIVE MEDIAN-UNBIASED ESTIMATE AND 90% CONFIDENCE INTERVAL FOR LARGEST AR ROOT

Figure 2 ROLLING MEDIAN-UNBIASED ESTIMATE AND 90% CONFIDENCE INTERVAL FOR LARGEST AR ROOT

essentially one throughout this sample. This is in fact the case; the correlation between the running mean of inflation and the rolling estimate of the largest AR root in Figure 2 over the same 12 years is −0.035.

2.3 INSTABILITY OF INFLATION AT HIGHER FREQUENCIES

Cogley and Sargent examine instability of inflation dynamics, both short- and long-run, via spectral estimates implied by their time-varying VAR. Here, I consider a more tightly parametrized approach and ask whether there appears to have been a break in the parameters of a univariate AR(5) model of the inflation rate. This is readily examined using the Quandt likelihood-ratio (or "sup-Wald") test for parameter stability. Although this test is designed around a single break, it is powerful against slow parameter evolution and multiple breaks as well. A technical issue is that the critical values need to hold when the largest root is one or nearly so; I handle this by using the critical values appropriate if the largest root is in fact one, taken from Banerjee, Lumsdaine, and Stock (1992), rather than the critical values appropriate when the largest root is well less than one. The test, implemented with conventional 15% trimming, fails to reject the hypothesis of parameter stability at the 10% significance level. However, using CPI inflation and different lag specifications can yield a significant break at the 10%, but not 5%, level, with the estimated break date in 1981. This evidence suggests that, on the whole, the inflation process has been stable, although there might have been some changes in its short-run dynamics between the first and the second half of the sample.

2.4 INSTABILITY OF THE PHILLIPS CURVE

Whether the backward-looking Phillips curve, interpreted as the relation between inflation, its lags, and current and past values of the unemployment rate, is unstable has attracted much attention. The evidence I provide here is borrowed from Staiger, Stock, and Watson (2001), who investigate the stability of the backward-looking Phillips relation of the type investigated by Gordon (1997, 1998).

A subdebate in this area has been whether the natural rate of unemployment should be estimated as the low-frequency component of the unemployment rate [the approach advocated by Hall (1999) and adopted by Cogley and Sargent] or whether it should be estimated off an estimated drift in the intercept of an empirical Phillips curve [the approach adopted by King, Stock, and Watson (1995), Gordon (1997, 1998), Staiger, Stock, and Watson (1997), and others].

Staiger, Stock, and Watson (2001) adopt Hall's and Cogley and Sargent's approach and estimate the natural rate by applying a low-

pass filter to the unemployment rate. Because the natural rate is esti-
mated using only the univariate unemployment rate, it is possible to
test separately for drift in intercept of the Phillips curve and for drift in
the slope coefficient; the NAIRU is the sum of the estimated natural
rate and the rescaled estimated intercept drift. Thus the NAIRU and
the natural rate are separately identified. Their conclusion is that in fact
these two series are very close to each other empirically, typically
within a few tenths of a percentage point of unemployment. The hy-
pothesis that there is no intercept drift in the Phillips curve, specified
as the deviation of the unemployment rate from its univariate long-run
trend, cannot be rejected at the 10% significance level. In practice,
then, there appears to be little difference between estimates of the
natural rate based on the Hall's and Cogley and Sargent's idea of the
long-run trend in the unemployment rate and the alternative approach
of estimating the time-varying NAIRU from intercept drift in the Phil-
lips curve.

Staiger, Stock, and Watson (2001) also test for drift in the slope of the
Phillips curve and cannot reject the null that the slope is stable.

Another way to see whether the Phillips curve has been stable is to see
how it has performed for forecasting. Interpreted broadly, the Phillips
relation links changes in the rate of inflation to economic activity, of
which the unemployment rate is but one measure. In their comparisons
of models for forecasting inflation, Stock and Watson (1999, 2001) con-
sider several versions of the backward-looking Phillips curve, each
based on different activity measures. They conclude that several activity
measures have produced reliable and useful inflation forecasts, at least
as measured by pseudo-out-of-sample forecast comparisons with bench-
mark autoregressive models. These include a composite index of real
economic activity constructed using a large number of income and out-
put measures, as well as simpler single measures such as the rate of
capacity utilization. Based on these broader measures of output, the
backward-looking Phillips curve has been a reasonably reliable and sta-
ble predictive relation over the past three decades.

3. Why Do the Bayesian and Frequentist Results Differ?

These conclusions are quite different than Cogley and Sargent's, and the
obvious question is, why? There are many differences between my meth-
ods and theirs: theirs are Bayesian and multivariate, mine are frequentist
and mainly univariate. I believe, however, that there are two main
sources of these differences: their prior leads them away from finding
persistence, and their specification, by forcing all the time variation to

occur through the dynamics rather than through the innovation variances, confuses changes in persistence with changes in volatility.

These views are informed by the recent study by Pivetta and Reis (2001), who compare the frequentist analysis of inflation persistence of the previous section, Cogley and Sargent's Bayesian method, and a more conventional time-varying parameter model of the type used by Brainard and Perry (2000). Although their analysis remain preliminary at the time of writing this comment, Pivetta and Reis' (2001) results suggest that Cogley and Sargent's importance sampling plays an important role in biasing (from a frequentist perspective) their estimates away from a unit root. This forces their posterior to have a low mean persistence, even if the true persistence (from a frequentist perspective) is quite large. The problem that Cogley and Sargent confront is a difficult one, and even among Bayesian econometricians there appears to be no consensus about the best way to place a prior on large autoregressive roots (see the special issue of *Econometric Theory* in 1994 on Bayesian approaches to unit-root inference and in particular the survey article by Uhlig, 1994).

The problem of confounding persistence and volatility is especially important, and Cogley and Sargent recognize this issue. Their persistence measures are based on the spectrum at frequency zero, but this can change either because the persistence has changed or because the entire spectrum has shifted, that is, the volatility of the process has changed. One does not need fancy tests to see that the volatility of the inflation process has changed greatly over the postwar period: the 1960s and 1990s were times of quiescent low inflation, the 1970s and early 1980s, of volatile high inflation. Because the integral of the spectrum is the variance, on using the height of the spectrum as a measure of persistence, quiescence becomes low persistence, and volatility becomes high persistence.

4. Implications and Conclusions

The evidence in Figures 1 and 2 suggests that inflation has been highly persistent for the past three decades, and stably so. My interpretation of the widespread view—that of Brainard, Perry, Taylor, Cogley, and Sargent—is that this confuses volatility with persistence. Inflation was low and stable in the 1960s and 1990s, but this does not mean that it was low and mean-reverting.

Whether or not the persistence of inflation has evolved, one implication of this discussion is that we need additional investigations of the statistical properties of Cogley and Sargent's method before adopting it for widespread use as a tool for data description.

Finally, let me turn to Taylor's warning, for here I agree with Cogley and Sargent. The fact is that many monetary economists believe inflation to have become less persistent, and this view must be reckoned with. To the extent that this view is held (correctly or not) by policymakers or advisors and to the extent that it encourages a revisionist perspective on the natural rate, then it does raise concerns about inadvertently repeating the inflationary mistakes of the past.

ADDITIONAL REFERENCES

Banerjee, A., R. L. Lumsdaine, and J. H. Stock. (1992). Recursive and sequential tests of the unit root and trend break hypotheses: Theory and international evidence. *Journal of Business and Economic Statistics*, 10, 271–288.

Brainard, W. C., and G. L. Perry. (2000). Making policy in a changing world. In *Economic Events, Ideas, and Policies: The 1960s and After*, G. L. Perry and J. Tobin (eds.). Brookings Institution Press.

Gordon, R. J. (1997). The time-varying NAIRU and its implications for economic policy. *Journal of Economic Perspectives* 11(1):11–32.

———. (1998). Foundations of the Goldilocks economy: Supply shocks and the time-varying NAIRU. *Brookings Papers on Economic Activity* 1998(2):297–333.

King, R. G., J. H. Stock, and M. W. Watson. (1995). "Temporal Instability of the Unemployment-Inflation Relation," Federal Reserve Bank of Chicago, *Economic Perspectives* May/June, 2–12.

Pivetta, F. and R. Reis. (2001). "The Persistence of Inflation in the U.S.," manuscript, Dept. of Economics, Harvard University.

Staiger, D., J. H. Stock, and M. W. Watson. (1997). The NAIRU, unemployment, and monetary policy. *Journal of Economic Perspectives* 11(Winter):33–51.

———, ———, and ———. (2001). Prices, wages and the U.S. NAIRU in the 1990s. Cambridge, MA: National Bureau of Economic Research. NBER Working Paper 8320.

Stock, J. H. (1991). Confidence intervals for the largest autoregressive root in U.S. economic time series. *Journal of Monetary Economics* 28:435–460.

———, and M. W. Watson. (1999). Forecasting inflation. *Journal of Monetary Economics* 44:293–335.

———, and ———. (2001). Forecasting output and inflation: The role of asset prices. Cambridge, MA: National Bureau of Economic Research. NBER Working Paper 8180.

Uhlig, H. (1994). What macroeconomists should know about unit roots: A Bayesian perspective. *Econometric Theory* 10:645–671.

Discussion

Tom Sargent responded to the discussants by saying that his view of events differed fundamentally from theirs. While they believed in conditional heteroscedasticity of shocks, he believed in changing decision rules. He explained that the authors were inspired by a graph of inflation over three centuries, which showed a clear break around 1970.

Rick Mishkin was sympathetic to the suggestion that what happened in the 1970s was that the Federal Reserve thought that the natural rate of unemployment was lower than it actually was. He was not so worried that the Solow–Tobin test would cause problems in the future, as advances since the 1970s in the understanding of the natural-rate hypothesis and in time-series econometrics are unlikely to go away. He also noted that there had been a substantial restructuring of monetary institutions since the 1970s, including increased central-bank independence and an increased emphasis on price stability. He was most worried about recidivism occurring because of policymakers underestimating the natural rate of unemployment, noting the wide confidence intervals on Jim Stock's estimates of the natural rate. Mishkin suggested that inflation targeting was the way to avoid a repeat of the 1970s.

Ken Rogoff noted that the view that Japan was stuck in a liquidity trap was a very powerful one in the policy literature. As a result, many policy economists indeed believe that output growth may be harmed if the rate of inflation wanders too close to zero. He also remarked that in countries other than the United States, there had obviously been a lot of institutional change since the 1970s, so it was hard to see how monetary policy could have remained stable.

Mark Gertler remarked that he and Richard Clarida had constructed measures of core inflation for Germany similar in spirit to those of Cogley and Sargent, using long-horizon forecasts to get core inflation. The striking difference between the United States and Germany was that although Germany suffered the same shocks as the United States, and policymakers had the same reasons to be confused, core inflation was flat and stationary in Germany. This finding suggested that there was something different about U.S. monetary policy in the 1970s. Gertler raised the possibility that the shift to nonborrowed reserves in 1979–1982 could have allowed shocks to have a greater impact, although the policy shift could also have been cover for an attempt to raise interest rates.

Chris Sims explained that the mere fact that he believed monetary policy was stable did not mean that he believed it was optimal. On recidivism, he believed that there were dangers in the inertia of orthodoxy.

Sargent agreed with Sims on the problems of identification in VARs. He said the problem was more profound than just partitioning contemporaneous correlations, as agents could have more information in their histories than was revealed by the histories of variables in the VAR. This fact generates time aggregation problems.