Persistence, Excess Volatility, and Volatility Clusters in Inflation

Michael T. Owyang

Three key features of the U.S. inflation time series have been identified by empirical studies. First, innovations in the level are persistent—that is, changes in the inflation rate generally endure. Second, volatility appears in clusters directly after changes in inflation. Following an innovation in the inflation rate, short periods of increased volatility are indicated by the presence of autoregressive conditional heteroskedasticity (ARCH) in the regression residuals. Third, periods with high mean inflation have a correspondingly high variance of inflation, and vice versa. Likewise, periods with low levels of inflation tend to be associated with low variability.

This paper presents a model in a single, integrated framework that offers one possible explanation for these facts about the U.S. inflation time series. In this model, the policymaker faces a trade-off between inflation and unemployment in the form of an expectations-augmented Phillips curve.¹ The Phillips curve is subject to two shocks: a persistent shock that follows a Markov process and a white noise shock. The magnitudes of both shocks are unobservable, forcing the policymaker to employ an ordinary least-squares (OLS) learning technology to determine the policy target. Agents have the same information set as the policymaker and form rational expectations of monetary policy. In addition to the Phillips curve shocks, an independent Markov process governs the policymaker's preferences; agents form their expectations after observing current policymaker preferences.

Learning and the interaction between the

Markov processes governing the position of the Phillips curve and the policymaker's preferences provide one possible explanation for the three stylized facts about the U.S. inflation time series. Changes in the variable that determines the position of the Phillips curve (henceforth called the *structural* variable) are persistent and directly determine the policymaker's target. Thus, regime shifts in the structural variable induce persistent shifts in inflation—the first of three stylized facts. The second, volatility clustering, is driven by the policymaker's learning mechanism. Once the economic fundamentals change, the policymaker resets the learning algorithm to determine the magnitude of the new shock. As the policymaker learns, new information each period does not lead to as large of an update of his estimate of the position of the Phillips curve. Thus, periods of volatility follow the shock, then drop off.

The third fact, the relationship between mean and variance, is a result of the process governing the policymaker's preferences. In addition to persistence effects, uncertainty in the policymaker's estimate of the position of the Phillips curve produces variability in the policy target. When the policymaker is accommodative, i.e., inflation is high, this uncertainty is amplified. However, when the policymaker is in an inflation-intolerant regime, uncertainty about fluctuations in the Phillips curve are not amplified, as the policymaker is less willing to trade off inflation for small gains in unemployment.

The model is estimated using Gibbs sampling. The estimation procedure will generate both a parameter vector for the model and posterior densities for each Markov process. Monte Carlo simulations using the estimated parameters reveal that three-state versions of both Markov processes (governing the Phillips curve and the policymaker's preferences) produce artificial data that exhibit the three characteristics described above.

FEATURES OF THE U.S. DATA

Figure 1 shows the U.S. monthly annualized inflation rate for the period 1947:01–1998:05.² The inflation data over this period exhibit three particular time-series characteristics:

Michael T. Owyang is an economist at the Federal Reserve Bank of St. Louis. A previous version of this paper was written at the University of California, San Diego, and appears in Owyang (2000). The author thanks Jim Hamilton, Garey Ramey, and Valerie Ramey for their input, as well as Takeo Hoshi, Neal Beck, Thomas Sargent, Robert Shimer, Leora Freidberg, Kevin Lansing, Glenn Rudebusch, and seminar participants at the Federal Reserve Banks of Dallas, St. Louis, and San Francisco, the Board of Governors of the Federal Reserve System, and Emory University for useful comments. The author also thanks Chang-jin Kim and Charles Nelson for assistance with the program. Abigail J. Chiodo provided research assistance.

¹ This literature stems from the seminal work of Kydland and Prescott (1977) and Barro and Gordon (1983). Models possessing policymaker preferences of similar form include Barro (1986), Cukierman and Meltzer (1986), and Ball (1992, 1995).

² The data are seasonally adjusted CPI data taken from Citibase.

Figure 1

U.S. Inflation

1. Changes in inflation are persistent.
2. Inflation series have volatility clusters and can be modeled as some form of ARCH series.
3. High/low levels of inflation are associated with relatively high/low variance and uncertainty.

The first of these phenomena has been considered in the empirical literature by Barsky (1987) and Fuhrer and Moore (1995). They find that innovations in the rate of inflation are largely permanent, causing persistent shifts in trend; in the presence of an aggregate shock, the inflation rate rises and stays high for an extended period. Consider the subperiod 1968-73, for example; here, an innovation in the inflation rate is associated with a largely permanent shift in trend.³

Table 1 shows the results from an analysis of the autocorrelations for monthly data over the sample period 1947:01-1998:05. Autocorrelations for the first eight lags range between approximately 0.4 and 0.6, providing strong evidence of the presence of serial correlation. Results indicate that inflation is largely persistent with lagged coefficients that are clearly significant, indicating that innovations that occur in any period spill over into subsequent periods.

Following Engle's (1982) original analysis of the U.K. inflation data, a number of papers have attempted to fit the U.S. inflation data to an ARCH model to test for the presence of volatility clusters. Kim (1993) tests ARCH against alternate specifications and finds that ARCH does not perform as well as an unobserved-component time series with Markov-switching heteroskedasticity. However, Baillie, Chung, and Tieslau (1996) employ an alternate specification to model the inflation process, using an autoregressive, fractionally integrated, moving-average version of a generalized autoregressive conditional heteroskedasticity model (ARFIMA-GARCH). They find evidence of persistence and mean reversion, as well as heteroskedasticity, in the inflation time series.

In light of this evidence, consider a GARCH(1,1) model for the variance of the form

$$\sigma_t^2 = \kappa + \chi e_{t-1}^2 + \delta \sigma_{t-1}^2,$$

in which the current conditional variance depends

³ This period coincides with the beginning of the two-tiered system for gold coverage in 1968, the closing of the gold window in 1971, and the end of the adjustable peg in 1973.

Table 1**Autocorrelations and AR(4) Regression for U.S. Time Series Inflation 1947:01–1998:05**

Autocorrelations			
Lag 1	0.611	Lag 5	0.473
Lag 2	0.575	Lag 6	0.465
Lag 3	0.501	Lag 7	0.463
Lag 4	0.466	Lag 8	0.475
Variable	Coefficient	Standard error	t Statistic
ΔCPI_{t-1}	0.392983	0.044755	8.780743
ΔCPI_{t-2}	0.217638	0.053232	8.780741
ΔCPI_{t-3}	0.153694	0.048365	3.177771
ΔCPI_{t-4}	0.180432	0.044039	4.097130
C	0.334763	0.114108	2.933746
ARCH(1)	0.141611	0.017654	8.021489
GARCH(1)	0.830818	0.020444	40.63829
Summary statistics			
R ²	0.382133	Mean inflation	4.099219
Adjusted R ²	0.376006	SD inflation	4.511847
SE of regression	3.564057	Akaike info criterion	2.553171
LM test			
F Statistic	62.33962	Probability	0.00000
Obs R ²	56.66900	Probability	0.00000
Variable	Coefficient	Standard error	t Statistic
C	6.959891	1.016461	6.847182
RESID^2(−1)	0.302882	0.038361	7.895545

on the lagged conditional variance and the lagged squared residual. A test for ARCH by employing a Lagrange multiplier (LM) test indicates that the null hypothesis of ARCH cannot be rejected. Table 1 contains the results from the LM test of an AR(4) model and the variance results of the GARCH(1,1) regression of the U.S. inflation time series. Both the ARCH and GARCH components of the variance equations are significant.

Ball and Cecchetti (1990) cite a relationship between the level of inflation and its variance, noting that an increase in the level of inflation is not only persistent but often associated with a corresponding increase in the variance and/or uncertainty of future inflation.⁴ Explanations for

this phenomenon have focused on three primary areas: changes in the expectations-augmented Phillips curve, temporary and permanent aggregate shocks, and idiosyncratic policy.

Using the U.S. quarterly gross national product (GNP) deflator and consumer price index (CPI) data divided into subperiods of various lengths, Ball and Cecchetti (1990) test the hypothesis that the level and variability of inflation are related. Their result is that the correlation between level and variance rises for the first few years and then begins to fall.

⁴ Logue and Willet (1976), Cosimano and Jansen (1988), Devereux (1989), and Evans (1991) also test this relationship. Evans pays particular attention to estimating inflationary uncertainty. Ball (1992) provides a theoretical model that attempts to explain this correlation.

Figure 2

Correlation Between Mean and Variance for Five-Year Subperiods

Ball and Cecchetti decompose inflation into a series of permanent shocks and temporary shocks—a trend stationary component and a white noise component. They show that permanent shocks have increasing variance with level and thus cause a rise in uncertainty when trend rises.

Consider the mean and uncertainty of inflation over five-year, non-overlapping subperiods and the correlation between sample means and standard deviations across these subperiods. The results are plotted in Figure 2. Examination of the results reveals a relationship between mean and standard deviation, with a correlation of about 0.23. Figure 3 plots the inflation means and standard deviations for each subsample. The magnitude of these increases in mean and variance varies with the sample period, but this indicates a relationship that might not be completely revealed using a simple correlation test. Note that only for the period 1975-79 does the variance fall when the subsample mean rises.

Model

To formulate an integrated model that provides one possible explanation for the previously mentioned time-series characteristics of U.S. inflation data, I propose a reduced-form macroeconomic model with the following features:

- a neoclassical rational expectations–augmented Phillips curve subject to a Markov shock to the natural rate of unemployment,

- a monetary policymaker with Markov-switching preferences for low inflation relative to low unemployment, and
- policymaker learning.

The policymaker faces a short-run tradeoff between inflation and unemployment embodied in the Phillips curve.⁵ Mankiw (2000) argues for the inclusion of both expectations and supply shocks to provide a complete and stable view of the economy. This model incorporates these features but assumes that the magnitude of these supply shocks is unobservable and must be learned.⁶ This model focuses on the actions of a policymaker under uncertainty and subject to shifts in preferences. During each period, events occur in the following order:

- The policymaker's preferences are determined and revealed to the public.
- Agents and the policymaker simultaneously set expectations and the policy target, respectively.
- The economic shocks then occur, and the policymaker and agents observe the realized inflation and unemployment rates.
- The policymaker's and agents' information sets are updated.

Consider an economy in which a monetary policymaker sets an inflation target and private agents form expectations about the realized inflation rate.⁷ Suppose that the policymaker faces an expectations-augmented Phillips curve of the form

$$(1) \quad u_t = k(\pi_t^e - \pi_t) + \eta_t + \varepsilon_t,$$

where u_t is the unemployment rate; π_t is the inflation rate; π_t^e are the period t inflation expectations; η_t is a persistent unemployment shock that follows

⁵ In order to abstract from interest rate properties, this paper does not consider directly the policymaker's instrument. In addition, I do not include smoothing as a policymaker objective. If interest rates were added as the policymaker's instrument, the paper could be nested as a special case of the model found, for example, in Clarida, Gali, and Gertler (1999). The backward-looking nature of this model, driven by the learning process in agent expectations, is similar to that proposed originally by Taylor (1981) and later by Fuhrer and Moore (1995) and Rudebusch and Svenson (1998).

⁶ A number of academic papers consider the effect of the Fed learning about the world. Sargent (1999) considers a model in which uncertainty generates paths between economic equilibria. Kasa (1999), Lansang (2001), Sack (1998), and Wieland (1998, 2000) also consider models in which the policymaker faces some degree of parameter or model uncertainty.

⁷ Owyang and Ramey (2001) use a similar model with adaptive expectations to measure monetary policy.

Figure 3**Mean and Variance of Five-Year Periods Over Time**

a three-state Markov-switching process with a vector of possible structural states $\vec{h} = (h_1, h_2, h_3)$, a transition kernel T_η , and period t state indicator vector S_t ; and $\varepsilon_t \sim N(0, \sigma_\varepsilon^2)$ is a white noise shock that occurs each period.

Shifts in η_t represent dramatic, persistent changes in the current economic environment. These shifts can be viewed as unexpected but highly visible events, such as wars or oil crises, in which the timing of the event is clearly observed.⁸ However, the new value of η_t that arises following a switch cannot be observed, but rather must be learned by the policymaker and agents.⁹ Let $\hat{\eta}_t = E_t[\eta_t | \Omega_t]$ denote the policymaker's period t estimate of η_t conditional on his information set $\Omega_t = (u_1, u_2, \dots, u_t, \pi_1, \pi_2, \dots, \pi_t)$.

The policymaker has preferences over inflation and unemployment, $L_t = \alpha_t u_t^2 + \pi_t^2$, with the relative weight on unemployment, α_t , governed by a Markov process with transition kernel T^α and possible states $\vec{\alpha} = (\alpha_1, \alpha_2, \alpha_3)$. A high α_t policymaker attaches more weight to output and sets a higher inflation target, conditional on the estimated state of the world, $\hat{\eta}_t$. The low α_t policymaker is an inflation hawk and sets a lower conditional inflation target. It can be shown that a policymaker who minimizes the current period value of his quadratic loss will form a short-run, discretionary inflation target that can be given by

$$(2) \quad \bar{\pi}_t = k\alpha_t \frac{(k\pi_t^e + \hat{\eta}_t)}{1 + k^2\alpha_t},$$

where $\hat{\eta}_t$ is the policymaker's estimate of the magnitude of the unemployment shock.¹⁰ Although the policymaker can set an inflation target, the

realized inflation rate is subject to some normally distributed noise, v_t .¹¹

For exposition, assume that agents form expectations rationally and with the same information set, Ω_t , as the policymaker. Agent expectations are then

$$(3) \quad \pi_t^e = \bar{\pi}_t.$$

Combining equations (2) and (3) gives

$$(2') \quad \bar{\pi}_t = k\alpha_t \hat{\eta}_t.$$

Equation (2') can be interpreted in terms of current and historical policy. Assuming that $\alpha_t \neq 0$, the policymaker's short-run target moves as a function of his belief about the state of the economy. As the balance of risks shifts toward higher perceived

⁸ It is convenient but not necessary to assume that the policymaker observes the timing of the event. Similar results obtain if the event is unobserved. If there is a switch but the policymaker believes there is none, the policymaker's inflation target will induce a poor unemployment outcome. If this persists for a number of periods, the policymaker can conclude that a switch has occurred. On the other hand, if the policymaker believes a switch occurred when it in fact had not, the policymaker resets the gain sequence. New information does not change the target on average but does induce more volatility.

⁹ Agents in the model, including the policymaker, are not assumed to know the structure of the underlying process. They are assumed to believe that the economy can take on a continuum of possible states.

¹⁰ Note that this formulation assumes the policymaker is optimizing over the current period only. A more forward-looking policymaker complicates the learning via feedback from the economy and is left as an extension. This formulation, however, does not presuppose a lack of a consistent long-run inflation target.

¹¹ This can be interpreted as the policymaker setting an intermediate interest rate target and realizing an inflation target with some error. For this application, I suppress the interest rate target and consider only the effect on inflation.

unemployment (i.e., higher $\hat{\eta}_t$), the policymaker will allow the inflation rate to wander while focusing attention on achieving a higher growth rate (and, thus, lower unemployment) and vice versa.

The effect of the switching process α_t on the inflation target can be historically exemplified by changes in Federal Reserve chairman, although the model does not restrict it to be so. Clarida, Gali, and Gertler (2000), among others, recognize the fundamental change in Fed objectives at the onset of the Volcker regime in October 1979.

Finally, since the magnitude of the unemployment shock is unknown, the policymaker must infer the state of the economy from the data. Since the timing of the shock is known, the policymaker can employ an OLS learning technology of the form

$$(4) \quad \hat{\eta}_t = \frac{1}{g_t} u_{t-1} + \frac{g_t - 1}{g_t} \hat{\eta}_{t-1},$$

where g_t is the number of periods since the last switch in η_t . The sequence is $g_1, g_2, \dots, g_t$ referred to as the *gain sequence*. The policymaker resets the gain sequence when a switch occurs. This is synonymous with discarding previous (and now useless) information gathered before the shock and attaching more weight to incoming information. As the policymaker accumulates more information, new information becomes less valuable and the weight attached falls.

Inflation Dynamics

The model specified by equations (1) through (4) simplifies into two time-series equations governing unemployment and inflation: the Phillips curve (equation (1)) and

$$(5) \quad \pi_t = \frac{g_t - 1}{g_t} \bar{\pi}_{t-1} + \frac{k\alpha_t}{g_t} (\eta_{t-1} + \varepsilon_{t-1} - v_{t-1}) + v_t,$$

where v_t is normally distributed inflation noise that occurs after the policymaker sets the inflation target and can be interpreted as control error.¹²

Equation (5) embodies the aforementioned time-series characteristics of the inflation time series. First, the autoregressive nature of equation (5) and the Markov structure of the shocks indicate that shocks to the inflation rate in the form of preference shocks to α_t or structural shocks to η_t are persistent, i.e., shifts in the inflation rate are lasting. This can be verified if the diagonal elements of the transition kernel sum to greater than 1. Second, a Markov shock to the Phillips curve in the form of

a shift in η_t will induce periods of volatility while the policymaker learns. Third, an accommodative policymaker (i.e., high α_t) will tend to induce more volatility in the inflation rate than the relatively hawkish (i.e., low α_t) policymaker.¹³

The first property implies that the equilibrium inflation path exhibits persistence whenever $\alpha_t > 0$, as a direct result of persistence in the processes determining α_t and η_t and indirectly through the effect of the learning rule on the policy choices. The second property is demonstrated in the following thought experiment. Consider an aggregate shock at time $t = 0$ that shifts the Phillips curve out through an increase in the parameter η_t . The policymaker would react to the shock by raising his inflation target and raising the mean level of inflation, thus causing a persistent change in the inflation rate. The policymaker, however, does not know the exact value of the target and cannot infer the magnitude of the shock because of unemployment noise, ε_t , and inflation noise, v_t . He must therefore implement a policy based on estimates of the new Phillips curve parameters, which he constructs using the learning technology. In the short term, the policymaker's estimate—and thus the inflation rate—is greatly affected by new information in the form of new observations. This result stems from the policymaker resetting the gain sequence, g_t , in equation (4). The weight on new information is increased because information prior to the incidence of the new shock is no longer valuable. Over time, the policymaker collects inflation-unemployment data and updates his estimate of the shock magnitude, causing the accuracy of his prediction to increase. Also, the accumulation of information allows the policymaker to decrease the weight of new information, and inflation becomes less variable.¹⁴

Suppose now that the policymaker's preferences change, switching from an accommodating regime to an anti-inflationary regime. The policymaker attaches more emphasis to minimizing the level of inflation and subsequently reduces the inflation

¹² Equation (5) elucidates a time-series path for inflation. A similar equation could be written for unemployment. However, the time-series characteristics for unemployment are less well documented, and I leave discussion of them to another paper. It is sufficient to say that, in this formulation, the unemployment time series would follow a Markov process with noise.

¹³ A formal presentation of these results is given in Appendix A.

¹⁴ This paper assumes a particular learning specification. However, these results are robust to permutations of the policymaker's learning technology provided that the learning mechanism implemented converges (see Marcat and Sargent, 1989).

Table 2

Estimated Parameters

Parameter	Estimated value	Parameter	Estimated value
a_3	1.2130 (0.3681)	h_3	6.1817 (1.2968)
a_2	0.4856 (0.2468)	h_2	5.3110 (0.9644)
a_1	0.0929 (0.1516)	h_1	4.3160 (0.9168)
σ_v	2.7434 (1.1105)	σ_ε	2.4194 (1.0503)
$\Pr[\alpha_t = a_3 \alpha_{t-1} = a_3]$	0.9663 (0.0191)	$\Pr[\eta_t = h_3 \eta_{t-1} = h_3]$	0.9908 (0.0398)
$\Pr[\alpha_t = a_2 \alpha_{t-1} = a_2]$	0.9706 (0.0189)	$\Pr[\eta_t = h_2 \eta_{t-1} = h_2]$	0.7645 (0.2057)
$\Pr[\alpha_t = a_1 \alpha_{t-1} = a_1]$	0.9058 (0.0768)	$\Pr[\eta_t = h_1 \eta_{t-1} = h_1]$	0.6874 (0.2133)

NOTE: Standard deviations across iterations are given in parentheses.

target for a given estimate, $\hat{\eta}_t$. Additionally, because the policymaker's preferences enter multiplicatively instead of additively, the adjustments made during the learning process become smaller. Thus, the policymaker sets a lower target with less variability—the third property.

ESTIMATION

The preceding model of inflation and unemployment is estimated using Gibbs sampling. Gibbs sampling uses an iterative filtering algorithm and a Monte Carlo algorithm to generate the ergodic density for the parameter vector conditional on the data. An outline of the sampling procedure appears in Appendix B. Seasonally adjusted monthly data for both series are taken from Citibase for the sample period 1947:01–1998:05. The inflation rate is taken as the annualized rate of change of the CPI.¹⁵ The parameter estimates are given in Table 2.

Here, $\vec{\alpha} = (a_1, a_2, a_3)$ reflects the policymaker's preferences and $\vec{h} = (h_1, h_2, h_3)$ describes unemployment shocks to the Phillips curve. The combination $(a_i$ and $\Pr[\alpha_t = a_i | \alpha_{t-1} = a_j]$ for $i, j = 1, 2, 3)$ defines the Markov process that determines the policymaker's preferences. Recall that α directly affects the policymaker's inflation target. The high value for the preference parameter, a_3 , indicates the most accommodative policymaker, while the low value, a_1 , represents a policymaker less willing to trade high inflation for lower unemployment; a_1 defines a regime in which the policymaker sets a nearly zero inflation target that does not respond much to changes in the state of the economy. On the other hand, a_2 and a_3 correspond to policymakers who are increasingly responsive to shocks to the structural variable.

Suppose a shock hits the economy in the form of an increase in the natural rate η_t . The accommodative policymaker ($\alpha_t = a_3$) responds to the shock by increasing his inflation target. A 1 percent increase in the natural rate, η_t , when the policymaker is in the a_3 state implies a 1.3 percent increase in the inflation target. Similarly, a 1 percent increase in the natural rate when the policymaker is in the a_1 state implies a 0.5 percent increase in the inflation target.

The parameters $(\Pr[\alpha_t = a_i | \alpha_{t-1} = a_i])$ for $i = 1, 2, 3)$ are the on-diagonal transition probabilities that determine the duration of the policymaker's preference regime. They also determine the expected number of regime changes over the sample period. The number of regime shifts predicted by the estimated transition probabilities are similar to the predictions in Romer and Romer (1989).¹⁶

Now consider the three-state process for the structural variable η_t . The process $(h_i, \Pr[\eta_t = h_i | \eta_{t-1} = h_j])$ for $i, j = 1, 2, 3)$ represents three shocks to the natural rate of unemployment and their transition probabilities. One interpretation of this process is that the state $\eta_t = h_2$ represents the “normal” state of the economy, while the other two states are the product of exogenous shocks. The low value, h_1 , can be interpreted as the natural rate of unemployment in an unusually productive state of the economy,

¹⁵ Results using the monthly personal consumption expenditures (PCE) index are not appreciably different. Postwar quarterly GDP deflator data are of insufficient length to run this type of estimation.

¹⁶ Romer and Romer find seven instances during this sample when the Fed reacted to explicitly reduce inflation. For a more detailed comparison of the type of model presented here and the Romer dates, see Owyang and Ramey (2001).

Table 3

Three-State Monte Carlo Results

	Simulated*	U.S. time series†
AR(1)	0.3114 (0.0454)	0.3841 (0.0438)
AR(2)	0.2326 (0.0401)	0.2238 (0.0517)
AR(3)	0.1886 (0.0436)	0.1604 (0.0468)
AR(4)	0.1720 (0.0389)	0.1739 (0.0436)
Constant	3.2497 (2.883)	0.2554 (0.1053)
ARCH	0.0400 (0.0488)	0.1410 (0.0170)
GARCH	0.6753 (0.2754)	0.8406 (0.0195)
Five-year correlation‡	0.2453 (0.3660)	0.3155
Sample mean	3.89 (0.21)	3.99
Variance	4.30 (0.61)	4.40
Skewness	0.92 (0.15)	0.85
Kurtosis	2.70 (0.21)	5.02

NOTE: Simulated data statistics are taken from 1000 iterations of 608 period samples. The first column contains results from the regression of inflation on four lags. GARCH(1,1) parameters are taken from the following model of the variance $\sigma_t^2 = \kappa + \chi u_{t-1}^2 + \delta \sigma_{t-1}^2, \chi, \delta > 0$.

*Standard deviations across samples are given in parentheses.

†Standard errors for the AR(4) regression are given in parentheses.

‡Indicates the correlation between mean and variance of five-year intervals.

perhaps caused by a positive technology shock. The high value, h_3 , represents the innovation to the natural rate after an adverse event such as an oil shock.¹⁷

Simulations

To evaluate how well the proposed model fits the data, the estimated parameters from the previous section are used in Monte Carlo simulations. The goal is to determine whether the model replicates the moment and variability characteristics of the U.S. inflation time series. Artificial data series are created according to equation (5), and AR(4) regressions are performed for comparison with the U.S. time series.¹⁸ The residuals are then tested for GARCH(1,1). Results of the simulations and subsequent statistical tests are shown in Table 3.

The coefficients on the autoregressive terms and variability characteristics taken from the regres-

sions on the artificial data are statistically consistent with those obtained from the U.S. time series. However, both the ARCH and GARCH parameters are comparatively low and more variable. This can be explained by the manner in which the simulations are generated. The estimated variability characteristics show dependency on the average duration of the regime. Short-lived regimes in η_t will reduce the variability characteristics of the artificial data, as the policymaker is unable to carry out the learning process. The policymaker resets the gain at the onset of each switch; however, if regimes tend to be of short duration, the policymaker never has a chance to reduce the gain. Variability characteristics are driven by the switches in the regime rather than the learning process. Long regimes in the η_t variable can cause similar problems. When there are few switches, the policymaker's gain remains at low levels, reducing the amount of variability induced by learning. Analysis of individual simulations reveals that, in iterations in which the policymaker reverts to an anti-inflationary regime for unusually long periods, the GARCH coefficient is near zero. Essentially, the variability parameters are biased downward in both cases of extremely high numbers of switches and extremely low numbers of switches. Incorporating switching processes into Monte Carlo experiments of this kind makes these results unavoidable, as some samples are bound to be outliers.

The artificial data also exhibit the mean-variability relationship found in the U.S. time series. The correlation between mean and variance for five-year intervals is highly variable in the simulated data, however. This correlation is also caused by the incorporation of Markov processes into Monte Carlo experiments. Recall that the mean-variability relationship depends on switches in the preference process. Thus, the statistic will be biased downward if the changes in the inflation level are driven primarily by switches in the structural process. If no switches in the α_t process occur, the mean-variance correlation can actually be negative. A similar result is true if there is a large number of preference switches, which can also have a tendency to understate the relationship or cause a negative correlation.

¹⁷ An alternate explanation for shifts in the natural rate involves the change in demographics over time. Shimer (1999) attributes recent changes in the natural rate to the "baby-boomers" phenomenon.

¹⁸ Series lengths are 750 observations. The first 142 observations are discarded to avoid issues associated with initial conditions.

CONCLUSION

The above model combines a reduced-form model of monetary policy with a Marcat-Sargent (1989) policymaker OLS learning formulation to provide one possible explanation for the three stylized facts about U.S. inflation. Results indicate that a bivariate, three-state Markov-switching model can generate these characteristics of the U.S. inflation time series. The model is able to produce inflation persistence, volatility clusters, and a correlation between level and variance and the parameter estimates that are similar to those of the actual U.S. inflation time series.

These inflation characteristics are generated by the interaction among the unobserved shock to the Phillips curve, policymaker learning, and the switches in the policymaker's preferences. Persistent shocks to either preferences or the Phillips curve translate, through the policymaker's decision rule, into persistent changes in the inflation target. Additionally, more inflation-tolerant policymakers tend to allow shocks to the economy to have a greater effect on the inflation target, causing the inflation target to become more variable when it is relatively high. Finally, clusters of volatility are a reflection of the policymaker's learning process. The uncertainty generated by the onset of a new shock makes the target more variable; however, as the policymaker learns, the target settles down.

What might this model, if correct, imply for the future of the U.S. economy? Has the economy entered a new technological regime? Can the economy sustain a low level of unemployment? Is a slowdown inevitable? Forecasters move rapidly to change their opinions at the onset of new information. Fluctuations in economic indicators can quickly alter the tenor of expectations. The Fed scrambles to adjust rates; economists on Wall Street adjust expectations. The result: more volatility while we all figure out the real state of the economy.

REFERENCES

- Baillie, Richard; Chung, Ching-Fan and Tieslau, Margie. "Analysing Inflation by the Fractionally Integrated ARFIMA-GARCH Model." *Journal of Applied Econometrics*, January/February 1996, 11(1), pp. 23-40.
- Ball, Laurence. "Why Does High Inflation Raise Inflation Uncertainty?" *Journal of Monetary Economics*, June 1992, 29(3), pp. 371-88.
- _____. "Time-Consistent Policy and Persistent Changes in Inflation." *Journal of Monetary Economics*, November 1995, 36(2), pp. 329-50.
- _____. and Cecchetti, Stephen. "Inflation and Uncertainty at Short and Long Horizons." *Brookings Papers on Economic Activity*, 1990, (1), pp. 215-54.
- Barro, Robert. "Reputation in a Model of Monetary Policy with Incomplete Information." *Journal of Monetary Economics*, January 1986, 17(1), pp. 3-20.
- _____. and Gordon, David. "Rules, Discretion and Reputation in a Model of Monetary Policy." *Journal of Monetary Economics*, July 1983, 12(1), pp. 101-21.
- Barsky, Robert. "The Fischer Hypothesis and the Forecastability and Persistence of Inflation." *Journal of Monetary Economics*, January 1987, 19(1), pp. 3-24.
- Carter, Chris K. and Kohn, Robert. "On Gibbs Sampling for State Space Models." *Biometrika*, August 1994, 81(3), pp. 541-53.
- Casella, George and George, Edward. "Explaining the Gibbs Sampler." *The American Statistician*, August 1992, 46(3), pp. 167-74.
- Clarida, Richard; Galí, Jordi and Gertler, Mark. "The Science of Monetary Policy: A New Keynesian." *The Journal of Economic Literature*, December 1999, 37(4), pp. 1661-707.
- _____.; _____. and _____. "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory." *Quarterly Journal of Economics*, February 2000, 115(1), pp. 147-80.
- Cosimano, Thomas and Jansen, Dennis. "Estimates of the Variance of U.S. Inflation Based Upon the ARCH Model: A Comment." *Journal of Money, Credit, and Banking*, August 1988, 20(3), pp. 409-21.
- Cukierman, Alex and Meltzer, Allan. "A Theory of Ambiguity, Credibility, and Inflation Under Discretion and Asymmetric Information." *Econometrica*, September 1986, 54(5), pp. 1099-128.
- Devereux, Michael. "A Positive Theory of Inflation and Inflation Variance." *Economic Inquiry*, January 1989, 27(1), pp. 105-16.
- Engle, Robert. "Autoregressive Conditional Heteroskedasticity

REVIEW

- with Estimates of the Variance of United Kingdom Inflation." *Econometrica*, July 1982, 51(4), pp. 987-1008.
- _____. "Estimates of the Variance of U.S. Inflation Based Upon the ARCH Model." *Journal of Money, Credit, and Banking*, August 1983, 15(3), pp. 286-301.
- Evans, Martin. "Discovering the Link Between Inflation Rates and Inflation Uncertainty." *Journal of Money, Credit, and Banking*, May 1991, 23(2), pp. 169-84.
- Fuhrer, Jeffery and Moore, George. "Inflation Persistence." *Quarterly Journal of Economics*, February 1995, 110(1), pp. 127-59.
- Hamilton, James. "Analysis of Time Series Subject to Changes in Regime." *Journal of Econometrics*, July/August 1990, 45(1-2), pp. 39-70.
- Kasa, Kenneth. "Will the Fed Ever Learn?" *Journal of Macroeconomics*, Spring 1999, 21(2), pp. 279-92.
- Kim, Chang-jin. "Unobserved-Component Time Series Models with Markov-Switching Heteroscedasticity: Changes in Regime and the Link Between Inflation Rates and Inflation Uncertainty." *Journal of Business and Economic Statistics*, July 1993, 11(3), pp. 341-49.
- _____. and Nelson, Charles. *State-Space Models with Regime Switching: Classical and Gibbs-Sampling Approaches with Applications*. Cambridge, MA: MIT Press, 1999.
- Kydland, Finn E. and Prescott, Edward C. "Rules Rather Than Discretion: The Inconsistency of Optimal Plans." *Journal of Political Economy*, June 1977, 85(3), pp. 473-91.
- Lansang, Kevin. "Learning About a Shift in Trend Output: Implications for Monetary Policy and Inflation." Working Paper No. 00-16, Federal Reserve Bank of San Francisco, 2001.
- Logue, Dennis and Willet, Thomas. "A Note on the Relation Between the Rate and Variability of Inflation." *Economica*, May 1976, 43(17), pp. 151-58.
- Marcet, Albert and Sargent, Thomas. "Convergence of Least Squares Learning Mechanisms in Self Referential Linear Stochastic Models." *Journal of Economic Theory*, August 1989, 48(2), pp. 337-68.
- Mankiw, N. Gregory. "The Inexorable and Mysterious Tradeoff Between Inflation and Unemployment." Working Paper No. W7884, National Bureau of Economic Research, September 2000.
- Owyang, Michael. "Three Essays on the Behavior of Monetary Policy." Ph.D. Dissertation, University of California, San Diego, 2000.
- _____. and Ramey, Garey. "Regime Switching and Monetary Policy Measurement." Working Paper No. 2001-03, University of California, San Diego, 2001.
- Romer, Christina D. and Romer, David H. "Does Monetary Policy Matter? A New Test in the Spirit of Friedman and Schwartz," in *National Bureau of Economic Research Macroeconomics Annual: 1989*. Cambridge, MA: MIT Press, 1989, pp. 121-70.
- Rudebusch, Glenn D. and Svenson, Lars E.O. "Policy Rules for Inflation Targeting." Working Paper No. W6512, National Bureau of Economic Research, April 1998.
- Sack, Brian. "Uncertainty, Learning and Gradual Monetary Policy." Working Paper No. 1998-34, Federal Reserve Board of Governors, August 1998.
- Sargent, Thomas. *The Conquest of American Inflation*. Princeton: Princeton University Press, 1999.
- Shimer, Robert. "Why Is the U.S. Unemployment Rate So Much Lower?" in *National Bureau of Economic Research Macroeconomics Annual: 1998*. Cambridge, MA: MIT Press, 1999, pp. 11-61.
- Taylor, John. "On the Relation Between the Variability of Inflation and the Average Inflation Rate." Carnegie-Rochester Conference Series on Public Policy, Autumn 1981, 15(0), pp. 57-85.
- _____. "Aggregate Dynamics and Staggered Contracts." *Journal of Political Economy*, February 1980, 88(1), pp. 1-23.
- Wieland, Volker. "Monetary Policy and Uncertainty About the Natural Unemployment Rate." Finance and Economics Discussion Paper Series, No. 98/22, Board of Governors of the Federal Reserve System, April 1998.
- _____. "Learning by Doing and the Value of Optimal Experimentation." *Journal of Economic Dynamics and Control*, April 2000, 24(4), pp. 501-34.

Appendix A

PROPOSITIONS AND PROOFS

Proposition 1. A shift in either α_t or η_t results in a persistent change in the inflation target $\bar{\pi}_t$.

Proposition 2. A one-time switch in the structural parameter causes the inflation variance to rise in periods following the switch.

Proof. Given equation (5) and the fact that the shocks are uncorrelated, write the variance of inflation, conditional on g_t , as

$$(A1) \quad \text{Var}(\pi_t | g_t) = \left(\frac{g_t - 1}{g_t} \right)^2 \text{var}(\bar{\pi}_{t-1} | g_t) + \left(\frac{k\alpha}{g_t} \right)^2 (\sigma_v^2 + \sigma_\varepsilon^2) + \sigma_v^2.$$

Further

$$\text{Var}(\hat{\eta}_t | g_t) = \frac{\sigma_\varepsilon^2}{g_t - 1}.$$

Combining yields

$$(A2) \quad \text{Var}(\pi_t | g_t) = \frac{(k\alpha_t)^2}{g_t} \sigma_\varepsilon^2 + \left(\frac{k\alpha_t}{g_t} \right)^2 \sigma_v^2 + \sigma_v^2.$$

For $t < t'$, the economy is in a steady state with $g_t = \infty$, and $\text{Var}(\pi_t | g_t) = \sigma_v^2$. Further, (A2) implies that $\text{Var}(\pi_t | g_t)$ is strictly decreasing in g_t , from which follows that the variance of inflation rises following t' . **Q.E.D.**

Proposition 3. The mean level of inflation is increasing in α . Additionally, when α is larger, the switch in η_t causes the variance of inflation to rise by a greater amount.

Proof. The first statement follows directly from (5). When α rises, the expected value of the second term on the right-hand side of (5) rises and the expected value of the next period target rises. In period t , the expected value of the second term is again larger than it was before the regime shift. The first term has also risen because the target rose that period. The conclusion follows by induction. Consider again (A1). It is easy to verify that the variance of inflation rises with α , which implies the latter result. **Q.E.D.**

Appendix B

THE GIBBS SAMPLER

The goal of the sampling routine is to estimate the conditional posterior distribution $p(\tilde{\beta}_T, \tilde{S}_T, \tilde{Z}_T, \omega | \tilde{y}_T)$, where $\tilde{y}_T = (y_1, y_2, \dots, y_T)$ is the vector of observables through time T and ω is the vector of parameters governing both Markov processes and the variances of the white noise shocks to the Phillips curve and the inflation rate.¹⁹ The Markov processes make the direct estimation of this distribution impractical. However, draws from $p(\tilde{\beta}_T, \tilde{S}_T, \tilde{Z}_T, \omega | \tilde{y}_T)$ can be made from an ergodic distribution of Markov simulations generated iteratively from the following conditional densities:

$$(B1) \quad p(\tilde{\beta}_T, \tilde{S}_T, \tilde{Z}_T, \omega | \tilde{y}_T) \begin{cases} p(\tilde{\beta}_T | \tilde{y}_T, \tilde{S}_T, \tilde{Z}_T, \omega) \\ p(\tilde{S}_T | \tilde{y}_T, \tilde{\beta}_T, \tilde{Z}_T, \omega) \\ p(\tilde{Z}_T | \tilde{y}_T, \tilde{\beta}_T, \tilde{S}_T, \omega) \\ p(\omega | \tilde{y}_T, \tilde{\beta}_T, \tilde{S}_T, \tilde{Z}_T) \end{cases}$$

The process of drawing from each of the above marginal distributions is accomplished by

a version of Carter and Kohn's (1994) multi-move algorithm, with one exception. The conditional distribution $p(\tilde{S}_T | \tilde{y}_T, \tilde{\beta}_T, \tilde{Z}_T, \omega)$ differs from the others in (B1) in that S_t depends on $\tilde{S}_{t-1}$ because of the policymaker's gain sequence. The multi-move algorithm generates the entire vector $\tilde{S}_T$ simultaneously. However, given S_t 's dependence on previous states, it cannot be generated in this manner. Thus, $\tilde{S}_T$ will be drawn from sequential sampling from

$$(B2) \quad p(S_t | \tilde{S}_{-t}, \tilde{y}_T, \tilde{\beta}_T, \tilde{Z}_T, \omega),$$

where $\tilde{S}_{-t} = (S_1, S_2, \dots, S_{t-1}, S_{t+1}, \dots, S_T)$. The joint density $p(\tilde{\beta}_T, \tilde{S}_T, \tilde{Z}_T, \omega | \tilde{y}_T)$ can then be estimated using the marginal densities (B1) and (B2).

¹⁹ More detailed analyses of the Gibbs sampler are available in Cassella and George (1992) and Kim and Nelson (1999).