Title
Regime Switching and Monetary Policy Measurement

Authors
Owyang, Michael T.
Ramey, Garey

Publication Date
2001

2001-03

UNIVERSITY OF CALIFORNIA, SAN DIEGO

DEPARTMENT OF ECONOMICS

REGIME SWITCHING AND MONETARY POLICY MEASUREMENT

BY

MICHAEL OWYANG

AND

GAREY RAMEY

DISCUSSION PAPER 2001-03
JANUARY 2001

REGIME SWITCHING AND MONETARY POLICY MEASUREMENT

Michael Owyang and Garey Ramey*

January 2001

Abstract. This paper applies regime switching methods to the problem of measuring monetary policy. Policy preferences and structural factors are specified parametrically as independent Markov processes. Interaction between the structural and preference parameters in the policy rule serves to identify the two processes. The estimates uncover policy episodes that are initiated by switches to "dove regimes," shown to Granger cause both NBER recessions and the Romer dates. These episodes imply real effects of monetary policy that are smaller than those found in previous studies.

1. Introduction

Beginning with Hamilton's (1989) study of business cycle dynamics, regime switching methods have proven extremely useful in a wide range of applications in macroeconomics and finance.[1] This approach also holds promise for the measurement of monetary policy, since policy is typically regarded in terms of fluctuations between persistent regimes involving stronger or weaker anti-inflationary postures. Commonly-used vector autoregression (VAR)

* Owyang: Federal Reserve Bank of St. Louis. Ramey: University of California, San Diego. Jim Hamilton and Valerie Ramey provided much important input into this paper, for which we are grateful. We also thank William Gavin, Daniel Thornton, Takeo Hoshi and Neal Beck for useful comments. Paige Skiba provided research assistance. Ramey thanks the NSF for financial support. Opinions expressed herein are those of the authors and do not represent or reflect the official opinions of the Federal Reserve Bank of St. Louis or the Federeal Reserve system.

[1] See Kim and Nelson (1999) for a survey of regime switching methods and applications.

methods for measuring policy cannot capture such persistent policy regimes, as these methods can identify only highly transitory policy shocks.[2]

This paper conducts an exploratory study of the use of regime switching for estimating monetary policy preferences. Our strategy is to avoid theoretical details by adopting a stylized model of policy determination that captures inflation/unemployment tradeoffs in a simple way. This approach allows for straightforward resolution of identification and estimation issues, and our results may be viewed as a first assessment of the usefulness of regime switching for monetary policy measurement.

The model posits that the policymaker is constrained by a standard expectations augmented Phillips curve. The Phillips curve contains a structural parameter that follows a two-state Markov process, reflecting periodic shifts in the natural rate of unemployment. The policymaker adopts an inflation target that embodies tradeoffs between inflation and unemployment, captured by a preference parameter that follows an independent two-state Markov process. The process switches between a "dove regime" in which the policymaker more readily accommodates increases in the natural rate, and a "hawk regime" in which there is less accommodation. The key identifying assumption is that a rise in the natural rate leads to a larger increase in the inflation target when the preference parameter is in the dove state, relative to the hawk state. This distinguishes the structural process probabilistically from the policy process, making it possible to estimate both processes.

Estimates of the model are obtained by means of Gibbs sampling using monthly data over the period 1965:3 to 1999:2. Highly persistent structural and policy processes are estimated, each having statistically distinct state values. Further, we obtain estimates of the posterior expected values of both the structural and policy parameters over the sample period, providing a picture of the evolution of structural and policy regimes. The policy

[2]See Christiano, Eichenbaum and Evans (1999) for a survey of the large literature that has utilized VAR methods to measure monetary policy. Within a structural VAR model, Bernanke and Mihov (1998) have utilized regime switching to measure shifts between the targeting of the federal funds rate and nonborrowed reserves. Recently, Sims (1999) has estimated a regime switching model of a federal funds rate reaction function.

process, in particular, displays three "dove episodes" occurring in the late 1960's, mid-1970's, and in an interval around 1980. These episodes correspond closely to the onset of NBER recessions, as well as to the dates identified by Romer and Romer (1989,1994) as reflecting policy tightening by the Federal Reserve. The three episodes follow a basic pattern: a switch to the dove regime first occurs, followed roughly a year later by a Romer date and then a recession. A switch back to the hawk regime occurs after another year. Switches toward the dove regime are shown to Granger cause both recessions and the Romer dates. This suggests that monetary policy regimes are driven by shifts toward looser policy, initiating a process of policy reversal that takes roughly two years. In other words, monetary policy is driven by persistent "dove shocks."

The implications of these policy episodes for output, prices and other variables are assessed by means of a VAR that treats the estimated posterior expected values of the structural and policy parameters as exogenous variables. Using the estimated VAR, we study the dynamic effects of a stylized policy episode in which the policy parameter switches to the dove regime for 24 periods, and then switches back to the hawk regime. The onset of the dove regime initiates a steady rise in prices, while output begins to decline after a year. Prices fall after the hawk regime is restored, and output bottoms out about a year later. The switch to the dove regime also induces a sharp rise in the federal funds rate, and the federal funds rate jumps upward again halfway through the dove episode.

Our dove shock shares a number of characteristics of a positive federal funds rate shock in the standard VAR model, including the increase in price levels following the shock. In the present case, however, there is no "price puzzle," since the shock actually represents a loosening of policy. Our policy episode generates a significantly smaller decline in output than that associated with federal funds shocks in standard VAR models. Moreover, the cumulative increase in unemployment associated with the restoration of the hawk regime, relative to the corresponding reduction of inflation (the so-called "sacrifice ratio"), is only 0.87, less than half the value found in previous studies. Overall, our results suggest that the real effects of monetary policy may be less significant than previously believed.

Section 2 presents the model, estimates are given in Section 3, and comparisons with NBER recessions and the Romer dates are carried out in Section 4. Implications of policy episodes and the sacrifice ratio are considered in Sections 5 and 6, respectively, and Section 7 concludes.

2. Model

We adopt a stylized model that focuses on policy tradeoffs between inflation and unemployment, in the vein of Kydland and Prescott (1977) and Barro and Gordon (1983a,b). The key innovation is that structural factors and policy preferences are specified parametrically as independent Markov processes. Unemployment is determined by an expectations-augmented Phillips curve:

$$u_t = k(\pi_t^e - \pi_t) + \eta_t + \varepsilon_{1t}, \tag{1}$$

where π_t^e denotes inflation expectations of private agents, π_t gives realized inflation, η_t denotes the structural parameter, and ε_{1t} is a white noise shock. Private agents form adaptive inflation expectations according to an autoregressive rule:

$$\pi_t^e = \gamma_1 \pi_{t-1} + \gamma_2 \pi_{t-2} + \varepsilon_{2t}, \tag{2}$$

where ε_{2t} is a white noise shock. The structural parameter η_t follows a two-state Markov process, taking on values $\overrightarrow{h} = (h_1, h_2)$, with $h_1 < h_2$. Note that η_t can be interpreted as the natural rate of unemployment. Let T^η indicate the matrix of transition probabilities for the structural process.

Realized inflation is determined by

$$\pi_t = \overline{\pi}_t + \varepsilon_{3t}, \tag{3}$$

where $\overline{\pi}_t$ indicates the inflation target set by a policymaker, and ε_{3t} is a white noise shock reflecting control error. We assume that the policy target responds to underlying policymaker preferences that embody desired tradeoffs between inflation and unemployment. Policy preferences and parameters of the Phillips curve interact in a manner that imposes restrictions on the policy target.

Figure 1 illustrates the nature of these restrictions. The negatively sloped lines in the figure give expected unemployment $\overline{u}_t$ as a function of the policy target $\overline{\pi}_t$, derived from (1) and (3). These lines serve as constraints on the policymaker. Under a dove policy regime, the policymaker's preferences place relatively greater weight on reducing unemployment. Thus, a rise in the natural rate from h_1 to h_2 shifts the policy outcome from point A to point B. Under a hawk regime, in contrast, reducing inflation receives relatively greater weight, and correspondingly the policy outcome shifts from point C to point D when the natural rate rises. Thus, a switch in the structural state induces larger changes in inflation relative to unemployment when the policymaker is in the dove regime, compared to the hawk regime.

To capture this basic relationship between preferences and structure, we adopt the following specification of the policy rule:

$$\overline{\pi}_t = \alpha_t(k\pi_t^e + \eta_t), \tag{4}$$

where α_t gives the policy preference parameter. Assume that α_t follows a two-state Markov process, taking on values $\overrightarrow{a} = (a_1, a_2)$, with $a_1 < a_2$. The matrix of transition probabilities is given by T^α. Note that the higher-valued state a_2 indicates the dove regime, since for $\alpha_t = a_2$ an increase in either the structural state or expected inflation leads to a larger rise in the inflation target.

Combining (3) and (4), we have

$$\pi_t = \alpha_t(k\pi_t^e + \eta_t) + \varepsilon_{3t}. \tag{5}$$

Equations (1), (2) and (5) comprise the complete model. The model has the important feature that, as a consequence of the restrictions implied by (4), the posterior distribution of the path of inflation, conditional on any realized path of the structural parameter η_t, will vary depending on the realized path of the preference parameter α_t. Because of this, the two parameters are identified.

3. Estimates

The model is estimated using Gibbs sampling; the appendix describes details of the estimation procedure.[3] We use monthly data on inflation and unemployment over the period 1965:3 to 1999:2, taken from Citibase.[4] The Gibbs sampler produces estimates of the model parameters k, γ_1 and γ_2, along with state values and transition matrices of the two Markov processes and the variances of the three white noise shocks. In addition, we obtain estimates of the posterior probabilities of the structural and preference states in each period, conditional on the full sample.

Parameter estimates are reported in Table 1. Of the estimated parameters, only the variance of the white noise shock in the inflation expectations function, $\sigma^2_{\varepsilon_2}$, and the weight on the second lag of inflation in that function, γ_2, are not significantly different from zero. Specifically, estimated values for each state of the structural and preference processes are significant at the 95 percent level. Moreover, for each process the estimated state values are distinct at the 95 percent level. The table also shows estimated values of diagonal elements of the transition matrices T^η and T^α. Observe that these estimates lie very close to unity, indicating that the structural and preference processes are highly persistent.

The estimates of h_1 and h_2 indicate that the economy fluctuates between a "low natural rate" state, in which the natural rate of unemployment is about 3.3 percent, and a "high natural rate" state, where the natural rate is 6.7 percent. Since the estimated values of the preference states a_1 and a_2 are positive, it follows that the policymaker accommodates increases in the natural rate under both the dove and hawk regimes.

Quantitative implications for inflation may be assessed by computing the steady-state expected inflation rates implied by the model, where steady states are determined by holding the state profile (h_i, a_i) fixed for all time and setting $\pi^e_t = E[\pi_t]$. Steady-state inflation rates for each state profile are reported in Table 2. Observe that inflation is higher in the high

[3]Casella and George (1992) and Kim and Nelson (1999) provide references on Gibbs sampling.

[4]Inflation data are the seasonally-adjusted, annualized rate of change of CPI-U. Unemployment is measured by the seasonally-adjusted, annualized rate of unemployment in the U.S.

natural rate state, compared with the low natural state, and it is higher in the dove regime, compared with the hawk regime. Moving to the high natural rate state produces a larger rise in inflation when policy preferences are in the dove regime, consistent with (4). Further, changes in the preference state have a larger quantitative effect on inflation than do changes in the structural state.

Note from Table 1 that $\gamma_1 + \gamma_2 \simeq 1$. Thus, according to (2), private agents' inflation expectations converge to rational expectations if the structural and policy states remain unchanged for a sufficiently long time. Our results imply that important departures from rational expectations occur only during transitions following regime switches. Moreover, γ_1 lies very close to unity, indicating that expectations adjust quite rapidly following a switch.

Posterior expected values for the structural and preference parameters can be computed by multiplying the estimated state values by the posterior state probabilities generated by the Gibbs sampler. For example, the posterior expected value of the structural parameter in period t is given by

$$E[\eta_t|\widetilde{y}_T, \widetilde{\varphi}] = h_1 \Pr[\eta_t = h_1|\widetilde{y}_T, \widetilde{\varphi}] + h_2 \Pr[\eta_t = h_2|\widetilde{y}_T, \widetilde{\varphi}],$$

where $\widetilde{y}_T$ indicates the inflation and unemployment data over the full sample, and $\widetilde{\varphi}$ gives the estimated parameter vector. Figure 2 reports estimated posterior expected values of the structural parameter, graphed along with the inflation and unemployment data. Observe that the sample period begins in the low natural rate state. A switch to the high natural rate state occurs near the end of 1970, coinciding with an oil price increase episode identified by Hamilton (1983). The high natural rate state persists until mid-1997, when a switch back to the low natural rate state is observed.[5]

Estimated posterior expected values for the policy preference parameter are shown in Figure 3. The hawk regime predominates for most of the sample, but there are three major dove episodes, occurring in the late 1960's, mid-1970's, and in an interval around 1980. As

[5]The estimates are roughly consistent with Shimer's (1998) argument that demographic factors induced an extended rise in the natural rate beginning in the early 1970's.

may be seen in the figure, inflation tends to be higher on average during the dove episodes, while switches back to the hawk regime coincide with lower inflation and sharp increases in unemployment. The next section offers further interpretations of the posterior expected policy parameter.

In summary, we obtain estimates of distinct, highly persistent processes for the structural and preference parameters. The estimated state values imply quantitatively important effects on the levels of inflation and unemployment generated by the model. Posterior expected values of the structural and policy parameters uncover switches that occur on a number of occasions over the sample period.

4. NBER Recessions and Romer Dates

The posterior expectation of the policy parameter, shown in Figure 3, indicates numerous shifts in the monetary policy regime. In this section we relate these shifts to NBER recessions and the Romer dates. The latter are dates, identified by Romer and Romer (1989,1994), at which the Fed declared an intent "to exert a contractionary influence on the economy in order to reduce inflation (1989, p. 134)."

Five NBER recessions and five Romer dates lie within the sample period we consider. Figure 4 graphs the recessions and Romer dates, along with the posterior expected value of the policy parameter. Observe that recessions tend to coincide with the latter parts of the dove episodes, with reversions to the hawk regime occurring near the end of recessions. Further, four of the five Romer dates occur shortly after switches to the dove regime. Reversion to the hawk regime occurs soon after these Romer dates. Based on the graph, there appears to be a close relationship between NBER recessions, the Romer dates, and the policy regime switches that we identify.

A more objective assessment of the relationship between these variables can be obtained by testing for Granger causality. For this purpose, we construct two variables from our posterior expected policy series, an "upward switch" variable, us_t, and a "downward switch"

variable, ds_t, defined as follows:

$$us_t = \max\{E[\alpha_t] - E[\alpha_{t-1}], 0\},$$

$$ds_t = \max\{-(E[\alpha_t] - E[\alpha_{t-1}]), 0\},$$

where $E[\alpha_t]$ indicates the posterior expected value of the policy parameter. We capture the onset of NBER recessions by means of a variable N_t having the value unity in the initial month of each of the five recessions in the sample, and zero for the remaining months. The Romer dates are indicated by a series R_t having a value of unity for the five Romer dates, and zero for the other months.

We first regress N_t on 24 lags of N_t and us_t, and then rerun the regression excluding the us_t lags. Table 3 reports the results of an F test for this exclusion restriction. Observe that the high F value of 1.968, and the associated p-value of 0.004, indicate that the lags of us_t are highly significant for predicting N_t. The table also reports the results of the reverse exercise, and the values $F = 0.848$ and $p = 0.674$ show that N_t does not help to predict us_t. This provides strong statistical evidence that the upward switches in our measured posterior expected policy series Granger cause NBER recessions. The reverse relationship holds when recessions are related to the downward switches: ds_t does not help predict N_t ($F = 0.259$, $p = 0.999$), while N_t is highly significant in predicting ds_t ($F = 5.150,\ p = 0.000$). Thus, NBER recessions Granger cause the downward switches.

Turning to the Romer dates, it can be observed in Table 3 that us_t helps to predict R_t, while the reverse is not true, meaning that the upward switches Granger cause the Romer dates. The relationship between the Romer dates and downward switches is less clear. Table 3 shows that ds_t is of no help in predicting R_t, while R_t has little explanatory power for ds_t. Thus, the data do not reveal any clear causal relationship between R_t and ds_t.

We conclude that the onset of dove regimes, as measured by upward switches in the posterior expectation of the policy parameter, can be viewed as a driving force behind both NBER recessions and the policy initiatives observed by the Romers. This suggests a new interpretation of policy episodes over the period since the mid-1960's. Beginning in the hawk

regime, a policy episode is initiated by a switch to the dove regime. Roughly a year after this switch, policymakers begin to express clear intent to tighten policy, based on fears of inflation. Close to this time a recession ensues. A little over a year later, the intent becomes realized in the form of a switch back to the hawk regime. In the next two sections we assess the implications of this stylized pattern for a range of macroeconomic variables.

5. Implications of Policy Episodes

In this section we consider the dynamic implications for output, prices and other variables of the policy episodes discussed in the preceding section. To begin, we estimate a standard VAR model consisting of five endogenous variables, including industrial production, the implicit price deflator for consumption expenditures, the federal funds rate, the ratio of nonborrowed reserves to total reserves, and M1.[6] Natural logs of all variables except the federal funds rate are taken, and each variable is regressed on 11 lags of the five endogenous variables, along with the current value and 11 lags of the posterior expected values of the structural and policy parameters (these are the series graphed in Figures 2 and 3).

Given the estimated VAR, we conduct the following policy experiment. We set the value of the posterior expected structure parameter equal to its unconditional expected value 5.85 for all periods, and compute the steady state associated with the hawk regime, where the posterior expected policy parameter is equal to 0.58 in each period. Beginning in this hawk steady state, the policy variable is increased to the dove regime value of 1.43 for 24 months, after which it is returned to the hawk regime value of 0.58 for all future periods. Thus, we simulate a policy episode initiated by a switch to the dove regime, followed by a switch back to the hawk regime. Observe from Figure 3 that the three identified dove episodes average about two years in length, so that the simulation can be viewed as a typical episode.

Results are given in Figure 5. In this figure, the dove regime begins in period 13 and ends in period 36. Output as measured by industrial production remains flat for the first year of the dove regime, while prices rise steadily. About halfway through the dove regime,

[6]We thank Charles Evans for providing these data.

output begins to decline sharply. Interestingly, this coincides roughly with the Romer dates occurring about halfway through the dove episodes depicted in Figure 4. Once the hawk regime takes hold in period 37, prices begin to fall, and output bottoms out about a year after restoration of the hawk regime.

Notice further that when the dove regime hits, there is an upward spike in the federal funds rate, while the nonborrowed reserve ratio declines and M1 begins a gradual upward movement. A little under halfway through the dove regime there is another upward spike of the federal funds rate, and the rate remains well above its hawk steady state value (by over five percentage points) for the remainder of the dove regime. After the second spike in the federal funds rate, the nonborrowed reserve ratio begins to increase, while the rise in M1 levels off. Once the hawk regime is restored, the federal funds rate gradually falls.

These findings suggest a number of interpretations. Note first of all that the dove shock initiating the policy episode generates a response much like that of a positive federal funds rate shock in the standard VAR policy analysis: an upward spike in the federal funds rate is followed by a gradual rise in prices, a downward movement in nonborrowed reserves, and a decline in output.[7] From our perspective, the increase in prices following the shock does not constitute a "price puzzle," since the shock actually represents a loosening of policy. Output responds with a lag in our setting because of lags in the reversal of policy.

The initial upward movement of the federal funds rate may be understood as a response to higher anticipated inflation. As policymakers respond to the increase in inflation, the federal funds rate rises again, and the nonborrowed reserve ratio begins to increase. Thus, our results confirm previous findings that the federal funds rate serves as an important instrument for implementing the anti-inflationary policy. Further, contractionary policy appears to entail a rise in the nonborrowed reserve ratio, contrary to the suggestion of Strongin (1995).

Finally, the overall policy episode leads to a smaller decline in output than that found in

[7]See Christiano, Eichenbaum and Evans (1999, p. 90) for a standard VAR analysis at monthly frequency that displays these characteristics. Our results depart from the VAR findings, however, in predicting a rise in M1 following the shock.

previous studies. Christiano, Eichenbaum and Evans (1999), for example, utilize standard VAR methods to show that a policy shock generating an 80 basis point rise in the federal funds rate reduces output by about 0.5 percent after two years. Our policy episode, in contrast, implies a much larger increase in the federal funds rate, accompanied by a smaller decline in output.

Our approach has the key advantage that we are able to measure shifts in the policy regime that induce persistent changes in policy, i.e., we measure a systematic component of policy. This allows us to ask whether monetary policy affects the economy only through unanticipated changes, or instead whether persistent changes in the policy regime can have an effect. To assess this issue, we carry out a policy simulation in which the policy variable takes on the dove value of 1.43 for a single period only, with the hawk value being maintained for all other periods. Thus, only the switch itself, which is the unanticipated part of the policy episode, can have an effect. Figure 6 compares this "one-shot" policy switch to the persistent 24-month dove episodes considered earlier. Observe that the one-shot policy episode has a tiny effect on output and prices when compared to the persistent episode; results are similar for the other variables. It follows that the systematic component of policy, captured here as a persistent policy state, has an important effect on real variables that goes beyond the effect of the unanticipated component.

6. Sacrifice Ratio

A range of studies have used Phillips curve frameworks to assess the unemployment effects of disinflationary policies.[8] These studies have constructed "sacrifice ratios" that measure the cumulative increase in unemployment associated with each percentage point of policy-induced inflation reduction. In particular, there is a rough consensus that a one percent reduction in inflation increases cumulative unemployment by at least two percentage points per year.

Our model can be used to measure the cost of disinflation by considering the effects of

[8]For example, see Okun (1978), Gordon and King (1982), Blinder (1987) and Ball (1994).

switches from the dove to the hawk regime. For this purpose, we estimate a VAR having two endogenous variables, the inflation and unemployment rates used in the original estimation. Each variable is regressed on 11 lags of the two endogenous variables, along with the current value and 11 lags of the posterior expected values of the structural and policy parameters. Using these estimates, we reconsider the policy experiment analyzed in the preceding section: beginning in the hawk steady state, the policy parameter switches to the dove regime for 24 periods, then switches back to the hawk regime.

Results are shown in Figure 7. During the dove episode, the inflation rate averages 6.74 percent higher than in the hawk steady state, while the unemployment rate goes down by about half a percentage point. After the hawk preferences are restored, inflation returns (somewhat erratically) to the hawk steady state level, while a large increase in unemployment ensues. In particular, unemployment peaks one year after the switch back to the hawk regime, at about two percentage points above the hawk steady state level.

The sacrifice ratio implied by these results can be calculated as follows. Beginning in the last six months of the dove regime, when unemployment rises above the hawk steady state value, the cumulative increase in years of unemployment over the next four years amounts to 5.88 percentage points. Since inflation falls by 6.74 percent as a result of the disinflation policy, the implied value of the sacrifice ratio is 0.87, or less than half of the consensus figure.

Our finding of a low sacrifice ratio can be comprehended in terms of our parameter estimates. In Phillips curve models, high costs of disinflation emerge when unemployment is highly sensitive to inflation surprises, or when inflation expectations are slow to adjust to a new disinflationary regime. However, our estimate of $k = 0.747$ indicates that unemployment is affected only slightly by inflation surprises. Moreover, since $\gamma_1 = 0.9772$, adjustment of expectations is extremely rapid. Both of these factors militate against disinflation costs and support the finding of a low sacrifice ratio.[9]

[9]Past estimates of structural models by Sargent (1976), Fair (1979) and Broadbent and Barro (1997) have found low values of the parameter relating unemployment to inflation surprises, although our estimate is even smaller. In a model incorporating both forward- and backward-looking expectations, Galí and Gertler (1999)

An alternative perspective on the costs of disinflation can be gained by considering the policy episode as a whole. Beginning in period 13, the switch to the dove regime and back to the hawk regime generates a cumulative increase in years of unemployment of 5.51 percent over six years. Thus, unemployment reductions obtained early in the dove phase are more than offset by the higher unemployment needed to restore the hawk steady state. The initial switch to the dove regime implies a highly unfavorable tradeoff, in that very large increases in inflation yield only slight reductions in unemployment.

7. Conclusion

In this paper we apply regime switching techniques to the measurement of monetary policy regimes. Using a stylized model of inflation/unemployment policy tradeoffs, we obtain estimates that reveal highly persistent processes of policy preferences and economic structure, switching between distinct states. The estimated posterior expected values of the policy parameter trace out episodes involving switches to a dove regime for about two years, followed by reversion to a hawk regime. The switches that initiate these episodes Granger cause both NBER recessions and the Romer dates, suggesting that incidents of monetary tightening might be best regarded as responses to earlier dove shocks. Our estimated policy episodes imply smaller effects on real variables than have been obtained in previous studies using different policy measures.

Our model may be extended to allow regime switches to depend on the duration of regimes or economic variables. The methodology can be applied to a broader set of variables that may influence policymaker preferences, including employment, output and financial market variables such as interest rates, and it may be applied to other countries. Policymaker objectives may be combined with policy instruments to create a synthetic analysis linking policy regimes with the particular instruments used to implement these regimes.

We have relied on a bare-bones structural model that has allowed us to obtain sharp estimates, but that also raises valid questions of robustness. A more theoretically complete

have found that backward-looking expectations are quantitatively unimportant, suggesting that expectations adjust rapidly in the aggregate.

model would incorporate explicit utility maximization by a forward-looking policymaker, lags in implementation of policy targets, and a richer structure of expectation formation by private agents. The econometric implementation of such a model represents a challenging and, in view of our results, potentially fruitful avenue for future research.

8. Appendix

Define state vectors S_t and Z_t such that $\eta_t = \overrightarrow{h} S_t$ and $\alpha_t = \overrightarrow{a} Z_t$. Let period t data be denoted by

$$y_t = \begin{bmatrix} \pi_t \\ u_t \end{bmatrix}.$$

$\widetilde{S}_T$ defines the vector of states $\widetilde{S}_T = (S_1, S_2, ..., S_T)$, where T indicates the length of the sample, and $\widetilde{Z}_T$, $\widetilde{\pi}^e_T$ and $\widetilde{y}_T$ are defined similarly. The vector of model parameters is given by $\widetilde{\varphi} = (a_1, a_2, h_1, h_2, \sigma^2_{\varepsilon_1}, \sigma^2_{\varepsilon_2}, \sigma^3_{\varepsilon_3}, k, \gamma_1, \gamma_2, T^\alpha, T^\eta)$.

It is convenient to rewrite the model (1), (3) and (5) in a state space representation. Define Y_t as

$$Y_t = \begin{bmatrix} u_t + k\pi_t - \eta_t \\ \pi_t - \alpha_t \eta_t \end{bmatrix}.$$

The model may be rewritten as follows:

$$Y_t = H_t \pi^e_t + e_t,$$

$$\pi^e_t = \mu_t + \varepsilon_{2t},$$

where

$$H_t = \begin{bmatrix} k \\ k\alpha_t \end{bmatrix}, \quad e_t = \begin{bmatrix} \varepsilon_{1t} \\ \varepsilon_{3t} \end{bmatrix},$$

and $\mu_t = \gamma_1 \pi_{t-1} + \gamma_2 \pi_{t-2}$.

The objective of the Gibbs sampler is to characterize the joint density $p(\widetilde{S}_T, \widetilde{Z}_T, \widetilde{\pi}^e_T, \widetilde{\varphi} | \widetilde{y}_T)$ using the ergotic distribution of a Markov simulation of the following conditional joint densities that are generated iteratively:

$$p(\widetilde{\pi}^e_T | \widetilde{y}_T, \widetilde{S}_T, \widetilde{Z}_T, \widetilde{\varphi}),$$
$$p(\widetilde{Z}_T | \widetilde{y}_T, \widetilde{S}_T, \widetilde{\pi}^e_T, \widetilde{\varphi}),$$
$$p(\widetilde{S}_T | \widetilde{y}_T, \widetilde{\pi}^e_T, \widetilde{Z}_T, \widetilde{\varphi}),$$
$$p(\widetilde{\varphi} | \widetilde{y}_T, \widetilde{S}_T, \widetilde{Z}_T, \widetilde{\pi}^e_T).$$

Samples from these densities are drawn at each step and used to generate the other densities, constituting a Markov chain. After an appropriate number of iterations, the ergotic distribution of this chain of conditional densities is the joint density $p(\widetilde{S}_T, \widetilde{Z}_T, \widetilde{\pi}^e_t, \widetilde{\varphi}|\widetilde{y}_T)$.[10]

Conditional Density of Inflation Expectations.

The conditional density $p(\widetilde{\pi}^e_T|\widetilde{y}_T, \widetilde{S}_T, \widetilde{Z}_T, \widetilde{\varphi})$ can be obtained by applying a Kalman filter modified for the presence of the two Markov processes that govern η_t and α_t. The Kalman filter produces the densities $p(\pi^e_t|\widetilde{y}_t, \widetilde{S}_t, \widetilde{Z}_t, \widetilde{\varphi})$ for all t. Given some initial conditions $\pi^e_{t-1|t-1}$ and $V_{t-1|t-1}$, the filter generates, for all t:

$$\pi^e_{t|t-1} = \mu_t,$$

$$V_{t|t-1} = \sigma^2_{\varepsilon_2},$$

$$\pi^e_{t|t} = \pi^e_{t|t-1} + (H_t V_{t|t-1} H_t' + R)^{-1} V_{t|t-1} H_t' (Y_t - H_t \pi^e_{t|t-1}),$$

$$V_{t|t} = V_{t|t-1}(I - (H_t V_{t|t-1} H_t' + R)^{-1} V_{t|t-1} H_t' H_t),$$

where

$$R = \begin{bmatrix} \sigma^2_{\varepsilon_1} & 0 \\ 0 & \sigma^2_{\varepsilon_3} \end{bmatrix}.$$

Then rewrite $p(\widetilde{\pi}^e_T|\widetilde{y}_T, \widetilde{S}_T, \widetilde{Z}_T, \widetilde{\varphi})$ as

$$p(\widetilde{\pi}^e_T|\widetilde{y}_T, \widetilde{S}_T, \widetilde{Z}_T, \widetilde{\varphi}) = p(\pi^e_T|\widetilde{y}_T, \widetilde{S}_T, \widetilde{Z}_T, \widetilde{\varphi}) \prod_{t=1}^{T-1} p(\pi^e_t|\widetilde{y}_t, \widetilde{S}_t, \widetilde{Z}_t, \widetilde{\varphi}, \pi^e_{t+1}).$$

The final iteration of the Kalman filter provides the first term. Elements of the second term are determined by the results of the Kalman filter and the following recursive conditional densities:

$$\pi^e_{t|t,\pi^e_{t+1}} = \pi^e_{t|t},$$

$$V_{t|t,\pi^e_{t+1}} = V_{t|t},$$

[10]Reference on the convergence of the sampling algorithm can be found in Gelfand and Smith (1990) and Geweke (1992).

where $Q = \sigma^2_{\tilde{\varepsilon}_2}$, $\pi^e_{t|t} = p(\pi^e_t|\widetilde{y}_t, \widetilde{S}_T, \widetilde{Z}_T, \widetilde{\varphi})$ and $V_{t|t}$ is the conditional variance as determined by the Kalman filter.

Conditional Densities of Policy and Structural States.

Recall that, conditional on $\widetilde{y}_T$, $\widetilde{S}_T$, $\widetilde{\pi}^e_T$ and $\widetilde{\varphi}$, (5) is linear in $\widetilde{Z}_T$. Given $p(Z_0|\widetilde{y}_0)$, a prior probability for the initial state, the Hamilton (1989) filter generates the conditional density $p(Z_T|\widetilde{y}_T, \widetilde{S}_T, \widetilde{\pi}^e_T, \widetilde{\varphi})$. Then, following Carter and Kohn (1994) and Kim and Nelson (1998), the density $p(\widetilde{Z}_T|\widetilde{y}_T, \widetilde{S}_T, \widetilde{\pi}^e_T, \widetilde{\varphi})$ is obtained from

$$p(\widetilde{Z}_T|\widetilde{y}_T, \widetilde{S}_T, \widetilde{\pi}^e_T, \widetilde{\varphi}) = p(Z_T|\widetilde{y}_T, \widetilde{S}_T, \widetilde{\pi}^e_T, \widetilde{\varphi}) \prod_{t=1}^{T-1} p(Z_t|\widetilde{y}_T, \widetilde{S}_T, \widetilde{\pi}^e_T, \widetilde{\varphi}, Z_{t+1}). \tag{6}$$

Each density $p(Z_t|\widetilde{y}_T, \widetilde{S}_T, \widetilde{\pi}^e_T, \widetilde{\varphi}, Z_{t+1})$ is generated from a filtering algorithm and Bayes' Law:

$$\begin{aligned} p(Z_t|\widetilde{y}_T, \widetilde{S}_T, \widetilde{\pi}^e_T, \widetilde{\varphi}, Z_{t+1}) &= \frac{p(Z_{t+1}|\widetilde{y}_t, Z_t, \widetilde{S}_t, \widetilde{\pi}^e_t, \widetilde{\varphi}) p(Z_t|\widetilde{y}_t, \widetilde{S}_t, \widetilde{\pi}^e_t, \widetilde{\varphi})}{p(Z_{t+1}|\widetilde{y}_t, \widetilde{S}_t, \widetilde{\pi}^e_t, \widetilde{\varphi})} \\ &= \frac{p(Z_{t+1}|Z_t) p(Z_t|\widetilde{y}_t, \widetilde{S}_t, \widetilde{\pi}^e_t, \widetilde{\varphi})}{\sum_{Z_t} p(Z_{t+1}|\widetilde{y}_t, Z_t, \widetilde{S}_t, \widetilde{\pi}^e_t, \widetilde{\varphi}) p(Z_t|\widetilde{y}_t, \widetilde{S}_t, \widetilde{\pi}^e_t, \widetilde{\varphi})} \\ &= \frac{p(Z_{t+1}|Z_t) p(Z_t|\widetilde{y}_t, \widetilde{S}_t, \widetilde{\pi}^e_t, \widetilde{\varphi})}{\sum_{Z_t} p(Z_{t+1}|Z_t) p(Z_t|\widetilde{y}_t, \widetilde{S}_t, \widetilde{\pi}^e_t, \widetilde{\varphi})}, \end{aligned}$$

where $p(Z_{t+1}|Z_t)$ is the transition probability and the filter determines the density $p(Z_t|\widetilde{y}_t, \widetilde{S}_t, \widetilde{\pi}^e_t, \widetilde{\varphi})$. The first inequality is simply an application of Bayes' Law. The final two inequalities arise from the Markov property of Z_t: in determining the density for Z_{t+1}, the only relevant information in the available set is the previous state Z_t. The numerator in (6) is calculated from the Hamilton filter as

$$\begin{aligned} p(Z_t|\widetilde{y}_t, \widetilde{S}_t, \widetilde{\pi}^e_t, \widetilde{\varphi}) &= \frac{f(y_t|\widetilde{y}_{t-1}, Z_t, \widetilde{S}_t, \widetilde{\pi}^e_t, \widetilde{\varphi}) p(Z_t|\widetilde{y}_{t-1}, \widetilde{S}_t, \widetilde{\pi}^e_t, \widetilde{\varphi})}{f(y_t|\widetilde{y}_{t-1}, \widetilde{S}_t, \widetilde{\pi}^e_t, \widetilde{\varphi})} \\ &= \frac{f(y_t|\widetilde{y}_{t-1}, Z_t, \widetilde{S}_t, \widetilde{\pi}^e_t, \widetilde{\varphi}) p(Z_t|\widetilde{y}_{t-1}, \widetilde{S}_t, \widetilde{\pi}^e_t, \widetilde{\varphi})}{\sum_{Z_t} f(y_t|\widetilde{y}_{t-1}, Z_t, \widetilde{S}_t, \widetilde{\pi}^e_t, \widetilde{\varphi}) p(Z_t|\widetilde{y}_{t-1}, \widetilde{S}_t, \widetilde{\pi}^e_t, \widetilde{\varphi})}, \end{aligned}$$

where

$$p(Z_t|\widetilde{y}_{t-1}, \widetilde{S}_t, \widetilde{\pi}^e_t, \widetilde{\varphi}) = \sum_{Z_{t-1}} p(Z_t|Z_{t-1}) p(Z_{t-1}|\widetilde{y}_{t-1}, \widetilde{S}_t, \widetilde{\pi}^e_t, \widetilde{\varphi}).$$

The density $p(Z_{t-1}|\widetilde{y}_{t-1}, \widetilde{S}_t, \widetilde{\pi}_t^e, \widetilde{\varphi})$ is taken from the previous iteration. The conditional density $p(\widetilde{S}_T|\widetilde{y}_T, \widetilde{\pi}_T^e, \widetilde{Z}_T, \widetilde{\varphi})$ can be generated using a similar process.

Conditional Density of the Parameter Vector.

The conditional densities for the elements of the parameter vector are generated by employing Bayesian OLS. The Bayesian posterior distribution for each element of $\widetilde{\varphi}$, conditional on all other elements of $\widetilde{\varphi}$, can be determined given a prior distribution. If model parameters, excluding the variances and transition probabilities, have prior distributions of the form $p_0(\varphi_i) \sim N(a_i, A_i\sigma^2)$, then their corresponding posterior conditional distributions are given by

$$p(\varphi_i|\varphi_{-i}, \widetilde{y}_T, \widetilde{S}_T, \widetilde{Z}_T, \widetilde{\pi}_T^e) \sim N(a_i^*, A_i^*\sigma^2), \tag{7}$$

where $a_i^* = (A_i^{-1} + X_i'X_i)^{-1}(A_i^{-1}a_i + X_i'Y_i)$, $A_i^* = (A_i^{-1} + X_i'X_i)^{-1}$, X_i is the appropriate regressor, and Y_i is a forecast error. For example, consider estimating the values of the vector $\overrightarrow{a}$ in (5). In this case, define $X_\alpha = B\widehat{Z}_T$ and $Y_\alpha = \widetilde{\pi}_T$, where the row t on-diagonal element of B is $\overrightarrow{\eta}T^\eta S_{t-1} + k\pi_t^e$ and the off-diagonal elements are zero, $\widehat{Z}_T = [\widetilde{z}_{1T}, \widetilde{z}_{2T}]$, and $\widetilde{z}_{iT}$ is a $T \times 1$ vector with representative element $z_{it} = 1$ iff $Z_i = i$. Other elements of the parameter vector can be estimated similarly. The values for the vector $\overrightarrow{h}$ can be generated from a posterior normal similar to (7) in which

$$X_\eta = \begin{bmatrix} \widehat{S}_T \\ \overrightarrow{a}\widetilde{Z}_T\widehat{S}_T \end{bmatrix}, \quad Y_\eta = \begin{bmatrix} \widetilde{u}_T - k(\widetilde{\pi}_T^e - \widetilde{\pi}_T) \\ \widetilde{\pi}_T - k\overrightarrow{a}\widetilde{Z}_T\widetilde{\pi}_T^e \end{bmatrix},$$

where $\widehat{S}_T$ is defined similarly to $\widehat{Z}_T$.

Given conditional priors for the variances of the form $\sigma_{\varepsilon_1}^2 \sim IG(\frac{\rho_0}{2}, \frac{R_0}{2})$ and $\sigma_{\varepsilon_i}^2 \sim IG(\frac{\lambda_0}{2}, \frac{L_0}{2})$, $i = 2, 3$, the posterior conditional probabilities are given by

$$\sigma_{\varepsilon_1}^2|\widetilde{y}_T, \widetilde{\pi}_T^e, \widetilde{S}_T, \widetilde{Z}_T, \widetilde{\varphi} \sim IG(\frac{\rho_0 + T}{2}, \frac{R_0 + \delta_1}{2}),$$

where

$$\delta_1 = (\widetilde{u}_T - k(\widetilde{\pi}_T^e - \widetilde{\pi}_T) - \overrightarrow{h}\widetilde{S}_T)'(\widetilde{u}_T - k(\widetilde{\pi}_T^e - \widetilde{\pi}_T) - \overrightarrow{h}\widetilde{S}_T);$$

and, for $i = 2, 3$:

$$\sigma_{\varepsilon_i}^2|\widetilde{y}_T, \widetilde{\pi}_T^e, \widetilde{S}_T, \widetilde{Z}_T, \widetilde{\varphi} \sim IG(\frac{\lambda_0 + T}{2}, \frac{L_0 + \delta_i}{2}),$$

where

$$\delta_i = (\widetilde{\pi}_t - \overrightarrow{a}\,\widetilde{Z}_t(\overrightarrow{h}\,\widetilde{S}_t + k\widetilde{\pi}_t^e))'(\widetilde{\pi}_t - \overrightarrow{a}\,\widetilde{Z}_t(\overrightarrow{h}\,\widetilde{S}_t + k\widetilde{\pi}_t^e)).$$

Transition Matrices.

Given a prior probability distribution for T_{ii}^a of the form $T_{ii}^a \sim \beta(u_{ii}, u_{ji})$ for $j \neq i$, the distribution for the transition probabilities T_{ii}^a is determined by

$$T_{ii}^a = \Pr[\alpha_t = a_i | \alpha_{t-1} = a_i] \sim \beta(u_{ii} + n_{ii}, u_{ji} + n_{ji}),$$

where n_{ii} is the number of periods that α_t remained in state i, and n_{ji} is the number of periods that α_t switched to state $j \neq i$ after beginning in state i. Then the other elements of T^α can be determined by $T_{ji}^\alpha = 1 - T_{ii}^\alpha$. A similar procedure is used to generate the elements of T^η.

References

[1] Albert, J., and Chib, S., 1993, "Bayes Inference via Gibbs Sampling of Autoregressive Time Series Subject to Markov Mean and Variance Shifts," *Journal of Business and Economic Statistics* 11, 1-15.

[2] Ball, L., 1994, "What Determines the Sacrifice Ratio?" in Mankiw, N.G., ed., *Monetary Policy*, Chicago: University of Chicago Press, 155-182.

[3] Barro, R., and Gordon, D., 1983a, "A Positive Theory of Monetary Policy in a Natural Rate Model," *Journal of Political Economy* 91, 589-610.

[4] Barro, R., and Gordon, D., 1983b, "Rules, Discretion and Reputation in a Model of Monetary Policy," *Journal of Monetary Economics* 12, 101-121.

[5] Bernanke, B., and Mihov, I., 1998, "Measuring Monetary Policy," *Quarterly Journal of Economics* 113, 869-902.

[6] Blinder, A., 1987, *Hard Heads, Soft Hearts*, Reading, MA: Addison Wesley.

[7] Broadbent, B., and Barro, R., 1997, "Central Bank Preferences and Macroeconomic Equilibrium," *Journal of Monetary Economics* 40, 17-43.

[8] Carter, C., and Kohn, R., 1994, "On Gibbs Sampling for State Space Models," *Biometrika* 81, 541-553.

[9] Casella, G., and George, E., 1992, "Explaining the Gibbs Sampler," *The American Statistician* 46, 167-174.

[10] Chib, S., 1993, "Bayes Regression and Autoregressive Errors: A Gibbs Sampling Approach," *Journal of Econometrics* 58, 275-294.

[11] Christiano, L., Eichenbaum, M., and Evans, C., 1999, "Monetary Policy Shocks: What Have We Learned and to What End?" in Taylor, J., and Woodford, M., eds., *Handbook of Macroeconomics, Volume 1*, Elsevier, 65-148.

[12] Fair, R., 1978, "An Analysis of the Accuracy of Four Macroeconometric Models," *Journal of Political Economy* 87, 701-718.

[13] Galí, J., and Gertler, M., 1999, "Inflation Dynamics: A Structural Econometric Analysis," *Journal of Monetary Economics* 44, 195-222.

[14] Gelfand, A. and Smith, A., 1990, "Sampling Based Approaches to Calculating Marginal Denisities," *Journal of the American Statistical Association 85*,398-409.

[15] Geweke, J., 1992, "Evaluating the Accuracy of Sampling-Based Approaches to the Calculation of Posterior Moments" in Bernardo, J., *et al* eds., *Bayesian Statistics 4*, Oxford: Oxford University Press, 169-193.

[16] Gordon, R., and King, S., 1982, "The Output Cost of Disinflation in Traditional and Vector AutoregressiveModels," *Brookings Papers on Economic Activity*, 205-242.

[17] Hamilton, J., 1983, "Oil and the Macroeconomy since World War II," *Journal of Political Economy* 91, 228-248.

[18] Hamilton, J., 1989, "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle," *Econometrica* 57, 357-384.

[19] Hamilton, J., 1990, "Analysis of Time Series Subject to Changes in Regime," *Journal of Econometrics* 45, 39-70.

[20] Kim, C.J., and Nelson, C., 1998, "Business Cycle Turning Points, a New Coincident Index, and Tests of Duration Dependence Based on a Dynamic Factor Model with Regime Switching," *Review of Economics and Statistics* 80, 188-201.

[21] Kim, C.J., and Nelson, C., 1999, *State-Space Models with Regime Switching: Classical and Gibbs-Sampling Approaches with Applications*, Cambridge, MA: MIT Press.

[22] Kydland, F., and Prescott, E., 1977, "Rules Rather Than Discretion: The Inconsistency of Optimal Plans," *Journal of Political Economy* 85, 473-491.

[23] Okun, A., 1978, "Efficient Disinflationary Policies," *American Economic Review* 68, 348-352.

[24] Romer, C., and Romer, D., 1989, "Does Monetary Policy Matter?: A New Test in the Spirit of Friedman and Schwartz," *NBER Macroeconomics Annual*, 121-170.

[25] Romer, C., and Romer, D., 1994, "Monetary Policy Matters," *Journal of Monetary Economics* 34, 75-88.

[26] Sargent, T., 1976, "A Classical Macroeconometric Model for the United States," *Journal of Political Economy* 84, 207-237.

[27] Shimer, R., 1998, "Why is the U.S. Unemployment Rate So Much Lower?" *NBER Macroeconomics Annual*, 11-61.

[28] Sims, C., 1999, "Drift and Breaks in Monetary Policy," draft.

[29] Strongin, S., 1995, "The Identification of Monetary Policy Disturbances: Explaining the Liquidity Puzzle," *Journal of Monetary Economics* 35, 463-97.

Parameter	Est. value	Parameter	Est. Value
a_1	0.5778 (0.0205)	h_1	3.2996 (0.1622)
a_2	1.4314 (0.0550)	h_2	6.6973 (0.1607)
γ_1	0.9772 (0.0184)	γ_2	0.0136 (0.0151)
$\sigma^2_{\varepsilon_1}$	1.5679 (0.1242)	$\sigma^2_{\varepsilon_2}$	0.5775 (0.7782)
$\sigma^2_{\varepsilon_3}$	7.6456 (0.6608)	k	0.0747 (0.0197)
$\Pr[\alpha_t = a_1 \mid \alpha_{t-1} = a_1]$	0.9823 (0.0098)	$\Pr[\eta_t = h_1 \mid \eta_{t-1} = h_1]$	0.9857 (0.0138)
$\Pr[\alpha_t = a_2 \mid \alpha_{t-1} = a_2]$	0.9519 (0.0251)	$\Pr[\eta_t = h_2 \mid \eta_{t-1} = h_2]$	0.9955 (0.0041)

Table 1. Estimated Parameters

Note: Standard deviations across iterations are given in parentheses.

State in t	Inflation Rate
(a_1, h_1)	1.99
(a_1, h_2)	4.04
(a_2, h_1)	5.29
(a_2, h_2)	10.73

Table 2. Steady State Inflation Rates

	Variable on us_t	us_t on variable	Variable on ds_t	ds_t on variable
Recession Dates	1.968 (0.004)	0.848 (0.674)	0.259 (0.999)	5.150 (0.000)
Romer Dates	3.581 (0.000)	0.327 (0.999)	0.273 (0.999)	1.016 (0.445)

Table 3. F Statistics for Granger Causality Tests

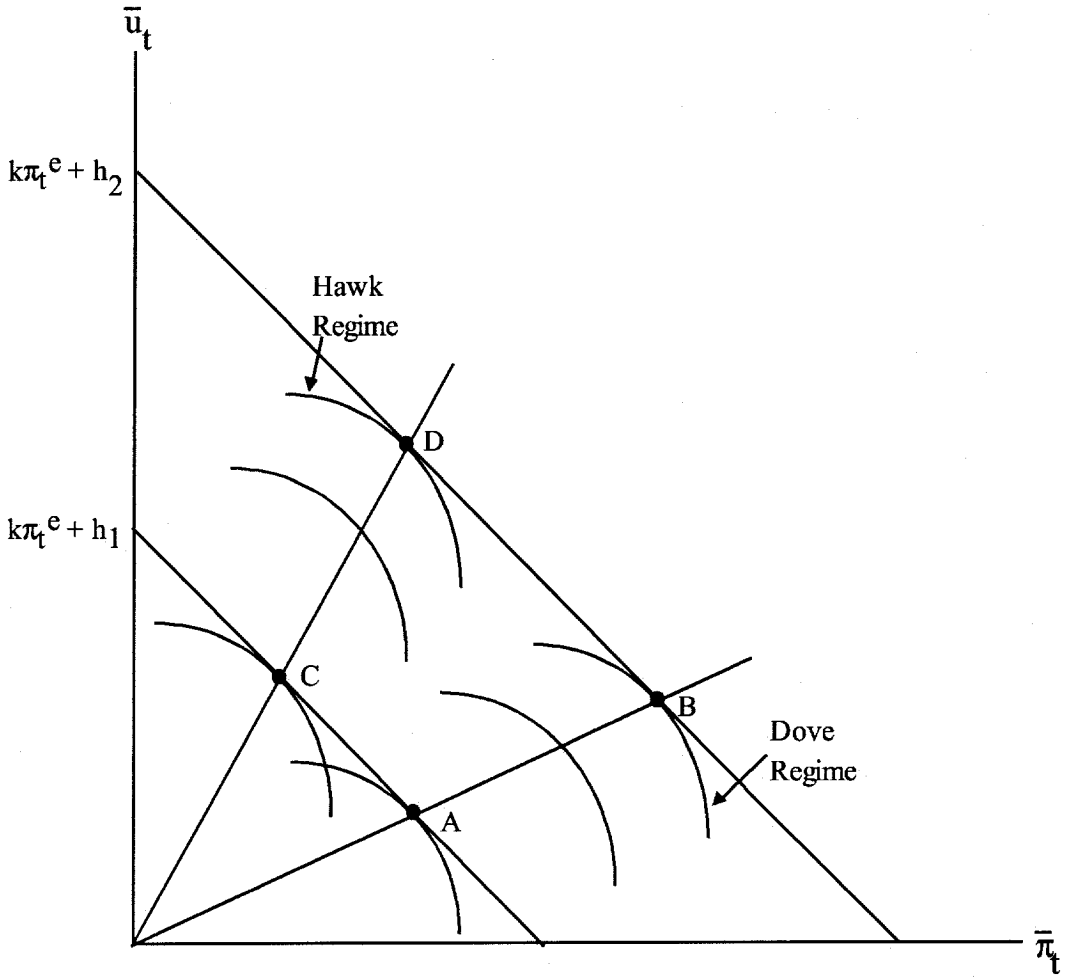

Figure 1

Posterior Expected Values of Preference Parameter

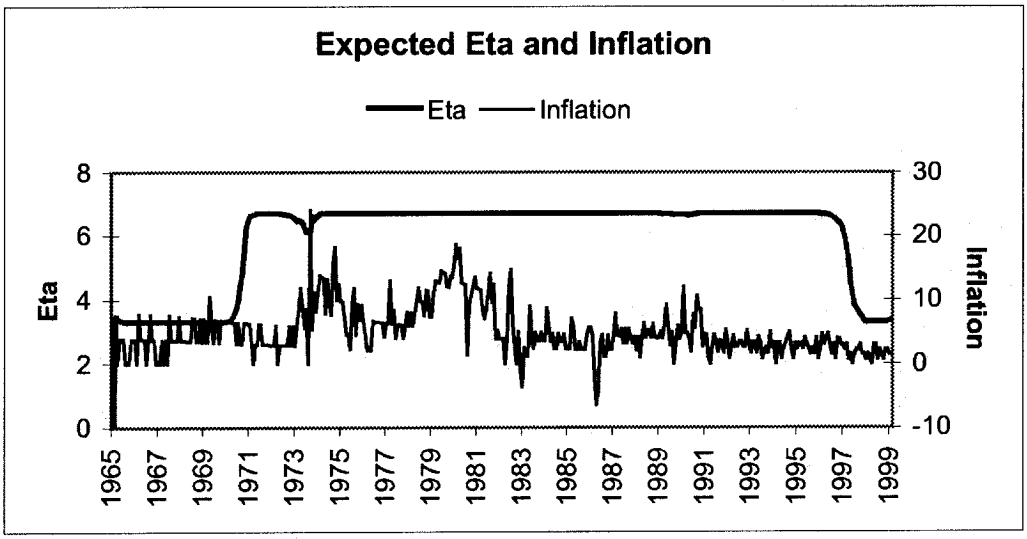

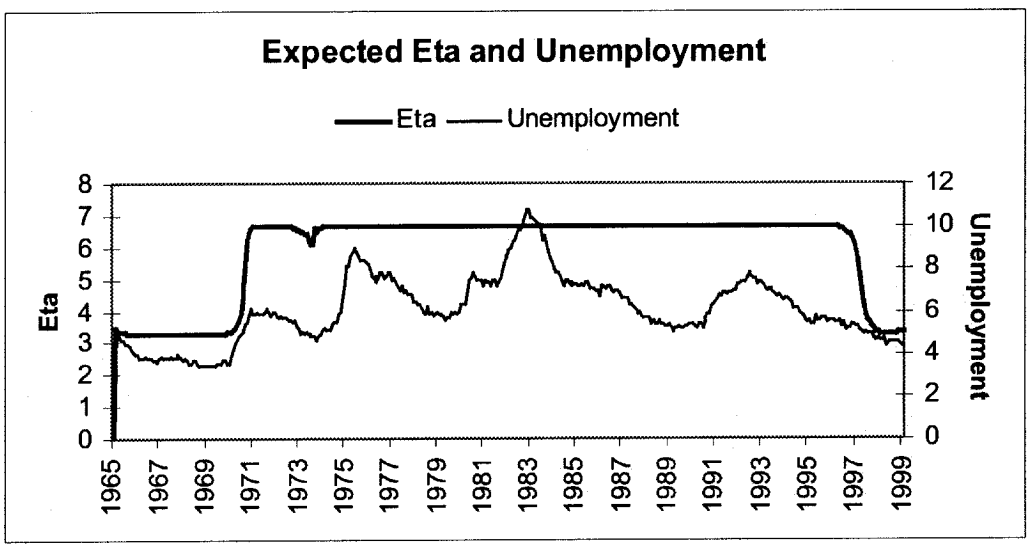

Figure 2

Posterior Expected Values of Structural Parameter

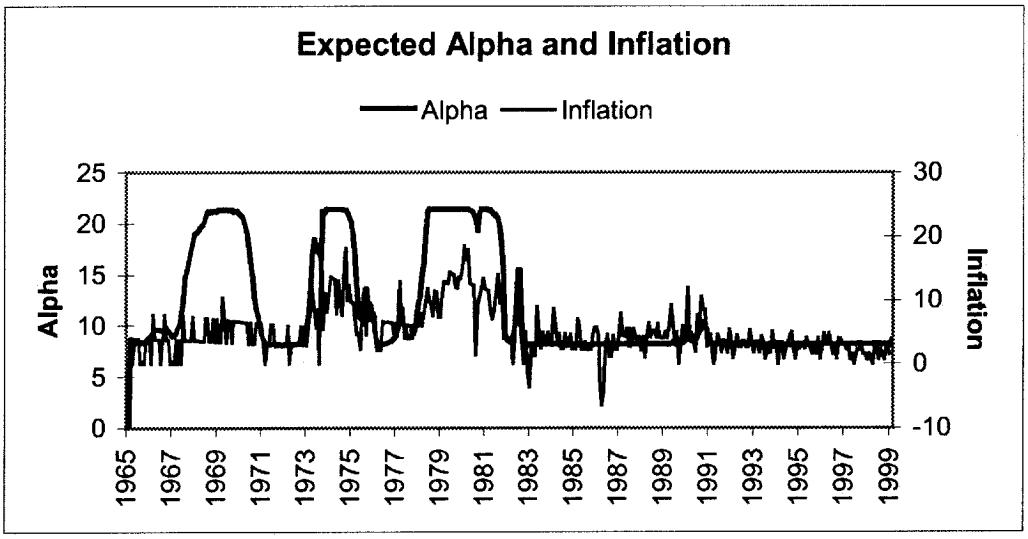

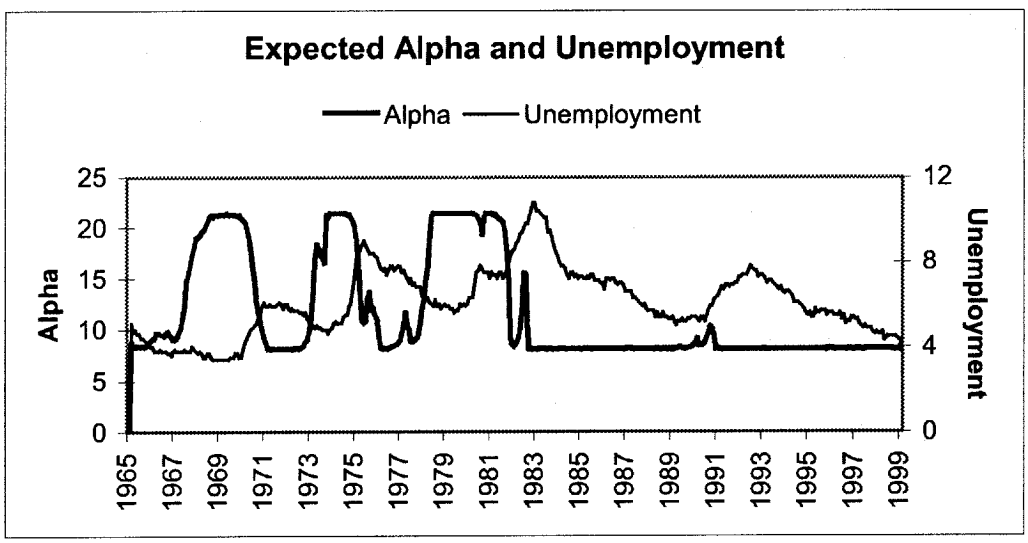

Figure 3

Posterior Expected Values of Policy Parameter

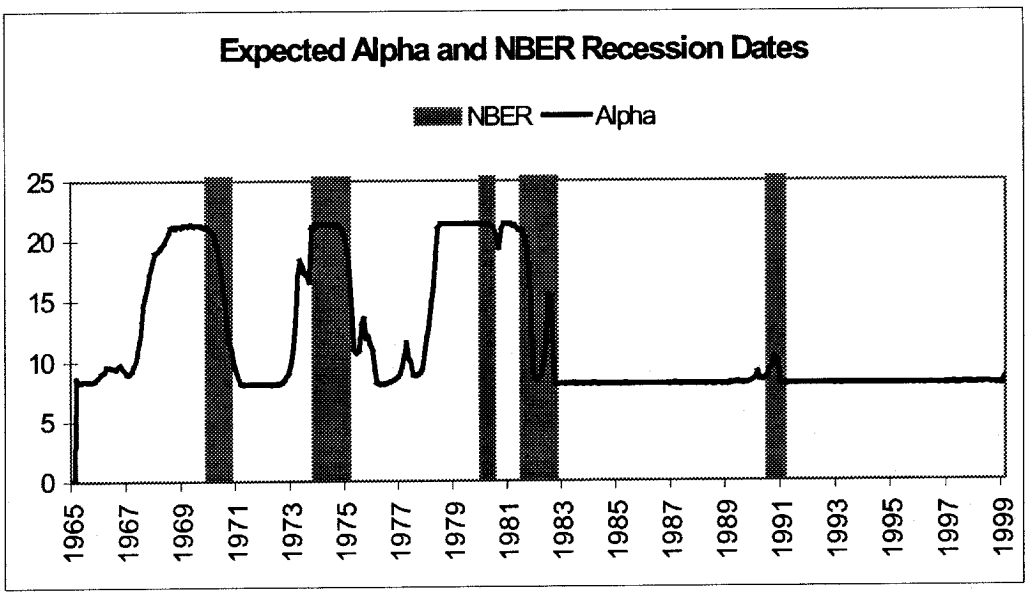

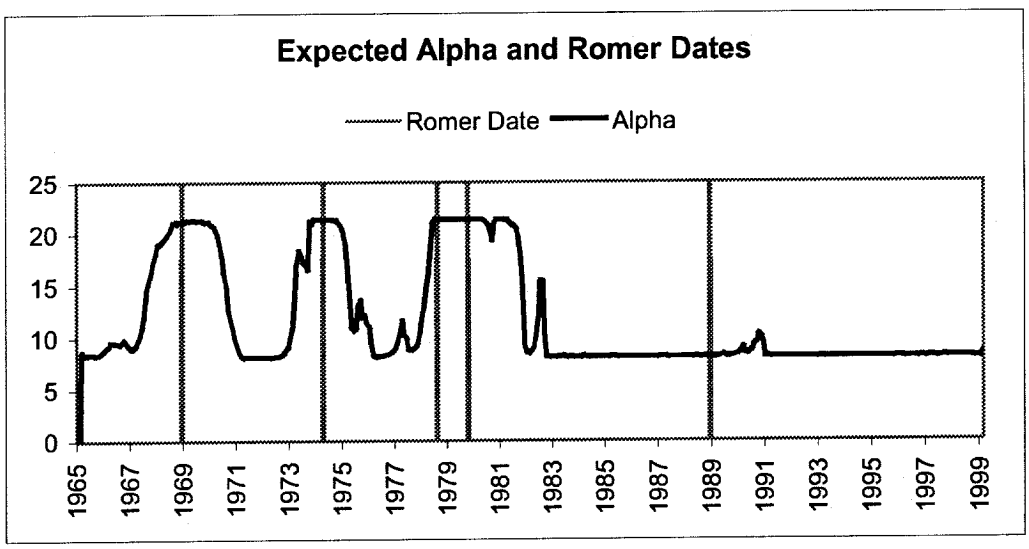

Figure 4

Comparison with NBER Recessions and Romer Dates

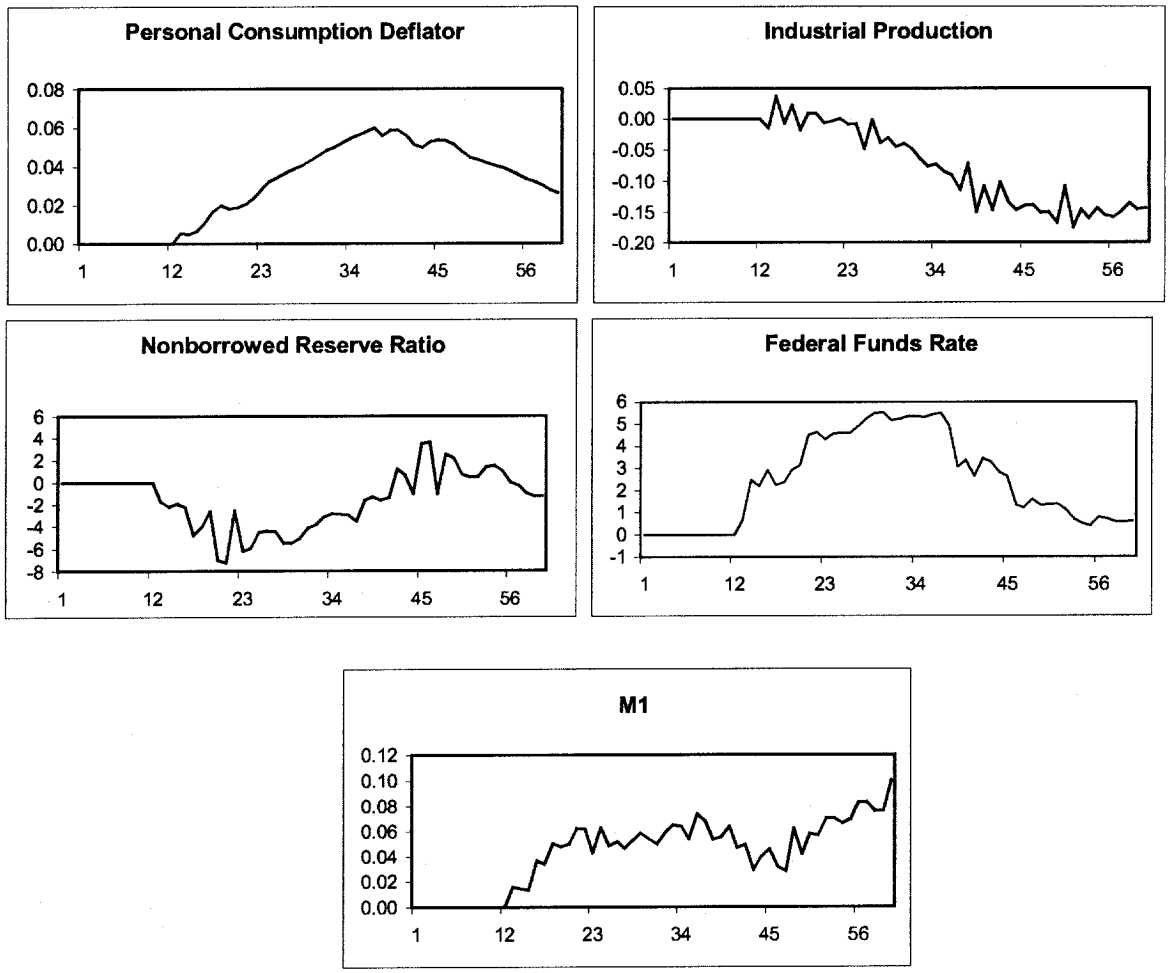

Figure 5

Effects of Inflationary Episode

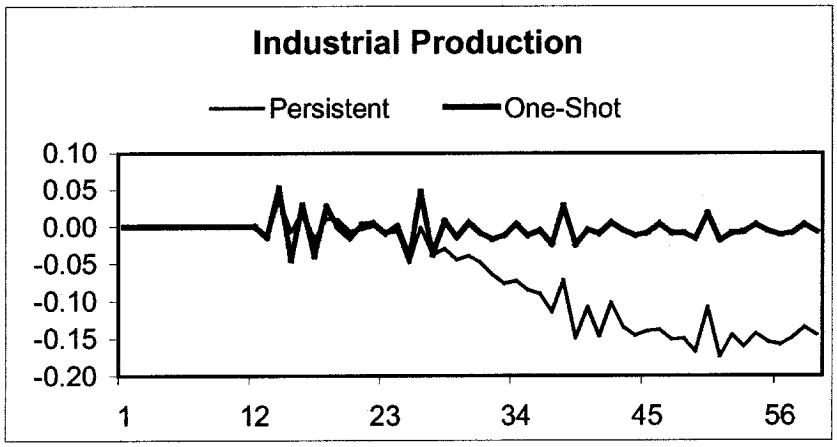

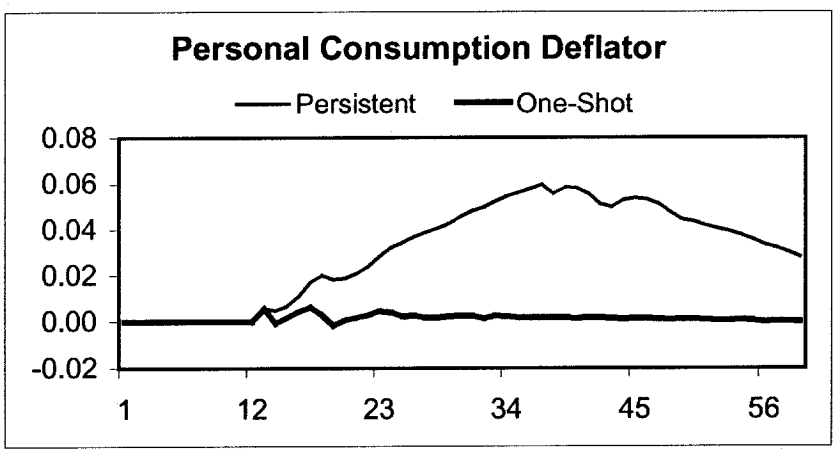

Figure 6

Persistent vs. One-Shot Switches

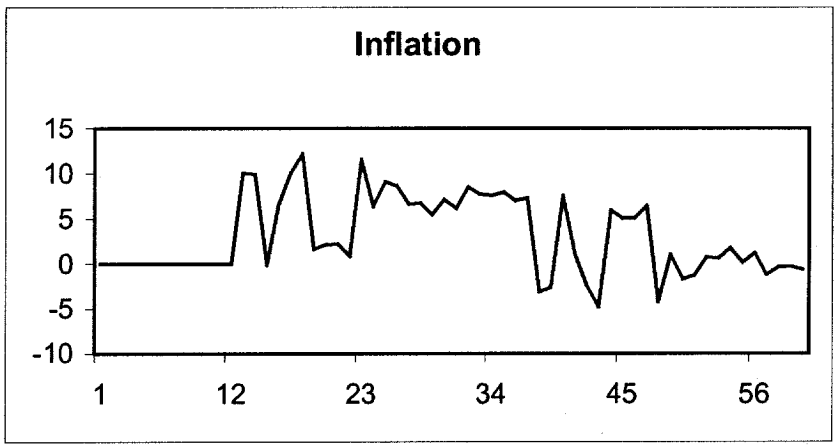

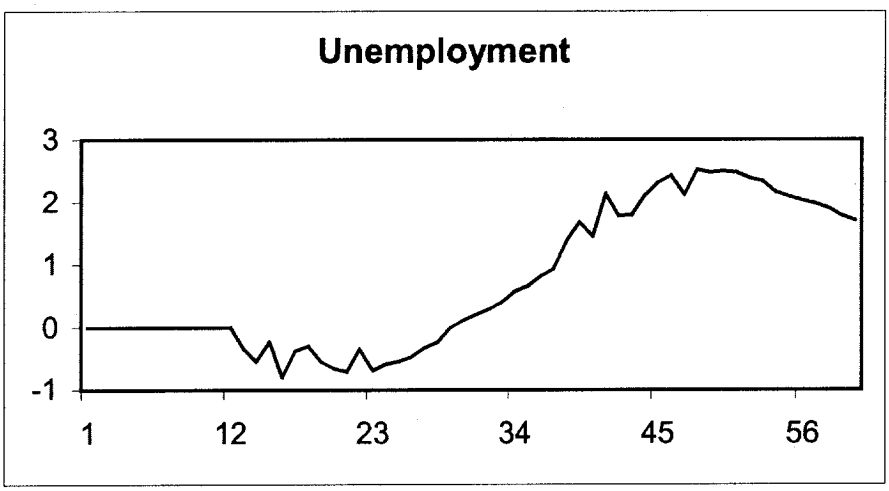

Figure 7

Effects of Inflation Episode