

Robust Estimation and Monetary Policy with Unobserved Structural Change

John C. Williams*

Federal Reserve Bank of San Francisco

July 27, 2004

Abstract

This paper considers the monetary policymaker's joint problem of model estimation and the design of a policy rule in the face of uncertainty regarding the process of structural change in the economy. Unobserved structural change is modeled through time variation in the natural rates of interest and unemployment. I show that certainty-equivalent optimal policies perform poorly when there is uncertainty about the natural rate processes. I then examine the properties of combined estimation methods and policy rules that are robust to this type of model uncertainty. I find that weighted sample means are robust estimators of natural rates for the purpose of setting policy. The optimal policy under uncertainty incorporates a significant degree of policy inertia and a muted response to the perceived unemployment gap; by comparison, the certainty-equivalent optimized policy in this model exhibits little policy inertia and a more aggressive response to the unemployment gap. This robust estimation/policy combination is shown to be highly effective at mitigating the effects of natural rate mismeasurement.

JEL Classification System: E52

* Correspondence: Economic Research, Federal Reserve Bank of San Francisco, 101 Market Street, San Francisco, CA 94105, Tel.: (415) 974-2240, e-mail: John.C.Williams@sf.frb.org. In preparation for the Conference on Models and Monetary Policy, Washington, D.C., March, 26-27, 2004. I thank Richard Dennis, Francis Diebold, Athanasios Orphanides, and Glenn Rudebusch for comments on earlier drafts. I also thank Kirk Moore and Judith Goff for outstanding research and editorial assistance, respectively. The opinions expressed are those of the author and do not necessarily reflect the views of the management of the Federal Reserve Bank of San Francisco.

The U.S economy has undergone substantial change over the past two decades. The IT revolution alone has transformed inventory management, fostered increased globalization of trade in goods and services, and improved the efficiency of labor and goods markets. These and other changes have had wide ranging effects on the economy: the magnitude of macroeconomic fluctuations has declined dramatically, estimates of the long-run growth rate of productivity have risen by more than a percentage points, and those of the natural rate of unemployment have declined by about a percentage point.¹ The implications of structural change for the conduct of monetary policy have attracted increased attention from researchers and policymakers, as evidenced by last year’s Jackson Hole conference entitled “Monetary Policy and Uncertainty: Adapting to a Changing Economy.” The goal of this paper is to examine issues related to the design of monetary policy when structural change is a regular feature of the economy and when there is considerable uncertainty about the precise nature of the underlying process of change.

I represent structural change by medium- and low-frequency variation in the natural rates of unemployment and interest. For the purpose of this paper, I define the natural rate of unemployment to be the unemployment rate consistent with a stable rate of inflation in the absence of transitory supply shocks; correspondingly, I define the natural rate of interest to be the real short-term interest rate consistent with the unemployment rate converging to its natural rate in the absence of transitory demand shocks. I focus on shifts in natural rates because of the relatively strong evidence that they vary over time; I leave the analysis of the monetary policy implications of changes in other aspects in the economy to future research.²

¹Kim and Nelson (1999) and McConnell and Quiros (2000) analyze the decline in variability in output in the United States. Orphanides and Williams (2002) and Edge, Laubach, and Williams (2004) document the evolution of estimates of the natural rate of unemployment and the long-run rate of productivity growth, respectively.

²For example, Rudebusch and Svensson (1999) find no evidence of a break in slope coefficients

If there is uncertainty about the true levels of the natural rates, but their true data generating processes (DGPs) were known, then, in a standard linear-quadratic framework, certainty equivalence obtains and the form and parameters of the optimal monetary policy would be the same as if there were no natural rate uncertainty.³ Similarly, Giannoni and Woodford (2004) emphasize that in a linear-quadratic framework the Euler equation describing the optimal monetary policy is invariant to the natural rate process. Nonetheless, implementation of the optimal policy requires that the policymaker form correct forecasts of endogenous variables, for which they need to know the true natural rate DGPs. However, the assumption that policymakers know the processes generating natural rates is highly unrealistic, as emphasized by Orphanides and Williams (2002).⁴ In addition, the processes underlying movements in natural rates themselves may change over time, as evidenced by shifts in innovation variances noted in Cogley and Sargent (2002) and others, further impairing policymakers' ability to infer the true model from the data. Thus, a key assumption of this approach is that uncertainty regarding the process of structural change is pervasive and that it is unlikely to vanish in the foreseeable future.

I assume that the policymaker must choose once and for all, without knowing the true DGPs for natural rates, both a method to estimate the natural rates in real time and a monetary policy rule. I consider two approaches to model uncertainty. In one,

in their model, while Estrella and Fuhrer (2003) find strong evidence of a break when testing jointly for intercepts— which are related to natural rates— and slope parameters. See also Kozicki and Tinsley (2001), who emphasize lower frequency variation in intercepts in estimated models. The evidence for change in other macroeconomic relationships, such as the autocorrelation of inflation, is the subject of ongoing research; see Bernanke and Mihov (1998), Cogley and Sargent (2001, 2002), Sims (2001), Boivin and Giannoni (2003), and Stock and Watson (2003).

³See Simon (1956), Theil (1958), Chow (1975) and Kalchbrenner and Tinsley (1976) for early analysis of certainty equivalence, and Swanson (2000), Svensson and Woodford (2003), and Woodford (2004), for more recent treatments. Stochastic natural rates are a form of additive uncertainty and therefore certainty equivalence applies for optimal policies and optimal filters. Certainty equivalence does not apply to uncertainty about slope parameters, as analyzed by Brainard (1967).

⁴See also Stock and Watson (1998), Lansing (2000) and Orphanides and van Norden (2002).

the policymaker has well-formed priors over the set of potential natural rate DGPs and chooses an estimation and policy rule combination that minimizes the expected loss—in terms of unconditional squared deviations of inflation, the unemployment rate, and the interest rate from target levels—integrating over the set of natural rate DGPs. The second approach corresponds to Knightian uncertainty, in which the policymaker does not have well-formed priors over the natural rate models. In this case, I follow the robust control literature and analyze the estimation and policy rule combinations that minimize the maximum loss over the set of potential natural rate DGPs. This method of examining robust monetary policy under model uncertainty follows the approach advocated by McCallum (1988) and implemented by Taylor (1999), Levin, Wieland, and Williams (1999, 2003), and others.⁵

I implement this approach by specifying three representative natural rate processes that make up the set of possible natural rate processes: a highly persistent AR(1), a long-memory or fractionally integrated process, and a Markov switching process. I choose these processes because they are all consistent with the data, but they differ in their implications for the specification of the optimal natural rate estimator. For each process, I also allow for uncertainty regarding the parameterization of the process, as evidenced by the wide range of estimates of natural rate innovation variances reported by Laubach and Williams (2003) and others in the case of the Kalman filter. To capture this form of parameter uncertainty, I include three different calibrations of each natural rate DGP.

A key contribution of the paper is the analysis of the joint problem of estimation and policy feedback when there is uncertainty about the processes underlying

⁵See, for example, Orphanides and Williams (2002), Laxton and Pesenti (2003), Levin and Williams (2003), Brock, Durlauf, and West (2003), and Onatski and N. Williams (2003). Cogley and Sargent (2003) extend this type of analysis to the case where the policymaker continuously updates his or her priors over models.

structural change. A number of researchers have examined the effects of natural rate mismeasurement on the performance and optimal specification of monetary policy rules, but most have treated natural rate mismeasurement as exogenous noise.⁶ In this paper, I directly examine the performance of real-time estimation strategies and policy rules where the true natural rates vary over time.⁷ Thus, the occurrence of natural rate misperceptions and their correlation with other variables arises endogenously and depends on both the estimation method and the policy rule.

I conduct the analysis using a variant of the Rudebusch-Svensson (1999) backward-looking model estimated on 50 years of postwar U.S. data. I focus on this model because Orphanides and Williams (2002) have shown that natural rate mismeasurement is relatively easy to overcome in forward-looking and hybrid models by specifying the policy rule in terms of changes of the interest rate reacting to inflation and the change in the unemployment rate. But such a strategy is far less effective in models of the Rudebusch-Svensson type, which have empirical support (Estrella and Fuhrer 2003).⁸ In addition, the model has been extensively studied in the monetary policy literature, facilitating the comparison of results from studies by Rudebusch (2001, 2002), Onatski and N. Williams (2003), Brock, Durlauf, and West (2003), and Levin and Williams (2003).

I find that there can be very large costs, especially in terms of inflation variability, to ignoring natural rate uncertainty. However, I also show that it is possible to design estimation and monetary policy rules that are robust to a variety of models

⁶See, for example, Orphanides et al (2000), Smets (2002), Orphanides (2002), Rudebusch (2001, 2002), and Orphanides and Williams (2002).

⁷The use of the term “real-time” to problems of this sort is due to Diebold and Rudebusch (1991).

⁸The assumption of adaptive expectations is not without cost, as this framework ignores the endogenous response of expectations, which can exacerbate the problems associated with policy errors induced by faulty estimates of model parameters, as discussed by Orphanides and Williams (2002, 2003) and others. The extension of the analysis of this paper to other models of expectations formation is left for future work.

of natural rate evolution. In the terminology of Levin and Williams (2003), such estimation/policy combinations display a high degree of fault tolerance in the face of model uncertainty about natural rates. I find that weighted sample means, where the weights on past data decline gradually, of the real interest rate and the unemployment rate yield very good estimates of their respective natural rates for use in the conduct of monetary policy. I also find that in the face of uncertainty about natural rates, the robust policy incorporates a greater deal of policy inertia and a more muted response to perceived unemployment gap than would be the case if the natural rates were known. By keeping its eye focussed primarily on the “inflation ball,” such a policy reduces the unavoidable policy “mistakes” resulting from natural rate mismeasurement.

1 The Model

I use a modified version of the Rudebusch-Svensson (1999) model for the analysis. Following Orphanides and Williams (2002), the model is specified in terms of the unemployment rate gap as opposed to the output gap specification of Rudebusch and Svensson, and allowance is made for time variation in the natural rates of interest and unemployment. Each unit of time corresponds to one quarter of one year.

1.1 Unemployment and Inflation Dynamics

The IS curve equation relates the unemployment rate, u_t , to its lags, its natural rate, u_t^* , and the lagged difference between the two-quarter average of the real federal funds rate, r_t , and its natural rate, r_t^* ,

$$u_t = (1 - \beta_1 - \beta_2)u_{t-1}^* + \beta_1 u_{t-1} + \beta_2 u_{t-2} + \beta_3 ((r_{t-1} + r_{t-2})/2 - r_{t-1}^*) + \epsilon_t, \quad (1)$$

where $\epsilon_t \sim N(0, \sigma_\epsilon^2)$ is a serially uncorrelated innovation. The real federal funds rate is defined to be the difference between the nominal federal funds rate and a measure of expected inflation assumed to equal the inflation rate over the past four quarters:

$$r_t \equiv i_t - \bar{\pi}_t,$$

where $\bar{\pi}_t$ denotes the four-quarter moving average of the inflation rate. Each period is one quarter of a year.

The Phillips curve equation relates the GDP price deflator inflation rate, π_t , to its own lags (with a unity sum imposed on the coefficients) and the lagged difference between the unemployment rate and its natural rate:

$$\pi_t = \gamma_1 \pi_{t-1} + (1 - \gamma_1) \frac{1}{3} \sum_{j=2}^4 \pi_{t-j} + \gamma_2 (u_{t-1} - u_{t-1}^*) + \eta_t, \quad (2)$$

where $\eta_t \sim N(0, \sigma_\eta^2)$ is a serially uncorrelated innovation. As noted above, the natural rates of interest and unemployment are time-varying and are therefore identified with time subscripts.

1.2 Modeling Natural Rates

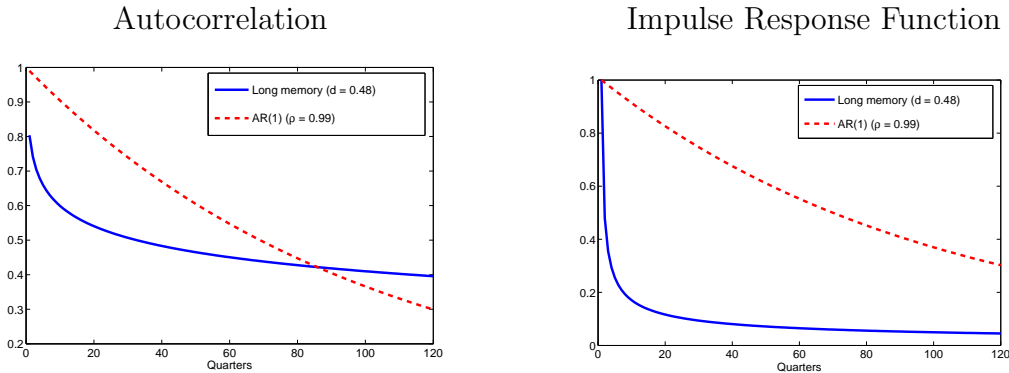
I consider three types of time series models for time-varying natural rates: a first-order autoregressive process, a fractionally integrated long-memory process, and a two-state Markov switching model. I focus on stationary processes for the natural rates.⁹ For each model, I consider three parameterizations of the variance of the process, as described below.

The first type of model is the standard first-order autoregressive process, or AR(1), according to which the natural rate z follows the law of motion:

$$z_t = (1 - \rho) \bar{z} + \rho z_{t-1} + \tau_t, \quad (3)$$

⁹Based on the ADF test, one can reject the null of nonstationarity of both the unemployment and real federal funds rate over 1950–2003 at the 5 percent level.

Figure 1: Characteristics of Long-Memory and AR(1) Processes



Notes: For these calculations, $\rho = 0.99$ for the AR(1) process (see equation 3), and $d = 0.48$ for the long-memory process (see equation 4).

where \bar{z} is the unconditional mean of z , $|\rho| < 1$, and τ_t is assumed to be a white noise innovation.

The second type of model is a fractionally integrated or “long-memory” model studied by Granger (1980) and Diebold and Rudebusch (1989). The corresponding law of motion is given by:

$$(1 - L)^d(z_t - \bar{z}) = \nu_t, \quad (4)$$

where $|d| < \frac{1}{2}$ and ν_t is a white noise innovation. I approximate this process by its binomial expansion, truncated after 5000 terms,

$$z_t = \bar{z} + \sum_{j=1}^{5000} (-1)^{j-1} \frac{\prod_{i=1}^j (d - i + 1)}{j!} (z_{t-j} - \bar{z}) \quad (5)$$

An interesting facet of, and possible justification for, the long-memory model highlighted by Granger (1980) is that it approximates the outcome from aggregating idiosyncratic AR(1) stochastic processes generated by differing values of the autocorrelation coefficient.

The fractionally integrated model differs from the AR(1) model in two important respects. First, the autocorrelation of an AR(1) decays geometrically, as shown by the

dashed line in Figure 1, while that of the long-memory process displays approximately hyperbolic decay, as shown by the solid line. Thus, the long-memory process is able to generate significant lower-frequency variability without imposing a very high degree of high frequency serial correlation. Second, the two processes differ markedly in their impulse responses. The impulse response function for an AR(1) declines geometrically, while that of the long-memory process falls rapidly for the first several periods, but then declines very gradually. Evidently, the long-memory process behaves like a combination of a weighted sum of two AR(1) processes, one with a relatively low root and the other with a root near unity.

The third process is a two-state Markov switching process as described by Hamilton (1989), in which with some probability, p , the natural rate shifts from the first state to the second state. I assume that the probability of switching states is the same for each state, so that the mean time spent in each state is the same and the unconditional mean of the natural rate is the average of the values in the “low” and “high” states.

Each of these three natural rate DGPs are characterized by two parameters, one describing the persistence of the natural rate and the other describing the variance of the innovations to the natural rate. Because I am interested in medium- and low-frequency variation in the natural rates, I assume values of $\rho = 0.99$, $d = 0.48$, and $p = 0.99$, which yield a high degree of lower-frequency persistence, as illustrated in Figure 1. I allow for uncertainty regarding the behavior of these processes by including three sets of values for the innovation variances, as discussed in the next section.

1.3 Monetary Policy

I assume that the monetary policymaker's objective is to minimize the expected unconditional squared deviations of the four-quarter inflation rate from its target rate, π^* , the unemployment rate from its natural rate, and the deviation of the nominal interest rate from the long-run target nominal interest rate, $i_t^* = r_t^* + \pi^*$. Specifically, the loss, \mathcal{L} is given by:

$$\mathcal{L} = E \{ (\pi - \pi^*)^2 + \lambda_u (u - u^*)^2 + \lambda_i (i - i^*)^2 \}, \quad (6)$$

where expectations are taken with respect to the innovations to the unemployment rate and inflation, $\{\epsilon_t\}_{t=0}^\infty$ and $\{\eta_t\}_{t=0}^\infty$, respectively, as well as to the natural rates of unemployment and interest, $\{u_t^*, r_t^*\}_{t=0}^\infty$. Throughout the remainder of the paper, I assume the following weights in the policymaker loss function: $\lambda_u = 1, \lambda_i = 0.5$; qualitatively, the results are not sensitive to moderate variations in these parameters. Furthermore, I assume that the inflation target is zero.

The expectation in the loss function takes into account both uncertainty about the realization of future innovations and uncertainty about the DGPs for the natural rates. Let S denote the set of the possible natural rate DGPs (which may differ across variables). Assume for the present purposes that the policymaker has well-defined prior beliefs regarding the distribution of $s \in S$, denoted by $F(s)$. Let $\mathcal{L}(s)$ denote the expected policymaker loss for the data generating process s . Then, for a given specification of monetary policy, the expected unconditional loss is given by:

$$\mathcal{L} = \int_S \mathcal{L}(s) dF(s). \quad (7)$$

In practice, I represent this expectation with a finite set of discrete elements, $\{s_i\}_{i=1}^{N_s}$ of S , weighted by ω_i :

$$\mathcal{L} = \sum_{i=1}^{N_s} \mathcal{L}(s_i) \omega_i, \quad (8)$$

where $\sum \omega_i = 1$. In the example studied in this paper, $N_s = 9$.

I assume that monetary policy is implemented by setting the federal funds rate according to a monetary policy rule taking the form of an augmented Taylor (1993) rule similar to that found to perform well under natural rate uncertainty in Orphanides and Williams (2002). Absent natural rate uncertainty, policy rules of this type yield macroeconomic performance very nearly equal to the fully optimal policy. In particular, the federal funds rate is set according to:

$$i_t = \theta_i i_{t-1} + (1 - \theta_i) \left\{ \hat{r}_{t-1}^* + (1 + \theta_\pi) \bar{\pi}_{t-1} + \theta_\pi \pi_{t-1} - \theta_u (u_{t-1} - \hat{u}_{t-1}^*) - \theta_{\Delta u} \Delta u_{t-1} \right\}, \quad (9)$$

where \hat{r}_t^* and \hat{u}_t^* are the policymaker's real-time estimates of the natural rates of interest and unemployment, respectively, and Δu_t is the first-difference in the unemployment rate. I abstract from the zero lower bound on interest rates.

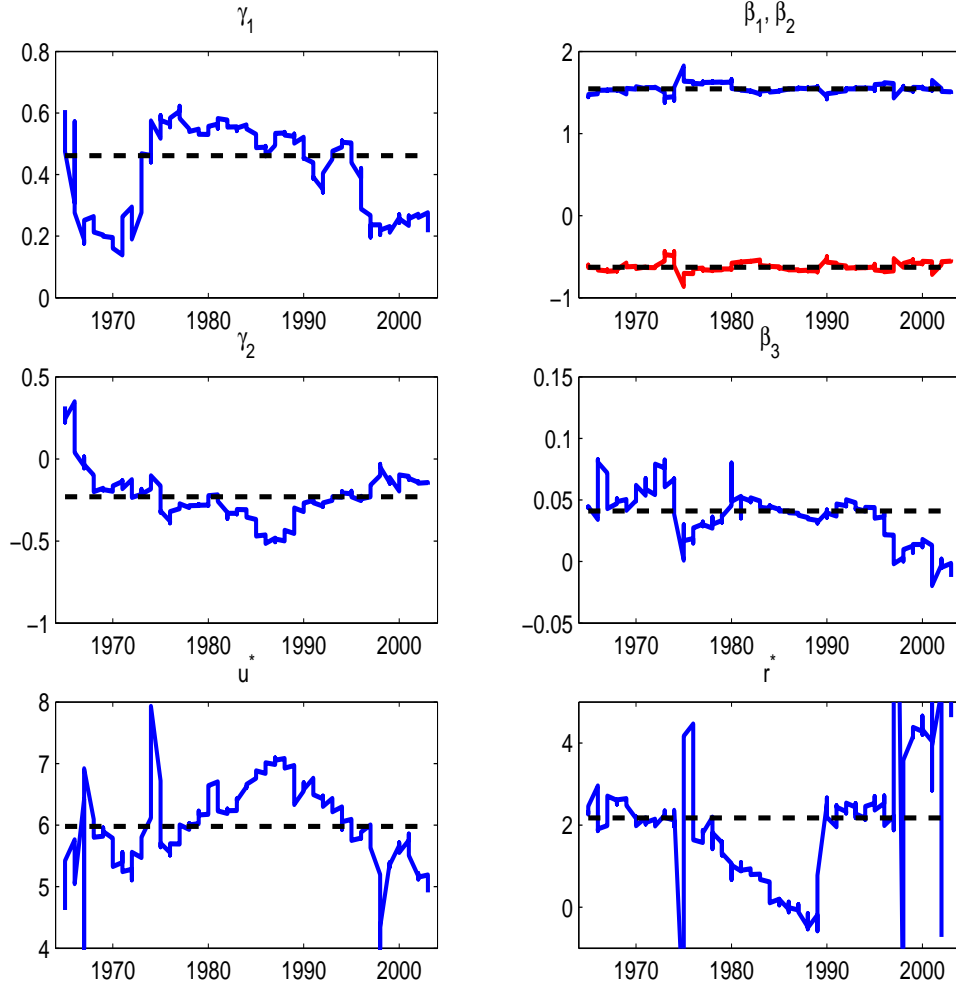
2 Model Estimation

In this section, I estimate the basic model and calibrate the set of DGPs for the natural rates of interest and unemployment.

2.1 Unemployment and Inflation Dynamics

If the natural rates of interest and unemployment were constant, OLS estimation of the parameters of the IS and Phillips curve equations would yield consistent estimates of these equations. However, if the natural rates change over time and are unobserved by the econometrician, the assumption of constant natural rates introduces mismeasurement in the right-hand-side variables, potentially biasing all parameter estimates. To mitigate these effects, I estimate the IS curve and Phillips curve equations using rolling regressions in which each sample contains only 15 years of data. In addition,

Figure 2: Rolling Regression Estimates



Notes: Sample length is 15 years. The dashed lines indicate the median estimates.

I proxy the true time-varying natural rate of unemployment with the Congressional Budget Office's (CBO) estimates. Given the sample of data of 1950-2003, I can thus run 156 regressions. I take as the model parameter estimate the median estimates from this set of 156 estimates.

Figure 2 plots the rolling regression estimates of the model parameters. In each case, the dashed line indicates the median estimate. The left-hand column of charts reports the estimates pertaining to the Phillips curve equation, and the right-hand

column reports the IS curve equation. For the Phillips curve equation, I plot estimates based on a specification that includes a constant but excludes the CBO estimate of the natural rate of unemployment; the bottom left panel shows the resulting rolling regression estimate of the natural rate of unemployment, given by the ratio of the estimated intercept divided by the estimate of γ_2 . For the IS curve equation, I use the CBO estimate of the natural rate of unemployment; the bottom right panel reports the resulting estimates of the natural rate of interest, given by the ratio of the estimated intercept divided by the negative of the estimate of β_3 .

The rolling regression estimates of the natural rates of unemployment and interest vary considerably over time, lending support for time variation in the true values of the natural rates or the difficulty in their real-time estimation, or both. Although not investigated in this paper, it is of interest to note that some of the other parameters exhibit some signs of time variation. The estimate of γ_1 , the coefficient on the first lag of inflation in the Phillips curve equation, also displays considerable time variation. This variation in estimates of γ_1 also obtains when the sum restriction on lagged inflation is not imposed.¹⁰ The slope of the IS curve displays a downward trend over the sample. In contrast, the estimates of the slope of the Phillips curve and the lags of the unemployment rate in the IS curve are relatively stable over the sample.

The median estimates from the rolling regressions yield the following two equations, which I use in the analysis that follows:

$$u_t = 0.09u_{t-1}^{*CBO} + 1.54u_{t-1} - 0.63u_{t-2} + 0.04 \left(\frac{r_{t-1} + r_{t-2}}{2} - 2.18 \right) + \epsilon_t, \quad (10)$$

$$\pi_t = 0.46\pi_{t-1} + 0.54 \frac{1}{3} \sum_{j=2}^4 \pi_{t-j} - 0.23(u_{t-1} - 5.96) + \eta_t. \quad (11)$$

¹⁰In fact, the sum restriction cannot be rejected for the sample from 1970 to 2003. Only during the 1960s is the sum of coefficients on lagged inflation well below unity; see Orphanides and Williams (2003) for a discussion of this issue.

These estimates are similar to those from full-sample estimation and conform to estimates from similar models, such as Rudebusch and Svensson (1999) and Orphanides and Williams (2002). A key difference is that full-sample estimation through 2003 yields a much lower value for β_3 , the slope of the IS curve, suggesting the possibility of bias owing to time variation of the natural rate of interest. The “sacrifice ratio” implied by the estimated Phillips curve is 2-1/4; that is, if the unemployment rate is 2-1/4 percentage points above its natural rate for one year, the inflation rate will eventually decline by 1 percentage point.

2.2 Natural Rates of Interest and Unemployment

I now describe the calibration of the three natural rate DGP models and the various parameterizations of each. As noted above, there exists a great deal of uncertainty regarding the parameters of any specific model for the natural rates. For example, the Kalman filter has been extensively used to estimate time-varying natural rates of interest and the unemployment.¹¹ A key finding in this literature is that the parameters describing the law of motion of natural rates are very imprecisely estimated (Laubach and Williams 2003), especially the innovation variance for the highly persistent component of natural rates. Thus, the data provide frustratingly little guidance on this key parameter.

I represent this uncertainty regarding the parameters describing natural rate processes by allowing for three parameterizations that span the set of values that are broadly consistent with the data. In one, the innovation variance for the natural rates is set to zero, corresponding to constant natural rates. In the second, the in-

¹¹See, for example, Staiger, Stock, and Watson (1997, 2002), Gordon (1998), Brainard and Perry (2000), and Laubach (2001), for Kalman filter estimates of the natural rate of unemployment. See Laubach and Williams (2003) and Orphanides and Williams (2002) for Kalman filter estimates of the natural rate of interest.

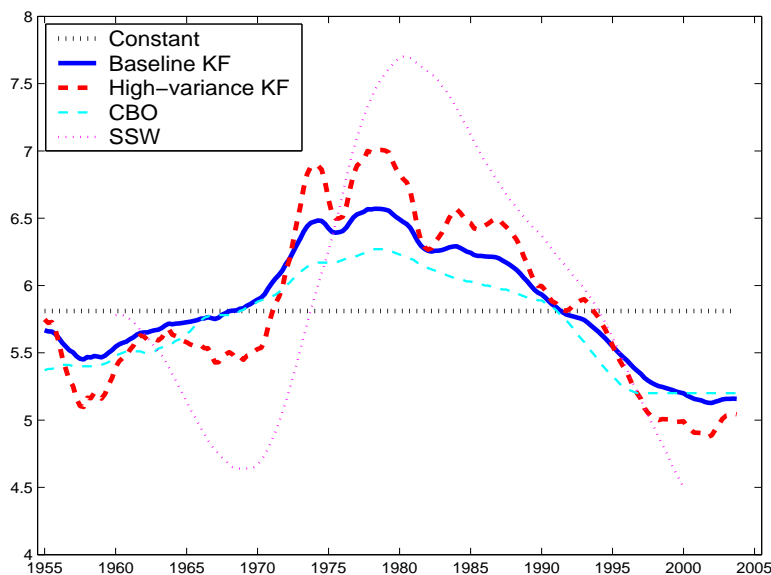
novation variance is set to the baseline value computed as described below. In the third, the innovation variance is set to a larger value that lies within the range of other published estimates. Note that in the high variance case, I am arguably being conservative in my approach in that even higher estimates of natural rate variability cannot be rejected by the data, according to some estimates. In the cases of a zero natural rate variance, the three DGPs collapse into one, so in the end there are seven unique alternative specifications of natural rate DGPs in all.

I follow the same basic procedure for calibrating the DGPs for both the natural rate of unemployment and interest. Starting with the natural rate of unemployment, I first estimate the natural rate using the Kalman filter applied to the Phillips curve equation, assuming that the natural rate follows a random walk. For this purpose, I use the Stock and Watson (1998) median-unbiased estimator. The sample is 1970-2003. The estimate of the standard deviation of the natural rate innovation, σ_τ , is 0.22. The resulting “smoothed” or “two-sided” estimates of the natural rate of unemployment are shown in Figure 3 and is comparable to, albeit slightly more variable than, the CBO estimates shown in the figure for comparison.

To capture the uncertainty regarding the innovation standard deviation, I consider two representative alternative values for σ_τ , 0 and 0.44. The value of zero corresponds to a constant natural rate, and the latter yields estimates of the natural rate that are in fact less variable than the baseline estimates reported in Staiger, Stock, and Watson (2002), as seen in the figure. In terms of in-sample fit of the inflation equation, the data cannot clearly distinguish between the baseline values and the two alternatives.

I use the values of the innovation standard deviations from the Kalman filter estimates to calibrate the three DGPs for the natural rate of unemployment. For the AR(1) model, I assume that $\rho = 0.99$ and set the standard deviation of the innovation to 0.22 for the baseline calibration. The resulting process for the natural rate of

Figure 3: Estimates of the Natural Rate of Unemployment



Notes: All estimates are smoothed or “two-sided” estimates. Baseline and high-variance KF refers to variations of the Kalman filter estimates computed for this paper; CBO refers to the current estimates from the Congressional Budget Office; SSW refers to estimates reported in Staiger, Stock, and Watson (2002).

unemployment has an unconditional standard deviation of 1.56 percentage points, as reported in Table 1. For the “high” variance version of the AR(1) model, I set the innovation standard deviation to 0.44; this yields an unconditional standard deviation of 3.12 percentage points. I do not formally estimate the long-memory process and the Markov-switching models for the natural rate of unemployment, but instead calibrate them to have the same unconditional variances as those of the AR(1) process. For the long-memory process model, I set $d = 0.48$. For the Markov-switching model, I set the common switching probability $p = 0.99$ and set the difference in the natural rate of unemployment between the states at 3.12 percentage points; for the high-variance variant, I set the difference between the states to 6.24 percentage points. Given this calibration, the unconditional variances are identical across the three processes for the baseline case and for the low and high-variance cases.

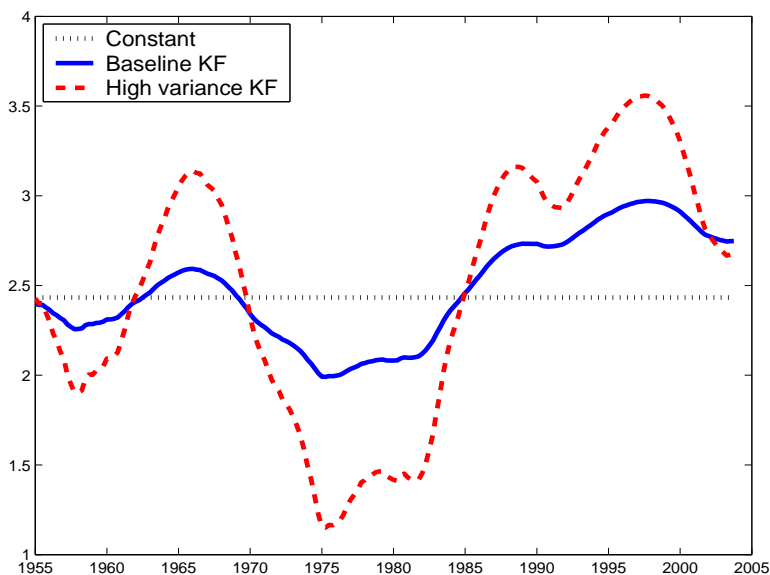
Table 1: Calibration of Natural Rate DGP Models

Natural Rate DGP	Innovation Standard Deviation				Natural Rate Stand. Dev.	
	ϵ	η	r^*	u^*	r^*	u^*
Zero variance (constant)	0.280	1.13	0.00	0.00	0.00	0.00
Baseline variance						
AR(1)	0.276	1.12	0.11	0.22	0.78	1.56
Long-memory	0.265	1.11	0.43	0.87	0.78	1.56
Markov switching	0.271	1.12	—	—	0.78	1.56
High variance						
AR(1)	0.257	1.10	0.22	0.44	1.56	3.12
Long-memory	0.210	1.03	0.84	1.75	1.56	3.12
Markov switching	0.253	1.07	—	—	1.56	3.12

The strategy for calibrating the DGPs for the natural rate of interest is the same as for the natural rate of unemployment. I use the same values of ρ , d , and p as before. Kalman filter estimation of the IS curve yields an innovation standard deviation of 0.11 percentage points, which I use for the baseline AR(1) process. This implies an unconditional standard deviation for the natural rate of interest of 0.78 percentage points. For the high-variance alternative, I set the innovation standard deviation to 0.22 percentage points, which is in the range of estimates reported by Laubach and Williams (2003). The resulting smoothed estimates are shown in Figure 4. For the two variants of the long-memory process, I set the innovation standard deviations so that they match those from AR(1) model. For the baseline Markov switching model, I set the difference between the two states at 1.56 percentage points, and for the high-variance alternative, I set the difference at 3.12 percentage points.

Finally, in order to make the seven DGPs equivalent in the sense of the implied overall variability of inflation and the unemployment rate gap, I adjust the standard deviations of the innovations to the Phillips curve and IS curve equations, as indicated

Figure 4: Estimates of the Natural Rate of Interest



Notes: All estimates are smoothed or “two-sided” estimates. Baseline and high-variance KF refers to variations of the Kalman filter estimates computed for this paper.

in Table 1. In the case of constant natural rates, I use the median estimate of the standard error of the regression for the fifteen-year rolling regressions used to estimate the parameters of the model, as described above. This procedure yields a Phillips curve innovation standard deviation of 1.13 percentage points and an IS curve innovation standard deviation of 0.28 percentage points. For the other natural rate DGPs, I set the Phillips curve and IS curve innovation standard deviations so that the median estimate of the standard errors of the regression of the Phillips curve and IS curve for rolling fifteen-year-sample regressions on simulated data yield the same estimated standard deviations. The resulting calibration of the innovation standard deviations are reported in Table 1.

3 Real-time Estimation of Natural Rates

I assume that the policymaker makes an ex ante commitment to one method of estimating the natural rates of both interest and unemployment. I additionally assume that the degree of uncertainty regarding the nature of time-variation in the natural rates variables is fixed. Thus, the policymaker can neither modify the method of natural rate estimation based on incoming data nor deduce the true DGPs.

I analyze two commonly used methods for estimating natural rates. The first method is the weighted sample mean, in which the estimate of the natural rate of interest (unemployment) equals a weighted sample mean of the real interest (unemployment) rate over the past n periods. In the case of constant weights, this method corresponds to the simple sample mean. Examples of weighted sample mean estimators of the current natural rate include the Hodrick-Prescott and Bandpass filters.¹²

The second method estimates the natural rates implied by the IS curve and Phillips curve equations. In each case, the dynamic equation is estimated with the natural rate term replaced by an intercept. In implementing this approach, I assume that the slope parameters are estimated jointly with the intercept. In estimating the natural rate of interest, estimates of the natural rate of unemployment are needed because the natural rate of unemployment appears in the IS curve. For this purpose, I use the policymaker's estimates of the natural rate of unemployment derived from the estimate Phillips curve equation. Note that inaccuracy in estimates of the natural rate of unemployment spills over to estimates of the natural rate of interest.

As in the case of the sample mean estimator, two variants of this second estimator method are common. In the first, used for example by Rudebusch (2001) in estimating

¹²See Hodrick and Prescott (1997), Baxter and King (1999), and Christiano and Fitzgerald (2002) for descriptions of these univariate filters. See Orphanides and Williams (2002) for a discussion of their real-time properties as estimators of natural rates.

the natural rate of interest, the natural rate is assumed to be constant over the sample period, and the natural rate estimate equals the estimated constant divided by the negative of the estimated coefficient on the unemployment rate gap (in the case of the Phillips curve) or the natural rate gap (in the case of the IS curve). In the second variant, OLS is replaced by weighted least squares, where the weights decline with the difference between the date of the past observation and that of the current quarter; Ball and Mankiw (2002) use a method in this class to estimate the natural rate of unemployment. In the case of geometrically declining weights, the latter method is identical to the constant-gain least squares commonly used in the learning literature (see, for example, Sargent 1999, Evans and Honkapohja 2001, and Orphanides and Williams 2004).

In the following, I will analyze the optimal choice of the single free parameter for each estimation method. Throughout, I assume that the maximum feasible choice of n is 200, consistent with the current availability of about 50 years of U.S. quarterly data on the unemployment rate and the inflation rate. For the methods that use weighted data, I assume that the weights decline geometrically, with the choice parameter being the decay factor, δ . For either method, increasing the decay factor provides better protection against time variation in the natural rate, but at the cost of increased sampling variation and resulting loss in precision if the natural rates are in fact constant.

4 Optimized Policy

In this section, I compute optimized policies assuming the policymaker knows the true DGPs for the natural rates. I start with the textbook case that the natural rates are constant and known. I then analyze the optimal estimators and policy rules

for the seven DGPs assuming the policymaker has only 200 observations on hand. To compute the policymaker loss under different policies, I perform model stochastic simulations that last 110,000 periods. I drop the first 10,000 periods to eliminate the effects of initial conditions and compute moments from the remaining 100,000 (25,000 years of) simulated observations.

4.1 Optimal Policies for Known Natural Rates

As a benchmark for comparison, I consider the performance of rules that are based on the belief that the natural rates are constant and known with certainty, as is common in much of the literature on monetary policy rules. The coefficients of the resulting optimized policy rule are given at the top of Table 2. Such a rule incorporates very little policy inertia ($\theta_i = 0.10$) and a large response to the unemployment gap ($\theta_u = 1.05$): these characteristics are consistent with the findings of Rudebusch and Svensson (1999) and Orphanides and Williams (2002) for backward-looking models with no natural rate uncertainty. The resulting unconditional standard deviations of inflation, the unemployment gap, and the interest rate gap, and the policymaker loss are given in the first row of the table. In this “ideal” world in which natural rates are constant and known, the policymaker loss under the optimized rule is 17.1.

I now evaluate the performance of the optimized policy rule that is designed assuming no natural rate uncertainty, but where in practice policymaking is based on real-time estimates of natural rates in situations where the natural rates may vary over time. I assume that the policymaker constructs real time estimates of the natural rates of interest and unemployment using one of the two estimators described above. For the present purpose, I assume that each period the policymaker estimates the natural rates of interest and unemployment using the most recent 50 years of simulated data, where all data are weighted equally. The interest rate is then set

Table 2: Policy Rule Optimized for Known Natural Rates
 $(\lambda_u = 1, \lambda_i = 0.5)$

Optimized policy: $\theta_i = 0.10, \theta_{\bar{\pi}} + \theta_{\pi} = 1.37, \theta_u = 1.05, \theta_{\Delta u} = 2.22$						
Natural Rate Model	Standard Deviation					Loss
	$\hat{u} - u^*$	$\hat{r} - r^*$	$\hat{\pi} - \pi^*$	$u - u^*$	$i - i^*$	\mathcal{L}^*
No Misperceptions (constant)	—	—	2.2	1.3	4.6	17.1
<i>Estimator: sample mean</i>						
Zero variance (constant)	0.4	1.0	2.3	1.3	4.8	18.7
Baseline variance						
AR(1)	1.5	1.2	2.6	1.4	4.9	20.9
Long-memory	1.3	1.1	2.4	1.6	4.8	19.7
Markov switching	1.6	1.3	2.6	1.4	4.9	20.9
High variance						
AR(1)	2.8	1.7	3.3	1.5	5.3	26.9
Long-memory	2.6	1.5	2.6	2.3	4.6	22.1
Markov switching	3.1	1.9	3.3	1.7	5.2	27.3
<i>Estimator: model equations</i>						
Zero variance (constant)	0.4	0.9	2.2	1.4	4.8	18.5
Baseline variance						
AR(1)	1.3	1.1	2.4	1.4	5.0	20.1
Long-memory	1.3	1.1	2.2	1.7	4.8	19.3
Markov switching	1.5	1.2	2.4	1.5	5.0	20.3
High variance						
AR(1)	2.4	1.6	2.8	1.6	5.2	23.8
Long-memory	2.5	1.5	2.3	2.3	4.6	21.3
Markov switching	2.7	1.8	2.8	1.8	5.2	24.5

according to the policy rule using the estimated natural rates as inputs. The middle portion of Table 2 reports the results for the case of the sample mean natural rate estimator for the seven different natural rate DGPs. The lower portion of the table reports the results when the estimates of natural rates are constructed based on the estimated intercepts from the IS and Phillips curve equations.

Increasing the variance of the natural rate innovations reduces the accuracy of the natural rate estimates. The first two columns show the unconditional standard

deviations of the difference between the real time and true values of the natural rates. For this policy rule, the estimator based on the IS and Phillips curve equations does a slightly better job of real-time estimation of both natural rates than does the sample mean.

Under the policy rule optimized assuming known natural rates, macroeconomic performance deteriorates modestly under the baseline calibrations of the natural rate DGPs, but declines much more significantly under the high variance alternatives. The policymaker loss in the case of the Markov switching model is over 40 to 60 percent higher than that implied by constant and known natural rates. With the greater variability in natural rates, the policymaker's estimates of natural rates become less accurate, adding persistent errors to the setting of policy. Experiments that consider the two sources of natural rate mismeasurement individually show that mismeasurement of the natural rate of unemployment, rather than interest, is the predominate source of the reduction in macroeconomic performance. In the cases of the AR(1) and Markov-switching processes, the rise in inflation variability accounts for much of the increase in the loss; in contrast, for the long-memory process, the variabilities of both the inflation rate and the unemployment rate gap contribute to the higher value of the loss. This latter result is due to the high-frequency component of the long-memory process that causes variability in the natural rate of unemployment that has little direct effect on inflation because these movements in the natural rate are short lived. Stabilization performance is modestly better when the policymaker derives estimates of the natural rates from estimated model equations rather than sample means.

4.2 Optimal Policies with Known Natural Rate DGPs

The preceding analysis illustrates that policies designed to be optimal when natural rates are known can perform poorly if natural rates vary over time and must be

estimated in real time. Some of the decline in macroeconomic performance associated with natural rate misperceptions is unavoidable: even with optimal estimates, natural rate mismeasurement will occur and this adds noise to the setting of policy. And, as shown above, the greater the variability in natural rates, the greater the natural rate mismeasurement on average, and the worse the macroeconomic performance. But, as I show in the following, appropriately designed estimation and monetary policy combinations can significantly reduce the stabilization costs associated with time varying natural rates.

I start by computing the optimized policies for each natural rate process, assuming that the true processes, but not the realized values of the natural rates, are known. These policies provide benchmarks for our analysis in the next section where I seek to find a single estimation and policy rule combination that is robust across the set of natural rate processes. Table 3 shows the optimized policies for each natural rate process and the resulting loss. The upper part of the table reports results for the weighted sample mean estimator. The first and second columns report the optimal choices for the decay factors used in estimating the natural rates of unemployment and interest, respectively. The next four columns report the coefficients of the optimized policy rule. The final column reports the policymaker loss, denoted by \mathcal{L}^* . The lower part of the table shows the results for the model equation-based estimator.

The performance of optimized policies is nearly the same for the two types of natural rate estimators.¹³ For both the sample mean and model-based estimators, the optimal choice of the decay factors is very small for all of the natural rate processes. For any given natural rate process, the coefficients of the policy rule are similar across the two natural rate estimators. The small differences in the performance of opti-

¹³I also tried rolling regression versions of these estimator with the sample length chosen as a free parameter and the decay parameters set to zero. These policies performed slightly worse in general than those that used the assumed “full” sample of 200 observations.

Table 3: Optimized Policies for Each Model
 $(\lambda_u = 1, \lambda_i = 0.5)$

	Policy Coefficients						Loss
	δ_u	δ_i	θ_i	$\theta_{\bar{\pi}} + \theta_{\pi}$	θ_u	$\theta_{\Delta u}$	\mathcal{L}^*
Natural Rate Model							
<i>Estimator: sample mean</i>							
Zero variance (constant)	0.000	0.000	0.1	1.5	1.1	2.3	18.6
Baseline variance							
AR(1)	0.011	0.000	0.4	1.5	0.7	2.9	19.8
Long-memory	0.002	0.000	0.2	1.5	0.9	2.5	19.4
Markov switching	0.002	0.001	0.5	1.5	0.6	3.1	19.9
High variance							
AR(1)	0.000	0.005	0.7	1.5	0.3	4.0	21.5
Long-memory	0.006	0.000	0.4	1.5	0.7	2.7	21.0
Markov switching	0.000	0.004	0.6	1.6	0.5	3.3	22.5
<i>Estimator: model equations</i>							
Zero variance (constant)	0.000	0.002	0.1	1.3	1.0	2.2	18.5
Baseline variance							
AR(1)	0.001	0.000	0.5	1.2	0.6	3.0	19.6
Long-memory	0.001	0.000	0.2	1.3	0.9	2.2	19.3
Markov switching	0.001	0.000	0.4	1.3	0.6	2.8	19.8
High variance							
AR(1)	0.001	0.002	0.7	1.3	0.3	3.8	21.3
Long-memory	0.001	0.000	0.4	1.2	0.7	2.6	20.9
Markov switching	0.002	0.003	0.6	1.3	0.4	2.8	22.3

mized policies between the types of estimators is striking given that the model-based estimators incorporate more information regarding the structure of the economy. In practice, such model-based methods are likely to be subject to misspecification, potentially causing them to perform worse than the sample mean estimator.

If the true natural rates are constant but unobserved, the optimized policy is nearly identical to the certainty-equivalent optimized policy; but, with time-varying natural rates, the optimized policies exhibit greater policy inertia and a more muted response to the unemployment rate gap. Note that certainty equivalence does not apply to this analysis owing to the violation of two necessary assumptions: first, the class of policy rules that I consider are not fully optimal and second, the estimators are not optimal in that they do not yield conditional mathematical expectations of the natural rates. In response to measurement error in the natural rate of interest, optimized policies reduce the direct response to the natural rate of interest through a larger degree of interest rate smoothing or policy inertia represented by θ_i . To compensate for the lack of accurate estimates of the natural rates optimized policies respond less aggressively to the perceived unemployment rate gap, but respond more strongly to the change in the unemployment rate. In so doing, these policies dramatically reduce the “cost” associated with natural rate mismeasurement. Interestingly, for a given natural rate estimator, the optimized policy responses to inflation differ relatively little across the natural rate processes. Thus, relative to the optimal policy in the case of no uncertainty, the optimal policies with uncertainty are biased towards combatting inflation relative to controlling variability in the unemployment gap and this bias is stronger under the high variance natural rate processes.

5 Robust Policies

I now analyze the choice of estimation and policy rule parameters that minimize the loss taking into account uncertainty regarding the natural rate processes. I first consider the case where the policymaker has priors across the natural rate processes; I then turn to the case of a min-max approach to model uncertainty.

5.1 Robust Policy with Priors across Natural Rate DGPs

I initially assume that the policymaker has well-formed priors over the seven natural rate processes. As noted above, I assume that the policymaker does not update these priors based on incoming data. I assume a weight of $1/4$ for the case of constant natural rates, $1/6$ on each of the baseline calibrations of the three natural rate processes (with the sum of the weights on the baseline calibrations equalling $1/2$), and $1/12$ on each equal weights on each of the three high-variance natural rate processes (with the sum of the weights on the high-variance processes equalling $1/4$). The weights are reported in the first column of Table 4.

For a given combination of a parameterized natural rate estimator and policy rule coefficients, I compute the loss in each model and sum the weighted losses to obtain the expected loss. I then numerically find the estimation and policy coefficients that minimize the expected loss. Table 4 reports the results. The coefficients of the optimized policy are reported at the top of the table. The second through fourth columns report the resulting unconditional standard deviations of the inflation rate, unemployment rate gap, and the nominal interest rate gap, respectively. The fifth column reports the resulting loss. The final column reports the minimum attainable loss (within the class of estimators and policy rule considered here) assuming the natural rate process is known, taken from Table 3, and denoted by \mathcal{L}^* .

Table 4: Robust Policy with Priors

($\lambda_u = 1, \lambda_i = 0.5$)						
Natural Rate Model	Weight	Standard Deviation			Loss	
	ω	$\bar{\pi}$	$u - u^*$	$i - r^*$	\mathcal{L}	\mathcal{L}^*
<i>Estimator: sample mean</i>						
Optimized policy: $\delta_u = 0.005, \delta_i = 0.000$						
Optimized policy: $\theta_i = 0.47, \theta_{\bar{\pi}} + \theta_{\pi} = 1.49, \theta_u = 0.60, \theta_{\Delta u} = 3.10$						
Zero variance (constant)	0.250	2.2	1.4	4.9	18.9	18.6
Baseline variance						
AR(1)	0.167	2.3	1.5	5.0	19.8	19.8
Long-memory	0.167	2.2	1.7	4.9	19.5	19.4
Markov switching	0.167	2.3	1.5	5.0	19.9	19.9
High variance						
AR(1)	0.083	2.6	1.6	5.1	22.0	21.5
Long-memory	0.083	2.2	2.3	4.6	21.0	21.0
Markov switching	0.083	2.6	1.8	5.1	22.9	22.5
Expected Loss					20.1	19.9
<i>Estimator: model equations</i>						
Optimized policy: $\delta_u = 0.001, \delta_i = 0.000$						
Optimized policy: $\theta_i = 0.43, \theta_{\bar{\pi}} + \theta_{\pi} = 1.26, \theta_u = 0.59, \theta_{\Delta u} = 2.75$						
Zero variance (constant)	0.250	2.2	1.4	4.9	18.8	18.5
Baseline variance						
AR(1)	0.167	2.3	1.5	5.0	19.7	19.6
Long-memory	0.167	2.2	1.7	4.8	19.4	19.3
Markov switching	0.167	2.3	1.5	5.0	19.8	19.8
High variance						
AR(1)	0.083	2.5	1.6	5.1	21.7	21.3
Long-memory	0.083	2.2	2.3	4.6	20.9	20.9
Markov switching	0.083	2.6	1.8	5.1	22.6	22.3
Expected Loss					19.9	19.8

The single optimized combination of estimators and policy rule coefficients does very well across all natural rate processes; indeed, the robust policy delivers performance nearly on par with the first-best policy for each natural rate process. By incorporating a significant degree of policy inertia and a muted response to the unemployment gap, this policy combination is very effective at protecting against the high-variance natural rate processes and accomplishes this at negligible cost in terms of performance in the case where the natural rates are constant. As before, the op-

timized rates of decay used in weighting past data in the natural rate estimators are very small. Finally, the optimized policy using the weighted sample mean estimators performs nearly as well as that using the model-based estimator. Although not reported here, this finding that a single estimation/policy combination is robust to natural rate model uncertainty generalizes to other parameterizations of the loss function.

5.2 Robust Policy with a Min-Max Objective

I now consider the case where the policymaker does not have well-formed priors over the different natural rate processes, but instead follows a min-max approach of choosing the estimation method and policy rule coefficients that minimize the maximum loss in any of the seven states of the world corresponding to different processes. This is the approach taken in the robust control literature (see Sargent 1999, Hansen and Sargent 2003). Table 5 reports the results from this experiment; as before, the values of the \mathcal{L}^* correspond to those from the weighted sample mean estimator reported in Table 3.

The min-max policy minimizes the “worst” state, that is, the high variance version of the Markov switching process. The min-max policy responds incorporates greater policy inertia, a larger response to inflation, and a smaller response to the perceived unemployment gap than the optimized policy assuming priors over the seven processes examined above. The decay parameter used in constructing the sample mean is nearly the same as before. The min-max policy does a better job of protecting against the high-variance natural rate processes at the cost of slightly worse performance in the zero and baseline variance processes.

Table 5: Robust Policy with Min-Max Loss

Natural Rate Model	$(\lambda_u = 1, \lambda_i = 0.5)$			Loss	
	Standard Deviation			\mathcal{L}	\mathcal{L}^*
	$\bar{\pi}$	$u - u^*$	$i - r^*$		
Estimator: weighted sample mean					
Optimized policy: $\delta_u = 0.000$, $\delta_i = 0.004$					
Optimized policy: $\theta_i = 0.60$, $\theta_{\bar{\pi}} + \theta_{\pi} = 1.60$, $\theta_u = 0.48$, $\theta_{\Delta u} = 3.28$					
Zero variance (constant)	2.1	1.5	5.1	19.5	18.6
Baseline variance					
AR(1)	2.1	1.6	5.1	20.1	19.8
Long-memory	2.1	1.8	5.0	20.1	19.4
Markov switching	2.1	1.6	5.1	20.2	19.9
High variance					
AR(1)	2.4	1.7	5.2	21.6	21.5
Long-memory	2.1	2.4	4.8	21.4	21.0
Markov switching	2.4	1.8	5.2	22.5	22.5
<i>Estimator: model equations</i>					
Optimized policy: $\delta_u = 0.002$, $\delta_i = 0.003$					
Optimized policy: $\theta_i = 0.59$, $\theta_{\bar{\pi}} + \theta_{\pi} = 1.26$, $\theta_u = 0.41$, $\theta_{\Delta u} = 2.85$					
Zero variance (constant)	2.1	1.5	5.0	19.1	18.5
Baseline variance					
AR(1)	2.2	1.5	5.0	19.9	19.6
Long-memory	2.1	1.8	4.9	19.7	19.3
Markov switching	2.2	1.6	5.0	20.0	19.8
High variance					
AR(1)	2.4	1.7	5.1	21.5	21.3
Long-memory	2.1	2.4	4.7	21.3	20.9
Markov switching	2.4	1.8	5.1	22.3	22.3

5.3 Simple Rules

The preceding analysis shows that the optimal degree of policy inertia and the responsiveness to inflation and the unemployment gap change once one accounts for uncertainty about the natural rates of unemployment and interest. For purposes of comparison with the literature, it is of interest to conduct the same analysis assuming that the policymaker is constrained to follow a version of the ‘‘Taylor Rule’’ (Taylor 1993) in which the interest rate is determined by the perceived natural rate of interest,

Table 6: Policy Evaluation Summary: Sample Mean Estimator
 $(\lambda_u = 1, \lambda_i = 0.5)$

Assumption	Optimal Policy Coefficients						Loss	
	δ_u	δ_i	θ_i	$\theta_{\pi} + \theta_{\pi}$	θ_u	$\theta_{\Delta u}$	Mean	Max.
<i>Optimized Taylor-style rule:</i>								
No uncertainty	0.000	0.000	—	1.34	1.28	—	23.3	31.6
Bayesian	0.021	0.000	—	1.56	1.07	—	21.5	25.1
Min-Max	0.012	0.000	—	1.70	0.81	—	21.9	24.4
<i>Optimized generalized Taylor rule:</i>								
No uncertainty	0.000	0.000	0.10	1.37	1.05	2.22	21.3	27.3
Expected loss	0.005	0.000	0.47	1.49	0.60	3.10	20.1	22.9
Min-Max	0.000	0.004	0.60	1.60	0.48	3.28	20.4	22.5

the inflation rate, and the perceived unemployment gap. In particular, I conduct the analysis exactly as before, but constrain $\theta_i = \theta_{\pi} = \theta_{\Delta u} = 0$. The upper part of Table 6 summarizes the results from this experiment; the lower part of the table summarizes the results from more complicated policy rules discussed in the earlier part of the paper. For this purpose, I focus on the case of the sample mean estimator. Each line reports the results for a particular assumption regarding the design of policy. The seventh column reports the expected loss assuming the priors over the natural rate DGPs described above; the final column reports the maximum loss across the DGPs.

As before, the policy optimized assuming no natural rate uncertainty performs relatively poorly when natural rates in fact vary over time. The optimized policy assuming priors over the various natural rate DGPs responds more aggressively to inflation and less so to the unemployment gap than the policy designed assuming no uncertainty. These differences are even more pronounced for the policy chosen to minimize the maximum loss. In both cases, the optimized decay factor in estimating the natural rate of unemployment is significantly higher than for the complicated

rules, while that for the natural rate of interest is zero. The Taylor-style rules perform nearly as well as the more complicated rules, consistent with the findings of Williams (2003).

6 Conclusion

This paper studies the policymaker's joint problem of model estimation and robust monetary policy in an environment in which there is uncertainty regarding the true process underlying movements in the natural rates of interest and unemployment. I show that the costs of ignoring natural rate uncertainty can be very large. Thus, there is a danger that policymakers could again fall into a pattern of mistakes such as occurred during the late 1960s and 1970s, when natural rate misperceptions contributed to the stagflation of that period, as argued in Orphanides and Williams (2003).

On the positive side, it is possible to design estimation and monetary policy rules that are remarkably robust to a variety of models of natural rate evolution. Weighted sample means perform well for the conduct of monetary policy in an environment where there is uncertainty regarding the data generating processes for natural rates. The robust policy incorporates a significant degree of policy inertia and muted response to the perceived unemployment gap. In contrast, in the model studied here, the certainty-equivalent policy is characterized by virtually no policy inertia and a relatively strong response to the unemployment gap. In addition, the robust policy's response to inflation is somewhat greater than that of the certainty-equivalent policy. As a result, the robust policy is tilted more strongly at the control of inflation relative to the unemployment gap than would be optimal if natural rates were known. This finding complements that of Orphanides and Williams (2004), who show that learning on the part of private agents calls for policies that react more strongly to deviations

of the inflation rate from its target.

The analysis can be extended in a number of fruitful ways, including the incorporation of private expectations and time variation in other model parameters. In addition, I have assumed that the policymaker does not update priors over the various natural rate data generating processes based on the incoming data. A natural extension would allow updating of beliefs about the various DGPs in the context of time variation in the innovation variances.

References

- Ball, Laurence, and N. Gregory Mankiw. "The NAIRU in Theory and Practice." NBER Working Paper 8940, May 2002.
- Baxter, Marianne, and Robert G. King. "Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series." *Review of Economics and Statistics*, 81(4), November 1999, 575-593.
- Bernanke, Benjamin S. and Ilian Mihov. "Measuring Monetary Policy." *Quarterly Journal of Economics*, 113, August 1998, 869-902.
- Boivin, Jean and Marc Giannoni. "Has Monetary Policy Become More Effective?" NBER Working Paper no. 9459, January 2003.
- Brainard, William C. "Uncertainty and the Effectiveness of Policy." *American Economic Review*, 57, 1967, 411-425.
- Brainard, William C. and George C. Perry. "Making Policy in a Changing World," in George L. Perry and James Tobin (eds.) *Economic Events, Ideas, and Policies: the 1960s and After*, Washington, DC: The Brookings Institution, 2000
- Brock, William A., Steven N. Durlauf, and Kenneth D. West. "Policy Evaluation in Uncertain Economic Environments." *Brookings Papers on Economic Activity*, 1:2003, 2003.
- Chow, Gregory C. *Analysis and Control of Dynamic Economic Systems*. New York: John Wiley & Sons, 1975.
- Christiano, Laurence J. and Terry J. Fitzgerald. "The Band Pass Filter." *International Economic Review*, 2002.
- Cogley, Timothy and Thomas Sargent. "Evolving Post-World War II U.S. Inflation Dynamics." in *NBER Macroeconomics Annual 2001*, 2001, 331-373.
- Cogley, Timothy and Thomas Sargent. "Drifts and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S." New York University, manuscript, August 2002.
- Cogley, Timothy and Thomas Sargent. "The Conquest of U.S. Inflation: Learning, Model Uncertainty, and Robustness." New York University, manuscript, November 2003.
- Diebold, Francis X. and Glenn D. Rudebusch, "Long Memory and Persistence in Aggregate Output," *Journal of Monetary Economics*, 24, September 1989, 189-209.
- Diebold, Francis X. and Glenn D. Rudebusch, "Forecasting Output with the Composite Leading Index: A Real-Time Analysis," *Journal of the American Statistical Association*, 86, September 1991, 603-610.

- Edge, Rochelle, Thomas Laubach and John C. Williams, “Learning and Shifts in Long-Run Productivity Growth.” Federal Reserve Bank of San Francisco Working Paper 04-04, March 2004.
- Estrella, Arturo and Jeffrey Fuhrer. “Monetary Policy Shifts and the Stability of Monetary Policy Models.” *Review of Economics and Statistics*, 85(1), February 2003, 94–104.
- Evans, George and Seppo Honkapohja. *Learning and Expectations in Macroeconomics*. Princeton: Princeton University Press, 2001.
- Giannoni, Marc P. and Michael Woodford. “Optimal Inflation Targeting Rules.” in Ben Bernanke and Michael Woodford (eds.) *The Inflation Targeting Debate*, Chicago: University of Chicago Press, 2004.
- Gordon, Robert J. “Foundations of the Goldilocks Economy: Supply Shocks and the Time-Varying NAIRU.” *Brookings Papers on Economic Activity*, 2, 1998, 297–333.
- Granger, Clive W. J. “Long Memory Relationships and the Aggregation of Dynamic Models Author.” *Journal of Econometrics*, 14(2), October 1980, 227–38.
- Hamilton, James “A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle.” *Econometrica*, 57, March 1989, 357–384.
- Hansen, Lars Peter, and Thomas J. Sargent, *Robust Control and Economic Model Uncertainty*, monograph, December 2003.
- Hodrick, Robert J., and Edward L. Prescott, “Post-war Business Cycles: An Empirical Investigation,” *Journal of Money, Credit, and Banking*, 29, 1997, 1–16.
- Kalchbrenner, J.H. and Peter A. Tinsley. “On the Use of Optimal Control in the Design of Monetary Policy.” Special Studies Paper No. 76, Board of Governors of the Federal Reserve System, 1976.
- Kim, Chang-Jin and Charles R. Nelson. “Has the U.S. Economy Become More Stable? A Bayesian Approach Based on a Markov Switching Model of the Business Cycle.” *Review of Economics and Statistics*, 81(4), 1999, 608–661.
- Kozicki, Sharon and Peter A. Tinsley. “Shifting Endpoints in the Term Structure of Interest Rates.” *Journal of Monetary Economics*, 47, June 2001, 613–52.
- Lansing, Kevin. “Learning about a Shift in Trend Output: Implications for Monetary Policy and Inflation.” Federal Reserve Bank of San Francisco Working Paper 2000-16, January 2000.
- Laubach, Thomas. “Measuring the NAIRU: Evidence from Seven Economies.” *Review of Economics and Statistics*, 83(2), May 2001, 218–231.
- Laubach, Thomas and John C. Williams. “Measuring the Natural Rate of Interest.” *Review of Economics and Statistics*, 85(4), November 2003, 1063–1070.
- Laxton, Douglas and Paolo Pesenti. “Monetary Policy Rules for Small, Open, Emerging Economies.” *Journal of Monetary Economics*, 50, June 2003.

- Levin, Andrew, Volker Wieland and John C. Williams. “Robustness of Simple Monetary Policy Rules under Model Uncertainty.” in John B. Taylor (ed.) *Monetary Policy Rules*, Chicago: University of Chicago, 1999.
- Levin, Andrew, Volker Wieland and John C. Williams. “The Performance of Forecast-Based Policy Rules under Model Uncertainty.” *American Economic Review*, 93(3) June 2003, 622–645.
- Levin, Andrew and John C. Williams. “Robust Monetary Policy with Competing Reference Models.” *Journal of Monetary Economics*, 50, June 2003, 945–975.
- McCallum, Bennett T. “Robustness Properties of a Rule for Monetary Policy.” *Carnegie-Rochester Conference Series on Public Policy*, 29, Autumn 1988, 175–203.
- McCallum, Bennett T. “Should Monetary Policy Respond Strongly to Output Gaps?” *American Economic Review*, 91(2), May 2001, 258–262.
- McConnell, Margaret and Gabriel Perez Quiros. “Output Fluctuations in the United States: What Has Changed Since the Early 1980s?” *American Economic Review*, 90(5), 2000, 1464–1476.
- Onatski, Alexei and Noah Williams. “Modeling Model Uncertainty.” *Journal of the European Economics Association*, 1(5), 2003.
- Orphanides, Athanasios. “Monetary Policy Rules and the Great Inflation.” *American Economic Association Papers and Proceedings*, 92(2), May 2002, 115-120.
- Orphanides, Athanasios, Richard Porter, David Reifschneider, Robert Tetlow and Frederico Finan. “Errors in the Measurement of the Output Gap and the Design of Monetary Policy.” *Journal of Economics and Business*, 52(1/2), January/April 2000, 117-141.
- Orphanides, Athanasios and Simon van Norden. “The Unreliability of Output Gap Estimates in Real Time.” *Review of Economics and Statistics*, 84(4), November 2002, 569–583.
- Orphanides, Athanasios and John C. Williams. “Robust Monetary Policy Rules with Unknown Natural Rate.” *Brookings Papers on Economic Activity*, 2:2002, 2002, 63–118.
- Orphanides, Athanasios and John C. Williams. “The Decline of Activist Stabilization Policy: Natural Rate Misperceptions, Learning, and Expectations.” Federal Reserve Bank of San Francisco Working Paper 2003-24, December 2003.
- Orphanides, Athanasios and John C. Williams. “Imperfect Knowledge, Inflation Expectations and Monetary Policy.” in Ben Bernanke and Michael Woodford (eds.) *The Inflation Targeting Debate*, Chicago: University of Chicago Press, 2004.
- Rudebusch, Glenn D. “Is the Fed Too Timid? Monetary Policy in an Uncertain World.” *Review of Economics and Statistics*, 83(2), May 2001, 203–17.
- Rudebusch, Glenn D. “Assessing Nominal Income Rules for Monetary Policy with Model and Data Uncertainty,” *Economic Journal*, 112, April 2002, 402–432.

- Rudebusch, Glenn D. and Svensson, Lars E. O. "Policy Rules for Inflation Targeting," in John B. Taylor, ed., *Monetary Policy Rules*, Chicago: University of Chicago Press, 1999, pp. 203–53.
- Sargent, Thomas J. *The Conquest of American Inflation*, Princeton: Princeton University Press, 1999.
- Simon, Herbert A. "Dynamic Programming Under Uncertainty with a Quadratic Criterion Function." *Econometrica*, 24(1), January 1956, 74–81.
- Sims, Christopher A. "Comment on Cogley and Sargent's 'Evolving Post-World War II U.S. Inflation Dynamics'." in *NBER Macroeconomics Annual 2001*, 2001, 373–379.
- Smets, Frank. "Output Gap Uncertainty: Does it Matter for the Taylor Rule?" *Empirical Economics*, 2002, 22(1), 113–29.
- Staiger, Douglas, James H. Stock, and Mark W. Watson (1997), "How Precise are Estimates of the Natural rate of Unemployment?" in: *Reducing Inflation: Motivation and Strategy*, ed. by Christina D. Romer and David H. Romer, Chicago: University of Chicago Press. 195-246.
- Staiger, Douglas, James H. Stock, and Mark W. Watson (2002), "Prices, Wages, and the U.S. NAIRU in the 1990s," in Krueger and Solow ed. (2002)
- Stock, James H. and Mark W. Watson. "Median Unbiased Estimation of Coefficient Variance in a Time-Varying Parameter Model," *Journal of the American Statistical Association*, 93, March 1998, 349–358.
- Stock, James H. and Mark W. Watson. "Has the Business Cycle Changed? Evidence and Explanations." prepared for Federal Reserve Bank of Kansas City 2003 Jackson Hole Symposium, August 2003.
- Svensson, Lars and Michael Woodford. "Indicator Variables for Optimal Policy." *Journal of Monetary Economics* 50, 2003, 691–672.
- Swanson, Eric T. (2000), "On Signal Extraction and Non-Certainty-Equivalence in Optimal Monetary Policy Rule," Finance and Economics Discussion Series, 2000-32, Board of Governors of the Federal Reserve System, June.
- Taylor, John B. "Discretion versus Policy Rules in Practice." Carnegie-Rochester Conference Series on Public Policy, 39, December 1993, 195–214.
- Taylor, John B. "The Robustness and Efficiency of Monetary Policy Rules as Guidelines for Interest Rate Setting by the European Central Bank." *Journal of Monetary Economics*, 43(3), 1999, 655-679.
- Theil, Henri. *Economic Forecasts and and Policy*. Amsterdam: North-Holland, 1958.
- Williams, John C. "Simple Rules for Monetary Policy." *Federal Reserve Bank of San Francisco Economic Review*, 2003, 1–12.
- Woodford, Michael. *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton University Press, 2004.