

Robust Estimation and Monetary Policy with Unobserved Structural Change

John C. Williams*

Federal Reserve Bank of San Francisco

March 2004

* 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 and Glenn Rudebusch for comments on earlier drafts. Kirk Moore provided outstanding research assistance. 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.

Abstract

This paper considers the joint problem of model estimation and implementation of monetary policy in the face of uncertainty regarding the process of structural change in the economy. We model unobserved structural change through time variation in the natural rates of interest and unemployment. We show that certainty equivalent optimal policies perform poorly when there is model uncertainty about the natural rate processes. We then examine the properties of combined estimation methods and policy rules that are robust to this type of model uncertainty. We find that weighted averages of sample means perform well as estimators of natural rates. The optimal policy under uncertainty responds more aggressively to inflation and less so to the perceived unemployment gap than the certainty equivalent policy. This robust estimation/policy combination is highly effective at mitigating the effects of natural rate mismeasurement.

JEL Classification System: E52

1 Introduction

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. The underlying rate of productivity growth may have risen by 1-1/2 percentage points. And the natural rate of unemployment appears to have declined by at least a percentage point. The implications of structural change for the conduct of monetary policy has 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 in an economy that regularly undergoes structural change and where there is considerable uncertainty as to the precise nature of the underlying process of change.

We represent structural change by medium- and low-frequency variation in the natural rates of interest and unemployment. For our purposes, the natural rate of unemployment is defined to be unemployment rate consistent with stable long-run inflation. Correspondingly, the natural rate of interest is defined to be the real short-term interest rate consistent with the unemployment rate equaling its natural rate in the long run. We focus on shifts in natural rates, rather than other changes in other aspects in the economy, because the empirical evidence is clearest for changes in natural rates.¹

If the true data generating processes for the natural rates were known,

¹For example, Rudebusch and Svensson (1999) find no evidence of a break in slope coefficients in their model, while Estrella and Fuhrer (2003) find strong evidence for a break when testing jointly for intercepts — which represent the natural rates — and slope parameters. See also Kozicki and Tinsley (2001). The evidence for change in other macroeconomic relationships, such as the autocorrelation of inflation, is the subject of ongoing debate; see Bernanke and Mihov (1998), Cogley and Sargent (2001, 2002), Sims (2001), Boivin and Giannoni (2003), and Stock and Watson (2003).

then certainty equivalence would obtain and the optimal policy rule would be the same as if there were no uncertainty.² However, the assumption that policymakers know the processes generating natural rates with certainty is surely highly unrealistic even as an approximation to reality, as emphasized by Orphanides and Williams (2002).³ Indeed, even if one were certain that one possessed the correct model of structural change, estimates of parameters of that model are likely to be very imprecise, as documented by Laubach and Williams (2003) in the case of the Kalman filter. Moreover, Sims (2001) and Cogley and Sargent (2002) provide evidence that the variance of the shock processes has changed over time; such time variation in higher moments implies that the parameters of the optimal filters are themselves changing over time, further impairing the ability to infer the true data generating process from the data. Thus, a key assumption of our approach is that the pervasive uncertainty regarding the process of structural change and that this uncertainty is unlikely to vanish in the foreseeable future.

We implement the idea of uncertainty regarding the processes generating the natural rates by considering a set of potential data generating processes (DGP) for natural rates. We then analyze policies that perform “well” across the set of 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.⁴ We consider three “reasonable” DGPs for natural rates that yield very different implications for the specification of the optimal filter. The DGPs are a highly persistent AR(1), a long-memory or fractionally-integrated

²See Simon (1956), Theil (1958), Chow (1975) and Kalchbrenner and Tinsley (1975) for early analysis of certainty equivalence, and Swanson (2000), Svensson and Woodford (2002), Woodford (2003), and Giannoni and Woodford (2004) for recent analysis. 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).

³See also Stock and Watson (1997) and Orphanides and van Norden (2003).

⁴See, for example, Orphanides and Williams (2002), Laxton and Pesenti (2003), Levin and Williams (2003), Brock, Durlauf, and West (2003), and Onatski and N. Williams (2004). Cogley and Sargent (2003) extend this type of analysis to the case where the policymaker continuously updates his or her priors over models.

process, and a Markov switching process. In addition, we consider three different calibrations of each DGP.

The primary contribution of the paper is the analysis of the joint problem of estimation and policy feedback when there is uncertainty about the data generating process underlying 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 of these papers have treated natural rate mismeasurement as exogenous noise.⁵ In this paper, we directly examine the performance of real-time estimation strategies and policy rules where the true natural rates vary over time.⁶ 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.

We assume that the policymaker must choose an estimation method and policy rule in advance without knowledge of which DGP is the true one. We evaluate the performance of a combination of estimation method and policy rule coefficients in terms of a standard loss function equal to a weighted sum of the unconditional variances of the inflation rate, the difference between the unemployment rate and its natural rate, and the difference between the interest rate and its natural rate. We consider two approaches to model uncertainty. In one, the policymaker has well-formed priors regarding the various DGPs. In that case, we analyze the estimation/policy combination that minimizes the expected loss, integrating over the various DGPs. The second case corresponds to Knightian uncertainty, in which the policymaker does not have priors over the models. In that case, we follow the robust control literature and analyze the estimation/policy combinations that minimize the maximum loss.

We conduct our analysis using a variant of the Rudebusch-Svensson (1999) backward-looking model estimated on 50 years of postwar U.S. data. We focus on this model because Orphanides and Williams (2002) have shown

⁵See, for example, Orphanides et al (2000), Smets (2002), Orphanides (2002), Rudebusch (2001, 2002), and Orphanides and Williams (2002).

⁶The use of the term “real-time” to problems of this sort is due to Diebold and Rudebusch (1991).

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 strong empirical support, as shown by Estrella and Fuhrer (2003).⁷ 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 (2002), Brock, Durlauf, and West (2003), and Levin and Williams (2003).

We find, consistent with many other studies, that there can be very large costs, especially in terms of inflation variability, of ignoring natural rate uncertainty. However, we also show that it is possible to design estimation and monetary policy rules that are robust to a variety of models 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. We find that weighted sample means, where the weights on past data decline gradually, of the real interest rate and the unemployment rate yield surprisingly good estimates of their respective natural rates. In the face of uncertainty about natural rates, the robust policy responds much more aggressively to inflation than under certainty. By keeping its eye on the “inflation” ball, such a policy automatically counteracts the unavoidable policy “mistakes” resulting from natural rate mismeasurement.

2 The Model

We use a version of the Rudebusch-Svensson (1999) model for our analysis. Note that expectations in this model are assumed to be adaptive and implicitly captured by the lags of the dependent variable. Following Orphanides and

⁷The 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), (2004), and others. The extension of the analysis of this paper to other models of expectations formation is left for future work.

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 we allow for time variation in the natural rates of interest and unemployment.

2.1 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 $\epsilon_t \sim N(0, \sigma_\epsilon^2)$. 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 to 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 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:

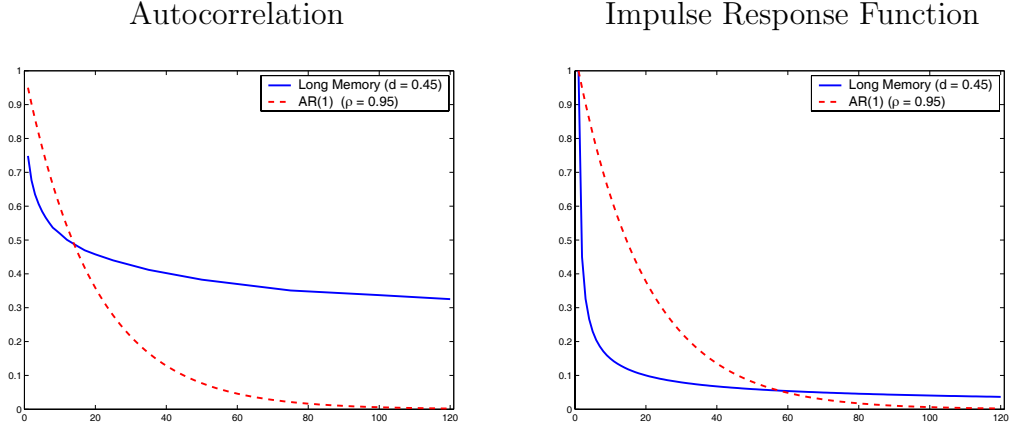
$$\pi_t = \gamma_1 \pi_{t-1} + \frac{1 - \gamma_1}{3} \sum_{j=2}^4 \pi_{t-j} + \gamma_2 (u_{t-1} - u_{t-1}^*) + \eta_t, \quad (2)$$

where $\eta_t \sim N(0, \sigma_\eta^2)$. As noted above, the natural rates of interest and unemployment are allowed to be time varying and are therefore identified with time subscripts. We describe the data generating process for these processes in the next section.

2.2 Modeling Natural Rates

As noted in the introduction, there is considerable uncertainty regarding the specification of the data generating processes (DGP) for the natural rates of interest and unemployment. We restrict ourselves to stationary processes for the

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



natural rates.⁸ We consider three types of time series models for time-varying natural rates: an AR(1), a fractionally-integrated long-memory process, and a two-state Markov switching model. We chose these models because they can generate highly persistent stationary series, but, importantly for our purposes, they imply different optimal filters. For each model, we consider three parameterizations of the variance of the process, as described below.

The first type of model is the standard AR(1), according to which the variable z follows the law of motion:

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

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). In this case, the 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. We approximate this process

⁸Based on the ADF test, one can reject the null of nonstationarity of the unemployment and real federal fund rate over 1950–2003 at the 5 percent level.

by its binomial expansion, truncated after 5000 terms,

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

Granger (1980) shows conditions under which such a long-term memory process approximates the aggregate process resulting from aggregating AR(1) processes over individuals with different values of ρ .

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 for the case of with $\rho = 0.95$, while that of the long-memory process displays approximately hyperbolic decay, as shown by the solid line (for $d = 0.45$). Thus, the long-term memory process is able to generate low frequency variability without resorting to near unit root behavior. Second, the two processes differ markedly in their impulse responses, as shown in the figure. The impulse response function (IRF) 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.

Finally, 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 its current state to the other. We 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 DGP are described by two parameters, one related to the persistence of the state and the second related to the variance of the innovations. Because we are interested in medium and low frequency variation in the natural rates, we assume values of $\rho = 0.99$, $d = 0.48$, and $p = 0.99$, which yield a high degree of persistence. We 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.

2.3 Monetary Policy

We assume that the objective of the monetary policymaker 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 natural rate of interest. Specifically, the loss, \mathcal{L} is given by:

$$\mathcal{L} = E \left\{ (\pi_t - \pi^*)^2 + \lambda_u (u_t - u_t^*)^2 + \lambda_i (i_t - \pi^* - r_t^*)^2 \right\}, \quad (6)$$

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 the natural rates of unemployment and interest, $\{u_t^*, r_t^*\}_{t=0}^\infty$. A loss function of approximately this functional form can be derived from a consumer welfare maximization criterion, as in Woodford (2003). In the following, we treat the weights, λ_u and λ_i , as fixed, and consider a range of values.⁹

The expectation in the loss function takes into account both uncertainty about the realization of future innovations, but also uncertainty about the data generating processes for the natural rates. Let S denote the set of such data generating processes (which may differ across variables). Assume for the present purposes that the policymaker has well-defined prior beliefs over 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, the unconditional loss is given by:

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

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

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

⁹In a micro-founded model, the weights λ_u and λ_i are functions of parameters describing technology and preferences. As discussed in Levin and Williams (2003), uncertainty about such parameters also implied uncertainty regarding these weights. In this paper, we ignore this link and hold the weights fixed.

where $\sum \omega_i = 1$. In the example studied in this paper, $N_s = 7$.

We assume that monetary policy is implemented by setting the federal funds rate according to a monetary policy rule taking the form of a modified Taylor (1993) rule. Rudebusch and Svensson (1999) show that such a rule yields performance very close to the first-best optimal policy under commitment in their model. And Levin and Williams (2003) shows that such rules are robust under model uncertainty. In particular, the federal funds rate is set according to:

$$i_t = \hat{r}_{t-1}^* + \bar{\pi}_{t-1} + \theta_1 \bar{\pi}_{t-1} - \theta_2 (u_{t-1} - \hat{u}_{t-1}^*), \quad (9)$$

where \hat{r}_t^* and \hat{u}_t^* are the policymaker's estimates of the natural rates of interest and unemployment, respectively. We assume that the inflation target is zero and abstract from the zero lower bound on interest rates. Note that the policymaker's estimates of the natural rates of interest and unemployment, are allowed to vary over time, as discussed below.


3 Model Estimation

In order to analyze the performance of filtering and monetary policies, we need to estimate the basic model and calibrate the set of data generating processes for the natural rates of interest and unemployment. We start with the estimation of the “slope” parameters of the IS and Phillips curve equations.

3.1 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 the parameters of the IS and Phillips curve equations. If, however, the natural rates are changing (and unobserved), omitted variable bias is likely to affect all model parameters. To reduce this effect, we used the Congressional Budget Office's (CBO) estimate of the natural rate of unemployment as a proxy for the true values and we estimated the two equations

Figure 2: Rolling Regression Estimates



using rolling regressions in which each sample contains only 15 years of data. Given our sample of data of 1950-2003, we were able to run 156 regressions. We take the median estimate of each parameter 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 first column of charts reports the estimates pertaining to the Phillips curve equation; the second column to the IS curve equation. For the Phillips curve equation, we plot a version of the model equation that includes a constant and 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 . In the case of the IS curve equation, we 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 vary considerably over time, lending some support either for time variation in the true processes or the difficulty in their real-time estimation. Interestingly, the estimate of γ_1 , the coefficient on the first lag of inflation in the Phillips curve equation also displays considerable time variation. Note that this occurs despite imposing the unit restriction on the sum of inflation lags in the Phillips curve equation as this characteristic also obtains when the sum restriction is not imposed.¹⁰ The slope of the IS curve displays a downward trend over the sample. Finally, the estimates of the slope of the Phillips curve and the lags of the unemployment rate in the IS curve are very stable over the sample.

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

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

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

These estimates are similar to those from full-sample estimation and conform with 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 resulting from the implicit assumption

¹⁰In fact, the sum restriction is clearly satisfied for the sample from 1970 to 2003. Only during the 1960s is the sum well below unity. See Orphanides and Williams (2003) for a discussion of this issue.

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

of a constant natural rate of interest. The implied “sacrifice ratio” 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.

3.2 Natural Rates of Interest and Unemployment

In addition to the uncertainty regarding the structure of the DGP for natural rates 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). In particular, the innovation variance for the highly persistent component of natural rates is estimated with little precision. Thus, the data provides frustratingly little guidance on this key parameter.

We now describe the calibration of the three natural rate DGP models and

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.


the various parameterizations of each. The results of this calibration are summarized in Table 1. For each DGP, we consider three parameterizations that span the set of values that are 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 innovation variance is set to the baseline value computed as described below. In the third parameterization, the innovation variance is set to a larger value that lies within the range of other published estimates. In the cases of a zero natural rate variance, the three DGPs collapse into one, so in the end we have seven alternative specifications of the DGP in all; we do not consider the various alternative combinations of these underlying processes for the two natural rates.

We follow the same basic procedure for calibrating both the natural rate of unemployment and interest. Starting with the natural rate of unemployment, we estimate the Phillips curve equation using the Kalman filter, assuming that the natural rate follows a random walk.¹² From this we extract our baseline estimate of the standard deviation of the innovation to the AR(1) model of the natural rate of unemployment, τ_t , of 0.22. The resulting “smoothed” estimate of the natural rate of unemployment is shown in Figure 3, and is similar to the CBO estimate.

To capture the uncertainty regarding the innovation standard deviation, we consider two representative alternative values of zero and 0.44. The value of zero corresponds to a constant natural rate and the latter yields an estimate of the natural rate similar, albeit smaller, to that of Stock and Watson (2001), 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. We use these values of the innovation standard deviations to calibrate our three DGP for the natural rate of unemployment. Of course, a zero innovation standard deviation implies a constant natural rate of unemployment, so in all three DGPs, the “zero” alternative is the same.

¹²For this purpose, we use the Stock and Watson (1998) median-unbiased estimator. The sample is 1970-2003.

Figure 3: Estimates of the Natural Rate of Unemployment




For the AR(1) model, we assume that $\rho = 0.99$. For the baseline calibration, we set the standard deviation of the innovation, τ_t , to 0.22, the baseline Kalman filter estimate. The resulting process for the natural rate of unemployment has an unconditional standard deviation of 1.56 percentage points. For the “high” variance version of the AR(1) model, we set the standard deviation of τ_t to 0.44; this yields an unconditional standard deviation of 3.12 percentage points.

We 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, we set $d = 0.48$. For the Markov-switching model, we set the common switching probability $p = 0.99$ and set the difference between the values at the states at 3.12 percentage points; for the high-variance variant, we set the difference between the states to 6.24 percentage points.

The strategy for calibrating the DGPs for the natural rate of interest is the same as for the natural rate of unemployment. We use the same values of ρ , d , and p as before. Kalman filter estimation of the IS curve yields an innovation

Figure 4: Estimates of the Natural Rate of Interest



standard deviation of 0.11 percentage points, which we use for the baseline AR(1) process (and again assume $\rho = 0.99$). This implies an unconditional standard deviation for the natural rate of interest of 0.78 percentage points. For the high-variance alternative, we 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, we again set $d = 0.48$ and set the innovation standard deviations so that they match those from AR(1) model. For the baseline Markov switching model, we again set $p = 0.99$ and set the difference between the two states at 1.56 percentage points, and for the high-variance alternative, the difference is set at 3.12 percentage points.

In order to make the seven DGPs equivalent in the sense of the implied overall variability of inflation and unemployment rate gap, we adjust the standard deviations of the innovation to the Phillips curve and IS curve equations, as indicated in Table 1. In the case of constant natural rates, we 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 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 each other DGP, we 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 yields the same estimated standard deviations (of 1.13 for the Phillips curve and 0.28 for the IS curve). The results are shown in the table.

4 Real-time Estimation of Natural Rates

We assume the policymaker makes an ex ante commitment to methods of estimating the natural rates of interest and unemployment. We additionally assume that the degree of uncertainty regarding the nature of time-variation in the natural rates variables is fixed. Thus, the policymaker is able to adapt to changes in the natural rates, but cannot deduce the true data generating process.

We 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 is simply the sample mean. Hodrick-Prescott and Bandpass filter estimates of the current level of the trend component of a series belong to this weighted sample mean class of estimators.¹³

The second method estimates the natural rates indirectly from 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, we assume the slope parameters are known with certainty, thus, we are overestimating the real-world precision of this estimator. In estimating

¹³See 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.


the natural rate of interest, estimates of the natural rate of unemployment are needed as they appear in the IS curve. Thus, inaccuracy in estimates of the natural rate of unemployment spills over to estimates of the natural rate of interest. As before, two variants of this estimator are common. In the first, used for example by Rudebusch (2001), 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 multiplying 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, ordinary least squares 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.¹⁴ In the case of geometrically declining weights, the latter method is identical to the steady-state Kalman gain for the simple model of a random walk natural rate discussed above and also is identical to constant-gain least squares commonly used in the learning literature (see, for example, Sargent 1999, Evans and Honkapohja 2002, and Orphanides and Williams 2004).

In the following, we will analyze the optimal choice of the free parameter for each estimation method. For both methods with constant weighting, the one free parameter is the sample length n . Throughout the following, we 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, we assume 200 observations are used and we assume that the weights decline geometrically, with the choice parameter being the decay factor, δ used in estimating both the natural rate of unemployment and the natural rate of interest.

For either method, shortening the optimal sample length or increasing the decay factor provides better protection against time variation in the natural rate, but carries the cost of increased sampling variation and resulting loss in

¹⁴Ball and Mankiw (2002) use a method in this class to estimate the natural rate of unemployment.

Figure 5: Accuracy of Natural Rate of Unemployment Estimators



precision. Figure 5 shows graphically the tradeoff associated with reducing the effective sample size. For this figure, the natural rate of unemployment is estimated from the Phillips curve equation assuming all slope parameters are known. The vertical axis is the root-mean-squared-error (RMSE) of estimates of the natural rate of unemployment, computed numerically. For the case of constant weights, the horizontal axis shows the sample length, in quarters, used in the rolling regressions. The lower solid curve shows the root mean squared error of estimates of the natural rate of unemployment for samples ranging from 10 years to 50 years in length. (This curve is simply a plot of $\frac{\sigma_\epsilon}{\gamma_2\sqrt{n}}$.) The middle solid curve shows the average RMSE for the third DGPs (AR(1)), long-memory process, and Markov switching) under the baseline calibration for different values of n . The minimum RMSE is achieved for samples of about 15 years. The upper thin solid line shows the average RMSE for the high-variance variants of the DGPs. In this case, the optimal length is shorter than 40 quarters. The ex ante best sample length depends on the relative weights one places on the various DGPs.

Based on accuracy, the constant-gain least-squares (CGLS) estimators out-

perform on average the rolling regression estimates. The lower dashed curve shows the RMSE of CGLS estimators that yield the same accuracy under a constant natural rate as the rolling regression for the sample size indicated on the horizontal axis (this corresponds to a decay factor of between 0.005 and 0.025). As seen in the figure, the curve generated by the CGLS estimator for the baseline calibration of the DGP lies under that generated by the rolling regression estimator, and the same is true for the high-variance calibration. These calculations are based on exogenous processes and are therefore invariant to the particular implementation of monetary policy in the model. This is not the case for the weighted sample mean estimators, however, the performance of which depends on the behavior of endogenous variables. We evaluate those estimators in the next section when we simulate the model.

5 Optimal Policy with a Known DGP

In this section, we compute optimized policies assuming the policymaker knows the correct DGP for the natural rates. We start with the textbook case that the natural rates are constant and known. We then analyze the optimal estimators and policy rules for the seven DGPs assuming the policymaker has only 200 observations at hand.

To compute the policymaker loss under different policies, we perform stochastic simulations of the model of 110,000 periods. We 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.

5.1 Optimal Policies for Known Natural Rates

As a benchmark for comparison, we consider the performance of rules designed based on the belief that the natural rates are constant and known with certainty. We assume one particular set of weights in the policymaker loss function: $\lambda_u = 1, \lambda_i = 0.5$; qualitatively, our results are not sensitive to moderate variations in these parameters, as discussed below. For the chosen policy

loss parameters, the optimized policy rule assuming known natural rates is characterized by $\theta_1 = 1.35$ and $\theta_2 = 1.25$.

We then evaluate the optimal policy based on the assumption of known natural rates using model simulations in which the natural rates are in fact generated by one of the seven DGPs described above. In implementing these policies we assume that the policymaker uses natural rate estimates constructed from 200 unweighted quarterly observations, either using the sample mean estimator or the estimator based on the dynamic equations. Table 2 reports the results for this experiment for the seven different natural rate DGPs; the upper part of the table refers to simulations in which the sample mean is used to estimate the natural rates and the lower part refers to simulations in which the natural rate is estimated using the dynamic equations (assuming the slope coefficients are known). Each row corresponds to one particular natural rate DGP. The first two columns report the simulated unconditional standard deviations of the real-time natural rate errors. The next three columns show the unconditional standard deviations of the four-quarter inflation rate, the difference between the unemployment rate and its natural rate, and the difference between the interest rate and its natural rate. The final column indicates the resulting value of the loss function for the specified loss function parameters.

Not surprisingly, increasing the variance of the natural rate innovations reduces the accuracy of the natural rate estimates. Given this policy rule, the estimator based on the IS and Phillips curve equations does a better job of real-time estimation of the natural rate of unemployment than the sample mean. The opposite is generally true for the natural rate of interest, though. The relatively poor performance of the natural rate of interest estimator based on the IS curve in part reflects the fact that this estimator incorporates estimates of the natural rate of unemployment, and measurement error of that natural rate adds noise to the estimates of the natural rate of interest.

Under the policy rule optimized assuming known natural rates, macroeconomic performance deteriorates modestly under the baseline calibrations of the natural rate DGPs, but declines dramatically under the high variance al-

Table 2: Policy Rule Optimized Assuming Known Natural Rates

Loss parameters: $\lambda_u = 1$, $\lambda_i = 0.5$						
Optimized policy rule parameters: $\theta_1 = 1.35$, $\theta_2 = 1.25$						
Natural Rate DGP	Standard Deviation					Loss
	$\hat{u} - u^*$	$\hat{r} - r^*$	$\bar{\pi}$	$u - u^*$	$i - r^*$	\mathcal{L}
<i>Estimator: Sample mean (unweighted)</i>						
Zero variance (constant)	0.8	1.0	2.5	1.5	5.0	20.5
Baseline variance						
AR(1)	1.5	1.2	2.7	1.5	5.1	22.9
Long-memory	1.5	1.2	2.6	1.7	5.0	21.7
Markov switching	1.6	1.3	2.7	1.5	5.1	22.5
High variance						
AR(1)	2.7	1.7	3.4	1.6	5.5	29.2
Long-memory	2.7	1.5	2.8	2.3	4.8	24.9
Markov switching	3.0	1.8	3.4	1.8	5.4	29.1
<i>Estimator: dynamic equations (unweighted)</i>						
Zero variance (constant)	0.4	0.8	2.3	1.4	4.8	19.1
Baseline variance						
AR(1)	1.4	2.1	2.5	1.7	4.9	24.6
Long-memory	1.3	1.4	2.3	1.6	4.3	21.0
Markov switching	1.5	1.9	2.9	1.5	5.1	23.6
High variance						
AR(1)	2.7	3.8	4.4	1.6	6.1	40.7
Long-memory	2.6	2.5	3.0	2.3	4.9	26.3
Markov switching	2.7	3.1	4.1	1.8	5.8	36.6

ternatives. The loss in the case of the Markov switching model is more than double that implied by constant natural rates. 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.

5.2 Optimal Policies with Known DGPs

Some of the decline in macroeconomic performance as the variance of the natural rate processes increases is unavoidable: more volatile unobserved natural rates are “bad” for monetary policy. But, as we show, mistakes in natural rate estimation is the primary culprit explaining the dramatic increases in the policymaker loss seen in Table 2.

To provide a benchmark of attainable performance under each natural rate DGP, we now compute the optimal policies for each DGP, assuming that the true DGP is known. In this model, certainty equivalence applies if the natural rate DGP is known with certainty. As a result, the best attainable outcome is that with the fully optimal policy under certainty using natural rate estimates generated by the optimal filter implied by the DGP. However, for purposes of comparison with our robust estimation and policy combination, we focus on the simple estimation techniques and the simple Taylor style policy rule as described above. Thus, certainty equivalence does not apply and the coefficients of the estimators and the policy rule will differ from those implied by optimal filtering and the certainty equivalent policy, respectively.

Table 3 shows the optimized policies for each DGP and the performance under that DGP. The upper part of the table reports results for the weighted sample mean estimator where the common decay factor for the two natural rate estimates and the policy rule parameters are simultaneously chosen to minimize the loss. The lower part of the table shows the CGLS estimator based on the intercepts of the IS and Philips curve equations, where again the decay factor is chosen as well as the policy rule coefficients. The results from the rolling regression versions to these two estimators are very similar to those shown here, but the weighted mean and the CGLS estimators perform better on average over the seven DGPs so, in the following, we focus on these

Table 3: Policy Rules Optimized Assuming Known DGPs

Loss parameters: $\lambda_u = 1, \lambda_i = 0.5$							
Natural Rate DGP	Coefficients			Standard Deviation			Loss \mathcal{L}^*
	δ	θ_1	θ_2	$\bar{\pi}$	$u - u^*$	$i - r^*$	
<i>Estimator: weighted sample mean</i>							
Zero variance	0.000	1.5	1.4	2.3	1.5	5.0	19.7
Baseline variance							
AR(1)	0.003	1.5	0.9	2.4	1.6	5.1	21.4
Long-memory	0.001	1.5	1.1	2.3	1.7	4.9	20.6
Markov switch.	0.002	1.5	0.8	2.5	1.6	5.1	21.5
High variance							
AR(1)	0.006	1.5	0.6	2.6	1.8	5.3	23.6
Long-memory	0.003	1.5	0.9	2.3	2.4	4.7	22.3
Markov switch.	0.004	1.6	0.6	2.6	1.9	5.3	24.6
<i>Estimator: dynamic equations</i>							
Zero variance	0.000	1.4	1.3	2.3	1.4	4.9	19.1
Baseline variance							
AR(1)	0.006	1.6	1.0	2.5	1.6	5.2	22.5
Long-memory	0.003	1.5	1.1	2.3	1.7	4.9	20.6
Markov switch.	0.002	1.6	0.9	2.5	1.6	5.2	22.2
High variance							
AR(1)	0.011	1.8	0.7	2.7	1.9	5.7	27.1
Long-memory	0.006	1.7	0.9	2.4	2.4	4.9	23.5
Markov switch.	0.003	2.0	0.8	2.8	2.0	5.6	27.3

estimators.¹⁵

For both estimators, the optimal choice of the decay is very small and well below the values implied by minimizing the accuracy of natural rate estimates for time-varying processes.

The policies that use the weighted sample mean outperform those based on CGLS estimates from the dynamic equations for the high variance alternatives and perform slightly worse in the case of a constant natural rates. Given

¹⁵Performance under the AR(1) and long-memory processes is slightly better using weighted sample means than rolling regression sample means, but under the Markov switching process, performance is better with the rolling regressions sample means over a truncated sample.

that we have exaggerated the accuracy of the CGLS estimator based on the dynamic equations by assuming away sampling error in the slope coefficients in these models, we conclude that the weighted sample mean is a more robust method of estimating natural rates. An important aspect of this finding is that policies that effectively stabilize inflation also improve accuracy of natural rate estimates from sample averages.

The optimized policies compensate for the lack of accurate estimates of the natural rates by responding more aggressively to inflation and less to the perceived unemployment rate gap, and in so doing dramatically reduce the “cost” associated with natural rate mismeasurement. 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. This bias is even stronger under the high variance DGP, and even is present to a small extent under the assumption of constant, but unknown, natural rates.

6 Robust Policies with DGP Uncertainty

A striking result implied by Table 3 is that a single estimation/policy combination is likely to perform well under all seven DGPs. The optimal policies for the six cases with time-varying natural rates feature small decay rates, combined with a response to inflation of between 1-1/2 and 2 and a response to the perceived negative unemployment gap of between 1/2 and 1. We now formalize this intuition by analyzing the choice of estimation and policy rule parameters that minimize the expected loss, where the expectation incorporates the uncertainty regarding the natural rate DGPs.

6.1 Robust Policy with Priors across DGPs

We initially assume that the policymaker has well-formed priors over the seven DGPs; we turn to the case where the policymaker does not have priors below. Following the logic that the policymaker must commit to estimation method and

Table 4: Robust Policy with DGP Uncertainty (with priors)

Loss parameters: $\lambda_u = 1$, $\lambda_i = 0.5$					
Estimator: weighted sample mean					
Optimal policy: $\theta_1 = 1.50$, $\theta_2 = 0.87$, $\delta = 0.003$					
Natural Rate DGP	Standard Deviation			Loss	
	$\bar{\pi}$	$u - u^*$	$i - r^*$	\mathcal{L}	\mathcal{L}^*
Zero variance (constant)	2.3	1.6	5.0	20.2	19.1
Baseline variance					
AR(1)	2.4	1.6	5.1	21.4	21.4
Long-memory	2.3	1.8	5.0	20.8	20.6
Markov switching	2.4	1.6	5.1	21.5	21.5
High variance					
AR(1)	2.8	1.7	5.2	24.6	23.6
Long-memory	2.3	2.4	4.7	22.3	22.3
Markov switching	2.9	1.9	5.2	25.2	24.6
Weighted average				21.7	21.2

policy in advance, we assume that the policymaker does not update these priors based on incoming data. As a benchmark case, we assume equal weights on the three baseline DGPs with the sum adding to 0.5, a weight of 0.25 on the zero variance DGP, and equal weights adding up to 0.25 on the high variance DGPs. For a given combination of parameterized estimation method and policy rule parameters, we compute the loss in each model and sum the weighted losses to obtain the expected loss. We then use a hill-climber routine to find the free parameters to minimize this loss. Table 4 reports the results. The final column reports the minimum attainable loss (within the classes of estimators and policy rule considered here) assuming the DGP is known, calculated above, and denoted \mathcal{L}^* . In this table and for the remainder of the paper, we focus exclusively on weighted sample mean estimation of natural rates because these policies outperformed those based on natural rates estimated based on dynamic equations.

As expected, the robust estimation/policy does very well for all DGPs. This combination is very effective at stabilizing the economy even under the

Table 5: Robust Policy with DGP Uncertainty (with priors)

Estimator: weighted sample mean					
Loss parameters: $\lambda_u = 0.5$, $\lambda_i = 0.25$					
Optimal policy: $\theta_1 = 1.81$, $\theta_2 = 0.99$, $\delta = 0.004$					
	Standard Deviation			Loss	
	$\bar{\pi}$	$u - u^*$	$i - r^*$	\mathcal{L}	\mathcal{L}^*
Natural Rate DGP					
Zero variance (constant)	2.1	1.6	5.2	12.7	12.3
Baseline variance					
AR(1)	2.3	1.7	5.3	13.5	13.5
Long-memory	2.1	1.9	5.2	13.0	12.8
Markov switching	2.4	1.6	5.1	13.7	13.6
High variance					
AR(1)	2.6	1.8	5.3	15.4	14.8
Long-memory	2.2	2.4	4.9	13.7	13.7
Markov switching	2.7	1.9	5.4	16.3	15.7
Weighted average				13.7	13.4
Loss parameters: $\lambda_u = 2$, $\lambda_i = 2$					
Optimal policy: $\theta_1 = 1.21$, $\theta_2 = 0.75$, $\delta = 0.001$					
Zero variance (constant)	2.5	1.5	4.9	58.4	57.1
Baseline variance					
AR(1)	2.7	1.5	5.0	61.5	61.5
Long-memory	2.5	1.7	4.8	59.2	58.6
Markov switching	2.4	1.7	5.2	62.5	61.7
High variance					
AR(1)	3.1	1.6	5.2	69.4	67.2
Long-memory	2.6	3.3	4.6	60.4	60.4
Markov switching	3.2	1.8	5.3	72.4	70.3
Weighted average				62.0	61.1

high variance DGPs, at negligible cost in terms of performance under the zero variance DGP. Again, the rate of decay used in weighting past data is surprisingly small.

This finding that a single estimation/policy combination is robust to natural rate model uncertainty generalizes to other parameterizations of the loss function. The upper part of Table 5 shows results for a case where policymakers place more weight on the goal of inflation stabilization ($\lambda_u = 0.5$, $\lambda_i = 0.25$); the lower part of the table shows results for a case where policymakers place relatively more weight on the goals of unemployment and interest rate stabi-

Table 6: Min-max Robust Policy with DGP Uncertainty

Loss parameters: $\lambda_u = 1, \lambda_i = 0.5$						
Estimator: weighted sample mean						
Optimal policy: $\theta_1 = 1.57, \theta_2 = 0.67, \delta = 0.004$						
Natural Rate DGP	Standard Deviation					Loss \mathcal{L}
	$\hat{u} - u^*$	$\hat{r} - r^*$	$\bar{\pi}$	$u - u^*$	$i - r^*$	
Zero variance (constant)	0.4	1.1	2.2	1.6	5.2	21.1
Baseline variance						
AR(1)	1.4	1.2	2.3	1.7	5.2	21.8
Long-memory	1.3	1.1	2.2	1.9	5.1	21.6
Markov switching	1.4	1.2	2.3	1.7	5.2	21.7
High variance						
AR(1)	2.5	1.5	2.6	1.8	5.3	23.7
Long-memory	2.5	1.5	2.2	2.5	4.8	22.7
Markov switching	2.9	1.8	2.6	1.9	5.2	24.0

lization ($\lambda_u = 2, \lambda_i = 2$). For comparison, the final column of the table shows the minimum feasible loss within this class of estimator and policy rule, \mathcal{L}^* , assuming the natural rate DGP were known.

As seen by comparing the final two columns of the table, the robust policy delivers performance nearly on par with the first-best policy for each DGP. As before, the decay factor used in estimating natural rates is very small, and the robust policies respond more aggressively to inflation and less so to the unemployment gap than would be optimal assuming the natural rates were known.

6.2 Robust Policy with a Min-Max Objective

We now consider the case where the policymaker does not have well-formed priors over the different DGPs, but instead follows a min-max approach of choosing the estimation method and policy rule coefficients to minimize the maximum loss in any of the 7 states of the world corresponding to different DGPs. This is the approach taken in the robust control literature (see Sargent 1999, Hansen and Sargent 2002). Table 6 results the results.

Finally, we explored the effects of modifying the parameters of the DGP, ρ , d , and p , and found that the results were not very sensitive to the values of these parameters.

The min-max policy minimizes the ‘worst’ state: the high variance version of the Markov switching process. The policy responds even more aggressively to inflation and is less responsive to the perceived unemployment gap than the robust policy assuming a priors over all seven DGPs examined above. The decay parameter used in constructing the sample mean is nearly the same as before. Under this policy, the loss under the time-varying natural rate model is only slightly higher than if policy were optimized for that particular DGP.

6.3 The Performance of More Complicated Policy Rules

The preceding analysis assumed that policy followed a simple Taylor-style policy rule and that the same estimation method and parameter were used for estimating both natural rates. Orphanides and Williams (2002) argue that policies that respond to the change in the unemployment rate, as well as the perceived unemployment rate gap, perform markedly better when there is uncertainty about the natural rate of unemployment. In addition, given the different processes for the natural rates and their observed counterparts, the use of the same parameterized estimator for each natural rate may be a costly restriction.

We now gauge the importance of these modifications to the estimation and policy rule specifications by allowing the decay factor to differ for the two processes and by specifying a generalized version of the policy rule given by:

$$i_t = \hat{r}_{t-1}^* + \bar{\pi}_{t-1} + \theta_1 \bar{\pi}_{t-1} - \theta_2 (u_{t-1} - \hat{u}_{t-1}^*) + \theta_4 (u_{t-1} - u_{t-2}) + \theta_5 \pi_{t-1}. \quad (12)$$

Note that we have added two new free parameters to the policy rule, one on the change in the unemployment rate and the second on the annualized inflation rate in the most recent quarter. We consider the weighted sample mean estimator with decay factors, δ_u and δ_r , applied to the unemployment rate and real interest rate, respectively. We assume again that the policymaker

Table 7: Robust Policy with DGP Uncertainty (with priors)

Estimator: weighted sample mean				
Loss parameters: $\lambda_u = 1$, $\lambda_i = 0.5$				
Optimal policy: $\theta_1 = 0.17$, $\theta_2 = 0.78$, $\theta_3 = -2.03$, $\theta_4 = 1.38$				
Optimal estimator: $\delta_u = 0.010$, $\delta_r = 0.000$				
Natural Rate DGP	Standard Deviation			Loss
	$\bar{\pi}$	$u - u^*$	$i - r^*$	\mathcal{L}
Zero variance (constant)	2.2	1.4	4.9	18.9
Baseline variance				
AR(1)	2.3	1.5	5.0	19.9
Long-memory	2.2	1.7	4.9	19.6
Markov switching	2.3	1.5	5.0	20.0
High variance				
AR(1)	2.6	1.6	5.1	22.2
Long-memory	2.2	2.3	4.6	21.1
Markov switching	2.7	1.8	5.1	23.2
Weighted average				20.2

has priors over the various models, as described above, and assume $\lambda_u = 1$ and $\lambda_i = 0.5$.

This generalization of the policy rule and the estimation method reduces average loss by about 7 percent relative to the simpler robust estimator/policy combination, but does not change the main conclusions of the paper. The optimal decay for estimating the natural rate of unemployment is larger than before, but still relatively small. The optimal decay for estimating the natural rate of interest is zero, reflecting the relatively low innovations variances for the natural rate of interest process and the relatively high overall variability of real interest rates. The overall response to inflation is stronger than that if the natural rates are assumed to be known and the response to the perceived unemployment gap is weaker. The sum of the responses to inflation for the robust policy is 1.56, relative to 1.38 for the same specification of the policy rule under the assumption of known natural rates; the response to the perceived unemployment gap in the robust policy is 0.78, relative to 1.07 if the natural rate is assumed to be known.

7 Conclusion

This paper studied the joint problem of model estimation and the robust monetary policy in an environment in which the policymaker is uncertain as to the true model of movements in the natural rates of interest and unemployment. We show that the costs of ignoring natural rate uncertainty can be very large. Thus, there is a danger that policymakers could fall into a similar pattern of mistakes 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) . However, 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, where the weights on past data decline very gradually, of the real interest rate and the unemployment rate yield surprisingly good estimates of the respective natural rates.

A key finding is that in the face of uncertainty about natural rates, the robust policy responds much more aggressively to inflation than would be optimal if natural rates were known. This finding reinforces 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, we have assumed that the policymaker does not update their priors over the various natural rate data generating processes based on the incoming data. An interesting extension of the model would allow such 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. "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.
- 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
- Gordon, Robert J. “Foundations of the Goldilocks Economy: Supply Shocks and the Time-Varying NAIRU.” *Brookings Papers on Economic Activity*, 2, 1998, 297–333.
- Giannoni, Marc P. and Michael Woodford. “Optimal Inflation Targeting Rules.” in Michael Woodford and Ben Bernanke (eds.) *Inflation Targeting*, Chicago: University of Chicago Press, 2004.
- 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.
- Kozicki, Sharon and Peter A. Tinsley. “Shifting Endpoints in the Term Structure of Interest Rates.” *Journal of Monetary Economics*, 47, June 2001, 613–52.
- 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. “Should Monetary Policy Respond Strongly to Output Gaps?” *American Economic Review*, 91(2), May 2001, 258–262.
- Onatski, Alexei and Noah Williams. “Modeling Model Uncertainty.” *Journal of the European Economics Association*, 1(5), 2003.
- Onatski, Alexei and Noah Williams. “Empirical and Policy Performance of a Forward-Looking Monetary Model.” Columbia University manuscript, March 2004.
- 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.) *Inflation Targeting*, 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.
- 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.
- Woodford, Michael. *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton University Press, 2004.