

Robust Monetary Policy with Imperfect Knowledge

Athanasios Orphanides
Board of Governors of the Federal Reserve System
and
John C. Williams*
Federal Reserve Bank of San Francisco
October 25, 2006

Abstract

We examine the performance and robustness properties of monetary policy rules in an estimated macroeconomic model in which the economy undergoes structural change and where private agents and the central bank possess imperfect knowledge about the true structure of the economy. Private agents rely on an adaptive learning technology to form expectations and update their beliefs based on incoming data. Policymakers follow an interest rate rule aiming to maintain price stability and to minimize fluctuations of unemployment around its natural rate but are uncertain about the economy's natural rates of interest and unemployment. We show that in this environment the scope for economic stabilization is significantly reduced relative to an economy under rational expectations with perfect knowledge. Furthermore, policies that would be optimal under perfect knowledge can perform very poorly when knowledge is imperfect. Efficient policies that take account of private learning and misperceptions of natural rates call for greater policy inertia, a more aggressive response to inflation, and a smaller response to the perceived unemployment gap than would be optimal if everyone had perfect knowledge of the economy. We show that such policies are quite robust to potential misspecification of private sector learning and the magnitude of variation in natural rates.

KEYWORDS: Monetary policy, natural rate misperceptions, rational expectations, learning.

JEL Classification System: E52

Correspondence: Orphanides: Federal Reserve Board, Washington, D.C. 20551, Tel.: (202) 452-2654, e-mail: Athanasios.Orphanides@frb.gov. Williams: 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. *This draft has been prepared for the Carnegie-Rochester Conference on Public Policy, Pittsburgh, November 10-11, 2006. The opinions expressed are those of the authors and do not necessarily reflect views of the Board of Governors of the Federal Reserve System or the management of the Federal Reserve Bank of San Francisco.

1 Introduction

To paraphrase Clausewitz, monetary policy is conducted in a fog of uncertainty. Our understanding of many key features of the macroeconomic landscape remains imperfect, and the landscape itself evolves over time. As emphasized by Taylor (1993) and McCallum (1999), a crucial requirement for a monetary policy rule is that its good performance be robust to various forms of model misspecification. In this view, it is not enough for a monetary policy rule to be optimal in one specific model, but instead it must be “stress tested” in a variety of alternative model environments before one can conclude with any confidence that the policy is likely to perform well in practice.¹ In this paper, we examine the performance and robustness properties of monetary policy rules in the context of fundamental uncertainty related to the nature of expectations formation and structural change in the economy. Our goal is to identify characteristics of policy rules that are robust to these types of imperfect knowledge, as well as to identify those that are not.

The first form of uncertainty facing the policymaker that we consider relates to the way in which agents form expectations. There is a growing literature that analyzes a variety of alternative models of expectations formation. The key conclusion we take from our reading of this literature is that there is a great deal of uncertainty regarding exactly how private expectations are formed. In particular, the standard assumption of rational expectations may be overly restrictive for monetary policy analysis, especially in the context of an economy undergoing structural change. But, the available evidence does not yet provide unequivocal support for any other single model of expectations formation. Therefore, fundamental uncertainty about the nature of expectations formation appears to be an unavoidable aspect of the policy environment facing central banks face today.

In this paper, we consider two popular alternative models of private expectations formation. Our approach can easily be extended to incorporate other alternative models of expectations as well, but for reasons of tractability, we leave this for future research. One

¹For past applications of this approach, see Levin, Wieland, and Williams (1999, 2003) and Levin and Williams (2003), who study the characteristics of monetary policy rules that are robust to model uncertainty related to macroeconomic dynamics.

model is rational expectations, which assumes that private agents know all the parameters of the model and forms expectations accordingly. This, of course, is the model used in much of the recent monetary policy rule literature. The second model is perpetual learning, where it is assumed that agents do not know the true parameters of the model, but instead continuously reestimate a forecasting model (see Sargent (1999) and Evans and Honkapohja (2001) for expositions of this model). This form of learning represents a relatively modest, and arguably realistic, deviation from rational expectations. An advantage of the perpetual learning framework is that it allows varying degrees of deviations in expectations formation relative to the rational expectations benchmark, which are characterized by variation in a single model parameter. As shown in Orphanides and Williams (2004, 2005a), the resulting process of perpetual learning on the part of economic agents introduces an additional layer of interaction between monetary policy, expectations, and economic outcomes.

The second source of uncertainty that we consider is unobserved structural change, which we represent in the form of low-frequency variation in the natural rates of unemployment and interest. The equilibrium of our model economy is described in terms of deviations from these natural rates. In particular, the inflation rate is in part determined by the unemployment gap, the deviation of the unemployment rate from its natural rate. Similarly, the unemployment rate gap is determined in part by the real interest rate gap, the difference between the real short-term interest rate and the real natural rate of interest. We assume that the central bank does not observe the true values of the natural rates and, indeed, is uncertain about the processes generating the natural rates.

Natural rate uncertainty presents a difficulty for policymakers who follow an interest rate rule with the goal of maintaining price stability and minimizing fluctuations of unemployment around its natural rate. With perfect knowledge of natural rates, the setting of policy would ideally account for the evolution of the economy's natural rates. But, if policymakers do not know the values of the natural rates of interest and unemployment when they make policy decisions, they must either rely on inherently imprecise real-time estimates of these rates for setting the policy instrument, or, alternatively, eschew natural

rates altogether and follow a policy rule that does not respond to natural rate estimates.

The evidence suggests there exists considerable uncertainty regarding the natural rates of unemployment and interest and ambiguity about how best to model and estimate natural rates, even with the benefit of hindsight.² Indeed, the measurement of the natural rate of output has been a key issue in U.S. monetary policy debates in both the 1970s and 1990s, and uncertainty about the natural rate of interest has been the topic of increasing discussion. The evidence indicates that substantial misperceptions regarding the economy's natural rates may persist for some time, before their presence is recognized. In the meantime, policy intended to be contractionary may actually inadvertently be overly expansionary, and vice versa. Moreover, in an environment where the private sector is learning, the learning process can interact with the policy errors and feed back to economic outcomes, as pointed out Orphanides and Williams, (2004, 2005a, 2005b) and Gaspar, Smets and Vestin (2005).

We examine the effects and policy implications of imperfect knowledge of expectations formation and unknown time-varying natural rates using a quarterly model of the U.S. economy estimated over 1981–2004. We first consider the performance and robustness characteristics of simple operational monetary policy rules under perfect and imperfect knowledge. We then analyze the characteristics and performance of policy rules optimized taking into account model uncertainty about expectations formation and natural rate uncertainty. We approach this problem of optimal policy under uncertainty from both a Bayesian and min-max perspectives and compare the results.

Our analysis yields several key findings. First, the scope for economic stabilization in our model with imperfect knowledge is significantly reduced relative to the economy under perfect knowledge (where private agents and the central bank know are assumed to know all features of the model). Second, monetary policies that would be optimal under

²See, for instance, Staiger, Stock, and Watson (1997), and Orphanides and Williams (2002) for documentation of the difficulties associated with the measurement of the natural rate of unemployment and real-time estimates of the unemployment gap; Orphanides and van Norden (2002), and van Norden (2002) for the related problem regarding the output gap; and Laubach and Williams (2002), Orphanides and Williams (2002) and Clark and Kozicki (2004) for the errors in real-time estimates of the natural rate of interest.

perfect knowledge can perform very poorly when knowledge is imperfect. Third, the optimal Bayesian policy under uncertainty performs very well across all of our model specifications and is therefore highly robust to the types of model uncertainty that we examine here. This policy features greater policy inertia, a larger response to inflation, and a smaller response to the perceived unemployment gap than would be optimal under perfect knowledge. Fourth, the resulting optimal Bayesian and min-max policies are closely related to policies that target the price level as opposed to the inflation rate, despite the fact that we assume that the objective is to stabilize the inflation rate not the price level.

The remainder of the paper is organized as follows. Section II discusses the problems for monetary policy caused by natural rate mismeasurement. Section III briefly describes the estimated macro model. Section IV describes the class of monetary policy rules that we study. Section V presents the models of expectations formation and natural rate estimation. Section VI provides details on the simulation methodology. Section 7 analyzes monetary policy under different models of expectations formation, but assuming constant natural rates. Section 8 explores the joint effects of alternative models of expectations and time-varying natural rates. Section 9 examines the optimal Bayesian and min-max policies. Section 10 concludes.

2 Natural Rates, Misperceptions, and Policy Errors

We start our analysis with an illustration of some of the difficulties presented by the evolution of the economy's natural rates. To highlight the role of natural rate misperceptions and the role of policy in propagating them in the economy, consider a generalization of the simple policy rule proposed by Taylor (1993). Let i_t denote the short-term interest rate employed as the policy instrument, (the federal funds rate in the United States), π_t the rate of inflation, and u_t the rate of unemployment, all measured in quarter t . The classic Taylor rule can then be expressed by

$$i_t = \hat{r}_t^* + \pi_{t-1} + \theta_\pi(\pi_{t-1} - \pi^*) + \theta_u(u_{t-1} - \hat{u}_t^*), \quad (1)$$

where π^* is the policymaker's inflation target and \hat{r}_t^* and \hat{u}_t^* are the policymaker's latest estimates of the natural rates of interest and unemployment, based on information available during period t . Note that in this formulation, we restrict attention to the operational version of the Taylor rule recognizing that, as a result of reporting lags, the latest available information about actual inflation and economic activity in period t regards the previous period, $t - 1$. Note also that here we consider a variant of the Taylor rule that responds to the unemployment gap instead of the output gap for our analysis, recognizing that the two are related by Okun's (1962) law.³ In his 1993 exposition, Taylor examined response parameters equal to 1/2 for both the inflation gap and the output gap. Using an Okun's coefficient of 2, this corresponds to setting $\theta_\pi = 0.5$ and $\theta_u = -1.0$.

The Taylor rule has been found to perform quite well in terms of stabilizing economic fluctuations, at least when the natural rates of interest and unemployment are accurately measured.⁴ However, historical experience suggests that policy guidance from this family of rules may be rather sensitive to misperceptions regarding the natural rates of interest and unemployment. The experience of the 1970s, discussed in Orphanides (2003) and Orphanides and Williams (2005b), offers a particularly stark illustration of policy errors that may result.

Following Orphanides and Williams (2002), we explore two dimensions along which the Taylor rule has been generalized that in combination offer the potential to mitigate the problem of natural rate mismeasurement. The first aims to mitigate the effects of mismeasurement of the natural rate of unemployment by partially (or even fully) replacing the response to the unemployment gap with one to the change in the unemployment rate.⁵ The second dimension we explore is incorporation of policy inertia, represented by the presence of the lagged short-term interest rate in the policy rule. Policy rules that exhibit a sub-

³In what follows, we assume that an Okun's law coefficient of 2 is appropriate for mapping the output gap to the unemployment gap. This is significantly lower than Okun's original suggestion of about 3.3. Recent views, as reflected in the work by various authors place this coefficient in the 2 to 3 range.

⁴See, e.g. the contributions in Taylor (1999), which are also reviewed in Taylor (1999b).

⁵This parallels a modification of the Taylor rule suggested by numerous researchers who have argued in favor of policy rules that respond to the growth rate of output rather than the output gap when real-time estimates of the natural rate of output are prone to measurement error. See, in particular, McCallum (2001), Orphanides (2003b), Orphanides et al. (2000), Leitimo and Lonning (2002), and Walsh (2003).

stantial degree of inertia typically improve the stabilization performance of the Taylor rule in forward-looking models.⁶ As argued by Orphanides and Williams (2002), the presence of inertia in the policy rule also reduces the influence of the estimate of the natural rate of interest on the current setting of monetary policy and, therefore, the extent to which misperceptions regarding the natural rate of interest affect policy decisions. To see this, consider the generalized Taylor rule of the form

$$i_t = \theta_i i_{t-1} + (1 - \theta_i)(\hat{r}_t^* + \pi_{t-1}) + \theta_\pi(\pi_{t-1} - \pi^*) + \theta_u(u_{t-1} - \hat{u}_t^*) + \theta_{\Delta u}(u_{t-1} - u_{t-2}). \quad (2)$$

The degree of policy inertia is measured by $\theta_i \geq 0$; cases where $0 < \theta_i < 1$ are frequently referred to as “partial adjustment”; the case of $\theta_i = 1$ is termed a “difference rule” or “derivative control” (Phillips 1954), whereas $\theta_i > 1$ represents superinertial behavior (Rotemberg and Woodford 1999). These rules nest the classic Taylor rule as the special case when $\theta_i = \theta_{\Delta u} = 0$.⁷

To see more clearly how misperceptions regarding the natural rates of unemployment and interest translate to policy errors it is useful to distinguish the real-time estimates of the natural rates, \hat{u}_t^* and \hat{r}_t^* , available to policymakers when policy decisions are made, from their “true” values u^* and r^* . If policy follows the generalized rule given by equation (2), then the “policy error” introduced in period t by misperceptions in period t is given by

$$(1 - \theta_i)(\hat{r}_t^* - r^*) + \theta_u(\hat{u}_t^* - u_t^*).$$

Although unintentional, these errors could subsequently induce undesirable fluctuations in the economy, worsening stabilization performance. The extent to which misperceptions regarding the natural rates translate into policy induced fluctuations depends on the parameters of the policy rule. As is evident from the expression above, policies that are relatively unresponsive to real-time assessments of the unemployment gap, that is, those with small θ_u

⁶See e.g. Levin et al. (1999, 2002), Rotemberg and Woodford (1999), Williams (2003), and Woodford (2003).

⁷Policy rules that allow for a response to the lagged instrument and the change in the output gap or unemployment rate have been found to offer a simple characterization of historical monetary policy in the United States for the past few decades in earlier studies, e.g. Orphanides and Williams (2003) and Orphanides (2003c).

minimize the impact of misperceptions regarding the natural rate of unemployment. Similarly, inertial policies with θ_f near unity reduce the direct effect of misperceptions regarding the natural rate of interest. That said, inertial policies also carry forward the effects of past misperceptions of the natural rates of interest and unemployment on policy, and one must take account of this interaction in designing policies robust to natural rate mismeasurement.

A limiting case that is immune to natural rate mismeasurement of the kind considered here is a “difference” rule, in which $\theta_i = 1$ and $\theta_u = 0$:

$$i_t = i_{t-1} + \theta_\pi(\pi_t - \pi^*) + \theta_{\Delta u}(u_t - u_{t-1}). \quad (3)$$

As Orphanides and Williams (2002), point out, this policy rule is as simple, in terms of the number of parameters, as the original formulation of the Taylor rule and is arguably simpler to implement in practice since does not require knowledge of the natural rates of interest or unemployment. However, because this type of rule ignores potentially useful information about the natural rates of interest and unemployment, its performance relative to the classic “level” Taylor rule and the generalized rule will depend on the degree of mismeasurement and the structure of the model economy, as we explore below.

3 An Estimated Model of the U.S. Economy

We examine the interaction of natural rate misperceptions, learning, and expectations for the design of robust monetary policy rules using a simple quarterly model motivated by the recent literature on micro-founded models incorporating habit formation in consumption and indexation in price-setting. (Woodford, 2004). The specification of the model is closely related to that in Gianonni and Woodford (2004), Smets (2002) and others.

3.1 The Structural Model

The model consists of the following two structural equations:

$$\pi_t = \phi_\pi \pi_{t+1}^e + (1 - \phi_\pi)\pi_{t-1} + \alpha_\pi(u_t - u_t^*) + e_{\pi,t}, \quad e_\pi \sim \text{iid}(0, \sigma_{e_\pi}^2), \quad (4)$$

$$u_t = \phi_u u_{t+1}^e + (1 - \phi_u)u_{t-1} + \alpha_u(r_t^e - r^*) + e_{u,t}, \quad e_u \sim \text{iid}(0, \sigma_{e_u}^2), \quad (5)$$

where π denotes inflation, u denotes the unemployment rate, u^* denotes the true natural rate of unemployment, r denotes the ex ante short-term real interest rate and r^* the natural real rate of interest.

The “Phillips curve” in this model (equation 4) relates inflation (measured as the annualized percent change in the GNP or GDP price index, depending on the period) during quarter t to lagged inflation, expected future inflation, and the unemployment gap during the current quarter. The parameter ϕ_π measures the importance of expected inflation on the determination of inflation, with $(1 - \phi_\pi)$ capturing the role of indexation. The unemployment equation (equation 5) relates the unemployment rate during quarter t to the expected future unemployment rate and one lag of the unemployment rate and the ex ante real interest rate gap. Here, $(1 - \phi_u)$ reflects the role of habit formation.

For our simulation analysis, we imposed the coefficients $\phi_\pi = \phi_u = 0.5$ on the lead-lag structure of the two equations. We opted to concentrate attention on this case to ensure that expectations are of comparable importance for the determination of the rates of inflation and unemployment in the model. These values for ϕ_π and ϕ_u are the largest allowable by the micro-founded theory developed in Woodford (2003) and are consistent with the empirical findings of Giannoni and Woodford (2004) and others.⁸ To estimate the remaining parameters, as in Orphanides and Williams, (2002) we rely on survey forecasts as proxies for the expectations variable which allows estimation of equations (4) and (5) with ordinary least squares. Specifically, we rely on the mean values of the forecasts provided in the Survey of Professional Forecasters. From this survey, we use the forecasts of the unemployment rate and three-month treasury bill rate as reported. For inflation, we rely on annualized log difference of the GNP or GDP price deflator, which we construct from the forecasts of real and nominal GNP or GDP which are reported in the survey. We posit that the relevant expectations are those formed in the previous quarter; that is, we assume that the expectations determining π_t and u_t are those collected in quarter $t - 1$. This

⁸We note that in the specification shown in equations (4) and (5), the data do not reject the value 0.5 for either ϕ_π or ϕ_u . The unrestricted point estimate of ϕ_π is in fact close to 0.5. However the unrestricted estimate of ϕ_u is noticeably lower.

matches the informational structure in the theoretical models (Giannoni and Woodford, 2004 and Woodford, 2003). Finally, to match the inflation and unemployment data as best as possible with these forecasts, we use first announced estimates of these series. Our primary sources for these data are the Real-Time Dataset for Macroeconomists and the Survey of Professional Forecasters, both currently maintained by the Federal Reserve Bank of Philadelphia (Zarnowitz and Braun (1993), Croushore (1993) and Croushore and Stark (2001)).

Using ordinary least squares, we obtain the following estimates for our model between 1981:4 and 2004:2, where the starting point of this sample reflects the availability of the Survey of Professional Forecasters data for the short-term interest rate.

$$\pi_t = 0.5 \pi_{t+1}^e + 0.5 \pi_{t-1} - 0.192 (u_t^e - u_t^*) + e_{\pi,t}, \quad \hat{\sigma}_{e_\pi} = 1.11 \quad (6)$$

(0.084)

$$u_t = 0.5 u_{t+1}^e + 0.5 u_{t-1} + 0.036 (\tilde{r}_t^e - r^*) + e_{u,t}, \quad \hat{\sigma}_{e_u} = 0.29 \quad (7)$$

(0.017)

The numbers in parentheses are the estimated standard errors of the corresponding regression coefficients. The estimated unemployment equation also includes a constant term that provides an estimate of the natural real interest rate, which is assumed to constant in estimating this equation. The estimated residuals show no signs of serial correlation in the price equation. Some serial correlation is suggested by the residuals of the unemployment equation, but for simplicity we ignore this serial correlation in evaluating the performance of monetary policies.

We model the natural rates as exogenous AR(1) processes independent of all other variables. We assume these processes are stationary based on the finding using the standard ADF test that one can reject the null of nonstationarity of both the unemployment rate and real federal funds rate over 1950–2003 at the 5 percent level. However to capture the near-nonstationarity of the series, we set the AR(1) coefficient to 0.99 and then calibrate the innovation variances to be consistent with estimates of time variation in the natural rates in postwar U.S. data. In particular, we set the innovation standard deviation of the natural

rate of unemployment to 0.07 and that of the natural rate of interest to 0.085. These values imply an unconditional standard deviation of the natural rate of unemployment (interest) of 0.50 (0.60), in the low end of the range of standard deviations of smoothed estimates of these natural rates suggested by various estimation methods (see Orphanides and Williams 2002 for details).

4 Monetary Policy

We evaluate the performance of monetary policies rules using a loss equal to the weighted sum of the unconditional variances of the inflation rate, the unemployment gap, and the change in the nominal federal funds rate:

$$\mathcal{L} = Var(\pi) + \lambda Var(\tilde{u}) + \nu Var(\Delta(i)), \quad (8)$$

where $Var(x)$ denotes the unconditional variance of variable x .⁹ We assume an inflation target of zero percent. As a benchmark for our analysis, we assume $\lambda = 4$ and $\nu = 0.25$. Based on an Okun's gap type relationship, the variance of the unemployment gap is about 1/4 that of the output gap, so this choice of λ corresponds to equal weights on inflation and output gap variability. We consider the sensitivity of our results to alternative specifications.

We complete the structural model by specifying a monetary policy rule according to which the federal funds rate is determined by a generalized Taylor Rule of the form:

$$i_t = \theta_i i_{t-1} + (1 - \theta_i)(\hat{r}_t^* + \pi_{t-1}) + \theta_\pi(\pi_{t-1} - \pi^*) + \theta_u(u_{t-1} - \hat{u}_t^*) + \theta_{\Delta u}(u_{t-1} - u_{t-2}), \quad (9)$$

where \hat{r}_{t-1}^* is the policymaker's real-time estimate of the natural rate of interest and \hat{u}_t^* is the real-time estimate of the natural rate of unemployment. We describe the policymaker's estimation of natural rates in the next section. As mentioned earlier, we used lagged data in the policy rule reflecting the lag in data releases. We focus on this class of four-parameter monetary policy rule because further increases in the number of terms in the policy the

⁹Taken literally, the structural model implies a second-order approximation to consumer welfare that is related to the weighted and discounted sum of expected variances of the *change* in the inflation rate and the *level* and *change* in the unemployment rate gap. For the present purposes, we use a standard specification of the loss used in much of the monetary policy evaluation literature.

rule yield relatively small reductions in the central bank loss. In particular, under rational expectations, the optimized four-parameter rule nearly replicates the first-best loss when $s=0$.

In the following we focus on different versions of this policy rule. In one, all four parameters are freely chosen. We also examine the two alternative simpler, 2-parameter rules that are nested by the generalized rule: The “level” variant, where we constrain θ_i and $\theta_{\Delta u}$ to be zero, and which is closer to the original Taylor rule; and the “difference” variant, where we impose the constraints $\theta_i = 1$ and $\theta_u = 0$.

5 Learning

We assume that private agents form expectations using an estimated forecasting model, and that the central bank forms estimates of the natural rates of interest and unemployment using simple time-series methods. Each period, both private agents and the central bank reestimate their respective models using constant-gain least squares that weighs recent data more heavily than past data. In this way, these estimates allow for time variation in the economy. Following Orphanides and Williams (2004), private agents reestimate their forecasting models each period using a constant gain algorithm that places more weight on recent observations.¹⁰ Given the structure of the model, agents need to forecast inflation, the unemployment rate, and the federal funds rate for up to two quarters into the future.

5.1 Perpetual Learning with Least Squares

Under perfect knowledge with no shocks to the natural rate of unemployment, the predictable components of inflation, the unemployment rate, and the funds rate each depend on a constant, one lag each of the inflation and the ex post real funds rate (the difference between the nominal funds rate and the inflation rate), and one or two lags of the unemployment rate, depending on whether the policy rule responds to just the lagged unemployment gap or also the change in the unemployment rate. We assume that agents

¹⁰See also Sargent (1999), Cogley and Sargent (2001), Evans and Honkapohja (2001), Gaspar and Smets (2002), and Gaspar, Smets and Vestin (2005) for related treatments of learning.

estimate forecasting equations for the three variables using a restricted VAR of the form corresponding to the reduced form of the RE equilibrium with constant natural rates. They then construct multi-period forecasts from the estimated VAR.

Consider the case where policy is described by the Taylor rule. To fix notation, let Y_t denote the 1×3 vector consisting of the inflation rate, the unemployment rate, and the federal funds rate, each measured at time t : $Y_t = (\pi_t, u_t, i_t)$; let X_t be the 5×1 vector of regressors in the forecast model: $X_t = (1, \pi_{t-1}, u_{t-1}, i_{t-1} - \pi_{t-1})$; let c_t be the 4×3 vector of coefficients of the forecasting model. This corresponds to the case of the Taylor rule. In the case of the generalized policy rule, the second lag of the unemployment rate also appears in X_t .

Note that we impose that the forecasting model include only the variables that appear with non-zero coefficients in the reduced form of the rational expectations solution of the model with constant natural rates. In principle, these zero restrictions may help or hinder the forecasting performance of agents in the model. In practice, allowing agents to include additional lags of variables in the forecasting model worsens macroeconomic outcomes. Thus, by imposing this structure, we are likely erring on the side of understating the costs of learning on macroeconomic performance.

Using data through period t , the least squares regression parameters for the forecasting model can be written in recursive form:

$$c_t = c_{t-1} + \kappa_t R_t^{-1} X_t (Y_t - X_t' c_{t-1}), \quad (10)$$

$$R_t = R_{t-1} + \kappa_t (X_t X_t' - R_{t-1}), \quad (11)$$

where κ_t is the gain.

Under the assumption of least squares learning with infinite memory, $\kappa_t = 1/t$, so as t increases, κ_t converges to zero. Assuming a constant natural rate of unemployment, as the data accumulate this mechanism converges to the correct expectations functions and the economy converges to the perfect knowledge rational expectations equilibrium. That is, in our model the perceived law of motion that agents employ for forecasting corresponds to the correct specification of the equilibrium law of motion under rational expectations.

As noted above, to formalize perpetual learning we replace the decreasing gain implied by the infinite memory recursion with a small constant gain, $\kappa > 0$.¹¹ With imperfect knowledge, expectations are based on the perceived law of motion of the inflation process, governed by the perpetual learning algorithm described above.

5.2 Calibrating the Learning Rate

A key parameter for the constant-gain-learning algorithm is the updating rate κ . To calibrate the relevant range for parameter we examined how well different values of κ fit either the expectations data from the Survey of Professional Forecasters, following Orphanides and Williams (2005b). To examine the fit of the Survey of Professional Forecasters (SPF), we generated a time series of forecasts using a recursively estimated VAR for the inflation rate, the unemployment rate, and the federal funds rate. In each quarter we reestimated the model using all historical data available during that quarter (generally from 1948 through the most recent observation). We allowed for discounting of past observations by using geometrically declining weights. This procedure resulted in reasonably accurate forecasts of inflation and unemployment, with root mean squared errors (RMSE) comparable to the residual standard errors from the estimated structural equations, (6) and (7). We found that discounting past data with values corresponding to κ in the range 0.01 to 0.04 yielded forecasts closest on average to the SPF than the forecasts obtained with lower or higher values of κ . In light of these results, we consider $\kappa = 0.02$ as a baseline value for our simulations, but also examine the robustness of policies to alternative values of this parameter. The value $\kappa = 0.02$ is also in line with the discounting reported by Sheridan (2003) as best for explaining the inflation expectations data reported in the Livingston Survey.

5.3 Policymaker's estimation of natural rates

Given the time variation in the natural rates, policymakers need to continuously reestimate these variables in real time. Based on the results of Williams (2004) that found that such a

¹¹In terms of forecasting performance, the “optimal” choice of κ depends on the relative variances of the transitory and permanent shocks, as in the relationship between the Kalman gain and the signal-to-noise ratio in the case of the Kalman filter.

procedure performed well and was reasonably robust to model misspecification, we assume that policymakers use a simple constant gain method to update their natural rates based on the observed rates of unemployment and ex post real interest rates. Thus, policymakers update their estimates of the natural rates of unemployment and interest as follows:

$$\hat{r}_t^* = \hat{r}_{t-1}^* + \zeta_r(i_{t-1} - \pi_{t-1} - \hat{r}_{t-1}^*), \quad (12)$$

$$\hat{u}_t^* = \hat{u}_{t-1}^* + \zeta_u(u_{t-1} - \hat{u}_{t-1}^*), \quad (13)$$

where ζ_r and ζ_u are the updating parameters. We set $\zeta_r = \zeta_u = \zeta = 0.005$, a lower value would imply far greater history of usable data than we possess while a higher value reduces natural rate estimate accuracy. We specify the updating equation for the perceived natural rate of unemployment exactly the same.

The model under imperfect knowledge consists of the structural equations for inflation, the unemployment gap, the federal funds rate (the monetary policy rule), the forecasting model, and the updating rule for the natural rates of interest and unemployment.

6 Simulation Methodology

As noted above, we measure the performance of alternative policies rules based on the central bank loss equal to the weighted sum of unconditional variances of inflation, the unemployment gap, and the change in the funds rate. In the case of rational expectations with constant and known natural rates, we compute the unconditional variances numerically as described in Levin, Wieland, and Williams (1999). In all other cases, we compute approximations of the unconditional moments using stochastic simulations of the model.

6.1 Stochastic Simulations

For stochastic simulations, the initial conditions for each simulation are given by the rational expectations equilibrium with known and constant natural rates. Specifically, all model variables are initialized to their steady-state values, assumed without loss of generality to be zero. The central bank's initial perceived levels of the natural rates are set to their true values, likewise equal to zero. Finally, the initial values of the C and R matrices

describing the private agents' forecasting model are initialized to their respective values corresponding to reduced-form of the rational equilibrium solution to the structural model assuming constant and known natural rates.

Each period, innovations are generated from Gaussian distributions with variances reported above. The innovations are assumed to be serially and contemporaneously uncorrelated. For each period, the structural model is simulated, the private agent's forecasting model is updated and a new set of forecasts computed, and the central bank's natural rate estimate is updated. We simulate the model for 41,000 periods and discard the first 1000 periods to mitigate the effects of initial conditions. We compute the unconditional moments from sample root mean squares from the remaining 40,000 periods (10,000 years) of simulation data.¹²

Private agents' learning process injects a nonlinear structure into the model that may generate explosive behavior in a stochastic simulation of sufficient length for some policy rules that would have been stable under rational expectations. One source of instability is due to the possibility that the forecasting model itself may become unstable. We take the view that in practice private forecasters reject unstable models. Each period of the simulation, we compute the maximum root of the forecasting VAR excluding the constants. If this root falls below the critical value of 0.995, the forecast model is updated as described above; if not, we assume that the forecast model is not updated and the matrices C and R are held at their respective previous period values.¹³

Stability of the forecasting model is not sufficient to assure stability in all simulations. For this reason, we impose a second condition that restrains explosive behavior. In particular, if the inflation rate, nominal interest rate, or unemployment gap exceed in absolute value six times their respective unconditional standard deviations (computed under the assumption of rational expectations and known and constant natural rates), then the vari-

¹²Based on simulations under rational expectations in which we can compute the moments directly, this sample size is sufficient to yield very accurate estimates of the unconditional variances. In addition, testing indicates that 1000 periods is sufficient to remove the effects of initial conditions on simulated second moments.

¹³We chose this critical value so that the test would have a small effect on model simulation behavior while eliminating explosive behavior in the forecasting model.

ables that exceeds this bound is constrained to equal the corresponding limit in that period. These constraints on the model are sufficient to avoid explosive behavior for the exercises that we consider in this paper and are rarely invoked for most of the policy rules we study, particularly for optimized policy rules. An illustrative example is the benchmark calibration of the model with monetary policy given by the Taylor Rule with $\theta_\pi = 0.5$ and $\theta_u = -1$, for which the limit on the forecasting model is binding less than 0.1 percent of the time, and that on the endogenous variables, only about 0.4 percent of the time.

7 Monetary Policy and Learning

We first consider the design of optimal monetary policy in the presence of learning by private agents but assuming that natural rates are constant and known by the policymaker. In this way we can more easily identify the private sector effects of learning in isolation. In the next section, we analyze the case of private learning with time varying natural rates that are unobserved by the policymaker.

7.1 The Effects of Learning under the Taylor Rule

To gauge the effects of learning for a given monetary policy rule, we consider macroeconomic performance under the Taylor Rule under alternative assumptions regarding the public's updating rate, κ . For these exercises, we assume that the policymaker knows the true values of the natural rates of interest and unemployment. Table 1 reports the performance of the Taylor Rule given by $\theta_\pi = 0.4$ and $\theta_{\bar{u}}$. The coefficient on the unemployment gap has the reverse sign is twice the size of the coefficient of 0.5 on the output gap in the standard Taylor rule, the latter modification reflecting the smaller variation in the unemployment gap relative to the output gap. The first row shows the outcomes under rational expectations. The second through fifth rows show the outcomes under learning for values of κ ranging from 0.01 to 0.04 (recall that 0.02 is our benchmark value).

The time variation in the coefficients of the forecasting model determining expectations induces greater variability and persistence in inflation and the unemployment gap. As shown in Table 1, the variability in these variables rises with the learning rate, κ , as does

their first-order unconditional autocorrelation.

In this model, the introduction of learning with constant natural rates induces nearly proportional increases in the variability of inflation and the unemployment gap. For example, in the case of $\kappa = 0.02$, the standard deviation of inflation is 32 percent higher than under rational expectations, and that of the unemployment gap is 33 percent higher. This holds true for other values of κ and stems from the fact that the model equations for inflation and the unemployment rate have identical lead-lag structures. It is worth noting that in other models, the two variables may be affected differently by learning.

The rise in persistence results from the effects of shocks on the estimated parameters of the forecasting model. Consider, for example, a positive shock to inflation. Upon reestimation of the forecasting model, a portion of the shock will pass through to the intercept of the inflation forecasting equation. This raises in the next period the value of expected inflation, which boosts inflation, and so on. If by chance another positive shock arrives, the estimated coefficient on lagged inflation in the forecasting model will be elevated, further raising the persistence of inflation.

A key aspect of learning is that its effects are especially felt in episodes when particularly large shocks or a series of positively correlated shocks occurs. Indeed, the impulse responses to iid shocks in this model are quantitatively little different from those in the model under rational expectations. However, with large or serially correlated shocks, the nonlinear nature of the learning process has profound effects. The unconditional moments thus represent an average of periods in which the behavior of the economy is approximately that described by the rational expectations equilibrium and relatively infrequent episodes in which expectations deviate significantly from that implied by rational expectations. Such “problem” episodes contribute importantly to the deterioration in macroeconomic performance reported in the table.

7.2 Optimized Taylor-style Rules

We now consider the optimal coefficients of the Taylor-style Rule under different assumptions regarding learning. As noted above, for this exercise we assume weights of four on

unemployment gap variability and 0.25 on interest rate variability. Figure 1 and Table 2 report summary results. The first two columns in the table report the optimized coefficients of the policy rules, the third through fifth columns report the standard deviations of the target variables, and the sixth column reports the associated loss, denoted by \mathcal{L}^* . The final column reports the loss under the policy rule optimized under rational expectations, denoted by \mathcal{L}^{RE} , evaluated under the alternative specifications of learning.

In the figure, each panel shows the loss associated with policies for a range of alternative parameters θ_π and θ_u , as shown in the two axes. The top left panel shows the loss under rational expectations. The remaining three panels show the corresponding loss for the same policies under learning.

As can be seen from the figure and table, the optimized Taylor-style rule under rational expectations performs very poorly when the public in fact is learning. If policy is given by the optimal policy assuming rational expectations, the loss under the benchmark value of $\kappa = 0.02$ is nearly 60 percent higher than under the optimized Taylor-style rule policy given in the third row of the table. The problem with the policy rule coefficients chosen assuming rational expectations is the relatively weak response to inflation. This mild response to inflation allows inflation fluctuations to feed into inflation expectations and back to inflation, driving the standard deviation of inflation to 2.8 percent for $\kappa = 0.02$.

A particular problem with the policy optimized assuming rational expectations is that it allows the autocorrelation of inflation to rise, prolonging the response of inflation expectations to any shock. For example, under the optimal policy assuming rational expectations, the first-order autocorrelation of inflation rises from 0.71 under rational expectations to 0.90 under learning with $\kappa = 0.02$ and to 0.93 with $\kappa = 0.04$. Interestingly, the autocorrelation of the unemployment gap is about the same under the policy optimized assuming rational expectations as it is for policies that take account of learning.

The efficient policy response with learning responds more aggressively to inflation relative to the optimal response under rational expectations. In contrast, under learning the response to the unemployment gap is less than or about equal to that under rational ex-

pectations. The stronger response to inflation dampens inflation variability and lowers the autocorrelation of inflation. Indeed, focussing on the outcomes under the optimal policies, the resulting autocorrelation of inflation is only modestly higher under learning than it is under rational expectations. Together, these effects reduce damaging fluctuation in the coefficients of agents' forecasting model. The loss under the optimized Taylor-style rule is 37 percent below that under the Taylor-style rule optimized under the assumption of rational expectations for $\kappa = 0.02$ and 46 percent lower for $\kappa = 0.04$.

The figure also highlights the robustness of the responsiveness to inflation in the rule exhibits an important asymmetry. While near the RE optimal policy the loss is extremely sensitive to θ_π under learning, a similar sensitivity is not evident for the higher values of θ_π that are optimal under learning.

We also conducted the same experiments for a number of alternative parameterizations of the loss function and the results are qualitatively the same as for the benchmark parameterization reported here.

8 Interaction of Learning and Time-varying Natural Rates

We now introduce time variation in natural rates to the model. The learning model of the agents is unchanged. We add the innovations to the natural rates and the central bank's equations for updating their natural rate estimates. Otherwise, the simulation experiments are conducted as above.

8.1 The Effects of Learning and Natural Rate Variation

Table 3 reports the results where monetary policy follows the Taylor Rule. The first set of rows under the heading "s = 0" reports the results where both natural rates are assumed to be constant and known by the policymaker; these results are identical to those reported in Table 1 and provide a point of reference for the results that incorporate time variation in the natural rates. The second set of rows under the heading "s=1" reports the results for the main calibration of the innovation variances. The third set of rows under the heading "s = 2" reports the results where we have doubled the standard deviation of the natural

rate innovations. The layout of the table is the same as Table 1 except that we have added columns reporting the standard deviations of natural rate misperceptions.

Under the benchmark calibration of the innovation variances, the standard deviation of central bank misperceptions of the natural rate of unemployment is 0.6 percentage points, while that of the natural rate of interest ranges between 0.9 and 1.2 percentage points. With higher innovation variances given by $s = 2$, the standard deviation of misperceptions of the natural rate of unemployment increases to about 1.1 percentage point, and of the natural rate of interest rises to between 1.4 to 1.7 percentage points. In all cases, these misperceptions are highly persistent, with first-order autocorrelation of about 0.99. Time varying natural rates inject serially correlated errors to the processes driving inflation, the unemployment rate, and the interest rate. The coefficients of private agents' forecasting model only gradually adjust to changes in the natural rates. Moreover, policymakers themselves are confused about the true level of natural rates and these misperceptions feed back into the coefficient estimates of agents' forecasting model. As a result, these shocks and the feedback through policy back into expectations cause a deterioration in macroeconomic performance. For a given rate of learning, the inclusion of time varying natural rates affects the standard deviations of inflation and the unemployment gap in about the same proportion. The introduction of time-varying natural rates also raises the autocorrelations of inflation and the unemployment rate. Under the Taylor Rule, the persistence of these series exceeds 0.85 for our benchmark calibration and exceeds 0.90 for the calibrations with greater natural rate variation and higher learning rates.

Table 4 reports the optimized Taylor-style rules with learning and time-varying natural rates. The format of the table parallels that of Tables 2 and 3. For comparison, the case of constant natural rates reported in Table 2 is given in the upper part of the table.

For a given rate of learning, time variation in natural rates raises the optimal policy response to inflation and lowers that to the perceived unemployment gap. For example, for $\kappa = 0.02$, the optimal coefficient on inflation rises from 0.53 to 1.07 to 1.21 for $s = 0, 1$, and 2, respectively, and that on the unemployment gap falls from 1.20 to 0.99 to 0.67. The

performance of the optimal Taylor-style rule assuming rational expectations, given in the final column, is truly abysmal in the model with both learning and time-varying natural rates.

Interestingly, for a given positive natural rate innovation variance, the optimal coefficients both on inflation and the unemployment gap are higher the greater is κ . With time-varying natural rates but a low rate of learning, the optimal policy is to dampen the response to the mismeasured unemployment gap and to concentrate on inflation. In this case, expectations help stabilize the unemployment gap even with a modest direct policy response to the gap, as discussed in Orphanides and Williams (2002). But, with a higher rate of learning, noise in the economy, including that related to time-varying natural rates, interferes with the public's understanding of the economy and expectations formation may no longer act as a stabilizing influence. In these circumstances, policy needs to respond relatively strongly to the perceived unemployment gap, even recognizing that this may amplify policy errors owing to natural rate misperceptions. Doing so helps stabilize unemployment expectations and avoids situations where private expectations of unemployment veer away from fundamentals.

Figure 2 presents a graphical summary of the role of time-varying natural rates under learning. The structure is similar to that in Figure 1. The top left panel shows the loss under rational expectations. The remaining three panels show the loss under learning with $\kappa = 0.02$ for different degrees of variation in the natural rates, $s = \{0, 1, 2\}$.

8.2 Optimized Difference Rule

The Taylor-style rule implicitly places a coefficient of one on the perceived natural rate of interest and $-\theta_u$ on the perceived natural rate of unemployment. As discussed in Orphanides et al (2000) and Orphanides and Williams (2002) in forward-looking models with natural rate misperceptions, an alternative specification of a policy rule that does not respond directly to perceived natural rates may perform better than the Taylor-style rule specification. In this subsection, we consider one such specification of a two-parameter policy rule in which θ_i is constrained to equal one, θ_u is constrained to equal zero, and

θ_{π} and $\theta_{\Delta u}$ are freely chosen to minimize the policymaker loss. We refer to policy rules with this specification as “difference” rules. Because the policy rule responds to the lagged first-difference of the unemployment rate, we expand private agents’ forecasting model to include the second lag on the unemployment rate. With this specification, the learning model is identical to the reduced form rational expectations solution of the model with constant natural rates. Table 5 reports the results. The losses resulting under the optimized difference rules are reported in the sixth column under the heading \mathcal{L}_D^* ; for comparison, the loss under the optimized Taylor-style rule is given in the final column of the table.

With time-varying natural rates, the optimized first-difference rule outperforms the optimized Taylor-style rules. The more volatile the natural rates are, the greater the performance advantage of the difference rules over the Taylor-style rules. With constant natural rates, the Taylor-style rules perform better than the difference rules, reflecting the fact that when policymakers have perfect knowledge of the natural rates of interest and unemployment, it pays to use this information in the setting of policy. We conclude that in an environment of imperfect knowledge, difference rules may provide a better simple benchmark for policy than the Taylor-style rule.

As in the case of the Taylor-style rule, both the existence of private sector learning and time variation in natural rates imply stronger optimal responses to inflation relative to rational expectations. The optimal coefficient on the change in the unemployment rate, however, is relatively insensitive to the learning rate and the degree of natural rate variation.

8.3 Optimized Generalized Rules

We now consider a more generalized form of the policy rule that combines elements of both the Taylor rule and the difference rule studied above. The specification is the same as in Orphanides and Williams (2002). The interest rate depends on the lagged interest rate, the lagged inflation rate and perceived unemployment gap, and the lagged change in the unemployment rate. We repeat the experiments described above. Table 6 reports the results. The loss from the optimal generalized policy rule is denoted by \mathcal{L}_4^* ; for comparison, the loss resulting from the optimized difference rule, denoted by \mathcal{L}_D^* , and the optimized

Taylor-style rule, denoted by \mathcal{L}_2^* , are reported in the final two columns of the table.

The optimized four-parameter rules perform significantly better than the optimized Taylor-style rules, especially in the presence of time varying natural rates, and outperforms the simple difference rule, particularly when natural rates are constant. This superior performance is related to three factors. First, this class of rule responds to more information, in particular the lagged funds rate, and thus has an advantage over the simple Taylor rule. Second, by incorporating a near-unity response to the lagged funds rate, the optimal generalized rules nearly completely remove the perceived natural rate of interest from influencing policy. Movements in the true natural rate of interest affect the economy, but there is no direct feedback of central bank misperceptions of the natural rate of interest to the economy. Third, by responding to the change in the unemployment rate as a proxy for the unemployment gap, this specification allows for a strong response to utilization variables without relying exclusively on imperfect measures of the gap.

9 Robust Policy

A striking feature of the results from the generalized policy rule is that the optimal coefficients of the generalized rule do not appear to be very sensitive to the rates of learning that we consider or the magnitude of variation in natural rates, as long as both elements are present. In all cases, the optimal coefficient on the lagged funds rate is near one. The coefficients on inflation and the unemployment gap vary, but are generally of approximately the same size. And the coefficient on the change in the unemployment rate is relatively similar across the different cases. These findings suggest that a single policy should be relatively robust to the alternative specification of the economy considered here.

To examine this more closely, we turn to an examination of robustness of a benchmark policy rule following the methodology in Levin, Wieland, and Williams (1999). An informative benchmark rule may be identified with the optimal policy rule corresponding to an agnostic Bayesian prior when the policymaker does not know which among a range range of models is a better representation of the economy. For our benchmark, we assume that

the policymaker is unsure about both the degree of structural change in the economy, as reflected in variation in natural rates, as well as about how expectations formed, that is whether they are rational or based on adaptive learning. Thus, we assume the policymaker has a flat prior on three possible values of $s = \{0, 1, 2\}$, and on four possible models for expectations, rational and learning with values $\kappa = \{0.01, 0.02, 0.03\}$. The policymaker's objective, then, is to identify the policy rule (8) that minimizes the expected loss (9) accounting for his agnostic prior over the correct model. Note that since he is uncertain about the presence of structural change, the policymaker updates his estimates of natural rates using his updating rules (12) and (13) to set policy.

The optimal Bayesian policy is:

$$i_t = 0.96i_{t-1} + (1 - 0.96)(\hat{r}_t^* + \pi_{t-1}) + 0.69(\pi_{t-1} - \pi^*) - 0.75(u_{t-1} - \hat{u}_t^*) - 2.58(u_{t-1} - u_{t-2}).$$

Note that this is rather similar to the optimal policy corresponding to the optimized rule for $\kappa = 0.02$ and $s = 1$ reported in table 6. Table 7 reports the performance of the economy when this benchmark policy rule is followed for the various alternative specifications of expectations formation and natural rate variation. The last two columns present a summary comparison. The fifth column \mathcal{L} , reports the loss associated with the specification listed in the first column when the optimal Bayesian rule is followed. The last column, \mathcal{L}_{Δ}^* reports the best-obtainable loss from a four-parameter rule optimized to that particular specification of the model, as given in Table 6.

The benchmark Bayesian rule performs very well across all different combinations of parameterizations of learning and natural rate fluctuations. In the parlance of Levin and Williams (2003), the model is reasonably fault tolerant once policy has accounted for some degree of learning and natural rate variation. The relative performance of this rule is actually poorest in the cases of little or no learning and constant natural rates. But, these are states of the world that are associated with the lowest loss so from a robustness perspective, the loss in efficiency in such situations is less worrisome than the outcomes corresponding to the larger losses that might occur under substantial variation in natural rates and learning. Remarkably, the relative performance of the benchmark rule is excellent

for all values of κ for both $s = 1$ and $s = 2$.

We also consider the optimal policy under a min-max criterion. The min-max policy minimizes the loss for the model where κ and s take their maximum values of 0.03 and 2, respectively. The results for this rule are reported in the column labeled $\mathcal{L}_{\Delta}^{\mathcal{M}}$ in Table 8. This rule does very well when κ and s are in the upper halves of their respective allowable ranges, but not very well these parameters are both in the lower halves of the ranges. Indeed, in only two of the twelve cases of parameter combinations that we consider does the min-max policy outperform the optimal Bayesian policy and even in those cases the loss is only very modestly lower.

To examine the robustness of simpler policies than the generalized rule which has four parameters, we also compute the optimal Bayesian and min-max level and difference rules which have only two parameters. In each case we employ the same flat prior over the alternative models of learning and natural rates. The resulting optimal Bayesian rules are:

$$i_t = \hat{r}_t^* + \pi_{t-1} + 1.05(\pi_{t-1} - \pi^*) - 0.75(u_{t-1} - \hat{u}_t^*)$$

$$i_t = i_{t-1} + 0.65(\pi_{t-1} - \pi^*) - 4.12(u_{t-1} - u_{t-2}).$$

In Table 9, we present a summary comparison of the optimal two-parameter rules and the optimal generalized rule. The optimal level rule performs uniformly worse than the optimal difference rule in this comparison. Given a choice among these simple alternatives, the difference rule proves clearly more robust in protecting against the uncertainties regarding expectations formation and natural rates. But the generalized rule, with its added flexibility, delivers better performance especially when s is small.

10 Conclusion

In an environment of imperfect knowledge regarding the potential for structural change in the economy and the formation of expectations, the scope for economic stabilization may be significantly reduced relative to an economy under rational expectations with perfect knowledge. Policies that appear to be optimal under perfect knowledge can perform very

poorly if they are implemented in such an environment. In our model economy, the presence of imperfect knowledge tends to raise the persistence of inflation, partly as a result of the persistent policy errors due to misperceptions of the natural rates and partly as a result of the learning process agents may rely upon to form expectations. This leads to a deterioration in economic performance, especially with regard to a policymaker's price stability objective. Policymakers who recognize the presence of these imperfections in the economy can adjust their policies and protect against this deterioration in economic outcomes.

Efficient policies that take account of private learning and misperceptions of natural rates appear to have two important characteristics. First, and arguably most important, these policies call for more aggressive responses to inflation that would be optimal under perfect knowledge. This tends to confirm the conventional wisdom that associates good central bank policy practice with policies that may appear to stress the role of maintaining price stability more than might appear warranted in simple models of the economy under perfect knowledge. Second, efficient policies exhibit a high degree of inertia in the setting of the interest rate. Indeed, simple difference rules which circumvent the need to rely on uncertain estimates of natural rates in setting policy, appear to be robust to potential misspecification of private sector learning and the magnitude of variation in natural rates. Importantly, it seems possible to design a simple policy rule that can deliver reasonably good macroeconomic performance even in an environment of considerable uncertainty regarding expectations formation and natural rate uncertainty.

References

- Bullard, James and Stefano Eusepi (2003), “Did the Great Inflation Occur Despite Policymaker Commitment to a Taylor Rule?” Federal Reserve Bank of St. Louis Working Paper 2003-012A, June.
- Clarida, Richard, Jordi Gali, and Mark Gertler (1999), “The Science of Monetary Policy,” *Journal of Economic Literature*, 37(4), 1661–1707, December.
- Clarida, Richard, Jordi Gali, and Mark Gertler (2000), “Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory,” *Quarterly Journal of Economics*, 147–180, February.
- Clark, Todd, and Sharon Kozicki (2004), “Estimating Equilibrium Interest Rates in Real Time,” mimeo.
- Cogley, Timothy and Sargent, Thomas (2001), “Evolving Post-World War II U.S. Inflation Dynamics,” in *NBER Macroeconomics Annual*.
- Collard, Fabrice and Harris Dellas (2004), “The Great Inflation of the 1970s,” European Central Bank Working Paper 336, April.
- Congressional Budget Office (2001), “CBO’s Method for Estimating Potential Output: An Update,” Washington, DC: Government Printing Office (August).
- Congressional Budget Office (2002), *The Budget and Economic Outlook: An Update*. Washington DC: Government Printing Office (August).
- Croushore, Dean (1993), “Introducing: The Survey of Professional Forecasters,” Federal Reserve Bank of Philadelphia *Business Review*, November/December, 3–13.
- Croushore, Dean and Tom Stark (2001), “A Real-Time Data Set for Macroeconomists,” *Journal of Econometrics* 105, 111–130, November.
- Cukierman, Alex and Francesco Lippi (2003), “Endogenous Monetary Policy with Unobserved Potential Output,” mimeo, December.
- Evans, George and Honkapohja, Seppo (2001), *Learning and Expectations in Macroeconomics*. Princeton: Princeton University Press.
- Fuhrer, Jeffrey C. and George R. Moore (1995), “Inflation Persistence,” *Quarterly Journal of Economics*, 110(1), 127–59.
- Gaspar, Vitor and Smets, Frank (2002), “Monetary Policy, Price Stability and Output Gap Stabilisation,” *International Finance*, 5(2), Summer, 193–211.
- Gaspar, Vitor and Vestin, David (2005), “Optimal Monetary Policy Under Adaptive Learning,” mimeo, March.
- Goodfriend, Marvin, and Robert King (1997), “The New Neoclassical Synthesis and the Role of Monetary Policy,” *NBER Macroeconomics Annual*, (12), 231–283.
- Laubach, Thomas and John C. Williams (2002), “Measuring the Natural Rate of Interest,” *Review of Economics and Statistics*, 85(4) 1063–1070..
- Levin, Andrew, Volker Wieland and John Williams (1999), “Robustness of Simple Monetary Policy Rules under Model Uncertainty,” in *Monetary Policy Rules*, John B. Taylor (ed.), Chicago: University of Chicago.

- Levin, Andrew, Volker Wieland and John Williams (2003), "The Performance of Forecast-Based Policy Rules under Model Uncertainty," *American Economic Review*, 93(3), 622–645, June.
- McCallum, Bennett T. (1999), "Issues in the Design of Monetary Policy Rules," in Taylor, John B. and Michael Woodford, eds., *Handbook of Macroeconomics*, Amsterdam: Elsevier.
- McCallum, Bennett T. and Edward Nelson (1999), "Performance of Operational Policy Rules in an Estimated Semiclassical Structural Model." in *Monetary Policy Rules*, John B. Taylor (ed.), Chicago: University of Chicago, 15–45.
- Okun, Arthur (1962), "Potential Output: Its Measurement and Significance," in *American Statistical Association 1962 Proceedings of the Business and Economic Section*, Washington, D.C.: American Statistical Association.
- Orphanides, Athanasios (2003a), "Monetary Policy Evaluation With Noisy Information," *Journal of Monetary Economics*, 50(3), 605-631, April.
- Orphanides, Athanasios (2003b), "The Quest for Prosperity Without Inflation," *Journal of Monetary Economics*, 50(3), 633-663, April.
- Orphanides, Athanasios (2003c), "Historical Monetary Policy Analysis and the Taylor Rule," *Journal of Monetary Economics*, 50(5), 983-1022, July.
- Orphanides, Athanasios (2004), "Monetary Policy Rules, Macroeconomic Stability and Inflation: A View from the Trenches," *Journal of Money, Credit and Banking*, 36(2), 151-175, April.
- Orphanides, Athanasios, Richard Porter, David Reifschneider, Robert Tetlow and Frederico Finan (2000), "Errors in the Measurement of the Output Gap and the Design of Monetary Policy," *Journal of Economics and Business*, 52(1/2), 117-141, January/April.
- Orphanides, Athanasios and Simon van Norden (2002), "The Unreliability of Output Gap Estimates in Real Time," *Review of Economics and Statistics*, 84(4), 569–583, November.
- Orphanides, Athanasios and John C. Williams (2002), "Robust Monetary Policy Rules with Unknown Natural Rates", *Brookings Papers on Economic Activity*, 2:2002, 63-118.
- Orphanides, Athanasios and John C. Williams (2004) "Imperfect Knowledge, Inflation Expectations and Monetary Policy," in *The Inflation Targeting Debate*, Ben Bernanke and Michael Woodford (eds.), Chicago: University of Chicago Press.
- Orphanides, Athanasios and John C. Williams (2005a), "Inflation scares and forecast-based monetary policy", *Review of Economic Dynamics*, 8(2), 498-527, April.
- Orphanides, Athanasios and John C. Williams (2005b), "The Decline of Activist Stabilization Policy: Natural Rate Misperceptions, Learning, and Expectations" *Journal of Economic Dynamics and Control*. 29(11), 1927-1950, November.
- Phillips, A. W. (1954), "Stabilisation Policy in a Closed Economy," *Economic Journal*, 290–323, June.
- Roberts, John M. (1997), "Is Inflation Sticky?" *Journal of Monetary Economics*, 39, 173–196.

- Rotemberg, Julio J. and Michael Woodford (1999), "Interest Rate Rules in an Estimated Sticky Price Model," in *Monetary Policy Rules*, John B. Taylor (ed.), Chicago: University of Chicago Press, 57–119.
- Sargent, Thomas J. (1999), *The Conquest of American Inflation*, Princeton: Princeton University Press.
- Sheridan, Niamh (2003), "Forming Inflation Expectations," Johns Hopkins University, mimeo, April.
- Smets, Frank (2000), "What Horizon for Price Stability," European Central Bank Working Paper No 24.
- 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.
- Taylor, John B. (1993), "Discretion versus Policy Rules in Practice," Carnegie-Rochester Conference Series on Public Policy, 39, December, 195–214.
- Taylor, John B. and Michael Woodford, eds., (1999), *Handbook of Macroeconomics*, Amsterdam: Elsevier.
- Taylor, John B., ed., (1999), *Monetary Policy Rules*, Chicago: University of Chicago.
- Taylor, John B. (1999b), "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), 655-679.
- van Norden, Simon (2002), "Filtering for Current Analysis" Bank of Canada Working Paper.
- Walsh, Carl (2003), "Implications of a Changing Economic Structure for the Strategy of Monetary Policy," in *Monetary Policy and Uncertainty: Adapting to a Changing Economy*, Federal Reserve Bank of Kansas City.
- Williams, John C. (2003) "Simple Rules for Monetary Policy," *Federal Reserve Bank of San Francisco Review*, 1-12.
- Woodford, Michael (2003), *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton University Press.
- Zarnowitz, Victor and Phillip A. Braun (1993), "Twenty-two Years of the NBER-ASA Quarterly Economic Outlook Surveys: Aspects and Comparisons of Forecasting Performance," in *Business Cycles, Indicators, and Forecasting*, James H. Stock and Mark W. Watson (eds.), Chicago: University of Chicago Press, 11–84.

Table 1
Performance of the Taylor Rule with Learning
(constant natural rates)

κ	Standard Deviation			First-order Autocorrelation		
	π	$u - u^*$	Δi	π	$u - u^*$	i
RE	1.48	0.54	1.96	0.64	0.77	0.60
0.01	1.68	0.63	1.97	0.72	0.83	0.67
0.02	1.95	0.72	1.99	0.79	0.86	0.75
0.03	2.13	0.79	2.03	0.81	0.88	0.78

Notes: $i_t = \hat{r}_t^* + \pi_{t-1} + 0.5 \times (\pi_{t-1} - \pi^*) - 1 \times (u_{t-1} - \hat{u}_t^*)$

Table 2
Optimized Taylor-style Rule with Learning
 (constant natural rates)

κ	Policy Rule Coefficient		Standard Deviation			First-order Autocorrelation		Loss	
	π	$u - u^*$	π	$u - u^*$	Δi	π	$u - u^*$	\mathcal{L}^*	\mathcal{L}^{RE}
RE	0.16	-1.37	1.61	0.48	1.52	0.71	0.75	4.0	4.0
0.01	0.53	-1.21	1.69	0.61	2.02	0.73	0.82	5.4	7.0
0.02	0.77	-1.20	1.73	0.72	2.34	0.74	0.85	6.4	10.2
0.03	0.89	-1.38	1.81	0.77	2.52	0.75	0.86	7.3	12.5

Notes: $i_t = \hat{r}_t^* + \pi_{t-1} + \theta_\pi(\pi_{t-1} - \pi^*) + \theta_u(u_{t-1} - \hat{u}_t^*)$

Table 3
The Taylor Rule with Learning
(time-varying natural rates)

<i>s</i>	Standard Deviation					First-order Autocorrelation		
	π	$u - u^*$	Δi	$u^* - \hat{u}^*$	$r^* - \hat{r}^*$	π	$u - u^*$	i
<i>RE</i>								
0	1.48	0.54	1.96	—	—	0.64	0.77	0.60
1	1.91	0.56	1.96	0.42	0.50	0.78	0.78	0.74
2	2.82	0.60	1.99	0.83	1.00	0.89	0.78	0.87
$\kappa = 0.01$								
0	1.68	0.63	1.97	—	—	0.72	0.83	0.67
1	2.12	0.84	1.99	0.60	1.17	0.82	0.90	0.82
2	2.72	1.14	1.99	1.07	1.71	0.89	0.94	0.89
$\kappa = 0.02$								
0	1.95	0.72	1.99	—	—	0.79	0.86	0.75
1	2.46	0.90	2.01	0.60	0.96	0.86	0.91	0.85
2	3.14	1.18	2.01	1.09	1.52	0.92	0.94	0.91
$\kappa = 0.03$								
0	2.13	0.79	2.03	—	—	0.81	0.88	0.78
1	2.61	0.94	2.03	0.60	0.92	0.88	0.91	0.86
2	3.34	1.20	2.03	1.09	1.46	0.92	0.94	0.91

Notes: $i_t = \hat{r}_t^* + \pi_{t-1} + 0.5 \times (\pi_{t-1} - \pi^*) - 1 \times (u_{t-1} - \hat{u}_t^*)$

Table 4
Optimized Taylor-style Rule with Learning
(time-varying natural rates)

<i>s</i>	Policy Rule Coefficient		Standard Deviation			Loss	
	π	$u - \hat{u}^*$	π	$u - u^*$	Δi	\mathcal{L}_2^*	\mathcal{L}_2^{RE}
<i>RE</i>							
0	0.16	-1.37	1.60	0.47	1.52	4.0	4.0
1	0.72	-0.65	1.62	0.60	2.23	5.3	20.4
2	1.02	-0.39	1.72	0.68	2.64	6.5	69.5
$\kappa = 0.01$							
0	0.53	-1.21	1.69	0.61	2.02	5.4	7.0
1	0.72	-0.57	1.81	0.89	2.27	7.8	14.0
2	1.04	-0.32	1.88	1.23	2.71	11.4	25.9
$\kappa = 0.02$							
0	0.77	-1.20	1.73	0.72	2.34	6.4	10.2
1	1.07	-0.99	1.82	0.93	2.75	8.7	21.3
2	1.21	-0.67	1.98	1.23	2.94	12.1	34.2
$\kappa = 0.03$							
0	0.88	-1.38	1.81	0.77	2.52	7.3	12.5
1	1.14	-0.97	1.89	0.98	2.87	9.5	24.7
2	1.42	-0.89	2.02	1.23	3.26	12.8	39.1

Notes: $i_t = \hat{r}_t^* + \pi_{t-1} + \theta_\pi(\pi_{t-1} - \pi^*) + \theta_u(u_{t-1} - \hat{u}_t^*)$

Table 5
Optimized Difference Rule with Learning
(time-varying natural rates)

<i>s</i>	Policy Rule Coefficient		Standard Deviation			Loss	
	π	Δu	π	$u - u^*$	Δi	\mathcal{L}_D^*	\mathcal{L}_2^*
<i>RE</i>							
0	0.31	-3.76	1.75	0.58	1.34	4.9	4.0
1	0.39	-3.72	1.74	0.63	1.37	5.0	5.3
2	0.52	-3.63	1.74	0.72	1.43	5.6	6.5
$\kappa = 0.01$							
0	0.49	-4.04	1.79	0.72	1.54	5.9	5.4
1	0.54	-3.94	1.78	0.87	1.54	6.8	7.8
2	0.68	-3.82	1.78	1.15	1.59	9.1	11.4
$\kappa = 0.02$							
0	0.55	-4.07	1.86	0.81	1.59	6.7	6.4
1	0.66	-4.20	1.83	0.94	1.70	7.6	8.7
2	0.84	-4.15	1.81	1.19	1.82	9.8	12.1
$\kappa = 0.03$							
0	0.76	-4.51	1.85	0.92	1.88	7.7	7.3
1	0.85	-4.54	1.83	1.02	1.95	8.5	9.5
2	0.93	-4.18	1.85	1.24	1.91	10.5	12.8

Notes: $i_t = i_{t-1} + \theta_\pi(\pi_{t-1} - \pi^*) + \theta_{\Delta u}(u_{t-1} - u_{t-2})$

Table 6
Optimized Generalized Rule with Learning
(time-varying natural rates)

<i>s</i>	Policy Rule Coefficient				Standard Deviation			Loss		
	<i>i</i>	π	$u - u^*$	Δu	π	$u - u^*$	Δi	\mathcal{L}_4^*	\mathcal{L}_D^*	\mathcal{L}_2^*
<i>RE</i>										
0	0.76	0.18	-0.97	-0.43	1.51	0.47	0.77	3.3	4.9	4.0
1	0.92	0.45	-0.66	-1.72	1.58	0.61	1.08	4.2	5.0	5.3
2	0.92	0.56	-0.41	-2.28	1.64	0.71	1.27	5.0	5.6	6.5
$\kappa = 0.01$										
0	0.85	0.42	-1.01	-1.23	1.60	0.61	1.12	4.4	5.9	5.4
1	0.95	0.52	-0.74	-1.89	1.60	0.83	1.21	5.7	6.8	7.8
2	1.00	0.68	-0.62	-2.50	1.64	1.12	1.41	8.2	9.1	11.4
$\kappa = 0.02$										
0	0.87	0.61	-1.09	-1.78	1.64	0.71	1.43	5.2	6.7	6.4
1	0.95	0.66	-0.78	-2.43	1.66	0.89	1.47	6.5	7.6	8.7
2	1.00	0.85	-0.60	-3.04	1.67	1.17	1.70	9.0	9.8	12.1
$\kappa = 0.03$										
0	0.86	0.64	-1.15	-2.19	1.77	0.75	1.59	6.0	7.7	7.3
1	0.99	0.84	-0.95	-3.01	1.71	0.95	1.80	7.3	8.5	9.5
2	1.00	0.97	-0.71	-3.27	1.74	1.21	1.91	9.8	10.5	12.8

Notes: $i_t = \theta_i i_{t-1} + (1 - \theta_i)(\hat{r}_t^* + \pi_{t-1}) + \theta_\pi(\pi_{t-1} - \pi^*) + \theta_u(u_{t-1} - \hat{u}_t^*) + \theta_{\Delta u}(u_{t-1} - u_{t-2})$

Table 7
Robustness of Optimal Bayesian Generalized Rule

<i>s</i>	Standard Deviation			Loss	
	π	$u - u^*$	Δi	\mathcal{L}_4^B	\mathcal{L}_4^*
<i>RE</i>					
0	1.43	0.62	1.48	4.1	3.3
1	1.51	0.64	1.48	4.4	4.2
2	1.71	0.67	1.49	5.2	5.0
$\kappa = 0.01$					
0	1.51	0.71	1.50	4.8	4.4
1	1.55	0.85	1.50	5.8	5.7
2	1.70	1.11	1.51	8.4	8.2
$\kappa = 0.02$					
0	1.59	0.78	1.52	5.5	5.2
1	1.64	0.90	1.52	6.4	6.5
2	1.84	1.14	1.53	9.1	9.0
$\kappa = 0.03$					
0	1.70	0.85	1.54	6.4	6.0
1	1.77	0.96	1.55	7.3	7.3
2	1.99	1.18	1.57	10.0	9.8

Notes: Each row shows the performance of the economy under alternative assumptions regarding the true mechanism for the formation of expectations (RE and learning with $\kappa = \{0.01, 0.02, 0.03\}$) and variation in natural rates ($s = \{0, 1, 2\}$) when the policymaker follows the optimal Bayesian policy:

$$i_t = 0.96i_{t-1} + (1 - 0.96)(\hat{r}_t^* + \pi_{t-1}) + 0.69(\pi_{t-1} - \pi^*) - 0.75(u_{t-1} - \hat{u}_t^*) - 2.58(u_{t-1} - u_{t-2}).$$

The parameters in this rule minimize the expected loss associated with the alternative assumptions shown under a uniform prior.

Table 8
Robustness of Min-Max Generalized Rule

s	Standard Deviation			Loss		
	π	$u - u^*$	Δi	\mathcal{L}_4^M	\mathcal{L}_4^B	\mathcal{L}_4^*
<i>RE</i>						
0	1.43	0.69	1.86	4.8	4.1	3.3
1	1.46	0.70	1.86	5.0	4.4	4.2
2	1.57	0.73	1.87	5.5	5.2	5.0
$\kappa = 0.01$						
0	1.48	0.77	1.88	5.5	4.8	4.4
1	1.50	0.89	1.88	6.3	5.8	5.7
2	1.57	1.14	1.87	8.5	8.4	8.2
$\kappa = 0.02$						
0	1.54	0.84	1.89	6.1	5.5	5.2
1	1.56	0.94	1.88	6.9	6.4	6.5
2	1.66	1.16	1.88	9.0	9.1	9.0
$\kappa = 0.03$						
0	1.61	0.91	1.90	6.8	6.4	6.0
1	1.65	1.00	1.90	7.6	7.3	7.3
2	1.75	1.20	1.91	9.8	10.0	9.8

Notes: Each row shows the performance of the economy under alternative assumptions regarding the true mechanism for the formation of expectations (RE and learning with $\kappa = \{0.01, 0.02, 0.03\}$) and variation in natural rates ($s = \{0, 1, 2\}$) when the policymaker follows the min-max generalized policy:

$$i_t = i_{t-1} + 0.97(\pi_{t-1} - \pi^*) - 0.71(u_{t-1} - \hat{u}_t^*) - 3.27(u_{t-1} - u_{t-2}).$$

The parameters in this rule minimize the maximum loss expected under any of the alternative assumptions regarding the formation of expectations and variation in natural rates considered.

Table 9
Robustness of Two-Parameter Rules

s	Min-Max		Bayesian	
	\mathcal{L}_T^M	\mathcal{L}_D^M	\mathcal{L}_T^B	\mathcal{L}_D^B
<i>RE</i>				
0	6.0	5.7	5.3	5.2
1	6.2	5.8	5.6	5.3
2	6.9	6.0	6.6	5.7
$\kappa = 0.01$				
0	6.8	6.4	6.6	6.0
1	8.5	7.1	8.1	6.8
2	11.8	9.3	12.2	9.1
$\kappa = 0.02$				
0	7.5	7.1	7.2	6.8
1	9.0	7.9	8.8	7.6
2	12.3	9.8	12.9	9.9
$\kappa = 0.03$				
0	8.2	8.0	8.1	7.7
1	9.7	8.6	10.1	8.5
2	12.8	10.6	13.3	10.7

Notes: Each row compares the loss incurred under alternative assumptions regarding the true mechanism for the formation of expectations and variation in natural rates for two-parameter optimal Bayesian and Min-Max policy rules. The corresponding Min-Max rules for the level, \mathcal{L}_T^M , and difference specifications, \mathcal{L}_D^M , are:

$$i_t = \hat{r}_t^* + \pi_{t-1} + 1.42(\pi_{t-1} - \pi^*) - 0.89(u_{t-1} - \hat{u}_t^*)$$

$$i_t = i_{t-1} + 0.93(\pi_{t-1} - \pi^*) - 4.18(u_{t-1} - u_{t-2})$$

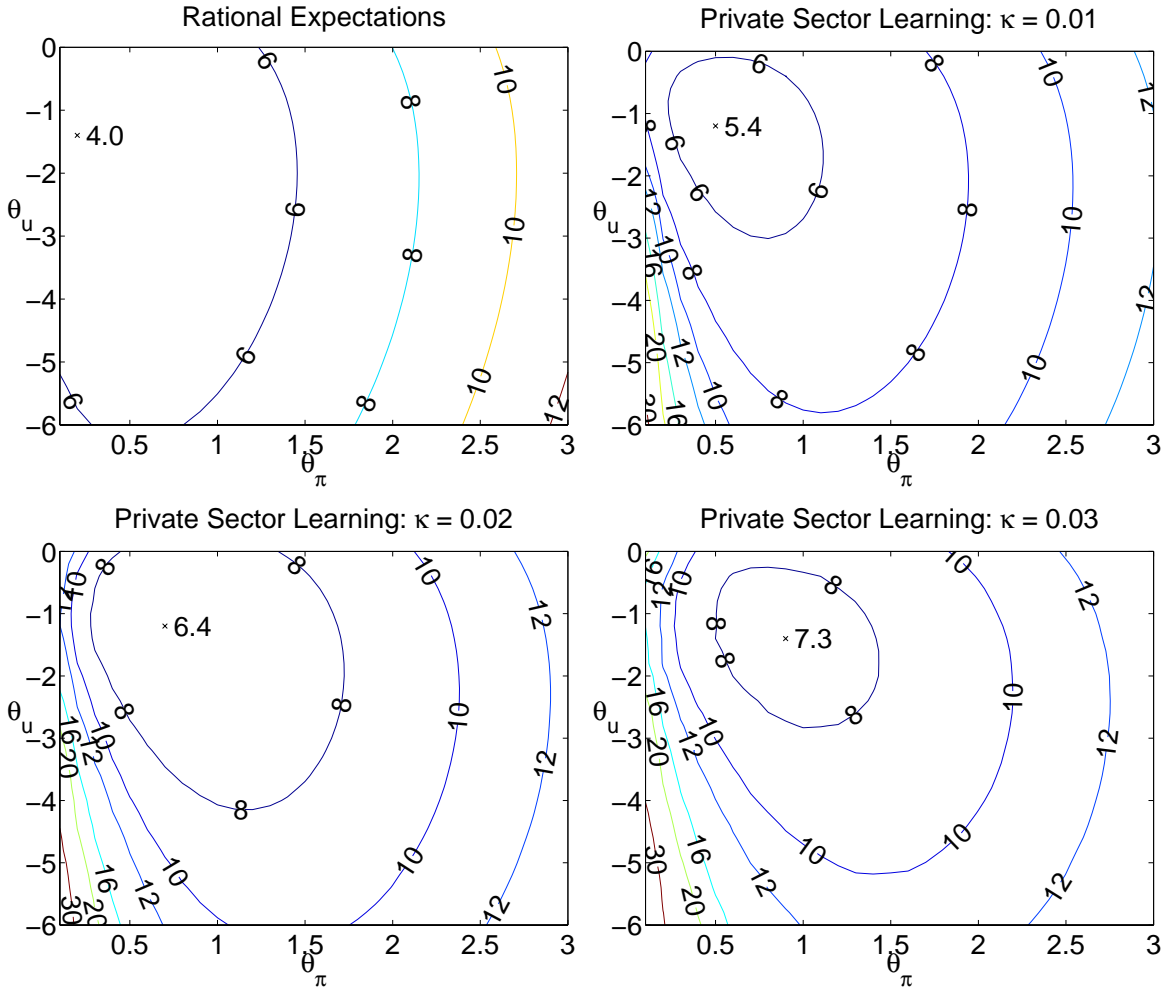
The corresponding Bayesian rules for the level, \mathcal{L}_T^B , and difference specifications, \mathcal{L}_D^B , are:

$$i_t = \hat{r}_t^* + \pi_{t-1} + 1.05(\pi_{t-1} - \pi^*) - 0.75(u_{t-1} - \hat{u}_t^*)$$

$$i_t = i_{t-1} + 0.65(\pi_{t-1} - \pi^*) - 4.12(u_{t-1} - u_{t-2}).$$

Figure 1

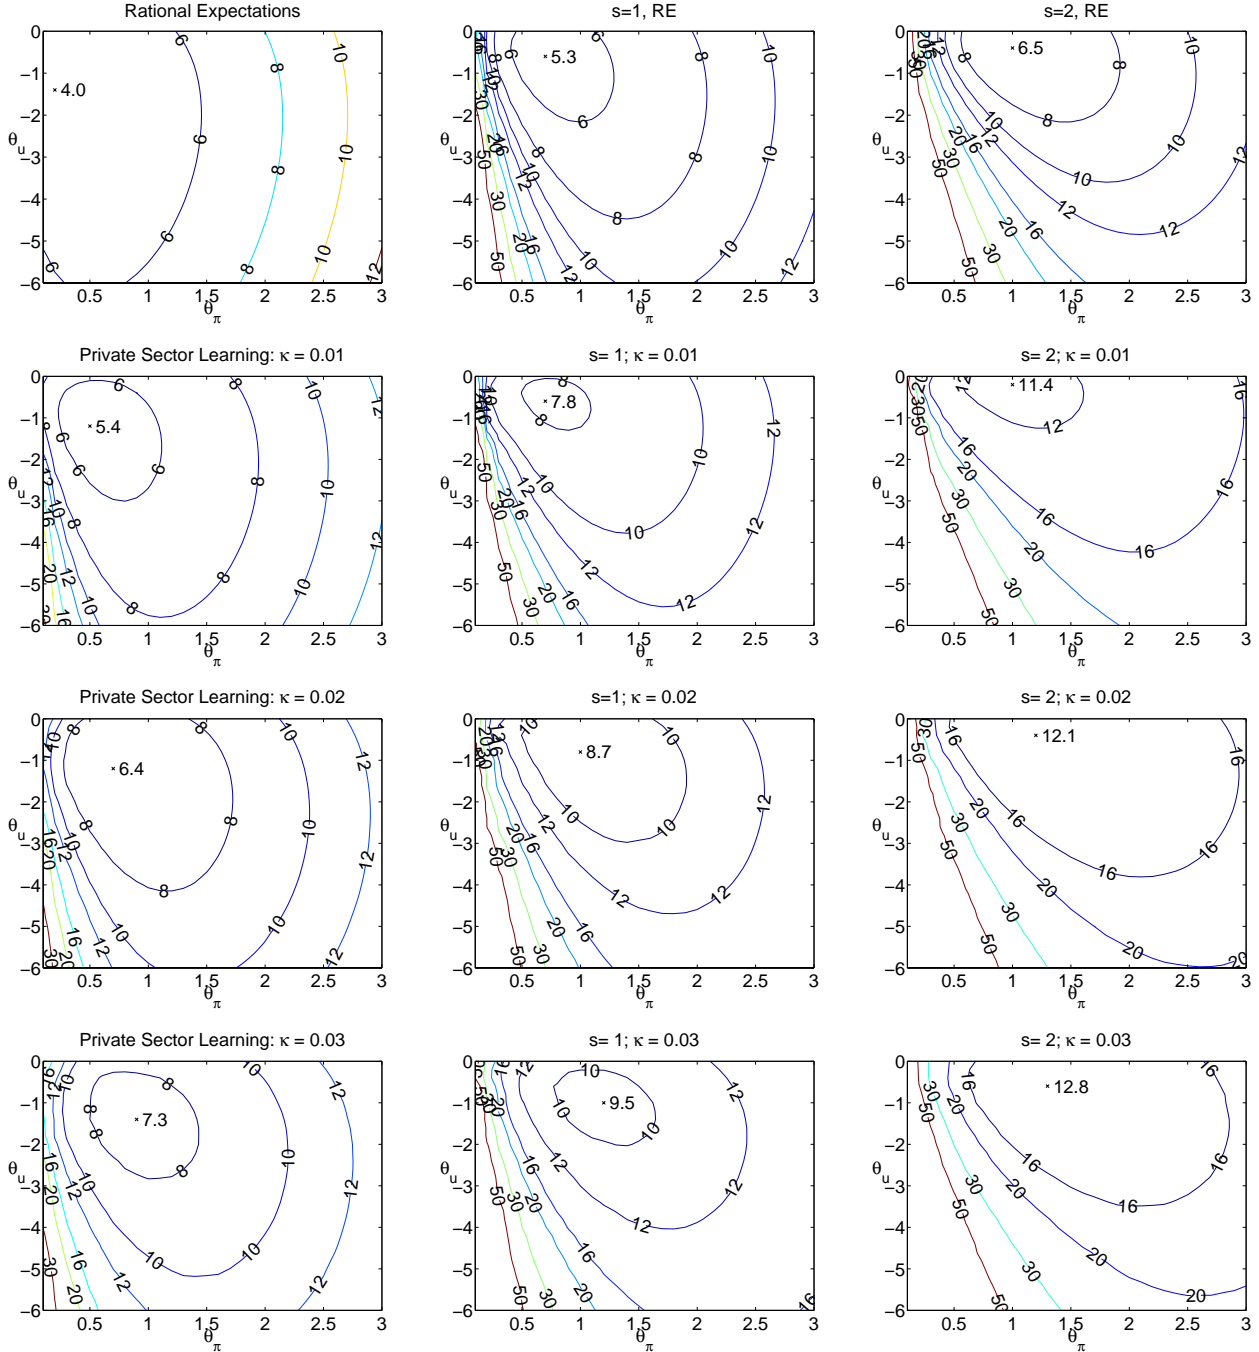
The Taylor Rule Under Learning



Notes: Each panel shows contours of the loss associated with the Taylor-style policy rule $i_t = \hat{r}_t^* + \pi_{t-1} + \theta_\pi(\pi_{t-1} - \pi^*) + \theta_u(u_{t-1} - \hat{u}_t^*)$ for the parameters θ_π and θ_u shown on the axes. The four panels correspond to alternative assumptions regarding expectations formation by the public. The top left panel corresponds to rational expectations with perfect knowledge. In all cases it is assumed that natural rates are constant.

Figure 2

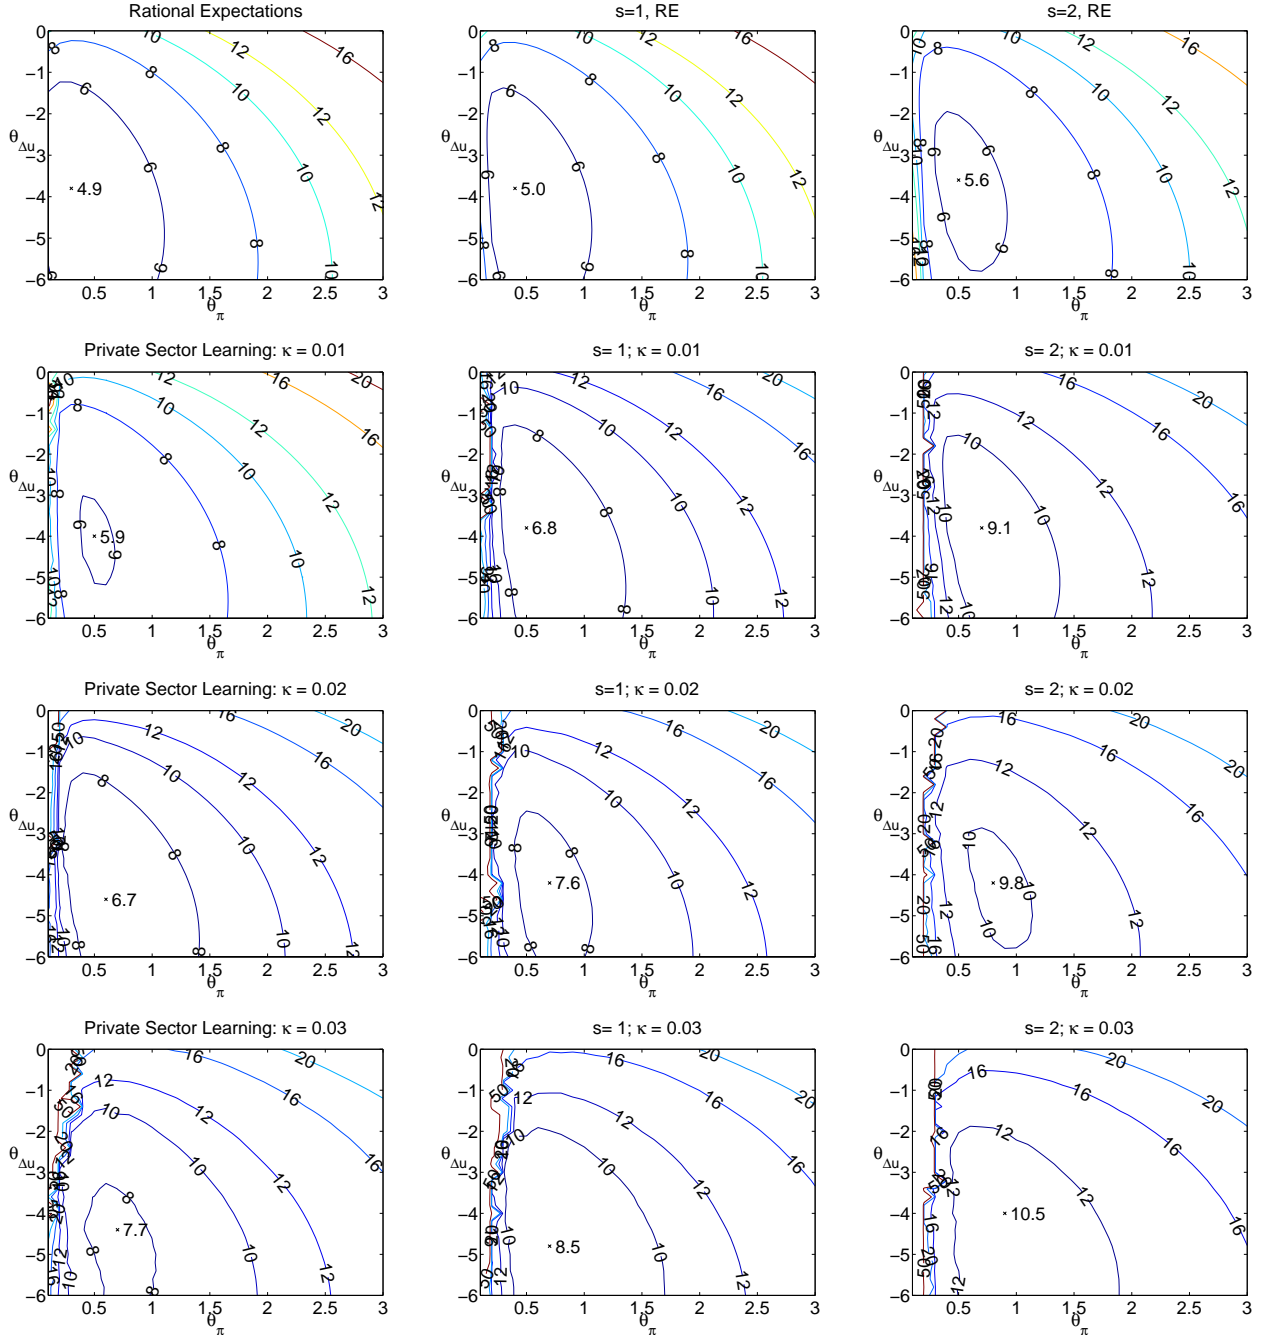
The Taylor Rule Under Learning with Time-Varying Natural Rates



Notes: Each panel shows contours of the loss associated with the Taylor-style policy rule $i_t = \hat{r}_t^* + \pi_{t-1} + \theta_\pi(\pi_{t-1} - \pi^*) + \theta_u(u_{t-1} - \hat{u}_t^*)$ for the assumptions regarding expectations formation and time-variation of the natural rates shown.

Figure 3

The Difference Rule Under Learning and Time-Varying Natural Rates



Notes: Each panel shows contours of the loss associated with the difference rule $i_t = i_{t-1} + \theta_{\pi}(\pi_{t-1} - \pi^*) + \theta_{\Delta u}(u_{t-1} - u_{t-2})$ for the assumptions regarding expectations formation and time-variation of the natural rates shown.