Imperfect Knowledge and the Pitfalls of Optimal Control Monetary Policy

Athanasios Orphanides
Central Bank of Cyprus
and
John C. Williams*
Federal Reserve Bank of San Francisco

November 2007

Abstract

This paper examines the robustness characteristics of optimal control policies derived under the assumption of rational expectations to alternative models of expectations formation and uncertainty about the natural rates of interest and unemployment. We assume that agents have imperfect knowledge about the precise structure of the economy and form expectations using a forecasting model that they continuously update based on incoming data. We also allow for central bank uncertainty regarding the natural rates of interest and unemployment. We find that the optimal control policy derived under the assumption of rational expectations performs relatively poorly when agents learn. These problems are exacerbated by natural rate uncertainty, even when the central bank's natural rates are efficient. We then examine to two types of simple monetary policy rules from the literature that have been found to be robust to model misspecification in other contexts. We find that these policies are robust to the alternative models of learning that we study and natural rate uncertainty and outperform the optimal control policy for empirically plausible parameterizations of the learning and natural rates models.

KEYWORDS: Rational expectations, robust control, model uncertainty, natural rate of unemployment, natural rate of interest.

JEL Classification System: E52

This is a preliminary draft (not for citation) prepared for the Central Bank of Chile Conference, November 15-26, 2007.

*We thank Jim Bullard, Richard Dennis, Andy Levin, Julio Rotemberg, and Mike Woodford for helpful comments and suggestions. The opinions expressed are those of the authors and do not necessarily reflect views of the Central Bank of Cyprus or the management of the Federal Reserve Bank of San Francisco. Correspondence: Orphanides: Central Bank of Cyprus, 80, Kennedy Avenue, 1076 Nicosia, Cyprus, Tel: +357-22714471, email: Athanasios.Orphanides@centralbank.gov.cy. 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.

1 Introduction

For nearly as long as macroeconomic models have existed, economists have proposed applying optimal control theory to the problem of monetary policy (see Chow (1976) for an early example). Support for this approach has waxed and waned in the past, reflecting in part swings in economists' confidence in macroeconometric models. Recently there has been renewed interest among academics and at central banks in applying optimal control to monetary policy, as spelled out in Svensson and Woodford (2003) , Svensson (2002), Woodford (2003), and Gianonni and Woodford (2005). Indeed, as described in Svensson and Tetlow (2006), analytical and computational advances now make it possible to operationalize this approach using the Federal Reserve Board's large-scale nonlinear macroeconomic model. One potential shortcoming of the optimal control approach is that it ignores uncertainty about the specification of the model. In principle, one could incorporate various types of uncertainty to the analysis of optimal policy. However, this is infeasible in practice, given current methods and computational power. As a result, existing optimal control policy analysis is done using a single reference model.

Given the prominence accorded to optimal control in the monetary policy literature and increasingly at central banks, it seems an especially propitious moment to examine the robustness properties of optimal control and other monetary policies when the reference model may be misspecified. The literature on monetary policy under uncertainty has tended to fall into one of two camps. The first, robust control, is closely related to optimal control, but allows local perturbations to the model structure of a general nature. Robust control methods of the type analyzed by Hansen and Sargent (2007), are best suited for relatively modest deviations from the reference model. The second approach, and the one that we follow in this paper, evaluates the performance of monetary policies across a set of possible non-nested models that embed substantive differences in structure, that is, moderate-sized deviations from a reference model.[1] This approach has been advocated by McCallum (1998)

[1]Svensson and Noah Williams (2006) have developed a methodology to compute optimal policy under model uncertainty using a Markov-switching framework; however, computing optimal policies under model uncertainty by this method is extremely computationally intensive and its application to real-world problems

and Taylor (1993), and has been implemented in numerous papers, including Taylor (1999), Levin, Wieland, and Williams (1999, 2003), Orphanides and Williams (2002), and Brock, Durlauf, and West (2005). A key finding in this literature is that properly calibrated simple possible rules that are generalizations of the Taylor Rule (Taylor 1993) can be very robust across a wide set of models. Moreover, optimal control policies can perform very poorly if the reference model is badly misspecified, as shown in Levin and Williams (2003).

In this paper, we reexamine the robustness of optimal control policies designed under the assumption of rational expectations to alternative models of expectations formation and uncertainty about the natural rates of interest and unemployment. The literature has tended to focus on issues of misspecification of the dynamics in structural equations. We abstract from these issues and assume that the basic structure of the central bank's reference model is correctly specified. Instead, we take seriously the information problems facing real-world agents, which may cause expectations to deviate from those implied by the model of the economy they inhabit (see Taylor (1975) for an early analysis of this issue and Sargent (2007) for a recent discussion). Evidence that survey measures of expectations are inefficient and display significant disagreement at each point in time, (see, for example, Mankiw, Reis, and Wolfers (2004) and Williams (2004)), call into question the assumption of rational expectations and suggest the need for monetary policies that are robust to deviations from rational expectations. We therefore assume that agents have imperfect knowledge of the precise structure of the economy and continuously learn by reestimating their forecasting models as new data become available. We consider various learning models that yield very good forecasts in our model economy.

We also allow for exogenous time variation in the natural rates of interest and unemployment that the central bank measures with error. We assume that the central bank has a good understanding of the process describing the evolution of these natural rates, but doe snot observe them directly. Instead, the central bank must estimate the natural rates using available data. We consider both the case where the central bank uses the optimal

is infeasible.

Kalman filter to estimate the natural rates, and the case where the central bank estimate of the key gain parameter of the Kalman filter is incorrect. Laubach and Williams (2003) and Clark and Kozciki (2005) document the imprecision in estimates of the gain parameter in the Kalman filter, making uncertainty about this key parameter a real-world problem for central bank estimates of natural rates.

We compare the performance of the optimal control policy to two types of simple monetary policy rules that have been found to be robust to model uncertainty of various types in the literature. The first is a forward-looking version of a Taylor-type policy rule, of the type that Levin, et al (2003) found to perform very well in a number of estimated rational expectations models of the U.S. economy. The second is rule proposed by Orphanides and Williams (2007) that differs from the first rule in that policy responds to the change in the measure of economic activity, rather than the level. This type of rule has been shown to be robust to mismeasurement of natural rates in the economy (Orphanides and Williams, 2002, 2007) and found to perform very well in a counterfactual analysis of monetary policy during 1996–2003 undertaken by Tetlow (2006). Both of these rules are strikingly parsimonious–they are characterized by only two free parameters.

We find that the optimal control policy constructed assuming rational expectations performs relatively poorly in our estimated model of the U.S. economy when agents do not possess perfect knowledge of the economy but instead must learn. The performance deteriorates further when we additionally allow for natural rate mismeasurement. The optimal control policy attempts to fine tune the economy very precisely, which works well when private expectations are perfectly aligned with those implied by rational expectations. But, when agents learn, expectations can deviate from those implied by rational expectations, and the finely-tuned optimal control policy can go astray. In particular, by implicitly assuming that inflation expectations are always well anchored, the optimal control policy responds insufficiently strongly to movements in inflation, which results in excessive variability of inflation.

In contrast, the two simple monetary policy rules that we study perform very well under learning and with natural rate mismeasurement. These rules clearly outperform the optimal control policy when agents learn. The relatively small advantage that the optimal control policy has over these robust rules when the model is correctly specified implies a small "insurance" payment to gain the sizable robustness benefits found here.

The remainder of the paper is organized as follows. Section 2 describes the model and reports the estimation results. Section 3 describes the central bank objective and the three alternative monetary policies. Section 4 describes the models of expectations formation. Section 5 discusses the simulation methods. Section 6 reports and analyzes the monetary policies under the assumption of rational expectations. Section 7 analyzes the performance of monetary policies with learning. Section 8 considers the robustness of the simple rules to alternative central bank forecasting models. Section 8 concludes.

2 An Estimated Model of the U.S. Economy

Our analysis is conducted using a simple quarterly model motivated by the recent literature on micro-founded models incorporating some inertia in inflation and output (see Woodford, 2003, for a fuller discussion). The specification of the model is closely related to that in Gianonni and Woodford (2005), Smets (2003), and others. The key difference is that instead of the output gap concept in these models, we employ the unemployment gap concept as the cyclical measure of real economic activity. The two concepts are closely related in practice by Okun's law and the properties of the model are largely invariant to this choice. In addition, the empirical problem of measuring the natural rate of unemployment—needed to define the unemployment gap—is essentially similar to the problem of measuring the level of potential output—needed to define the output gap.

2.1 The Model

The structural model consists of two equations that describe the behavior of the unemployment rate and the inflation rate. In addition, there are equations describing the time series properties of the exogenous shocks. To close the model, the short-term interest rate is set by the central bank, as described in the next section.

The "IS curve" equation is motivated by the Euler equation for consumption with adjustment costs or habit:

$$u_t = \phi_u u^e_{t+1} + (1 - \phi_u)u_{t-1} + \alpha_u \left(i^e_t - \pi^e_{t+1} - r^*_t\right) + v_t, \tag{1}$$

$$v_t = \rho_v v_{t-1} + e_{v,t}, \quad e_v \sim \mathrm{N}(0, \sigma^2_{e_v}). \tag{2}$$

We specify the IS equation in terms of the unemployment rather than output to facilitate the estimation of the equation using real-time data. This equation relates the unemployment rate, u_t, to the unemployment rate expected in the next period, one lag of the unemployment rate, and the difference between the expected ex ante real interest rate—equal to the difference between the nominal short-term interest rate, i_t, and the expected inflation rate in the following period, π_{t+1}—and the natural rate of interest, r^*_t. The unemployment rate is subject to a shock, v_t, that is assumed to follow an AR(1) process with innovation variance $\sigma^2_{e_v}$. The AR(1) specification for the shocks is based on the evidence of serial correlation in the residuals of the estimated unemployment equation, as discussed below.

The "Phillips curve" equation is motivated by the New Keynesian Phillips curve with indexation:

$$\pi_t = \phi_\pi \pi^e_{t+1} + (1 - \phi_\pi)\pi_{t-1} + \alpha_\pi(u_t - u^*_t) + e_{\pi,t}, \quad e_\pi \sim \mathrm{N}(0, \sigma^2_{e_\pi}). \tag{3}$$

It relates inflation, π_t, (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, denoted by π^e_{t+1}, and the difference between the unemployment rate, u_t, and and the natural rate of unemployment, u^*_t, during the current quarter. The parameter ϕ_π measures the importance of expected inflation on the determination of inflation, while $(1 - \phi_\pi)$ captures the effects of inflation indexation. The "mark-up" shock, $e_{\pi,t}$, is assumed to be a white noise disturbance with variance $\sigma^2_{e_\pi}$.

In the model simulations, we abstract from time variation in the natural rates of interest and unemployment and assume for convenience that these variables are constant. We further assume that they known by the central bank. See Orphanides and Williams (2007) for analysis of time-varying natural rates in a model with learning.

We model the low frequency behavior of the natural rates of unemployment and interest as exogenous AR(1) processes independent of all other variables:

$$u_t^* = (1-\rho_{u^*})\bar{u}^* + \rho_{r^*} u_{t-1}^* + e_{u^*,t}, \quad e_{u^*} \sim \mathrm{N}(0, \sigma_{e_u^*}^2), \tag{4}$$

$$r_t^* = (1-\rho_{r^*})\bar{r}^* + \rho_{u^*} r_{t-1}^* + e_{r^*,t}, \quad e_{r^*} \sim \mathrm{N}(0, \sigma_{e_r^*}^2). \tag{5}$$

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. The unconditional mean values of the natural rates are irrelevant for our analysis and so we set them both to zero in our analysis.[2]

2.2 Model Estimation and Calibration

We estimate the IS curve and Phillips curve equations using forecasts from the Survey of Professional Forecasters (SPF) as proxies for the expectations that appear in the equations.[3] We assume that expectations are formed in the previous quarter; that is, we assume that the expectations affecting inflation and unemployment in period t are those collected in quarter $t-1$. This matches the informational structure in many theoretical models (see Woodford, 2003, and Giannoni and Woodford, 2005). To match the inflation and unemployment data as best as possible with these forecasts, we use first announced estimates of these series, obtained from the Real-Time Dataset for Macroeconomists maintained by the Federal Reserve Bank of Philadelphia. In estimating the inflation equation, we use the Congressional Budget Office (2001) estimates of the natural rate of unemployment as proxies for the true values. The data sample used in estimation of the model runs from 1969:4 to 2004:2, where the starting date is the first sample point in the SPF.

[2]Because we e ignore the zero lower bound on nominal interest rates as well as any other potential source of nonlinear behavior in the structural model, the unconditional means of variables are irrelevant. Inclusion of the zero bound would severely complicate the analysis and is left for future work.

[3]Specifically, we use the mean forecasts of the unemployment rate and three-month treasury bill rate. We construct inflation forecasts using the annualized log difference of the GNP or GDP price deflator, which we construct from the reported forecasts of real and nominal GNP or GDP. The Survey is currently maintained by the Federal Reserve Bank of Philadelphia. See Croushore (1993) and Croushore and Stark (2001) for details on the survey methodology.

The estimation results are reported below, with standard errors indicated in parentheses. We estimate the IS curve equation using least squares with AR(1) residuals. Unrestricted estimation of the IS curve equation yields a point estimate for ϕ_u of 0.39, with a standard error of 0.15. This estimate is below the lower bound of 0.5 implied by theory; however, the null hypothesis of a value of 0.5 is not rejected by the data.[4] We therefore impose $\phi_u = 0.5$ in estimating the remaining parameters of the equation. Note that the estimated equation also includes a constant term (not shown) that provides an estimate of the natural real interest rate, which is assumed to constant for the purpose of estimating this equation.

$$u_t = 0.5\, u^e_{t+1} + 0.5\, u_{t-1} \underset{(0.022)}{+0.056} (\tilde{r}^e_t - r^*) + v_t, \tag{6}$$

$$v_t = \underset{(0.085)}{0.513} v_{t-1} + e_{v,t}, \quad \hat{\sigma}_{e_v} = 0.30, \tag{7}$$

$$\pi_t = 0.5\, \pi^e_{t+1} + 0.5\, \pi_{t-1} \underset{(0.087)}{-0.294} (u^e_t - u^*_t) + e_{\pi,t}, \quad \hat{\sigma}_{e_\pi} = 1.35, \tag{8}$$

Unrestricted estimation of the Phillips curve equation yields a point estimate for ϕ_π of 0.51, just barely above the lower bound implied by theory.[5] For symmetry with our treatment of the IS curve, we impose the $\phi_\pi = 0.5$ and estimated the remaining parameters using OLS. The estimated residuals for this equation show no signs of serial correlation in the price equation (DW = 2.09), consistent with the assumption of the model.

As discussed in Orphanides and Williams (2002), there is considerable uncertainty regarding the magnitude and persistence of low-frequency fluctuations in the natural rates of unemployment and interest.[6] We do not estimate a model of natural rates in this paper; instead, we calibrate the parameters of the AR(1) processes based on estimates fund elsewhere in the literature. To capture the highly persistent movements in natural rates we set the autocorrelation parameters, ρ_{u^*} and ρ_{r^*}, to 0.99.

[4]This finding is consistent with the results reported in Giannoni and Woodford (2005), who in a similar model, find that the corresponding coefficient is constrained to be at its theoretical lower bound.

[5]For comparison, Giannoni and Woodford (2005) find that the corresponding coefficient is constrained to be at its theoretical lower bound of 0.5.

[6]Recent papers that estimate the natural rates of unemployment and interest include Staiger, Stock, and Watson (1997), Laubach (2001), Laubach and Williams (2003), Clark and Kozicki (2005).

We consider alternative calibrations of the variances of the innovations to the natural rate processes, indexed by the parameter s. In particular, we consider three cases. In the first case, denoted $s = 0$, the variances are assumed to equal zero; that is, the natural rates are constant over time. In the second case, denote by $s = 1$,we calibrate the innovation variances to be consistent with estimates of time variation in the natural rates in postwar U.S. data. Specifically, 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. In the third case, denoted by $s = 2$, we double the standard deviations of the innovations to the natural rate processes, consistent with low end of the range of standard deviations of smoothed estimates of these natural rates suggested by various estimation methods.

3 Optimal Control Monetary Policy

We evaluate the performance of alternative monetary policies under model uncertainty. The monetary policy instrument is the nominal short-term interest rate. We assume that the central bank observes all variables from all previous periods, including private-sector forecasts, when making the current period policy decision. We further assume that the central bank has access to a commitment technology; that is, we study policy under commitment.

The central bank's objective is to minimize 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:

$$\mathcal{L} = Var(\pi - \pi^*) + \lambda Var(u - u^*) + \nu Var(\Delta(i)), \tag{9}$$

where $Var(x)$ denotes the unconditional variance of variable x. We assume an inflation target of zero percent. As a benchmark for our analysis, we assume $\lambda = 4$ and $\nu = 1$. 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.

The optimal control policy is that which minimizes the loss subject to the equations describing the economy. This is constructed, as is typical in the literature and practice, assuming that the policymaker knows the true parameters of the structural model and assumes all agents use rational expectations and the central bank's knows the natural rates of unemployment and interest. Note that for the optimal control policy, as well as the simple policy rules described below, we use lagged information in the determination of the interest rate, reflecting the lag in data releases. The optimal control policy is described by s set of equations that describe the first-order optimality condition for policy and the behavior of the Lagrange multipliers associated with the constraints on the optimization problem implied by the structural equations of the model economy.

Because we are interested in describing the setting of policy in a potentially misspecified model, it is useful to represent the optimal control policy in an equation that relates the policy instrument to observable variables, rather than in terms of Lagrange multipliers that depend on the model. However, there are infinitely many such representations, a subset of which do not yield a determinate rational expectations equilibrium. We consider several alternative representations that are closely related to those that have been studied in the literature.

In the first representation, which we denote "OC," the optimal control policy is described by a feedback rule where the setting of policy depends on the observed past values of the inflation rate, the unemployment gap (the difference between the unemployment rate and the natural rate of unemployment), and the interest rate gap (the difference between the ex post real interest rate and the natural rate of interest). We find that this representation yields a determinate rational expectations equilibrium. We find that including three lags of these variables is sufficient to mimic the optimal control outcome assuming naturals are known. In the following, we focus on this three-lag specification. Note that this formulation implicitly assumes that the central bank uses the structural model with rational expectations to generate forecasts.

The second representation of the optimal control policy is a form of a forecast-targeting policy similar to that proposed by Svensson and Woodford (2003). In principle, this form of the optimal control policy requires the inclusion of infinitely many leads of the objective variables. However, Gianonni and Woodford (2005) show that this policy can be well approximated by including only a few leads of the target variables. As discussed below, we find that a specification in which the policy instrument depends on the first three leads of the inflation rate and the unemployment rate and three lags of the policy instrument yields outcomes under rational expectations nearly identical to those under the optimal control policy. We denote this representation of the optimal control policy by "OC-FT."

3.1 Central Bank Estimation of Natural Rates

Given the time variation in the natural rates, the central bank needs to have real-time estimates of natural rates. We assume that the central bank does not observe the natural rates and must instead estimate them in real time.

We assume that the central bank uses the Kalman filter to estimate both natural rates. Given the assumptions of the model, this is the optimal filter. In particular, the real-time estimate of the natural rate of interest is given by:

$$\hat{u}_t^* = 0.99\,\hat{u}_{t-1}^* + \lambda_u \left\{ \left[\pi_t - (\phi_\pi\,\pi_{t+1}^e + (1-\phi_\pi)\,\pi_{t-1} + \alpha_\pi\,u_t^e)\right]/(-\alpha_\pi) - 0.99\,\hat{u}_{t-1}^* \right\}, \quad (10)$$

where λ_u is the Kalman gain associated that depends on the relative variances of the innovations to inflation and the natural rate of unemployment. The term multiplied by λ_u is the "surprise" inflation, conditional on the prior estimate of the natural rate of unemployment, scaled so that it is in unemployment rate units. The corresponding equation for the central bank estimate of the natural rate of interest is given by:

$$\hat{r}_t^* = 0.99\,\hat{r}_{t-1}^* + \lambda_r \left\{ \left[u_t - (\phi_u\,u_{t+1}^e + (1-\phi_u)\,u_{t-1} + \alpha_u\,(i_t^e - \pi_{t+1}^e))\right]/(-\alpha_u) - 0.99\,\hat{r}_{t-1}^* \right\}, \quad (11)$$

Note that this specification of the updating rules corresponds to the optimal filters for our model. It implicitly assumes that the central bank perfectly knows the specification and slope parameters of the structural model, including the laws of motion of the natural rate of unemployment. As a result, our assumptions represent a best case for the central bank with respect its ability to estimate natural rates. In other work, we examine the implications of model uncertainty regarding the data generating processes for natural rates (Orphanides and Williams 2005, 2007). The optimal values of the gains and the associated unconditional standard deviations of the natural rate mismeasurement are reported in Table 1.

In the following, we allow for some uncertainty regarding the estimation of natural rates. In particular, as noted above, estimates of the Kalman gain tend to be very imprecise. We therefore consider the possibility that the central bank uses incorrect estimates of λ_r and λ_u. Specifically, we examine the effects of the central bank using the optimal gain based on a particular value of s when the true data generating process is given by a different value of s between 0 and 2. For example the central bank may assume that natural rates are highly variables ($s = 2$) in estimating natural rates, when in fact they are constant ($s = 0$).

4 Expectations Formation

Because we are interested in robustness of monetary policies to uncertainty about how expectations are formed, we consider several different models of expectations formation. One model is rational expectations, where private agents are assumed to know all features of the model including the realized values of natural rates. We assume the model with rational expectations is the central bank's reference model that it uses to compute optimal monetary policies. The remaining models that we study involve real-time perpetual learning on the part of private agents. The models differ in the particular perceived laws of motion (PLM) of the economy that agents assume for their forecasting model.

4.1 Perpetual Learning

In the models of learning that we consider, we assume that private agents and, in some cases, the central bank form expectations using an estimated reduced-form forecasting model. Specifically, following Orphanides and Williams (2005a), we posit that private agents engage in perpetual learning, that is they reestimate their forecasting model using a constant-gain least squares algorithm that weights recent data more heavily than past data.[7] This approach to modeling learning allows for the possible presence of time variation in the economy, including the natural rates of interest and unemployment. It also implies that agents' estimates are always subject to sampling variation, that is, the estimates do not eventually converge to fixed values.

We assume agents forecast inflation, the unemployment rate, and the short-term interest rate using a unrestricted vector autoregression model (VAR) containing lags of these three variables and a constant. VAR models are well-suited for our purposes. First, variants of VARs are commonly used in real-world macroeconomic forecasting, making this a reasonable choice on realism grounds. Second, the rational expectations equilibrium of our model implies a reduced-form VAR of this form.

We consider three alternative specifications of the VAR used for forecasting, with lag lengths of one, two, and three quarters. The VAR with three lags nests the reduced-form of the model under the assumptions of rational expectations. In particular, under this assumption, the minimum state space reduced-form of the equilibrium implied by the Phillips and IS curves includes two lags each of the inflation rate and interest rate and three lags of the unemployment rate. The monetary policy rule may imply additional states for the economy, depending on the specification of the rule. For the rules that we consider, the three-lag VAR nests the reduced-form of the rational expectations equilibrium with constant natural rates.[8] We also consider VARs with shorter lag lengths to capture the possibility that agents do not know the true reduced-form structure of the economy. In addition, we know from the forecasting literature that parsimonious VARs can perform better at forecasting in small samples, so agents may optimally choose under-parameterized VARs to improve forecast accuracy.

At the end of each period, agents update their estimates of their forecasting model using

[7]See also Sargent (1999), Cogley and Sargent (2001), and Evans and Honkapohja (2001) for related treatments of learning.

[8]Time-varying natural rates add additional variables to the reduced-form representation of the economy, but as shown below a VAR(3) is a very close approximation to the reduced-form of the economy in that case.

data through the current period. To fix notation, let Y_t denote the 1×3 vector consisting of the inflation rate, the unemployment rate, and the interest rate, each measured at time t: $Y_t = (\pi_t, u_t, i_t)$. For a VAR with l lags, let X_t be the $(3 \cdot l + 1) \times 1$ vector of regressors in the forecast model: $X_t = (1, \pi_{t-1}, u_{t-1}, i_{t-1}, \ldots, \pi_{t-l}, u_{t-l}, i_{t-l})$. Let c_t be the $(3 \cdot l + 1) \times 3$ vector of coefficients of the forecasting model. Using data through period t, the coefficients of the forecasting model can be written in recursive form:

$$c_t = \quad c_{t-1} + \kappa R_t^{-1} X_t (Y_t - X_t' c_{t-1}), \tag{12}$$

$$R_t = \quad R_{t-1} + \kappa (X_t X_t' - R_{t-1}), \tag{13}$$

where κ is the gain. With these estimates in hand, agents construct multi-period forecasts needed for their decisons.

For some specifications of the VAR, R_t may not be full rank. For example, if policy follows the LWW rule and agents form expectations using a VAR(3), then R will be less than full rank under rational expectations. To avoid this problem, in each period of the model simulations, we check the rank of R_t. If it is less than full rank, we assume that agents apply a standard Ridge regression (Hoerl and Kennard, 1970), where R_t is replaced by $R_t + 0.00001 * I(k)$, and k is the dimension of R.

4.2 Calibrating the Learning Rate

A key parameter in the learning model is the private agent updating parameter, κ. Estimates of this parameter tend to be imprecise and sensitive to model specification, but tend to lie between 0 and 0.04.[9] We take 0.02 to be a reasonable benchmark value for κ, a value that implies that the mean age of the weighted sample is about the same as for standard least squares with a sample of 25 years. Given the uncertainty about this parameter, we report results for values of κ between 0.01 (equivalent in mean sample age to a sample of about 50 years) to 0.03 (equivalent in mean sample age to a sample of about 16 years). For comparison, we also report results for the case of $\kappa = 0$. In this case, agents do not update the coefficients of the forecast model. Instead, the coefficient values are fixed at the

[9] See Orphanides and Williams (2005), Milani (2005), Sheridan (2003), and Branch and Evans (2006).

initial values that are set as explained in the next section. We discuss below the forecasting performance of private agent's forecasts under alternative learning rates.

5 Simulation Method

In the case of rational expectations with constant and known natural rates, we compute model unconditional moments 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.

5.1 Stochastic Simulations

For the stochastic simulations, we initialize all model variables to their respective steady-state values, which we assume to be zero. The initial conditions of C and R are set to the steady-state values implied by the forecasting PLM in the rational expectations equilibrium with known natural rates.

Each period, innovations are generated from independent Gaussian distributions with variances reported above. The private agent's forecasting model is updated each period and a new set of forecasts computed, as are the central bank's natural rate estimates. We simulate the model for 44,000 periods and discard the first 4000 periods to eliminate the effects of initial conditions. We compute the unconditional moments from the remaining 40,000 periods (10,000 years) of simulated data.

5.2 The Projection Facility

Private agents' learning process injects a nonlinear structure into the model that may cause the model display explosive behavior in a simulation. In simulations where the model is beginning to display signs of explosive behavior, we follow Marcet and Sargent (1989) and stipulate modifications to the model that curtail the explosive behavior.

One potential source of explosive behavior is that the forecasting model itself may become explosive. We take the view that in practice private forecasters reject explosive models. Correspondingly, in 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 1, 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 values from the previous period.[10] This constraint is typically encountered in less than one percent of the simulation periods; however, in the case of of a high updating rate ($\kappa = 0.03$) and large natural rate variation ($s = 2$), this constraint can be encountered up to three percent of the time.

This constraint on the forecasting model is insufficient to assure that the model economy does not exhibit explosive behavior in all simulations. For this reason, we impose a second condition that eliminates explosive behavior. In particular, the inflation rate, nominal interest rate, and the unemployment gap are not allowed to exceed in absolute value six times their respective unconditional standard deviations (computed under the assumption of rational expectations and known natural rates) from their respective steady-state values. This constraint on the model is invoked extremely rarely in the simulations. However, in some instances, the projection facility is invoked a very high percentage of the time. This occurs because the model gets stuck at the bound of allowable interest rate variability. When that occurs, we adjust the bound on allowable interest rate variability modestly and find a simulation where the projection facility is invoked very rarely. This bound has a relatively small effect on simulation outcomes otherwise, so making this modification has little effect on our results.

6 Performance of the Optimal Control Policy

We first evaluate the performance of the two representations of the optimal control policy in the model assuming rational expectations. We then examine the performance when agents learn.

[10]We 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.

7 Outcomes under Rational Expectations

The impulse responses to the two inflation and unemployment shocks are nearly identical for these two policies under rational expectations, as shown in Figure 1. The implied central bank losses are likewise nearly identical at 6.593. The optimal policy possesses two key features. First, it generates noticeable secondary cycles associated with a a very high degree of policy inertia. Second, the response of the nominal interest rate to the inflation shock is very muted, with the interest rate rising only about 30 basis points in response to a 1.3 percentage point shock to inflation. Under rational expectations, this gradual and mild policy response is sufficient to bring inflation under control due to the fact that expectations are the future course of policy is perfectly understood by the public. As we will see, when these assumptions fail, this approach to policy can have unfortunate consequences.

8 Outcomes under Learning

We now turn to the performance of the different policies when agents learn. We start by evaluating the forecast performance of the various PLMs. We then turn to policy rule evaluation.

8.1 Forecast Model Selection

Table 2 shows the root-mean-squared one-step-ahead forecast errors in the model simulations with constant natural rates under the OC policy for different values of the learning parameters, κ, and the three specifications of the PLM. The first four rows show the results when the public forms expectations using the three-lag VAR, the second four rows show the results when the public forecasts using a two-lag VAR, and the final four rows show the results when the public uses a one-lag VAR. For each case, we report the forecast performance of the various VARs.

Overall, for inflation and unemployment forecasts, all three VARs do about equally well. In fact, the under-parameterized VARs with one and two lags tend to do slightly better than the three-lag VAR, when agents are learning $\kappa > 0$. The interest rate forecast are better with the three-lag VAR, reflecting the fact that the interest rate depends on variables not included in the one- and two-lag VARs.[11]

We now examine whether each forecast model is self-confirming, by which we mean that agents residing in an economy where all other agents used model X for forecasting would also choose model X to forecast. We do this by simulating the model assuming all agents use a VAR with l lags and then compute the forecast errors of the alternative forecast models. The off-diagonal blocks in Table 2 reports the results for the case of the optimal control policy.

Under the OC policy, the VAR(3) is self-confirming in that the RMS forecast errors are equal or larger for the alternative models. That said, the forecasting performance for inflation and unemployment of the VAR(2) model is nearly indistinguishable from that of the VAR(3) model, suggesting agents would on average be close to indifferent between the two models. Interestingly, the VAR(2) model is close to self-confirming as well. If all agents use the VAR(2) for forecasting, forecasts of the unemployment rate and inflation from the VAR(3) and VAR(1) would on average be about as accurate or worse than those form the VAR(2). The VAR(3) does slightly better at forecasting the interest rate. The VAR(1) does not appear to be self-confirming. The other models do better at forecasting unemployment and the interest rate, and only slightly worse at forecasting inflation when all agents use the VAR(1) for forecasting.

Table 3 reports the forecast accuracy of the three VARs in model simulations with constant and known natural rates under the OC policy assuming our benchmark value of $s = 1$ and assuming that the central bank estimates natural rates using the optimal values of the Kalman gains. This table reports only the forecast accuracy for the VAR that is actually used in determining expectations in the model simulations (and corresponds to the diagonal blocks of Table 2). The results are very similar to those for the case of constant

[11] The three-lag VAR encompasses the optimal control policy, so under that policy the interest rate forecasts errors would be zero if it were not for the effects when the projection facility on excessive variability of interest rates is invoked. We experimented with a version that imposes a much more relaxed restriction on interest rate fluctuations. With this modification, the interest rate forecast errors were zero and the performance of the rule was nearly the same as reported in the paper.

natural rates. One difference is that the forecast errors for the short-term interest rate are higher than with constant natural rates. The introduction of natural rate mismeasurement introduces serially correlated monetary policy shocks to the model. These additional shocks interfere with the ability of private agents' to forecast future policy actions.

In summary, all three VAR models seem to be reasonable for forecasting; and on this basis it is hard to dismiss any of them. The one exception is that interest rate forecasts are generally better in the VAR with three lags, but the forecasting accuracy of the other variables often suffers slightly in the VAR when agents are leaning. The VAR(3) and VAR(2) are close to self-confirming; however, the VAR(1) is not. Overall, we view the VAR(1) as the least plausible model.

8.2 Outcomes with Learning

We now examine the performance of the two representations of the optimal control policy computed under the assumption of rational expectations in an environment where agents are learning. We first consider the OC policy, then look at the OC-FT policy.

The behavior of the economy with learning under the OC policy is seen in the impulse responses to inflation and unemployment shocks, shown in Figure 2. For the simulations underlying this figure, we assume that private agents use the three-lag VAR with $\kappa = 0.02$ in forming expectations and that natural rates vary with $s = 1$. We assume that the central bank knows the value of s and uses the optimal Kalman gains. In the model with learning, the impulse response to a shock depends on the initial conditions. We therefore show the distribution of IRFs taken over the unconditional joint distribution of the c and R matrices and the endogenous variables in the model, as described in Orphanides and Williams (2007). Note that these are not confidence bands per se, but only reflect the effects of differing initial conditions on the response to a shock.

When agent learns, the OC policy does not effectively contain movements in inflation. Under rational expectations, the optimal control policy is characterized by a relatively modest rise in interest rates, but still manages to engineer a reduction of inflation through a period of below-target inflation starting about a year after the onset of the shock. However, with learning, the range of responses of inflation to both shocks is very large, indicating that this policy is effective at containing inflation only when agents' expectations formation is close to that implied by the rational expectations equilibrium.

Macroeconomic performance under the OC policy deteriorates with learning, with the magnitude in fluctuations in all three objective variables increasing in the updating rate, κ. Table 4 reports the results from these experiments assuming constant natural rates. The upper part of the table reports results where agents use a three-lag VAR in forming forecasts. The first row in this part of the table reports the results where agents do not learn, but instead hold fixed the coefficients of their forecasting model. Because the three-lag VAR nests the reduced-form of the rational expectations equilibrium, this case corresponds to rational expectations.[12]

The effects of learning under the OC policy are quite large: In the benchmark case of $\kappa = 0.02$, and agents use a three-lag VAR for forecasting, the central bank loss is double what it would be absent learning. The main problem with the optimal control rule is that it is designed to stabilize inflation in an a "perfect" world of rational expectations. Under learning, the modest policy responses to outbreaks of inflation or deflation are insufficient to keep inflation and inflation expectations under strict control. In effect, this policy is designed to "fine tune" policy responses, an approach that breaks down when the assumed structure of the economy turns out to be incorrect.

If agents use under-parameterized VARs for forecasting but do *not* learn, performance is somewhat worse than under rational expectations. Evidently, in this model, the optimal control policy works best if expectations are perfectly aligned with those implied by the policy. Interestingly, with these VAR forecasting models, the deleterious effects of learning are generally smaller in the case of the three-lag VAR. The parsimony of these forecasting models may minimize random fluctuations in the VAR coefficients that tend to plague larger-scale VARs.

[12] Note that the simulated moments reported here differ slightly from those computed analytically and reported in the previous section. These differences reflect the fact that a simulation of 40,000 periods is not sufficient to match unconditional moments *exactly*.

The combination of natural rate mismeasurement and private sector learning causes macroeconomic performance to deteriorate further under the OC policy. Table 5 reports the results for the OC policy when private agents use a three-lag VAR for forecasting for the three values of s. In each case, we assume that central bank knows the value of s and uses the corresponding optimal Kalman gains. As noted above, natural rate mismeasurement by the central bank introduces serially correlated monetary policy shocks to the model. These shocks directly increase aggregate variability and by adding additional noise to private forecasts have an indirect deleterious effect on macroeconomic performance.

The deterioration in macroeconomic performance due to natural rate mismeasurement is much more pronounced if the central bank's Kalman gains are too low. Table 6 reports the outcomes for the three values of s under different assumptions regarding the central bank's estimate of s and thereby the Kalman gains. The costs of over-filtering the data are smaller than those of under-filtering.

We now turn to the performance of the OC-FT policy with learning. With a forecast-based representation of optimal policy, we face of choice of how the central bank makes its forecast that it uses in setting policy. We consider two alternatives. In the first, we assume that the central bank computes its forecast using the structural model assuming that private agents do the same. We refer to this as"RE forecasts." The left half of Table 7 reports the results for this case. The outcomes are nearly the same as for the OC policy. This finding is not surprising. If the OC- and OC-FT policies yielded *identical* outcomes under rational expectations, then if the central bank uses the RE equilibrium of the model to generate forecasts, the outcomes will be identical regardless of how private agents form expectations. This equivalence is due to the fact that the leads in the optimal policy equation can be replaced with the forecasts implied by the reduced-form of the rational expectations equilibrium. In practice, however, the "OC" and "OC-FT" policies yield equilibria under RE that differ ever so slightly, so the outcomes under learning also differ somewhat.

In this model, forecast-targeting optimal control policies that use private-sector forecasts perform very poorly even with constant and known natural rates. The results for this case are shown in the right-hand portion of Table 7. The outcomes under the FT-OC rule using private sector forecasts are uniformly very poor under learning.[13] Indeed, the results are generally much worse than if the central bank used the forecasts implied by the structural model assuming rational expectations. We tried alternative forecast-based specifications of the optimal control policy and the results were qualitatively the same as those reported for the OC-FT policy. We also experimented with using alternative VARs to generate the central bank forecasts and again the results were qualitatively the same. In the following, we focus on the OC policy in analyzing optimal control policies.

9 Simple Rules

So far, we have documented that macroeconomic outcomes deteriorate significantly when the public is learning and the central bank follows a policy that would be optimal if expectations were rational. Of course, this finding alone does not demonstrate that optimal control policies are inadvisable, unless it can be shown that other policies are more robust to alternative models of expectations formation. We therefore consider two alternative monetary policies that have been recommended in the literature for being robust to various forms of model uncertainty.

The first rule is a version of the forecast-based policy rule proposed by Levin, Wieland, and Williams (2003). We refer to this as the "LWW" type of policy rule; according to this rule, the short-term interest rate is determined as follows:

$$i_t = i_{t-1} + \theta_\pi(\bar{\pi}^e_{t+3} - \pi^*) + \theta_u(u_{t-1} - \hat{u}^*_{t-1}), \tag{14}$$

where $\bar{\pi}^e_{t+3}$ is the forecast of the four-quarter change in the price level and u^* is the natural rate of unemployment which we take to be constant and known. Because this policy rule features characterizes policy in terms of the first-difference of the interest rate, it does not rely on estimates of the natural rate of interest, as does the standard Taylor Rule (1993).

[13]Note that in the case of the VAR(3) with $\kappa = 0.03$, the projection facility is invoked very often. For that case, we do not report the results. We modified the conditions for invoking the projection facility and find simulations that do not invoke the projection facility often. Macroeconomic performance was very poor, similar to the other results reported in the table.

The second rule we consider is that proposed by Orphanides and Williams (2007) for its robustness properties in the face of natural rate uncertainty.

$$i_t = i_{t-1} + \theta_\pi(\bar{\pi}^e_{t+3} - \pi^*) + \theta_{\Delta u}(u_{t-1} - u_{t-2}). \tag{15}$$

A key feature of this policy is the absence of any measures of natural rates in the determination of policy.

We choose the parameters of these simple rules to minimize the loss under rational expectations and constant natural rates using a hill-climbing routine.[14] The resulting optimized LWW rule is given by:

$$i_t = i_{t-1} + 1.05\ (\bar{\pi}^e_{t+3} - \pi^*) - 1.39\ (u_{t-1} - \hat{u}^*_t). \tag{16}$$

The optimized OW rule is given by:

$$i_t = i_{t-1} + 1.74\ (\bar{\pi}^e_{t+3} - \pi^*) - 1.19\ (u_{t-1} - u_{t-2}). \tag{17}$$

In the following, we refer to these specific parameterizations of these two rules simply as the "LWW" and "OW" rules. Table 8 reports the outcomes under the optimal control policy, the LWW rule, and the OW rule under rational expectations and constant natural rates. The outcomes under the OC and "OC-FT" policies are shown for comparison.

Under rational expectations and constant natural rates, the optimal control policies yield a loss only modestly lower than that under the LWW rule, a result consistent with the findings in Williams (2003) and Levin and Williams (2003) for other models. The small differences in outcomes between an OC policy and the LWW rule is illustrated in Figure 3, which plots the impulse responses to the two shocks for the OC policy, LWW rule, and the OW rule, under the assumption of rational expectations. The impulse responses under the LWW rule mimic very closely those of the optimal control policy. The only noticeable difference is seen in the responses to the inflation shock. The LWW rule prescribes a sharper initial rise in the nominal short-term interest rate and the unemployment rate than the

[14] If we allow for time-varying natural rates that are *known by all agents*, the optimized parameters of the LWW and OW rules under rational expectations are nearly unchanged. The relative performance of the different policies is also unaffected.

optimal control policy. Despite this, the optimal control policy manages to bring inflation down slightly more quickly owing the expectation of some overshooting of inflation past the target under the optimal control policy.

The difference between the loss under the optimal control policy and the OW rule is somewhat larger than for the LWW rule. In response to the inflation shock, the OW policy acts aggressively to bring inflation back to target, at the cost of larger rise in the unemployment rate. In response to the unemployment shock, this policy, which fails to take into account the level of the unemployment rate, brings the unemployment rate back to target too slowly, causing inflation to fall further below the target.

In contrast to the optimal control policy, the LWW and OW rules perform very well with agents learn. Table 9 compares the performance of these rules to that of the OC policy under learning with known constant natural rates. As in the case of the optimal control policy, the central bank losses are generally larger with the simple rules under learning than it would be absent learning, and the losses with learning are greatest when agents use the three-lag VAR for forecasting. The good performance of the LWW rule is seen clearly in the impulse responses to the shocks shown in Figure 4 for the case of the three-lag VAR and $\kappa = 0.02$; the assumptions are the same as in Figure 2. For both shocks, the range of responses of inflation is much narrower than for the optimal control policy. Thus, the LWW rule consistently brings inflation back to target quickly following a shock to inflation and contains the response of inflation to the unemployment shock. This tighter control of inflation does not come at a cost of a wider range of unemployment responses. The range of responses of the unemployment rate to the two shocks is comparable to those under the optimal control policy. As in the case of the LWW rule, the OW rule effectively contains the inflation responses to the two shocks, as seen in Figure 5 which shows the distribution of IRFs under learning for the OW policy rule. Indeed, it does even better at controlling inflation than the LWW rule, but at a cost of greater variability of the other target variables. As a result, the LWW performs somewhat better in terms of the central bank loss than the OW rule for all learning models that we consider.

With constant natural rates, the LWW rule outperforms the optimal control policy for learning rates of 0.01 and above, reflecting the much better stabilization of inflation under the LWW rule. This result holds regardless of the version of the VAR used for forecasting. The relative performance is seen in Figure 6, which shows the outcomes for values of κ between 0 and 0.03 for the optimal control policy (the solid line), the LWW rule (the dashed line), and the OW rule (the dashed-dotted line) when agents forecast using the three-lag VAR. For very low values of κ, the LWW rule yields slightly higher variability of all three objective variables than the optimal control policy. But, with higher values of κ, the LWW rule responds more effectively to inflation and keeps inflation, and thereby inflation expectations, well contained. It achieves this while allowing somewhat higher variability in the unemployment rate and the change in the interest rate. The results for the OW rule are similar and this rule outperforms the optimal control policy for learning rates of slightly above 0.01 and higher.

The simple rules significantly outperform the OC policy when agents learn and natural rates are mismeasured. Table 10 reports the results with time-varying natural rates. The addition of time-varying natural rates does not change the qualitative results regarding the lack of robustness of optimal control policies relative to the two simple rules that we study. Rather, it amplifies the effects that we found from introducing learning. Figure 7 shows the results for various values of κ for the case of s =1. Both the LWW and OW rules outperform the OC policy for all values learning rates κ above 0.01. Similar results obtain in the case of $s = 2$.

Finally, because these simple rules do not rely much or at all on natural rate estimates, they are robust to misspecification in estimation of natural rates. Table 11 reports the results for the LWW rule for different values of s allowing for the possibility that the central bank's estimate for s and thereby the Kalman gains is incorrect. The results for the OW rule are shown for comparison. For the OW rule this misspecification is irrelevant and the results are invariant to the central bank's estimate of s. For the LWW rule, policy does respond to the perceived natural rate of unemployment and performance generally deteriorates somewhat if the gain is misspecified. Indeed, in the case of $s = 2$ and the central bank erroneously believes natural rates are constant, the OW rule outperforms the LWW for values of κ of 0.02 and above. The performance of the OC policy (reported in Table 6) is much worse than the two simple rules when the Kalman gains are incorrect, reflecting the effects of the greater mismeasurement of the natural rates on policy and thereby the economy.

10 Conclusion

Current techniques for determining optimal control and robust control monetary policies rely on the assumption that the policymaker possess a very good reference model. This assumption is not tenable given the large degree of model uncertainty. This paper has focused on one facet of this uncertainty associated with expectations formation. The main finding is that optimal control policies are not robust to this form of model uncertainty in the estimated model that we study. Of course, our finding does not imply that there does not exist a reference model for which the optimal control policy is robust to the alternative models of expectations formation that we studied here, but it does provide a general warning about the potential pitfalls of optimal control policies when the reference model is misspecified. We also find that mismeaasurement of time-varying natural rates of interest and unemployment exacerbate the problems associated with learning for the optimal control policy. In contrast to optimal control policies, we find that simple rules that have been found to be robust to other types of model uncertainty are also robust to uncertainty about how expectations are formed and natural rate mismeasurement.

Until feasible methods are developed that allow for the derivation of optimal monetary policy under a realistic range of model uncertainty including models with learning and natural rate mismeasurement, the alterative approach of "stress testing" parsimonious policy rules across a wide set of models provides a practical and productive method of learning which characteristics of monetary policies are robust and which are fragile.[15] Of course,

[15] Gaspar, Smets and Vestin (2006) analyze optimal monetary policy in a very simple model with learning. Because the model with learning is nonlinear, they apply dynamic programming techniques that are infeasible

robustness of any policy cannot be “proved,” because the policy may perform poorly in an alternative model that has yet be considered. As Carlson and Doyle (2002) warn “They are robust, yet fragile, that is, robust to what is common or anticipated but potentially fragile to what is rare or unanticipated.” Recognition of this, of course, implies the need for more research into the robustness properties of all monetary policy strategies.

for the type of model studied in this paper.

References

Branch, William A. and George W. Evans (2006), "A Simple Recursive Forecasting Model," *Economics Letters* 91, 158166.

Brock, Williams A., Steven Durlauf and Kenneth West (2007), "Model Uncertainty and Policy Evaluation: Some Theory and Empirics," *Journal of Econometrics*," 127 (2), February, 629-664.

Carlson, JM and and J Doyle (2002), "Complexity and Robustness," *Proceedings of the National Academy of Sciences*, 99, Suppl. 1, February, 2538-2545.

Chow, Gregory C. (1976), "Control Methods for Macroeconomic Policy Analysis." *The American Economic Review*, 66 (2), May, 340-345.

Clark, Todd and Sharon Kozicki (2005), "Estimating Equilibrium Interest Rates in Real Time," *North American Journal of Economics and Finance,* 16(3), 395-413.

Cogley, Timothy and Sargent, Thomas (2001), "Evolving Post-World War II U.S. Inflation Dynamics," in *NBER Macroeconomics Annual*.

Congressional Budget Office (2001), "CBO's Method for Estimating Potential Output: 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 (2005), "Endogenous Monetary Policy with Unobserved Potential Output," *Journal of Economic Dynamics and Control,* 29(11), pp. 1951-83.

Evans, George and Honkapohja, Seppo (2001), *Learning and Expectations in Macroeconomics.* Princeton: Princeton University Press.

Gaspar, Vitor, Frank Smets, and David Vestin (2006). "Adaptive Learning, Persistence, and Optimal Monetary Policy" ECB Working paper, June.

Giannoni, Marc P. and Michael Woodford (2005). "Optimal Inflation Targeting Rules," in *The Inflation Targeting Debate*, Ben Bernanke and Michael Woodford (eds.), Chicago: University of Chicago Press.

Hansen, Lars Peter, and Thomas J. Sargent (2007), *Robustness*, forthcoming, Princeton University Press.

Hoerl, A.E. and R.W. Kennard (1970), "Ridge regression: Biased Estimation of Nonorthogonal Problems," *Tecnometrics*, 12, 69–82.

Laubach, Thomas (2001),"Measuring the NAIRU: Evidence from Seven Economies," *Review of Economics and Statistics*, 83, May, 218-231.

Laubach, Thomas and John C. Williams (2003). "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.

Mankiw, N. Gregory, Ricardo Reis, and Justin Wolfers (2004), "Disagreement about Inflation Expectations," in Mark Gertler and Kenneth Rogoff (ed.) *NBER Macroeconomics Annual 2003*, 18, Cambridge, Mass.: MIT Press.

Marcet, Albert and Thomas J. Sargent (1989), "Convergence of Least Squares Learning Mechanisms in Self-Referential Linear Stochastic Models," *Journal of Economic Theory*, 48 (2), August, 337-368.

McCallum, Bennett T. (1988). "Robustness Properties of a Rule for Monetary Policy." *Carnegie-Rochester Conference Series on Public Policy*, 29, Autumn, 173-203.

Milani, Fabio (2005). "Expectations, Learning, and Macroeconomic Persitence." mimeo, Univeristy of California, Irvine.

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 (2005a) "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 (2007) "Inflation Targeting Under Imperfect Knowledge," in *Monetary Policy Under Inflation Targeting*, Frederic Mishkin and Klaus Schmidt-Hebbel (eds.), Santiago: Central Bank of Chile.

Sargent, Thomas J. (1999), *The Conquest of American Inflation*, Princeton: Princeton University Press.

Sargent, Thomas J. (2007), "Evolution and Intelligent Design," mimeo, New York University, September.

Sheridan, Niamh (2003), "Forming Inflation Expectations," Johns Hopkins University, mimeo, April.

Smets, Frank (2003), "Maintaining Price Stability: How Long Is the Medium Term?" *Journal of Monetary Economics*, September 2003, 50(6), 1293-1309.

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.

Svensson, Lars .E.O. (1999), "Inflation Targeting as a Monetary Policy Rule." *Journal of Monetary Economics*, 43, 607-654.

Svensson, Lars E. O. (2002), "Inflation Targeting: should it be modeled as an instrument rule or a targeting rule?" *European Economic Review*, 46(4/5), 771-180.

Svensson, Lars E. O. and Robert Tetlow (2005,) "Optimum Policy Projections," International Journal of Central Banking,1: 177-207.

Svensson, Lars E. O. and Noah Williams (2007), "Monetary Policy with Model Uncertainty: Distribution Forecast Targeting," mimeo, Princeton University.

Svensson, Lars E. O. and Michael Woodford (2003), "Optimal Indicators for Monetary Policy," *Journal of Monetary Economics*, 46, 229-256.

Taylor, John B. (1975), "Monetary Policy during a Transition to Rational Expectations," *Journal of Political Economy*, 83(5), October, 1009–1021.

Taylor, John B. (1993), "Discretion versus Policy Rules in Practice," Carnegie-Rochester Conference Series on Public Policy, 39, December, 195–214.

Taylor, John B. (1999), "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.

Tetlow, Robert J. (2006), "Real-time Model Uncertainty in the United States: 'Robust' policies put to the test," mimeo,Federal Reserve Board, May 22.

Williams, John C. (2003) "Simple Rules for Monetary Policy," *Federal Reserve Bank of San Francisco Review*, 1-12.

Williams, John C. (2004), "Discussion of 'Disagreement about Inflation Expectations'," in Mark Gertler and Kenneth Rogoff (ed.) *NBER Macroeconomics Annual 2003*, 18, Cambridge, Mass.: MIT Press, 2004, 257-268.

Woodford, Michael (2003), *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton University Press.

Table 1
Natural Rate Estimation

	Optimal Kalman Gain		Unconditional Standard Deviations			
s	λ_r	λ_u	r^*	u^*	$r^* - \hat{r}^*$	$u^* - \hat{u}^*$
0	0	0	0	0	0	0
1	0.0027	0.0083	0.60	0.50	0.56	0.42
2	0.0086	0.0219	1.21	0.99	1.00	0.68

Table 2
Forecast Accuracy with Constant Natural Rates (RMSE)
(OC Policy with Optimal Kalman Gains; Constant Natural Rates)

	Forecast Model								
	VAR(3)			VAR(2)			VAR(1)		
κ	π	u	i	π	u	i	π	u	i
True Forecasting Model VAR(3)									
0.00	1.35	0.30	0.00	1.35	0.30	0.23	1.37	0.35	0.35
0.01	1.38	0.31	0.01	1.38	0.31	0.24	1.41	0.36	0.36
0.02	1.43	0.33	0.09	1.44	0.34	0.26	1.51	0.40	0.40
0.03	1.48	0.35	0.14	1.50	0.36	0.31	1.58	0.43	0.43
True Forecasting Model VAR(2)									
0.00	1.35	0.30	0.00	1.35	0.30	0.23	1.36	0.35	0.43
0.01	1.38	0.31	0.01	1.37	0.31	0.24	1.38	0.36	0.44
0.02	1.41	0.32	0.05	1.39	0.32	0.25	1.44	0.38	0.46
0.03	1.45	0.33	0.09	1.42	0.33	0.26	1.49	0.41	0.48
True Forecasting Model VAR(1)									
0.00	1.36	0.31	0.00	1.36	0.31	0.26	1.36	0.33	0.44
0.01	1.38	0.30	0.00	1.37	0.30	0.23	1.36	0.33	0.46
0.02	1.41	0.31	0.02	1.39	0.31	0.23	1.37	0.33	0.47
0.03	1.44	0.32	0.05	1.41	0.31	0.24	1.38	0.33	0.47

Table 3

Forecast Accuracy with Time-Varying Natural Rates (RMSE)

(OC Policy with Optimal Kalman Gains; $s = 1$)

	Standard Deviation		
κ	π	u	i
	VAR(3)		
0.00	1.36	0.30	0.10
0.01	1.39	0.31	0.07
0.02	1.45	0.34	0.17
0.03	1.50	0.37	0.19
	VAR(2)		
0.00	1.36	0.30	0.31
0.01	1.38	0.31	0.28
0.02	1.40	0.32	0.28
0.03	1.42	0.33	0.30
	VAR(1)		
0.00	1.37	0.35	0.78
0.01	1.37	0.33	0.55
0.02	1.38	0.33	0.54
0.03	1.39	0.34	0.54

Table 4

Performance of OC Policy with Constant and Known Natural Rates

	Standard Deviation			Loss
κ	π	$u - u^*$	Δi	$\mathcal{L}$
		VAR(3)		
0.00	1.84	0.68	1.20	6.65
0.01	2.14	0.76	1.32	8.63
0.02	2.75	0.92	1.57	13.39
0.03	3.15	1.04	1.79	17.45
		VAR(2)		
0.00	1.83	0.68	1.22	6.71
0.01	2.06	0.74	1.29	8.14
0.02	2.42	0.86	1.47	10.93
0.03	2.76	0.97	1.66	14.12
		VAR(1)		
0.00	1.94	0.78	1.43	8.27
0.01	2.15	0.75	1.36	8.75
0.02	2.46	0.84	1.48	11.06
0.03	2.70	0.92	1.59	13.21

Table 5

Performance of OC Policy with Time-Varying Natural Rates

(VAR(3) Forecasts; Optimal Kalman Gains)

	Standard Deviation			Loss
κ	π	$u-u^*$	Δi	$\mathcal{L}$
$s=0$				
0.00	1.84	0.68	1.20	6.65
0.01	2.14	0.76	1.32	8.63
0.02	2.75	0.92	1.57	13.39
0.03	3.15	1.04	1.79	17.45
$s=1$				
0.00	1.85	0.81	1.24	7.56
0.01	2.27	0.88	1.34	10.07
0.02	2.98	1.07	1.66	16.27
0.03	3.38	1.20	1.86	20.61
$s=2$				
0.00	1.89	0.98	1.28	9.02
0.01	2.44	1.06	1.39	12.39
0.02	3.16	1.26	1.75	19.38
0.03	3.71	1.43	2.08	26.20

Table 6

Central Bank Loss under OC Policy with Incorrect Kalman Gain

(VAR(3) Forecasts)

	Central Bank Estimate of s		
κ	0	1	2
	True value: $s = 0$		
0.00	6.65	6.93	7.53
0.01	8.63	8.27	9.33
0.02	13.39	13.58	12.33
0.03	17.45	17.35	16.62
	True value: $s = 1$		
0.00	7.56	7.56	8.14
0.01	12.95	10.07	9.74
0.02	19.28	16.27	13.33
0.03	22.61	20.61	17.94
	True value: $s = 2$		
0.00	10.10	8.98	9.02
0.01	28.51	19.09	12.39
0.02	34.35	27.24	19.38
0.03	37.30	31.74	26.20

Table 7
Performance under Rational Expectations

	Standard Deviation			Loss
Policy	π	$u - u^*$	Δi	$\mathcal{L}$
OC	1.83	0.67	1.20	6.59
OC-FT	1.83	0.67	1.20	6.59
LWW rule	1.87	0.69	1.23	6.93
OW rule	1.83	0.73	1.39	7.40

Table 8

Performance of OC-FT Policy under Learning with Constant Natural Rates

	CB: RE Forecasts				CB: VAR(3) Forecasts			
	Standard Deviation			Loss	Standard Deviation			Loss
κ	π	$u-u^*$	Δi	$\mathcal{L}$	π	$u-u^*$	Δi	$\mathcal{L}$
				VAR(3)				
0.00	1.84	0.68	1.21	6.66	1.84	0.68	1.21	6.66
0.01	2.23	0.76	1.33	9.04	4.33	1.17	2.71	31.54
0.02	2.86	0.92	1.59	14.10	5.43	1.65	4.40	59.67
0.03	3.26	1.04	1.79	18.09	*	*	*	*
				VAR(2)				
0.00	1.83	0.68	1.23	6.74	1.86	0.70	2.63	12.34
0.01	2.10	0.75	1.30	8.32	3.98	1.01	2.32	25.25
0.02	2.43	0.84	1.45	10.84	5.04	1.32	3.31	43.37
0.03	2.81	0.95	1.63	14.13	5.32	1.53	4.08	54.33
				VAR(1)				
0.00	1.95	0.79	1.44	8.38	3.57	2.42	10.67	149.88
0.01	2.17	0.75	1.35	8.78	4.74	1.84	4.54	56.65
0.02	2.44	0.82	1.44	10.72	5.28	1.97	4.81	66.53
0.03	2.75	0.92	1.59	13.49	5.55	1.96	4.84	69.65

Notes: The symbol "*" indicates that the projection facility is invoked more than 10 percent of the simulation periods.

Table 9

Performance of Simple Rules under Learning with Constant Natural Rates

	OC	LWW Rule				OW Rule			
	Loss	Standard Deviation			Loss	Standard Deviation			Loss
κ	$\mathcal{L}$	π	$u-u^*$	Δi	$\mathcal{L}$	π	$u-u^*$	Δi	$\mathcal{L}$
				VAR(3)					
0.00	6.65	1.88	0.69	1.24	6.97	1.84	0.73	1.39	7.43
0.01	8.63	1.93	0.80	1.37	8.17	1.90	0.86	1.56	8.97
0.02	13.39	1.99	0.91	1.58	9.78	1.96	0.97	1.75	10.66
0.03	17.76	2.08	1.04	1.79	11.82	2.05	1.09	1.98	12.87
				VAR(2)					
0.00	6.71	1.88	0.69	1.24	6.97	1.84	0.73	1.39	7.45
0.01	8.14	1.93	0.79	1.35	8.01	1.89	0.82	1.54	8.63
0.02	10.93	1.97	0.89	1.47	9.21	1.94	0.93	1.69	10.08
0.03	14.12	2.04	0.98	1.65	10.75	2.01	1.03	1.88	11.81
				VAR(1)					
0.00	8.27	1.89	0.73	1.36	7.56	1.84	0.76	1.41	7.66
0.01	8.75	1.87	0.76	1.28	7.46	1.83	0.80	1.49	8.11
0.02	11.06	1.91	0.85	1.42	8.53	1.89	0.91	1.68	9.72
0.03	13.21	1.96	0.95	1.58	9.95	1.97	1.04	1.87	11.68

Table 10

Performance of Simple Rules under Learning with TV Natural Rates

(VAR(3) Forecasts; Optimal Kalman Gains)

	OC	LWW Rule				OW Rule			
	Loss	Standard Deviation			Loss	Standard Deviation			Loss
κ	$\mathcal{L}$	π	$u-u^*$	Δi	$\mathcal{L}$	π	$u-u^*$	Δi	$\mathcal{L}$
				s=0					
0.00	6.65	1.88	0.69	1.24	6.97	1.84	0.73	1.39	7.43
0.01	8.63	1.93	0.80	1.37	8.17	1.90	0.86	1.56	8.97
0.02	13.39	1.99	0.91	1.58	9.78	1.96	0.97	1.75	10.66
0.03	17.45	2.09	1.04	1.78	11.83	2.05	1.09	1.98	12.87
				s=1					
0.00	7.56	1.90	0.82	1.30	8.02	1.86	0.84	1.40	8.23
0.01	10.07	1.94	0.92	1.41	9.15	1.93	0.95	1.60	10.16
0.02	16.27	1.99	1.02	1.57	10.62	1.99	1.08	1.78	11.84
0.03	20.61	2.08	1.15	1.79	12.80	2.09	1.18	2.03	14.03
				s=2					
0.00	9.02	1.92	1.01	1.43	9.85	1.95	1.26	1.41	12.15
0.01	12.39	1.95	1.10	1.47	10.85	1.99	1.17	1.65	12.16
0.02	19.38	2.00	1.18	1.62	12.18	2.06	1.28	1.86	14.27
0.03	26.20	2.07	1.27	1.81	13.97	2.13	1.36	2.07	16.21

Table 11

Central Bank Loss with Incorrect Kalman Gain

(VAR(3) Forecasts)

	LWW Rule			OW Rule
	Central Bank Estimate of s			
κ	0	1	2	
	True value: $s = 0$			
0.00	6.97	7.11	7.46	7.43
0.01	8.17	8.08	8.17	8.97
0.02	9.78	9.62	9.61	10.66
0.03	11.83	11.37	11.23	12.87
	True value: $s = 1$			
0.00	8.03	8.02	8.75	8.23
0.01	9.66	9.15	9.00	10.16
0.02	10.98	10.62	10.37	11.84
0.03	13.53	12.80	11.97	14.03
	True value: $s = 2$			
0.00	11.12	9.78	9.85	12.15
0.01	13.19	11.56	10.85	12.16
0.02	14.65	13.07	12.18	14.27
0.03	17.18	15.29	13.97	16.21

Figure 1

Impulse Responses under Rational Expectations

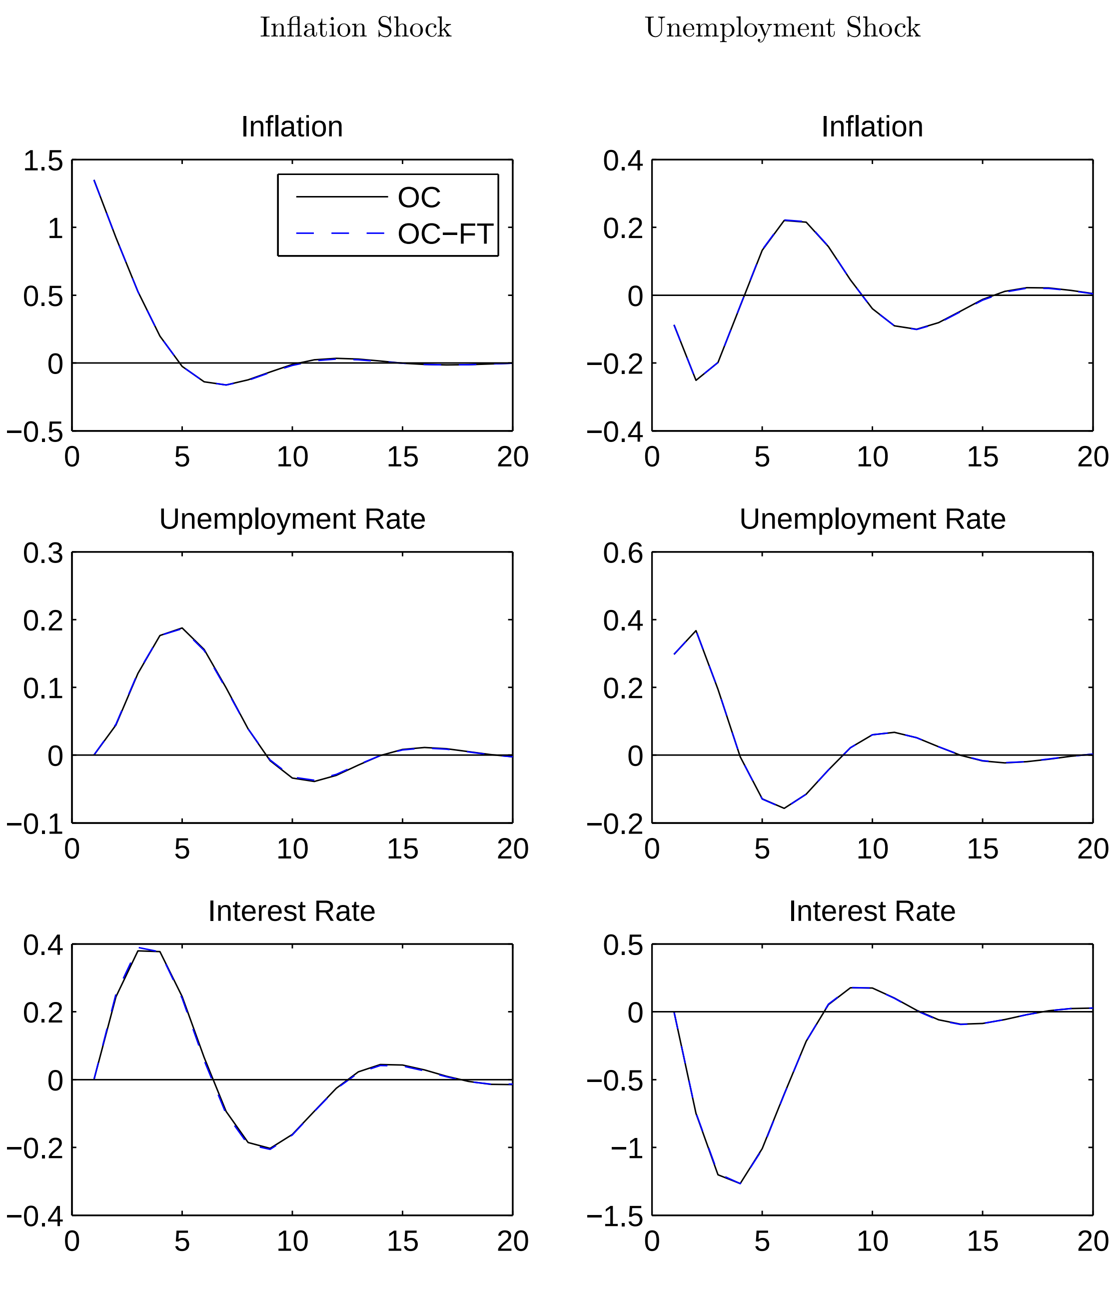

Notes: The first column of charts plots the impulse responses to a one standard deviation innovation to the inflation shock, e_π. The second column plots the impulse responses to a one standard deviation innovation to the unemployment shock, e_v.

Figure 2

OC Policy: Impulse Responses with Learning ($\kappa = 0.02$)

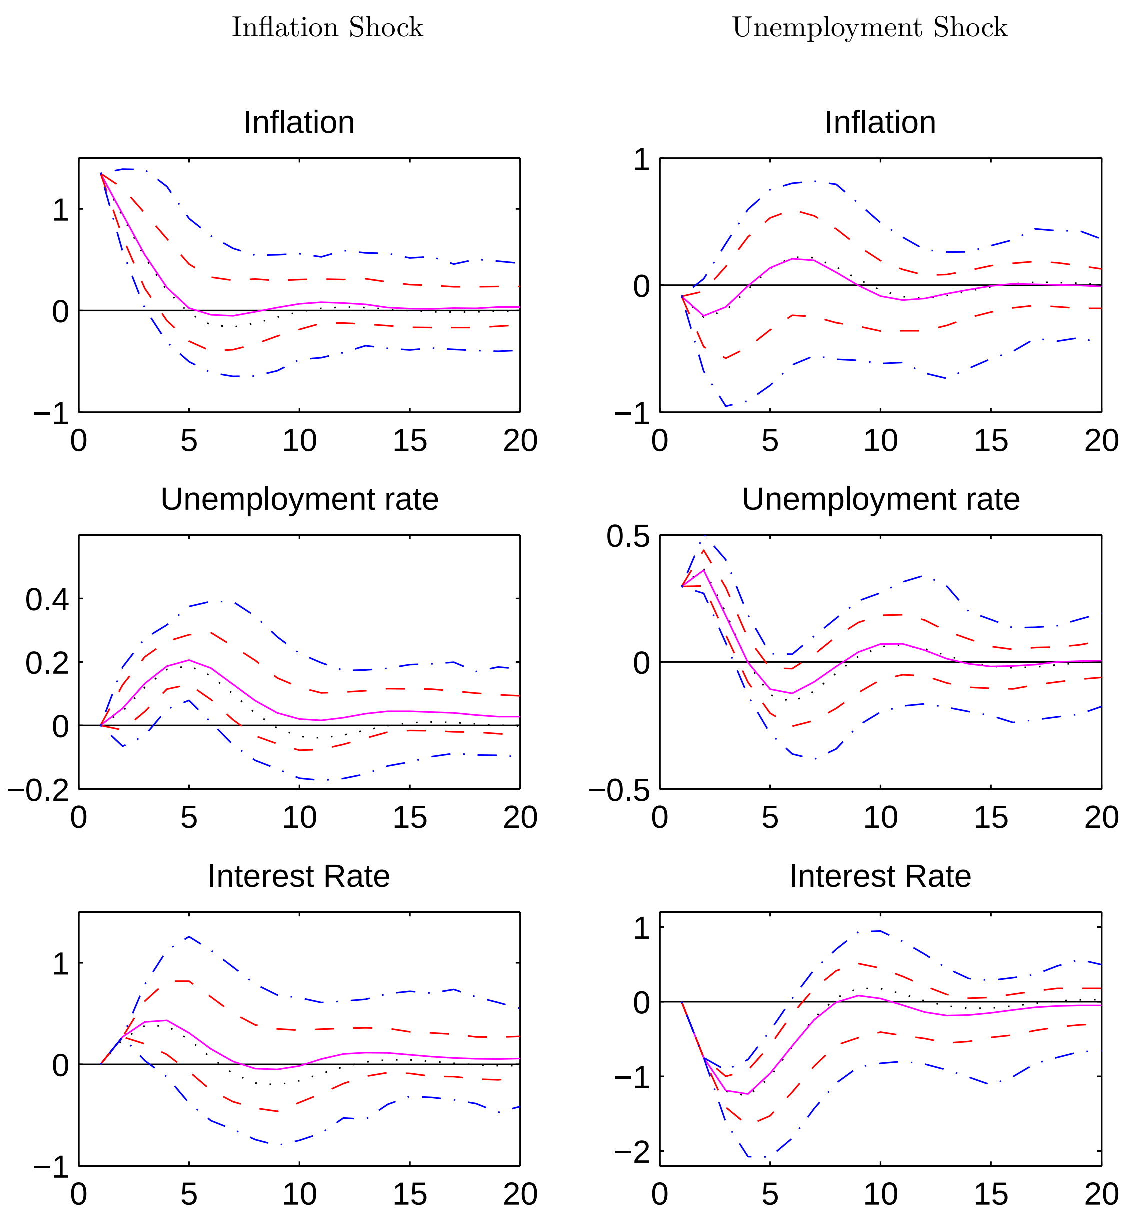

Notes: The left columns shows the responses to an inflation shock; the right column shows those to an unemployment shock. In each panel, the dotted line plots the impulse responses under rational expectations. The solid lines show the median responses under learning. The dashed lines show the 70 percent bands of the responses with learning; the dashed-dotted lines show the 90 percent bands. For these simulations, we assume agents use the three-lag VAR to form expectations with $\kappa = 0.02$ and that natural rates and known and constant.

Figure 3

Impulse Responses under Rational Expectations

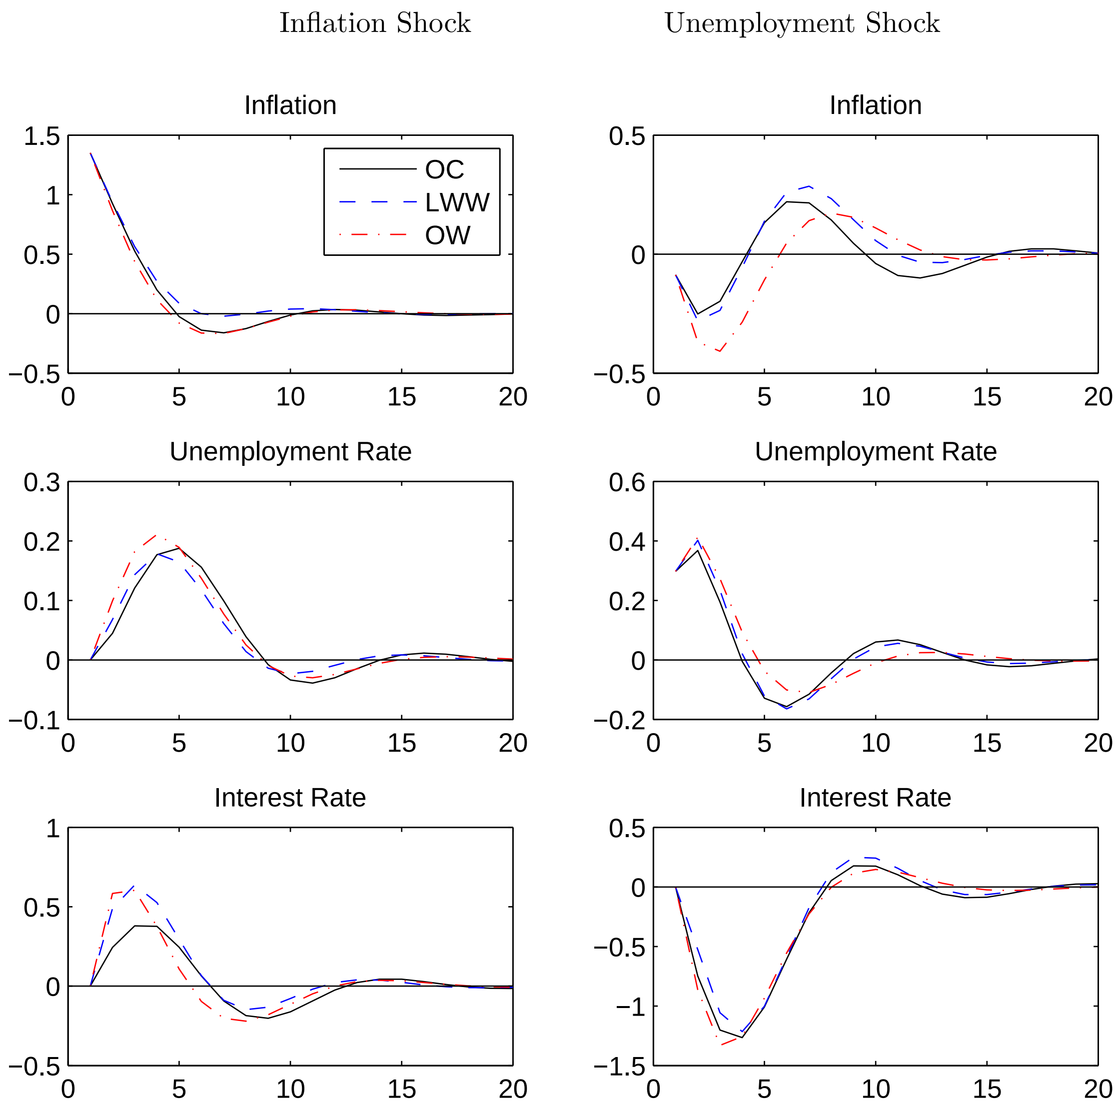

Notes: The first column of charts plots the impulse responses to a one standard deviation innovation to the inflation shock, e_π. The second column plots the impulse responses to a one standard deviation innovation to the unemployment shock, e_v.

Figure 4

LWW Rule: Impulse Responses with Learning ($\kappa = 0.02$)

Inflation Shock　　　　Unemployment Shock

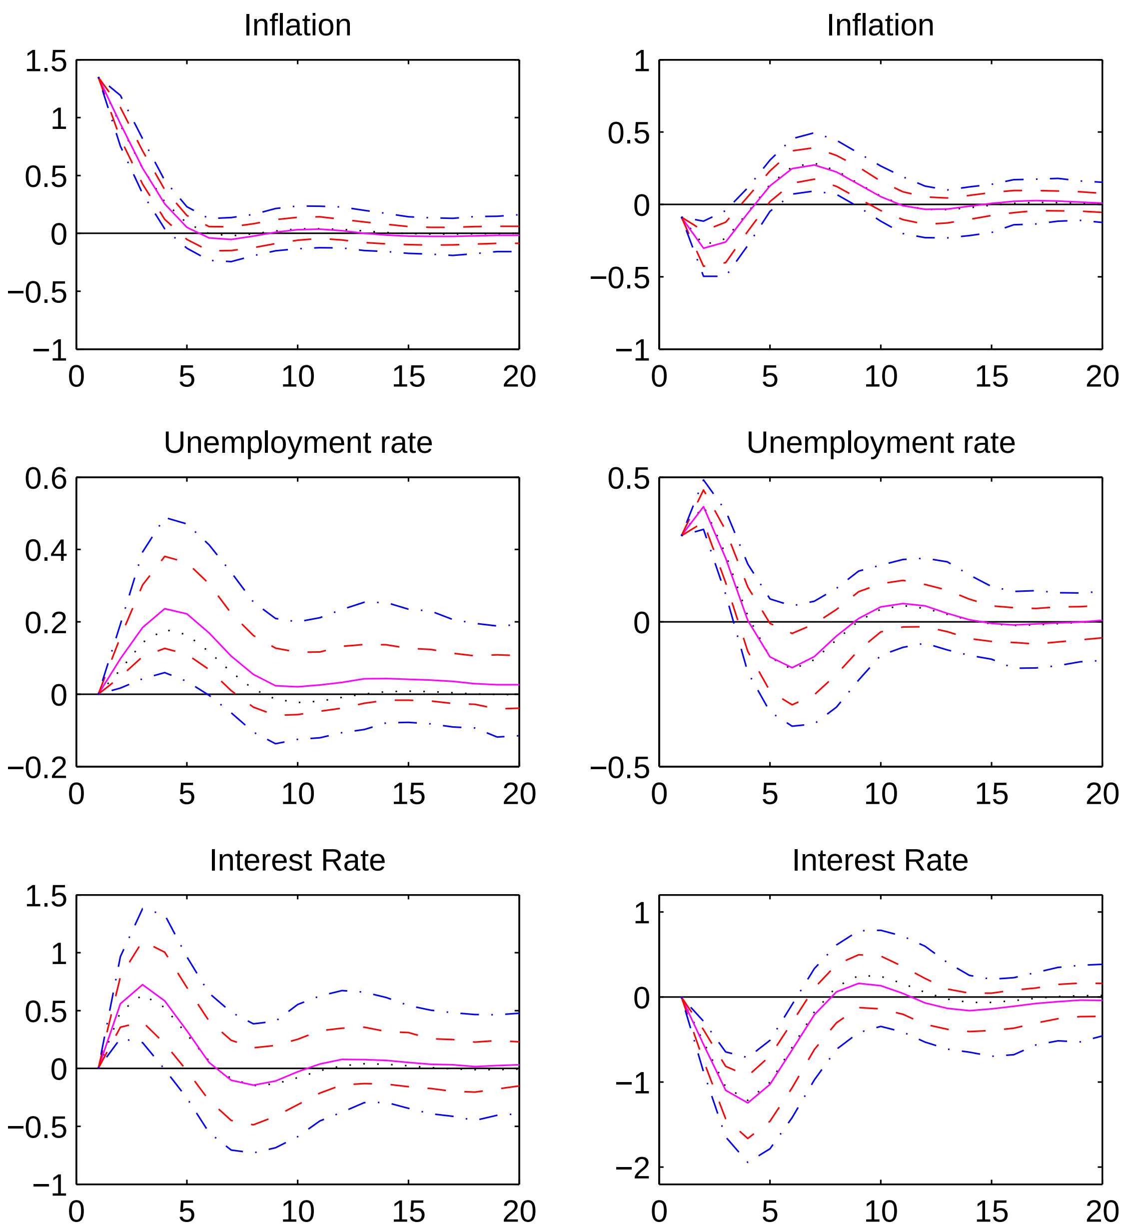

Notes: The left columns shows the responses to an inflation shock; the right column shows those to an unemployment shock. In each panel, the dotted line plots the impulse responses under rational expectations. The solid lines show the median responses under learning. The dashed lines show the 70 percent bands of the responses with learning; the dashed-dotted lines show the 90 percent bands. For these simulations, we assume agents use the three-lag VAR to form expectations with $\kappa = 0.02$ and that natural rates and known and constant.

Figure 5

OW Rule: Impulse Responses with Learning ($\kappa = 0.02$)

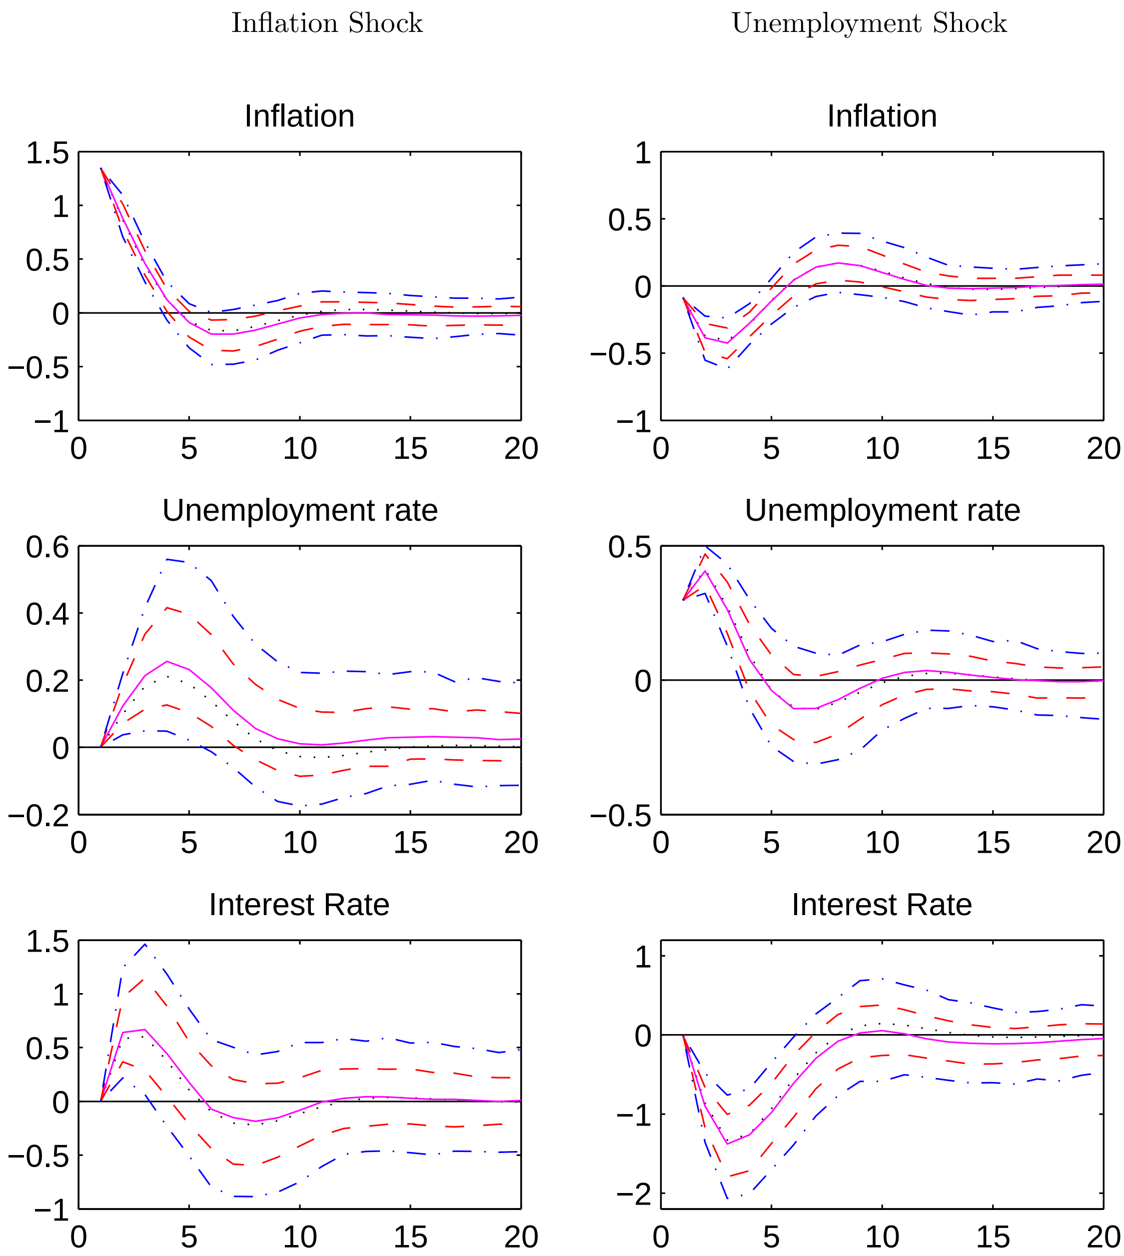

Notes: The left columns shows the responses to an inflation shock; the right column shows those to an unemployment shock. In each panel, the dotted line plots the impulse responses under rational expectations. The solid lines show the median responses under learning. The dashed lines show the 70 percent bands of the responses with learning; the dashed-dotted lines show the 90 percent bands. For these simulations, we assume agents use the three-lag VAR to form expectations with $\kappa = 0.02$ and that natural rates and known and constant.

Figure 6

Robustness to Learning: Constant Natural Rates

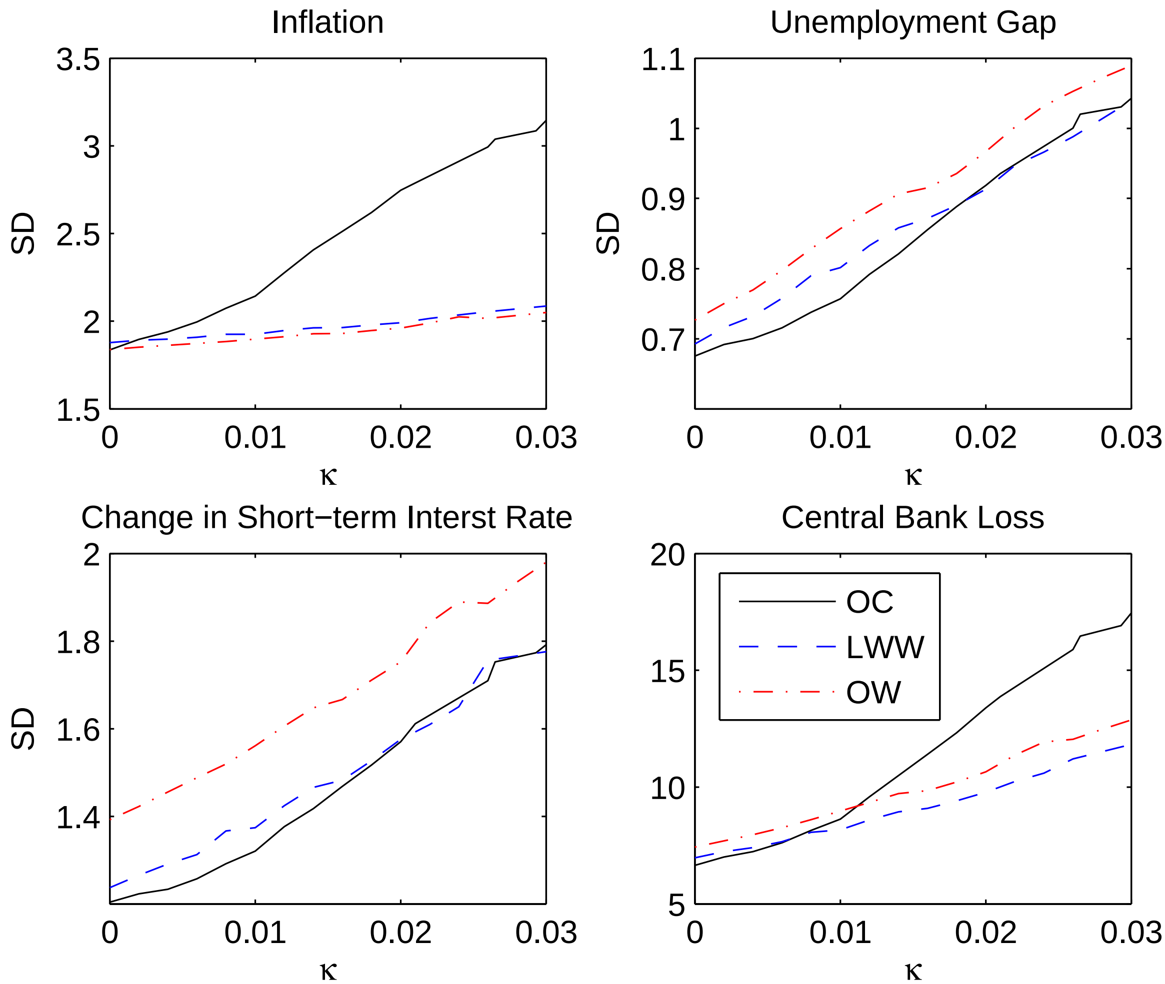

Notes: In each panel, each line plots the asymptotic standard deviation or expected loss that obtain under the specified monetary policy for alternative learning rates, κ, indicated on the horizontal axis. Natural rates are assumed to be constant and known.

Figure 7

Robustness to Learning: Time-varying Natural Rates (s=1)

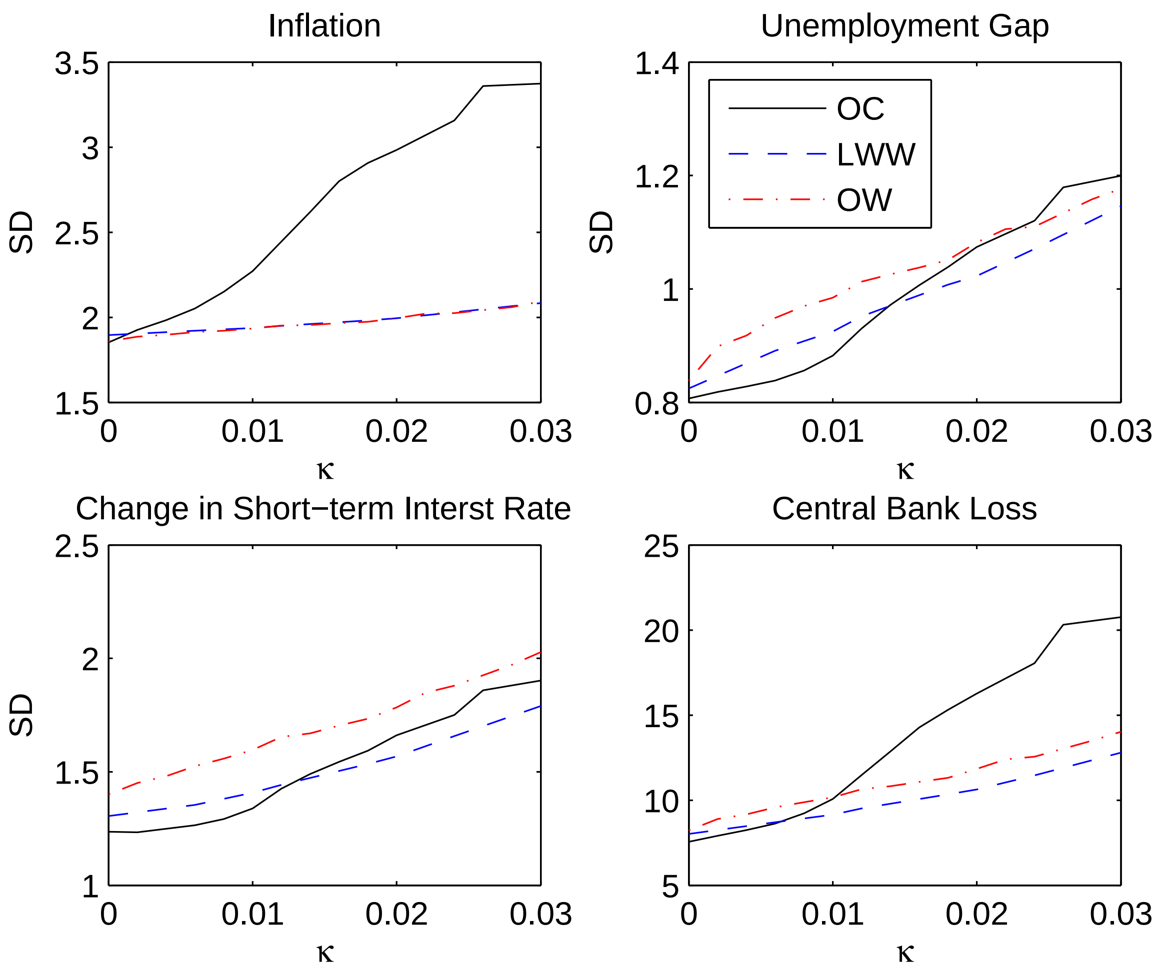

Notes: In each panel, each line plots the asymptotic standard deviation or expected loss that obtain under the specified monetary policy for alternative learning rates, κ, indicated on the horizontal axis.