Learning, Expectations Formation, and the Pitfalls of Optimal Control Monetary Policy

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

October 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. 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 find that the optimal control policy derived under the assumption of rational expectations performs relatively poorly when agents learn. 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 outperform the optimal control policy for empirically plausible parameterizations of the learning models.

KEYWORDS: Rational expectations, learning, robust control, model uncertainty.

JEL Classification System: E52

This is a preliminary draft (not for citation) prepared for the FRB Dallas Conference, "John Taylor's Contributions to Monetary Theory and Policy," October 12–13, 2007.

*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 Giononni 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, say by providing a distribution of forecasts for each policy path, to the analysis of optimal policy. But, in practice, this is infeasible 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.[1] The second approach, and the one that we follow in this paper, evaluates the performance of monetary policies across a set of possible

[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 is infeasible.

non-nested models that embed substantive differences in structure, that is, moderate-sized deviations from a reference model. This approach has been advocated by McCallum (1998) 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 practice, one may view potential misspecification of the reference model to be less egregious than the kind studied in Levin and Williams, and therefore the potential costs of following the optimal control policy to be correspondingly smaller.

In this paper, we reexamine the robustness of optimal control policies designed under the assumption of rational expectations to alternative models of expectations formation. 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 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 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. 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 $\sigma_{e_v}^2$. 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:

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

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_{t+1}^e, 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 $\sigma_{e_\pi}^2$.

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.

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.[2]

[2] 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.

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, 2002) 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.

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.[3] 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.

$$u_t = 0.5\, u^e_{t+1} + 0.5\, u_{t-1} + \underset{(0.022)}{0.556} \; (\tilde{r}^e_t - r^*) + v_t, \tag{4}$$

$$v_t = \underset{(0.085)}{0.513} \; v_{t-1} + e_{v,t}, \quad \hat{\sigma}_{e_v} = 0.30, \tag{5}$$

$$\pi_t = 0.5\, \pi^e_{t+1} + 0.5\, \pi_{t-1} \; - \underset{(0.087)}{0.294} \; (u^e_t - u^*_t) + e_{\pi,t}, \quad \hat{\sigma}_{e_\pi} = 1.35, \tag{6}$$

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.[4] For symmetry with our treatment of the IS curve, we impose the $\phi_\pi = 0.5$ and estimated the remaining parameters

[3] 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.

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

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.

3 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:

$$\mathcal{L} = Var(\pi - \pi^*) + \lambda Var(u - u^*) + \nu Var(\Delta(i)), \tag{7}$$

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 standard in the literature, assuming that the policymaker knows the true parameters of the structural model and assumes all agents use rational expectations (see Sargent, 2007, for a particularly clear description of the methodology). We represent the optimal control policy in the model simulations by a feedback rule where the setting of policy depends on the observed past values of inflation, the unemployment rate, and the interest rate. 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.

We 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:

$$i_t = i_{t-1} + \theta_\pi(\bar{\pi}^e_{t+3} - \pi^*) + \theta_u(u_{t-1} - u^*), \tag{8}$$

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.

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

$$i_t = i_{t-1} + \theta_\pi(\bar{\pi}^e_{t+3} - \pi^*) + \theta_{\Delta u}(u_{t-1} - u_{t-2}). \tag{9}$$

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

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. 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.[5] This approach to modeling learning allows for the possible presence of time variation in the economy. 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. 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 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

[5] See also Sargent (1999), Cogley and Sargent (2001), Evans and Honkapohja (2001), Gaspar and Smets (2002), and Gaspar, Smets and Vestin (2006) for related treatments of learning.

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:

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

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

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.[6] We take 0.02 to be a reasonable benchmark value for κ. Given the uncertainty about this parameter, we report results for values of κ up to and including 0.03. In the extreme case of $\kappa = 0$, agents do not update the coefficients of their forecast model; the coefficient values are fixed at the 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.

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

5 Simulation Method

In the case of rational expectations 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.

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. 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.[7] This constraint is typically encountered in less than one percent of the simulation

[7]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.

periods; however, in the case of of a high updating rate ($\kappa = 0.03$), this constraint can be encountered up to two 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) from their respective steady-state values. This constraint on the model is invoked extremely rarely in the simulations.

6 Results under Rational Expectations

We first evaluate the performance of the optimal control and two simple rules in the model assuming rational expectations. The parameters of the simple rules are chosen to minimize the loss using a hill-climbing routine. The resulting optimized LWW rule is given by:

$$i_t = i_{t-1} + 1.05\,(\bar{\pi}^e_{t+3} - \pi^*) - 1.39\,(u_{t-1} - \hat{u}^*_t). \tag{12}$$

The optimized OW rule is given by:

$$i_t = i_{t-1} + 1.74\,(\bar{\pi}^e_{t+3} - \pi^*) - 1.19\,(u_{t-1} - u_{t-2}). \tag{13}$$

In the following, we refer to these specific parameterizations of these two rules simply as the "LWW" and "OW" rules. Table 1 reports the outcomes under the optimal control policy, the LWW rule, and the OW rule under rational expectations.

The optimal control policy yields 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 the optimal control policy and the LWW rule is illustrated in Figure 1, which plots the impulse responses to the two shocks for the optimal control, 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 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.

7 Results 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.

7.1 Forecast Model Selection

Table 2 shows the root-mean-squared one-step-ahead forecast errors in the model simulations for the three monetary policies, different values of the learning parameters, κ, and the three specifications of the PLM. In each case, all agents are assumed to forecast with the specified model and updating parameter.

For inflation and unemployment forecasts, all three VARs do about equally well. In fact, the under-parameterized VARs with one and two lags do slightly better than the three-lag VAR, when agents are learning $\kappa > 0$. These results are insensitive to the monetary policy in place. The interest rate forecast are better with the three-lag VAR and the optimal policy, reflecting the fact that simple rules use VAR-generated forecasts, which are not variables in the forecasting model itself. So, the simple rules are at a disadvantage here in that they imply additional forecast errors relative to the optimal control policy.[8]

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 somewhat in the VAR when agents are leaning.

7.2 Optimal Control Policy

We now examine the performance of the optimal control policy computed under the assumption of rational expectations in an environment where agents are learning. Table 3 reports the results from these experiments. 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.[9]

Macroeconomic performance under the optimal control policy deteriorates with learning, with the magnitude in fluctuations in all three objective variables increasing in the updating rate, κ. The effects of learning under this 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.

[8]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.

[9]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 behavior of the economy under learning for the case of the three-lag VAR and $\kappa = 0.02$ is seen in the impulse responses, shown in Figure 2. 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. 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. This expectation is no longer assured under learning; indeed, the median inflation response is above the target in nearly every quarter. In addition, 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.

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.

7.3 Simple Rules

We now repeat the same policy evaluation using the LWW and OW rules; Table 4 reports the results. As in the case of the optimal control policy, the central bank losses are generally larger 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 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 3, 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). Under rational expectations, the LWW rule yields slightly higher variability of all three objective variables than the optimal control policy. But, under learning, 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 superior 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 median responses under learning are very close to the responses under rational expectations. More importantly, 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 about the same as under the optimal control policy.

The OW rule outperforms the optimal control policy for learning rates of 0.02 and higher. As in the case of the LWW rule, the OW rule effectively contains the inflation responses to the two shocks, as seen in Figure 4 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 than the OW rule for all learning models that we consider, under the assumption that natural rates are constant.

8 Central Bank Forecasts

Up to this point, we have assumed that the central bank that follows the LWW and OW policy rules uses the same forecasts of inflation in setting policy as private agents use in making decisions. Given that inflation forecasts are important state variables for determining macroeconomic outcomes under learning, this assumption is subject to the criticism that we are giving the central to much information about how private expectations are formed, which contributes to the excellent performance of these rules in model simulations.

To address this issue, we now evaluate the performance of simple rules assuming that the central bank uses a different reduced-form forecasting model than used by private agents. Specifically, we examine the outcomes assuming that the central bank uses either a VAR(1), VAR(2), or a VAR(3) in forming forecasts. For this exercise, we assume that private agents use a VAR(3) in forming expectations. The results are reported in Table 5. For comparison, the outcomes under the optimal control policy, which use forecasts from the structural model assuming rational expectations, are shown in the first column.

The qualitative results are not sensitive to the specific forecasting model used by the central bank for forecasting. For example, the performance of the LWW and OW rules is in nearly all cases as good or *better* than occur when the central bank uses a VAR(2) to forecast, as opposed to the VAR(3) used by private agents. The performance of these rules suffers somewhat if the central bank forms forecasts using a VAR(1), but even then the superior performance of the simple rules over the optimal control policy obtains for learning rates of 0.02 and above.

9 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 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. 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. In a companion paper, we show that the combination of learning by agents and natural rate mismeasurement by the central bank causes a sharp degradation in the performance of the optimal control policy relative to the simple rules analyzed here.

Until feasible methods are developed that allow for the derivation of optimal monetary policy under a realistic range of model uncertainty, 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. Of course, 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.

References

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

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.

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).

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.

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

Gaspar, Vitor and Smets, Frank (2002), "Monetary Policy, Price Stability and Output Gap Stabilisation," *International Finance*, 5(2), Summer, 193–211.

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.

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.

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 (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 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 (2005b), "Inflation scares and forecast-based monetary policy", *Review of Economic Dynamics*, 8(2), 498-527, April.

Orphanides, Athanasios and John C. Williams (2005c), "The Decline of Activist Stabilization Policy: Natural Rate Misperceptions, Learning, and Expectations" *Journal of Economic Dynamics and Control*. 29(11), 1927-1950, November.

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.

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.

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.

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 Jour- nal of Central Banking,1: 177-207.

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.

Wallich,Henry C. (1976), "Application of Optimal Control to Problems of Economic Stabilization: Discussion," *American Economic Review*, 66 (2), 356-59.

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.

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.

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 under Rational Expectations

Policy	Standard Deviation π	Standard Deviation $u - u^*$	Standard Deviation Δi	Loss $\mathcal{L}$
Optimal control	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 2

Forecast Accuracy (RMSE

	Optimal Control			LWW Rule			OW Rule		
κ	π	u	i	π	u	i	π	u	i
				VAR(3)					
0.00	1.35	0.30	0.00	1.35	0.30	0.00	1.35	0.30	0.00
0.01	1.38	0.31	0.01	1.38	0.31	0.24	1.38	0.31	0.37
0.02	1.43	0.33	0.09	1.41	0.32	0.44	1.41	0.32	0.62
0.03	1.48	0.36	0.16	1.44	0.34	0.57	1.44	0.34	0.87
				VAR(2)					
0.00	1.35	0.30	0.23	1.35	0.30	0.00	1.35	0.30	0.00
0.01	1.37	0.31	0.24	1.37	0.31	0.22	1.37	0.31	0.36
0.02	1.39	0.32	0.25	1.38	0.32	0.34	1.39	0.31	0.54
0.03	1.42	0.33	0.26	1.40	0.33	0.49	1.41	0.33	0.74
				VAR(1)					
0.00	1.36	0.33	0.44	1.36	0.33	0.11	1.35	0.33	0.36
0.01	1.36	0.33	0.46	1.36	0.33	0.23	1.36	0.33	0.56
0.02	1.37	0.33	0.46	1.37	0.33	0.34	1.37	0.34	0.80
0.03	1.38	0.33	0.47	1.38	0.34	0.45	1.38	0.34	0.99

Table 3

Performance of Optimal Control Policy under Learning

	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	3.17	1.05	1.81	17.76
		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 4
Performance of Simple Rules under Learning

	LWW Rule				OW Rule			
	Standard Deviation			Loss	Standard Deviation			Loss
κ	π	$u-u^*$	Δi	$\mathcal{L}$	π	$u-u^*$	Δi	$\mathcal{L}$
				VAR(3)				
0.00	1.88	0.69	1.24	6.97	1.84	0.73	1.39	7.43
0.01	1.93	0.80	1.37	8.17	1.90	0.86	1.56	8.97
0.02	1.99	0.91	1.58	9.78	1.96	0.97	1.75	10.66
0.03	2.09	1.04	1.78	11.83	2.05	1.09	1.98	12.87
				VAR(2)				
0.00	1.88	0.69	1.24	6.97	1.84	0.73	1.39	7.45
0.01	1.93	0.79	1.35	8.01	1.89	0.82	1.54	8.63
0.02	1.97	0.89	1.47	9.21	1.94	0.93	1.69	10.08
0.03	2.04	0.98	1.65	10.75	2.01	1.03	1.88	11.81
				VAR(1)				
0.00	1.89	0.73	1.36	7.56	1.84	0.76	1.41	7.66
0.01	1.87	0.76	1.28	7.46	1.83	0.80	1.49	8.11
0.02	1.91	0.85	1.42	8.53	1.89	0.91	1.68	9.72
0.03	1.96	0.95	1.58	9.95	1.97	1.04	1.87	11.68

Table 5

Central Bank Loss with Alternative CB Forecast Models

(Private agents forecast with VAR(3))

	Optimal	LWW Rule			OW Rule		
κ	Control	VAR(3)	VAR(2)	VAR(1)	VAR(3)	VAR(2)	VAR(1)
0.00	6.65	6.97	6.97	7.07	7.43	7.43	7.18
0.01	8.63	8.17	7.92	8.96	8.97	8.63	9.38
0.02	13.39	9.78	9.43	11.58	10.66	10.21	12.47
0.03	17.76	11.83	11.11	14.25	12.87	12.27	15.34

The columns under the LWW and OW Rule headings correspond to different central bank forecast models; i.e., VAR(1), VAR(2), or VAR(3). In all cases, private agents use a VAR(3) to compute forecasts.

Figure 1

Impulse Responses under Rational Expectations

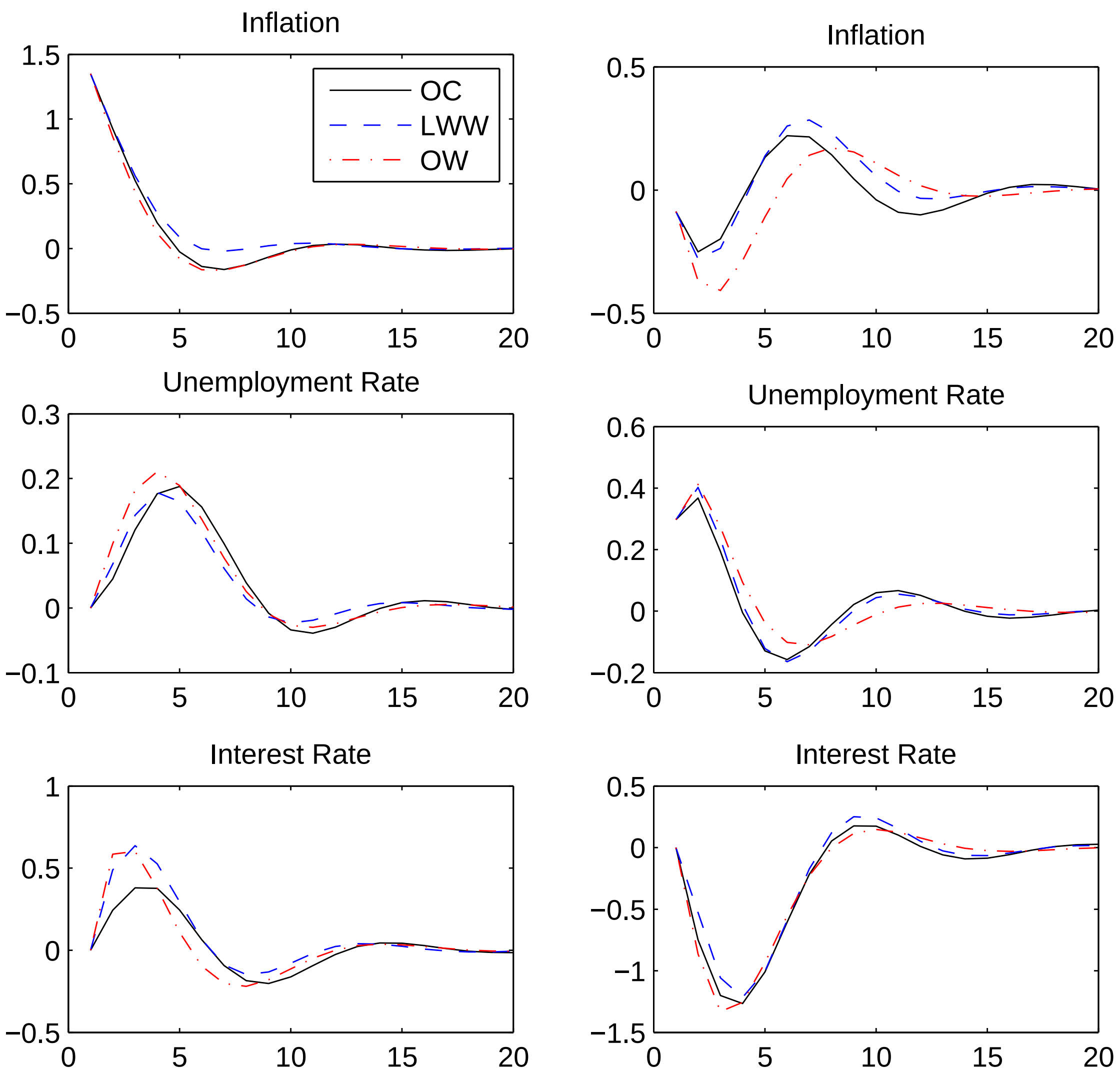

Notes: The first column of charts plots the impulse responses to a one standard deviation innovation to the inflation shock, e_π. The second column plots the impulse responses to a one standard deviation innovation to the unemployment shock, e_v.

Figure 2

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

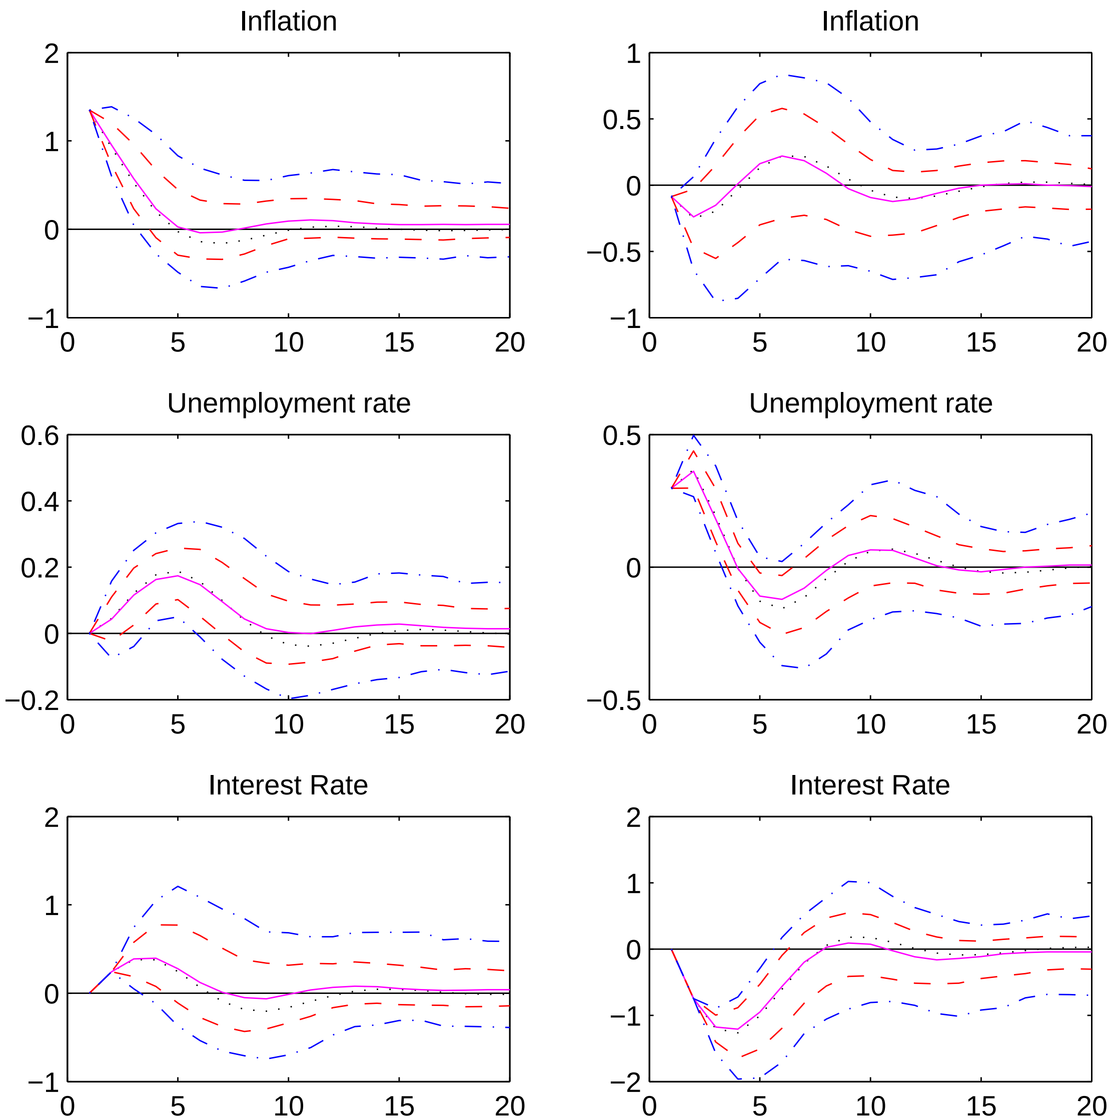

Notes: The left columns shows the responses to an inflation shock; the right column shows those to an unemployment shock. In each panel, the dotted line plots the impulse responses under rational expectations. The solid lines show the median responses under learning. The dashed lines show the 70 percent bands of the responses with learning; the dashed-dotted lines show the 90 percent bands.

Figure 3

Robustness to Learning

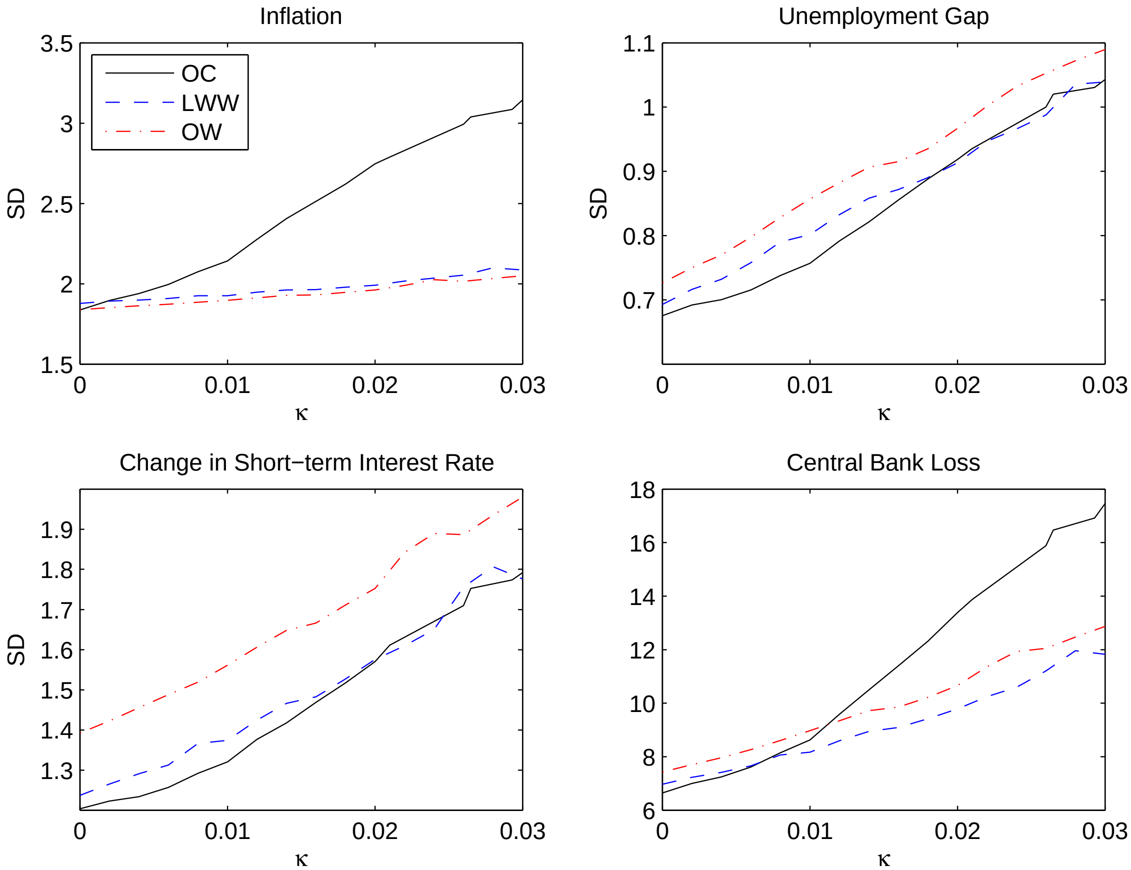

Notes: In each panel, each line plots the asymptotic standard deviation or expected loss that obtain under the specified monetary policy for alternative learning rates, κ, indicated on the horizontal axis.

Figure 4

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

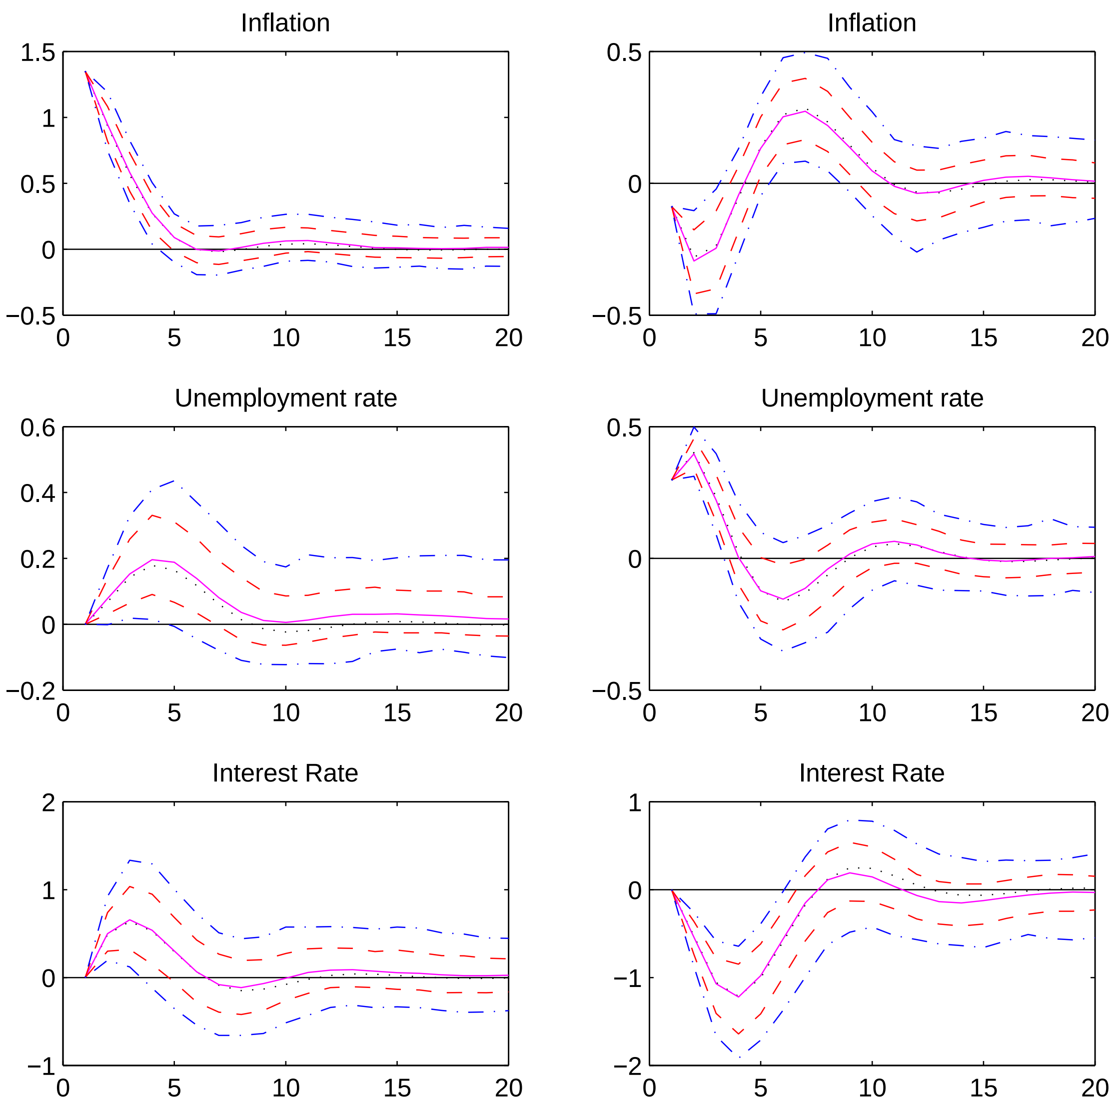

Notes: The left columns shows the responses to an inflation shock; the right column shows those to an unemployment shock. In each panel, the dotted line plots the impulse responses under rational expectations. The solid lines show the median responses under learning. The dashed lines show the 70 percent bands of the responses with learning; the dashed-dotted lines show the 90 percent bands.

Figure 5

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

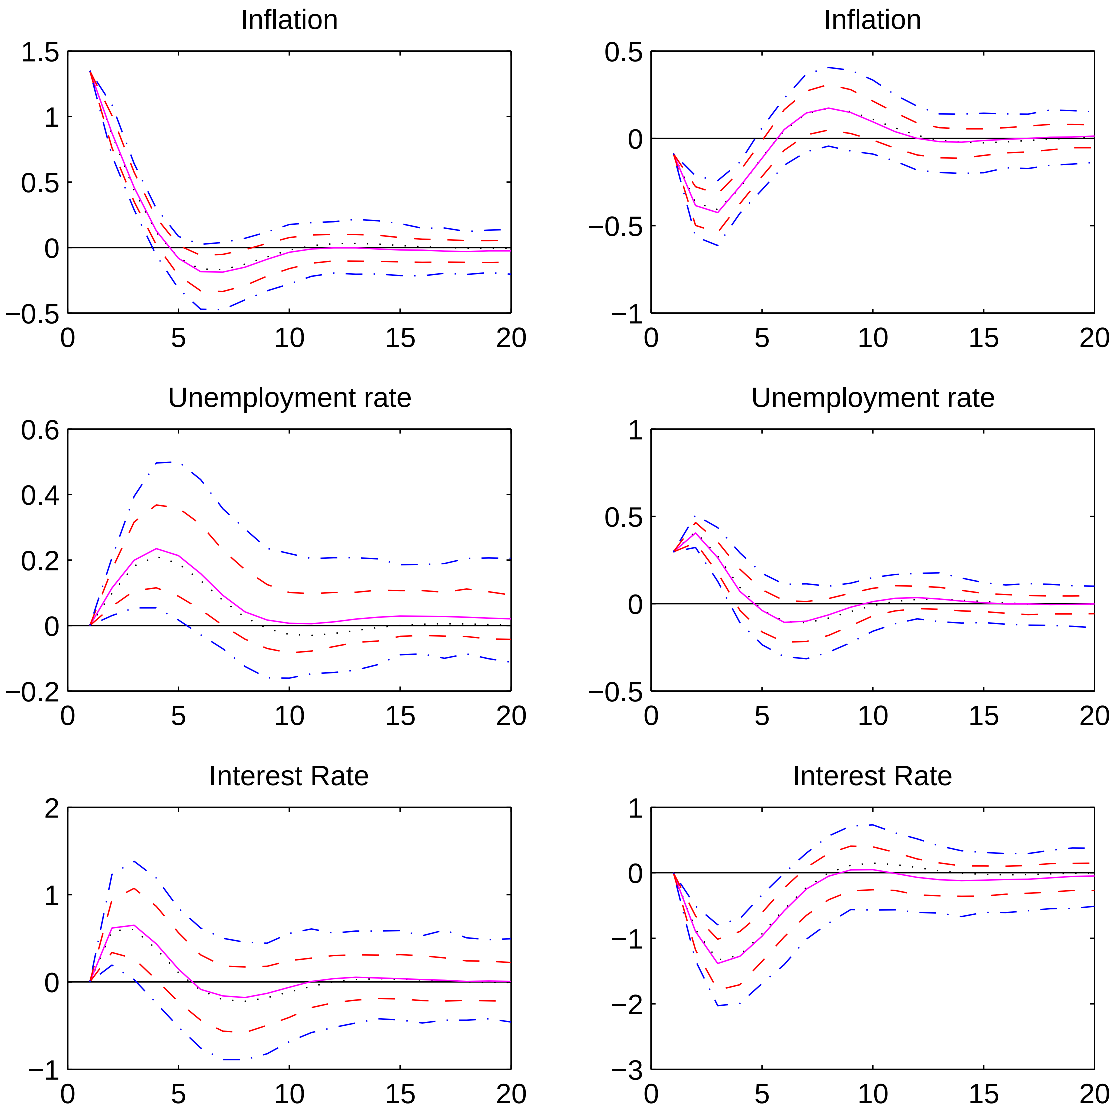

Notes: The left columns shows the responses to an inflation shock; the right column shows those to an unemployment shock. In each panel, the dotted line plots the impulse responses under rational expectations. The solid lines show the median responses under learning. The dashed lines show the 70 percent bands of the responses with learning; the dashed-dotted lines show the 90 percent bands.