

CENTRE FOR APPLIED MACROECONOMIC ANALYSIS

The Australian National University



CAMA Working Paper Series

July, 2008

LEARNING, EXPECTATIONS FORMATION, AND THE PITFALLS OF OPTIMAL CONTROL MONETARY POLICY

Athanasios Orphanides
Central Bank of Cyprus

John C. Williams
Federal Reserve Bank of San Francisco

CAMA Working Paper 17/2008
<http://cama.anu.edu.au>

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

April 2008

Abstract

This paper examines the robustness characteristics of optimal control policies derived under the assumption of rational expectations to alternative models of expectations. 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 can perform poorly when expectations deviate modestly from rational expectations. We then show that the optimal control policy can be made more robust by deemphasizing the stabilization of real economic activity and interest rates relative to inflation in the central bank loss function. That is, robustness to learning provides an incentive to employ a “conservative” central banker. We then examine 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 empirically plausible parameterizations of the learning models and perform about as well or better than optimal control policies.

KEYWORDS: Rational expectations, robust control, model uncertainty.

JEL Classification System: E52

*This paper was prepared for the FRB Dallas Conference, “John Taylor’s Contributions to Monetary Theory and Policy,” October 12–13, 2007. We thank Jim Bullard, Richard Dennis, Bob Hall, Andy Levin, Julio Rotemberg, David Stockman, Volker Wieland, Mike Woodford, Hakan Yilmazkuday, and an anonymous referee 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, the Governing Council of the European Central Bank, 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 been drawn to the idea of applying optimal control theory to the problem of monetary policy (see Chow, 1976, for an early example). Support for the use of optimal control for policy has waxed and waned, reflecting in part swings in economists' confidence in macroeconomic models. Recently, interest among academics and at central banks has re-emerged, as spelled out in contributions by Svensson (2002), Svensson and Woodford (2003), Woodford (2003), Giannoni and Woodford (2005), and others. As described in Svensson and Tetlow (2005), analytical and computational advances now make it possible to operationalize this approach using the Federal Reserve Board's large-scale nonlinear macroeconomic model and other models.

One potential shortcoming of the optimal control approach is that it ignores uncertainty about the specification of the model. Although in principle one could incorporate various types of uncertainty to the analysis of optimal policy, 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. Levin and Williams (2003) found that optimal control policies can perform very poorly if the central bank's reference model is badly misspecified. 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 more closely the robustness properties of optimal control and other monetary policies when the reference model may be misspecified.

In this paper, we study the robustness of optimal control policies derived under the assumption of rational expectations to potential model misspecification with respect to how private agents form expectations.¹ The existing literature in this area, including Orphanides and Williams (2007a,b), has focused exclusively on simple monetary policy rules

¹We follow the methodology advocated by McCallum (1988) and Taylor (1993) and implemented in numerous papers, including Taylor (1999), Levin, Wieland, and Williams (1999, 2003), Orphanides and Williams (2002, 2007a,b), and Brock, Durlauf, and West (2007). An alternative approach, and one that we view as complementary to ours, is robust control (Hansen and Sargent, 2007). These methods are best suited to situations where deviations from the reference model are thought to be relatively modest.

that are generalizations of the Taylor Rule (Taylor 1993) and has not analyzed optimal control policies. We take seriously the informational problems facing economic 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, Williams, 2004, and D'Amico and Orphanides, 2006) suggests that that the process of expectations formation is an important source of model uncertainty. As an alternative to rational expectations, we assume that agents learn about the economy by reestimating a forecasting model as new data become available. We consider a range of learning models, each of which yields good forecasts in our model economy and represents a relatively modest departure from the rational expectations benchmark.

We find that the optimal control policy derived assuming rational expectations can perform poorly in our estimated model of the U.S. economy when agents do not possess perfect knowledge of the economy but instead must learn. To evaluate the optimal policy, we assume the central bank aims to minimize the weighted unconditional variances of the deviation of the inflation rate from its target, the deviation of the unemployment rate from the natural rate, and the first difference of the nominal interest rate. The optimal control policy brings inflation gradually back to target while smoothing swings in the unemployment rate and interest rate. This policy works extremely well when private expectations are perfectly aligned with those implied by rational expectations; however, if agents are learning, expectations can deviate from those implied by rational expectations, and the finely-tuned optimal control policy can go awry. 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 highly persistent and large deviations of the inflation rate from its target.

We show that the lack of robustness of the optimal policy derived under rational expectations can be corrected by assigning the central bank a policy objective that places

significantly less weight on stabilization of the real economy and interest rates relative to inflation than that implied by society’s true objective function. Thus, robustness to learning provides a rationale for hiring a “conservative” central banker in the parlance of Rogoff (1985). We find that the optimal bias to the weights in the loss function are quite large, implying that the central bank should act as if it put up to ten times greater weight on inflation than implied by society’s true loss function. The optimal control policy derived under the biased objective function responds more strongly to inflation, anchoring inflation expectations and performing well under the models of learning that we consider.

We then 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 in the literature. The first is a forward-looking version of a Taylor-type rule, which Levin, et al (2003) found to perform very well in a number of estimated rational expectations models. The second is a rule proposed by Orphanides and Williams (2007a) that differs from the first rule in that policy responds to the change in the measure of economic activity, rather than to the level. This type of rule has been shown to perform very well in a counterfactual analysis of monetary policy during 1996–2003 undertaken by Tetlow (2006). Both rules are parsimonious in that they have only two free parameters. We find that these two simple monetary policy rules, where the parameters are chosen to be optimal under rational expectations, perform very well under learning. Indeed, these rules outperform the optimal control policies derived under rational expectations with the society’s true loss when agents learn, and yield comparable performance to the optimal control policy derived under the biased objective function.

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 optimal control policy. Section 4 describes the models of expectations formation. Section 5 discusses the simulation methods. Section 6 reports and analyzes economic outcomes for optimal control policies under rational expectations and under learning. Section 7 examines the performance of simple monetary policy rules, and section 8 concludes.

2 An Estimated Model of the U.S. Economy

Our analysis uses 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 Giannoni and Woodford (2005), Smets (2003), and others. The main difference is that we substitute the unemployment gap for the output gap to facilitate the estimation of the model using real-time data. The two concepts are closely related in practice by Okun’s law, and the key properties of the model are largely unaffected by this choice.

2.1 The Model

The structural model consists of two equations that describe the behavior of the unemployment rate and the inflation rate and 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:

$$u_t = \phi_u u_{t+1}^e + (1 - \phi_u)u_{t-1} + \alpha_u (i_t^e - \pi_{t+1}^e - r^*) + v_t, \quad (1)$$

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

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, r_t^e —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^* . 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:

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

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 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 are constant. We further assume that they are known by the central bank. (See Orphanides and Williams, 2007a and 2007b, 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 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

²Specifically, we use the mean forecasts of the unemployment rate and the 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.

values. The data sample used in estimation of the model runs from 1968:4 to 2004:2, where the starting date is the first sample point in the SPF.³

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 we assume to be constant:

$$u_t = 0.5 u_{t+1}^e + 0.5 u_{t-1} + 0.056 (r_t^e - r^*) + v_t, \quad (4)$$

(0.022)

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

(0.085)

$$\pi_t = 0.5 \pi_{t+1}^e + 0.5 \pi_{t-1} - 0.294 (u_t^e - u_t^*) + e_{\pi,t}, \quad \hat{\sigma}_{e_\pi} = 1.35, \quad (6)$$

(0.087)

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 $\phi_\pi = 0.5$ and estimate 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.

3 Optimal Control Monetary Policy

We assume that the monetary policy instrument is the nominal short-term interest rate and that the central bank's objective is to minimize a loss equal to the weighted sum of the

³Expectations for the Treasury bill rate were not collected in the first few years of the sample. When these are not available, we use the expectations of the three-month rate implied by the slope of the term structure under the expectations hypothesis.

⁴This 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.

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

unconditional variances of the inflation rate, the unemployment gap, and the change in the nominal federal funds rate. Specifically,

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

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. We also examine alternative parameterizations of the loss function.

In setting policy in every period, 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 optimal control policy is the policy that minimizes the loss subject to the constraints imposed by the equations describing the economy. This policy is constructed, as is standard in the literature, assuming that the policymaker knows the true parameters of the structural model and that all agents use rational expectations. Note that for the optimal control policy, as well as for the simple policy rules described below, we use lagged information in the determination of the interest rate, reflecting the lag in data releases. In principle, the optimal control policy is described by a 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 with an equation that relates the policy instrument to observable variables, rather than in terms of Lagrange multipliers that depend on the model. There are infinitely many such representations, some of which do not yield a determinate rational expectations equilibrium. We consider two alternative representations, each of which yields a unique determinate rational expectations equilibrium

that very closely mimics that under the fully optimal policy.

In the first representation, which we denote “OC,” the optimal control policy is approximated by a feedback rule where the setting of policy depends on the observed past values of inflation, the unemployment rate, and the interest rate. In our model, the fully optimal policy can be nearly perfectly approximated by an equation where the interest rate is determined by three lags of the interest rate, inflation rate, and the unemployment rate. For the benchmark parameterization of the weights in the loss function, the resulting OC policy is given by:

$$i_t = 1.17i_{t-1} + 0.03i_{t-2} - 0.28i_{t-3} + \quad (8)$$

$$0.17\pi_{t-1} + 0.03\pi_{t-2} + 0.01\pi_{t-3} - 2.47u_{t-1} + 2.12u_{t-2} - 0.32u_{t-3}.$$

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) and we denote it “OC-FT.” In principle, this form of the optimal control policy requires including infinitely many leads of the objective variables. However, Giannoni and Woodford (2005) show that this policy can be approximated well by including only a few leads of the target variables. 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. Abstracting from the constant, the OC-FT policy is given by:

$$i_t = 3.85i_{t-1} - 1.39i_{t-2} - 0.53i_{t-3} + \quad (9)$$

$$3.78\pi_{t+1}^e - 1.15\pi_{t+2}^e - 3.51\pi_{t+3}^e + 21.16u_{t+1}^e - 40.32u_{t+2}^e + 11.31u_{t+3}^e.$$

In implementing this policy, we assume that the central bank uses the forecasts consistent with reduced form of the structural under the assumption of rational expectations, as explained below.

4 Expectations Formation

Because we are interested in the 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 with certainty all features of the estimated model. We assume that the model with rational expectations is the central bank's reference model, that is, the model it uses to compute optimal monetary policies. In the remaining models that we study, agents are assumed to engage in real-time perpetual learning. 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 (2005), 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.⁶ 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 that agents forecast inflation, the unemployment rate, and the short-term interest rate using an 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 by macroeconomic forecasters, 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

⁶See also Sargent (1999), Cogley and Sargent (2001), and Evans and Honkapohja (2001) for related treatments of learning.

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 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: $X_t = (1, \pi_{t-1}, u_{t-1}, i_{t-1}, \dots, \pi_{t-l}, u_{t-l}, i_{t-l})$. 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:

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

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

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

For some specifications of the VAR, R_t may not be full rank. For example, if policy follows a policy rule where the nominal interest rate depends on the first lag of the inflation and unemployment rates and agents form expectations using a VAR(3), then R_t 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 0.04.⁷ We take 0.02 to be a reasonable benchmark value for κ , a value implying 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, and instead, the coefficient values are fixed at the initial values that are set as explained in the next section.

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 absence of learning. Under VAR(3) expectations, this is identical with the rational expectations equilibrium.

In each period, innovations are generated from independent Gaussian distributions with variances reported above. The private agents' 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.

⁷See Sheridan (2003), Orphanides and Williams (2005), Branch and Evans (2006), and Milani (2007).

5.2 The Projection Facility

Private agents' learning process injects a nonlinear structure into the model that may cause it to 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 the existence of a "projection facility," that is, 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.⁸ This constraint is typically encountered in less than 1 percent of the simulation periods; however, in the case of a high updating rate ($\kappa = 0.03$), this constraint can be encountered up to 2 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 Performance of the Optimal Control Policy

We first evaluate the performance of the OC and OC-FT policies in the model assuming rational expectations. We then examine the performance when agents learn.

⁸We 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.

6.1 Outcomes under Rational Expectations

The impulse responses to the two shocks are virtually indistinguishable for the OC and OC-FT policies under rational expectations and the implied central bank losses are likewise nearly identical. The dotted lines in Figure 1 show the impulse responses to inflation and unemployment shocks for the OC policy under rational expectations. The optimal policies possess two key features. First, they generate noticeable secondary cycles associated with 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 relatively modest policy response is sufficient to bring inflation under control because 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.

6.2 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 evaluate the performance of the optimal control policies.

6.2.1 Forecast Model Selection

Table 1 reports the root-mean-squared one-step-ahead forecast errors in the model simulations 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.

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 forecasts 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.⁹

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 1 report the results.

Under the OC policy, the VAR(3) is self-confirming in that the alternative forecasting models yield equal or larger root-mean-squared (RMS) forecast errors than the VAR(3) forecasts. That said, the forecasting performance for inflation and unemployment in the VAR(2) model is nearly indistinguishable from that in the VAR(3) model, suggesting that, on average, agents would be close to indifferent between the two models. Interestingly, the VAR(2) model is close to self-confirming as well. That is, if all agents use the VAR(2) for forecasting, forecasts of the unemployment rate and inflation from the VAR(3) and VAR(1) would be about as accurate as, or worse than, on average, those from 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.

In short, all three VAR models seem to be reasonable for forecasting; therefore, on this basis it would be hard to dismiss any of them out of hand. 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.

6.2.2 Policy evaluation with learning

We now examine the performance of the two representations of the OC and OC-FT policies computed under the assumption of rational expectations in an environment where agents

⁹The three-lag VAR encompasses the OC 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 relaxing the restriction on the size of permissible 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.

are learning.

The behavior of the economy with learning under the OC policy for the case of the three-lag VAR and $\kappa = 0.02$ is seen in the impulse responses, shown in Figure 1. In the model with learning, the impulse response to a shock depends on the initial conditions. We therefore show the distribution of impulse response functions (IRF) 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 (2007a). 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 agents learn, the OC policy does not effectively contain movements in inflation. Under rational expectations, the OC 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 very 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 . Columns two through five of Table 2 report the results from these experiments. The upper third of the table reports results where agents use a three-lag VAR in forming forecasts. The first row in this third 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.

The effects of learning under the OC policy are quite large. In the benchmark case of $\kappa = 0.02$, and when agents use a three-lag VAR for forecasting, the central bank loss is twice as large as obtains absent learning. The main problem with the OC policy is that it is designed to stabilize inflation in 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 tight control, as discussed further below.

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.

We now turn to the performance of the OC-FT policy with learning. With a forecast-based representation of optimal policy, we need to specify how the central bank makes its forecast that it uses in setting policy. Consistent with the assumption that the central bank takes its structural model to be true, we posit that the central bank computes its forecast using the structural model assuming that private agents do the same. Results are reported in the final four columns of Table 2, where it can be seen that 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, because we assume that the central bank uses the rational expectations equilibrium of the model to generate forecasts, the outcomes would be identical regardless of how private agents form expectations. This equivalence arises because the leads in the optimal policy equation can be replaced with the forecasts implied by the reduced form of the rational expectations equilibrium. In our approximation, however, the OC and OC-FT policies yield slightly different equilibria under rational expectations, so the outcomes under learning differ as well.

6.3 Alternative Weights in the Loss Function

We now examine the sensitivity of our results to the values of the weights in the loss function, λ and ν . Table 3 reports the outcomes for OC policies derived under a selection of alternative sets of loss function weights. For these computations we assume that private agents use the VAR(3) for forecasting. As in the benchmark calibration of the loss, in all cases reported here, the loss increases with the learning parameter κ .

The deleterious effects of learning on stabilization performance under the OC policy are smaller when the central bank loss puts less weight on unemployment and interest rate variability in the loss. For example, in the case of the benchmark values of $\lambda = 4$ and $\nu = 1$, the standard deviation of the inflation rate is 0.9 percentage point higher for $\kappa = 0.02$ than for $\kappa = 0$ (as seen in Table 2); but, if $\lambda = 1$, the corresponding increase in the standard deviation of the inflation rate is only 0.7 percentage point (as seen in Table 3). The effects of learning on the variability of unemployment and the change in the interest rate are similar for these two values of λ . Likewise, if the weight on interest rate variability, ν , is reduced, the effects of learning on performance are smaller, as seen by comparing the outcomes for $\nu = 0.5$ and $\nu = 1$ when $\lambda = 1$. The combination of large weights on both unemployment and interest rate variability greatly magnifies the effects of learning on macro performance under the OC policy, as illustrated by the case of $\lambda = 16$ and $\nu = 2$ reported in the table.

OC policies optimized with relatively low weights on unemployment and interest rate variability are characterized by larger responses to the inflation rate, which promotes the stabilization of inflation and inflation expectations when agents learn. When both of these weights are relatively low, that is, the central bank places a high relative weight on inflation stabilization, the OC policy acts to return inflation back to target very rapidly after a shock and, as a result, the public's expectation of inflation remains near the target and the effects of learning on the economy are muted. In contrast, when the weights on unemployment and interest rate variability are relatively high, the OC policy attempts to engineer very gradual return of inflation to target following a shock. The resulting persistent deviations of inflation from target can confuse the public, who have imperfect knowledge about the economy and, in particular, the central bank's objective.

6.4 Biasing the Central Bank Loss Function

We now consider the problem of the design of optimal policy that acknowledges uncertainty about how agents form expectations. One approach would be to compute the optimal control policy in the model with learning. For example, Gaspar et al (2006) compute optimal nonlinear policies in a very simple model. This approach has two serious shortcomings

that make it of limited use in practice. First, it requires one to specify the “true” model of learning, which we view as unrealistic. Second, deriving the optimal control policy in a real-world empirical monetary model typically used by central banks remains computationally prohibitive due to the nonlinear nature of models with learning.

We therefore consider the alternative approach of searching for a “biased” central bank loss function such that the implied OC policy derived assuming rational expectations performs well under learning for the true social loss function. This approach applies existing methods of computing optimal policies under rational expectations and is therefore feasible in practice. In particular, for a given value of κ , we searched for the values of $\tilde{\lambda}$ and $\tilde{\nu}$ such that the OC policy derived used the loss $\tilde{\mathcal{L}} = Var(\pi - \pi^*) + \tilde{\lambda}Var(u - u^*) + \tilde{\nu}Var(\Delta(i))$ minimizes the true social loss (which we assume to be given by the benchmark values of $\lambda = 4$ and $\nu = 1$), under learning.¹⁰ We used a grid search to find the optimal weights (up to one decimal place) for the central bank loss. The results are reported in Table 4.

The optimal weights for the central bank loss on unemployment and interest rate variability are significantly smaller than the true weights in the social loss and this downward bias is increasing in the learning rate κ . Thus, in the presence of learning, it is optimal to assign the central bank a loss that places much greater relative weight on inflation stabilization than the true social loss, that is, to employ a “conservative” central banker in the terminology of Rogoff (1985).¹¹ In addition, the relative weight on interest rate variability is reduced by a greater proportion than that on unemployment variability.

For the case of $\kappa = 0.02$, the optimal weights are given by $\tilde{\lambda} = 1.2$ and $\tilde{\nu} = 0.2$. The resulting OC policy is given by:

$$i_t = 1.11i_{t-1} - 0.05i_{t-2} - 0.20i_{t-3} + \tag{12}$$

$$0.21\pi_{t-1} + 0.19\pi_{t-2} + 0.01\pi_{t-3} - 3.14u_{t-1} + 2.34u_{t-2} - 0.34u_{t-3}.$$

Note that this policy is characterized by a nearly twice as large direct response to the

¹⁰Note that this approach can be generalized to allow for including additional variables in the loss function. We leave this to future research.

¹¹Orphanides and Williams (2005), using a very simple theoretical model, similarly found that a central bank loss function biased toward stabilizing inflation (relative to output) was optimal when private agents learn.

inflation rate compared to the OC policy derived for the benchmark loss (and reported in equation 8), reflecting the greater relative weight on inflation stabilization for the biased central bank loss. This policy also responds somewhat more aggressively to unemployment and exhibits less intrinsic policy inertia (as measured by the sum of the coefficients on the lagged interest rate), reflecting the smaller weight on interest rate variability. In the case of $\kappa = 0.03$, the optimal weights on unemployment and interest rate variability are about one tenth as large as those in the true social loss; that is, the relative weight on inflation is about 10 times as large as the true social loss.

7 Simple Rules

To provide a metric against which to measure the performance of the optimal control policies, we examine the performance under learning of two alternative monetary policy rules that have been recommended in the literature for their robustness 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:

$$i_t = i_{t-1} + \theta_\pi(\bar{\pi}_{t+3}^e - \pi^*) + \theta_u(u_{t-1} - u^*), \quad (13)$$

where $\bar{\pi}_{t+3}^e$ is the three-quarter-ahead 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 (2007a) for its robustness properties in the face of natural rate uncertainty; we refer to it as the OW rule:

$$i_t = i_{t-1} + \theta_\pi(\bar{\pi}_{t+3}^e - \pi^*) + \theta_{\Delta u}(u_{t-1} - u_{t-2}). \quad (14)$$

Note that this policy rule is related to the elastic price standard proposed by Hall (1984), whereby the central bank aims to maintain a stipulated relationship between the forecast of the unemployment rate and the price level. This policy rule is also closely related to the

first difference of a modified Taylor-type policy rule in which the forecast of the price level is substituted for the forecast of the inflation rate.¹²

We chose the parameters of these simple rules to minimize the benchmark loss under the assumption of rational expectations using a hill-climbing routine.¹³ The resulting optimized LWW rule is given by:

$$i_t = i_{t-1} + 1.05 (\bar{\pi}_{t+3}^e - \pi^*) - 1.39 (u_{t-1} - u^*). \quad (15)$$

The optimized OW rule is given by:

$$i_t = i_{t-1} + 1.74 (\bar{\pi}_{t+3}^e - \pi^*) - 1.19 (u_{t-1} - u_{t-2}). \quad (16)$$

In the following, we refer to these specific parameterizations of these two rules simply as the LWW and OW rules.

The OC policy derived using the true social loss yield a loss only modestly lower than that associated with 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 OC policy and the LWW rule are illustrated in Figure 2, which plots the impulse responses to the two shocks for the OC policy, LWW rule, and the OW rule, under 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 OC policy manages to bring inflation down slightly more quickly owing to effects of the expectation of overshooting of inflation past the target.

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 a larger rise in the

¹²See Aoki and Nikolov (2004) and Gaspar et al (2007) for evaluations of policies that target or respond to the price level, rather than the inflation rate, under uncertainty.

¹³We focus here on the rules optimized using the true value of the weights in the loss function. The performance under learning of the LWW and OW rules optimized under rational expectations are relatively insensitive to the weights in the central bank loss function, so the gains from biasing the central bank's loss function are correspondingly modest.

unemployment rate. In response to the unemployment shock, this policy, which fails to take into account the level of the unemployment rate relative to its target (the natural rate of unemployment), brings the unemployment rate back to target too slowly, causing inflation to fall further below the target.

In contrast to the OC policy, the LWW and OW rules perform very well when agents learn. The good performance of the LWW rule is seen clearly in the impulse responses to the shocks shown in Figure 3 for the case of the three-lag VAR and $\kappa = 0.02$. 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.¹⁴ 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 OC 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 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.

Table 5 compares the performance of these rules to that of the OC policy under learning. As in the case of the OC policy, the central bank losses are generally larger with the simple rules under learning than they 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 OC 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 5,

¹⁴Ferrero (2007) finds that policy rules that respond aggressively to the public's forecast of inflation accelerates the convergence of least squares learning to the rational expectations equilibrium. Although this result is not directly applicable to our analysis which is based on constant-gain learning, this mechanism may contribute to our finding of a narrow distribution of impulse responses under both the LWW and OW rules. The relationship between convergence under least-squares learning and behavior of under constant-gain learning is a fruitful area of further research.

which shows the outcomes for values of κ between 0 and 0.03 for the OC policy (the thick solid lines), the OC policy with the optimal biased central bank loss given by equation 12 (the thin solid lines), the LWW rule (the dashed lines), and the OW rule (the dashed-dotted lines). Under rational expectations, the LWW rule yields slightly higher variability of all three objective variables than the OC policy. But, because the LWW rule responds more strongly to inflation, it keeps inflation, and thereby inflation expectations, well contained when agents learn. 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 OC policy for learning rates of slightly above 0.01 and higher.

The OC policy with the biased central bank loss optimized for $\kappa = 0.02$ performs worse than the LWW rule and about the same as the OW rule for the full range of values of κ . Although not shown in the figure, the same is true for the the OC policy with the biased central bank loss optimized for $\kappa = 0.03$. Although the biased weights in the loss function does improve the performance of the OC policy, it still performs no better than a simple policy rule. Evidently, the extra fine-tuning that the OC policy is able to achieve under rational expectations is of little value in an environment characterized by imperfect knowledge.

7.1 Central Bank Forecasts

Up to this point, we have assumed that the central bank that follows the LWW or the OW policy rule implements the policy using the same forecasts of inflation that private agents use in making decisions. In fact, there are several measures of inflation expectations from surveys of households and economists and derived from bond yields. At any point in time, these various measures differ and policymakers must judge from the set of measures where inflation expectations lie. As a result, policymakers' estimates of private inflation expectations are contaminated by measurement error.

Given that inflation forecasts are important state variables for determining inflation and unemployment in this model and that the central bank's reading of private inflation expectations is subject to mismeasurement, we now evaluate the performance of simple

rules assuming that the central bank does not use the same forecasting model that private agents use. Specifically, we examine the outcomes assuming that the central bank bases its policy decisions on the average of the inflation forecasts generated by the three VAR models, with one, two, and three lags. We assume that the central bank uses the same value of κ as private agents. This use of average forecasts allows us to capture the presence of heterogeneity in public expectations as well as to introduce what could be interpreted as systematic measurement error emanating from the model's dynamics into the central bank's forecast. The resulting RMS of the difference between the central bank and private one-quarter-ahead forecasts ranges between 0.1 and 0.3 percentage points.

The qualitative results are not sensitive to this change in assumption regarding how the central bank forecasts inflation. The results from these experiments are reported in Table 6. The outcomes under the OC policy are shown for comparison. In fact, in nearly all cases, the performance of the LWW and OW rules in these circumstances is about as good as, or *better* than it is when the central bank uses the same VAR to forecast as private agents.

8 Conclusion

Current techniques for determining optimal control and robust control monetary policies rely on the assumption that the policymaker possesses a very good reference model, an assumption that we find untenable in practice. This paper focuses on one facet of model misspecification and uncertainty associated with the formation of expectations. Our main finding is that optimal control policies are not in general robust to this form of model uncertainty, providing a warning about the potential pitfalls of optimal control policies. We find that the lack of robustness of optimal control policies can be partially mitigated by assigning to the central bank an objective that places more weight on inflation stabilization relative to unemployment and interest rate stabilization.

We show that simple policy rules that are robust to other types of model uncertainty are also robust to uncertainty about how expectations are formed. These rules perform as well or better than optimal control policies derived under rational expectations, even

after biasing the central bank loss to maximize performance under learning. 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, the alternative approach of “stress testing” parsimonious policy rules across a wide set of models provides a practical and productive method of identifying which characteristics of monetary policies are robust and which are fragile.¹⁵ Of course, robustness of any policy cannot be definitively “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 suggests a need for continuing research into the robustness properties of all monetary policy strategies.

¹⁵Svensson and Williams (2007) propose a method to compute optimal policy under model uncertainty using a Markov-switching framework. Computing optimal policies under model uncertainty with this method is extremely computationally intensive and its application to real-world problems remains infeasible.

References

- Aoki, Kosuke and Kalin Nikolov (2004), "Rule-based Monetary Policy under Central Bank Learning," in *NBER International Seminar on Macroeconomics 2004*, Richard H. Clarida, Jeffrey A. Frankel, Francesco Giavazzi, and Kenneth D. West (eds.), Cambridge, MA: The MIT Press.
- Branch, William A. and George W. Evans (2006), "A Simple Recursive Forecasting Model," *Economics Letters*, 91, 158-166.
- Brock, William 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, J.M. and J. Doyle (2002), "Complexity and Robustness," *Proceedings of the National Academy of Sciences*, 99, Suppl. 1, February, 2538-2545.
- Cateau, Gino (2007), "Price-level or Inflation-Targeting under Model Uncertainty," mimeo, Bank of Canada.
- Chow, Gregory C. (1976), "Control Methods for Macroeconomic Policy Analysis." *The American Economic Review*, 66(2), May, 340-345.
- Cogley, Timothy and Thomas J. Sargent (2001), "Evolving Post-World War II U.S. Inflation Dynamics," in *NBER Macroeconomics Annual 2001*, Ben S. Bernanke and Kenneth S. Rogoff (eds.), Cambridge, MA: The MIT Press.
- 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, November, 111-130.
- D'Amico, Stefania and Athanasios Orphanides (2006), "Uncertainty and Disagreement in Economic Forecasting," Board of Governors of the Federal Reserve System, mimeo, August.
- Evans, George and Seppo Honkapohja (2001), *Learning and Expectations in Macroeconomics*, Princeton: Princeton University Press.
- Ferrero, Giuseppe (2007), "Monetary Policy, Learning and the Speed of Convergence," *Journal of Economic Dynamics and Control*, 31(9), September, 3006-3041.
- 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.
- Hall, Robert E. (1984), "Monetary Strategy with an Elastic Price Standard," *Price Stability and Public Policy: A Symposium Sponsored by the Federal Reserve Bank of Kansas City*, Federal Reserve Bank of Kansas City, 137-159.

- Hansen, Lars Peter, and Thomas J. Sargent (2007), *Robustness*, Princeton: Princeton University Press.
- Hoerl, A.E. and R.W. Kennard (1970), "Ridge Regression: Biased Estimation of Nonorthogonal Problems," *Technometrics*, 12, 69–82.
- Levin, Andrew T., Volker Wieland and John C. Williams (1999), "Robustness of Simple Monetary Policy Rules under Model Uncertainty," in *Monetary Policy Rules*, John B. Taylor (ed.), Chicago: University of Chicago Press.
- Levin, Andrew T., Volker Wieland and John C. Williams (2003), "The Performance of Forecast-Based Policy Rules under Model Uncertainty," *American Economic Review*, 93(3), June, 622–645.
- Levin, Andrew T. and John C. Williams (2003), "Robust Monetary Policy with Competing Reference Models," *Journal of Monetary Economics*, 50(5), July, 945–975.
- Mankiw, N. Gregory, Ricardo Reis, and Justin Wolfers (2004), "Disagreement about Inflation Expectations," in *NBER Macroeconomics Annual 2003*, in Mark Gertler and Kenneth Rogoff (ed.), Cambridge, MA: The 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 (2007), "Expectations, Learning, and Macroeconomic Persistence," *Journal of Monetary Economics*, 54(7), October, 2065-2082.
- 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 (2005), "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 (2007a), "Inflation Targeting Under Imperfect Knowledge," in *Monetary Policy Under Inflation Targeting*, Frederic Mishkin and Klaus Schmidt-Hebbel (eds.), Santiago: Central Bank of Chile.
- Orphanides, Athanasios and John C. Williams (2007b), "Robust Monetary Policy with Imperfect Knowledge," *Journal of Monetary Economics*, 54, August, 1406-1435.
- Rogoff, Kenneth (1985), "The Optimal Degree of Commitment to an Intermediate Monetary Target," *Quarterly Journal of Economics*, 100, 1169-90.
- 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. (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 (2006), "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 *NBER Macroeconomics Annual 2003*, Mark Gertler and Kenneth Rogoff (eds.), Cambridge, MA: The MIT Press, 2004, 257-268.
- Woodford, Michael (2003), *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton: Princeton University Press.

Table 1
Forecast Accuracy Horse Race: Root-Mean-Squared Forecast Errors
(OC Policy; $\lambda = 4, \nu = 1$)

κ	Forecast Model								
	VAR(3)			VAR(2)			VAR(1)		
	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

Notes: The upper third of the table reports the simulated root-mean-squared forecast errors when the central bank follows the OC policy and private agents use forecasts from a three-lag VAR. The middle and lower parts of the table report the corresponding results when decisions in the economy are based on forecasts from a two-lag VAR and one-lag VAR, respectively.

Table 2
Performance of OC Policy under Learning ($\lambda = 4, \nu = 1$)

κ	OC Policy				OC-FT Policy			
	Standard Deviation			Loss	Standard Deviation			Loss
	π	$u - u^*$	Δi	\mathcal{L}	π	$u - u^*$	Δi	\mathcal{L}
VAR(3)								
0.00	1.84	0.68	1.20	6.65	1.84	0.68	1.21	6.66
0.01	2.14	0.76	1.32	8.63	2.23	0.76	1.33	9.04
0.02	2.75	0.92	1.57	13.39	2.86	0.92	1.59	14.10
0.03	3.15	1.04	1.79	17.45	3.26	1.04	1.79	18.09
VAR(2)								
0.00	1.83	0.68	1.22	6.71	1.83	0.68	1.23	6.74
0.01	2.06	0.74	1.29	8.14	2.10	0.75	1.30	8.32
0.02	2.42	0.86	1.47	10.93	2.43	0.84	1.45	10.84
0.03	2.76	0.97	1.66	14.12	2.81	0.95	1.63	14.13
VAR(1)								
0.00	1.94	0.78	1.43	8.27	1.95	0.79	1.44	8.38
0.01	2.15	0.75	1.36	8.75	2.17	0.75	1.35	8.78
0.02	2.46	0.84	1.48	11.06	2.44	0.82	1.44	10.72
0.03	2.70	0.92	1.59	13.21	2.75	0.92	1.59	13.49

Notes: The left side of the table reports outcomes under the OC policy; the right side of the table reports results under the OC-FT policy. The upper third of the table reports the outcomes when private agents use forecasts from a three-lag VAR. The middle and lower parts of the table report the corresponding results when private agents use forecasts from a two-lag VAR and one-lag VAR, respectively.

Table 3
Performance of OC Policies with Alternative Loss Functions

κ	Standard Deviation			Loss
	π	$u - u^*$	Δi	\mathcal{L}
$\lambda = 4, \nu = 1$				
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
$\lambda = 1, \nu = 1$				
0.00	1.83	0.79	1.05	5.06
0.01	2.09	0.89	1.17	6.54
0.02	2.52	1.05	1.38	9.39
0.03	2.92	1.22	1.59	12.56
$\lambda = 0, \nu = 1$				
0.00	1.84	0.86	0.99	4.38
0.01	2.11	0.98	1.12	5.71
0.02	2.52	1.16	1.33	8.11
0.03	2.86	1.33	1.51	10.48
$\lambda = 1, \nu = 0.5$				
0.00	1.75	0.74	1.26	4.41
0.01	1.96	0.82	1.37	5.44
0.02	2.29	0.96	1.57	7.42
0.03	2.62	1.09	1.78	9.64
$\lambda = 16, \nu = 1$				
0.00	1.93	0.53	1.60	10.81
0.01	2.50	0.61	1.75	15.26
0.02	3.21	0.71	1.98	22.26
0.03	3.39	0.76	2.16	25.49
$\lambda = 16, \nu = 2$				
0.00	1.97	0.59	1.32	12.87
0.01	2.59	0.70	1.52	19.13
0.02	3.60	0.85	1.82	31.24
0.03	3.86	0.96	2.05	38.09

Notes: Agents are assumed to use the VAR(3) in forecasting.

Table 4
Optimal Weights $\tilde{\lambda}$ and $\tilde{\nu}$ for Central Bank Loss under Learning
(Society's true loss: $\lambda = 4, \nu = 1$)

κ	Optimal		Standard Deviation			True Loss
	$\tilde{\lambda}$	$\tilde{\nu}$	π	$u - u^*$	Δi	\mathcal{L}
0.00	4.0	1.0	1.84	0.68	1.20	6.65
0.01	2.3	0.4	1.95	0.73	1.51	8.20
0.02	1.2	0.2	2.04	0.83	1.81	10.23
0.03	0.5	0.1	2.03	0.95	2.19	12.54

Notes: The central bank is assumed to follow the OC policy derived from minimizing the central bank loss given by $\tilde{\mathcal{L}} = Var(\pi - \pi^*) + \tilde{\lambda}Var(u - u^*) + \tilde{\nu}Var(\Delta(i))$. The true loss is given by $\mathcal{L} = Var(\pi - \pi^*) + 4 \cdot Var(u - u^*) + Var(\Delta(i))$. Agents are assumed to use the VAR(3) in forecasting.

Table 5
Performance of Simple Rules under Learning

κ	OC	LWW Rule				OW Rule			
	Loss	Standard Deviation			Loss	Standard Deviation			Loss
	\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	2.09	1.04	1.78	11.83	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

Notes: The upper third of the table reports the simulated outcomes under the OC, LWW, and OW rules when private agents and the central bank use forecasts from a three-lag VAR. The middle and lower parts of the table report the corresponding results when private agents and the central bank use forecasts from a two-lag VAR and one-lag VAR, respectively.


Table 6
Central Bank Loss when Central Bank uses Model-Averaged Forecasts

	OC	LWW	OW
	Policy	Rule	Rule
Private forecasting model: VAR(3)			
0.00	6.65	7.00	7.33
0.01	8.63	8.14	8.83
0.02	13.39	9.63	10.52
0.03	17.76	11.27	13.03
Private forecasting model: VAR(2)			
0.00	6.71	6.95	7.35
0.01	8.14	8.01	8.53
0.02	10.93	9.15	9.81
0.03	14.12	10.54	11.65
Private forecasting model: VAR(1)			
0.00	8.27	7.55	7.84
0.01	8.75	7.45	8.08
0.02	11.06	8.55	9.73
0.03	13.21	9.75	11.42

Notes: The upper third of the table reports the simulated outcomes under the OC, LWW, and OW rules when private agents use forecasts from a three-lag VAR, but the central bank uses an average of forecasts from the one-, two-, and three-lag VARs. The middle and lower parts of the table report the corresponding results when private agents use forecasts from a two-lag VAR and one-lag VAR, respectively, and the central bank uses an average of forecasts from the one-, two-, and three-lag VARs.

Figure 1


Optimal Control Policy: Impulse Responses under OC Policy



Notes: The left column shows the responses to a one standard deviation innovation to the inflation shock; the right column shows those to a one standard deviation innovation to the unemployment shock. In each panel, the dotted line plots the impulse responses under rational expectations. The solid lines show the median responses under learning $\kappa = 0.02$; the dashed lines show the 70 percent bands of the responses with learning and the dashed-dotted lines show the 90 percent bands.

Figure 2


Impulse Responses under Rational Expectations



Notes: The left column plots the impulse responses to a one standard deviation innovation to the inflation shock, e_π . The right column plots the impulse responses to a one standard deviation innovation to the unemployment shock, e_u .

Figure 3


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



Notes: The left column shows the responses to a one standard deviation innovation to the inflation shock; the right column shows those to a one standard deviation innovation to the unemployment shock. Agents use the VAR(3) for forecasting. 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 4

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



Notes: The left column shows the responses to a one standard deviation innovation to the inflation shock; the right column shows those to a one standard deviation innovation to the unemployment shock. Agents use the VAR(3) for forecasting. 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
Robustness to Learning



Notes: In each panel, each line plots the asymptotic standard deviation or expected loss that obtains under the specified monetary policy for alternative learning rates, κ , indicated on the horizontal axis. The thin solid lines report the outcomes under the OC policy derived assuming $\tilde{\lambda} = 1.2$ and $\tilde{\nu} = 0.2$. Agents are assumed to use the VAR(3) in forming forecasts.