

DISCUSSION PAPER SERIES

No. 4844
CEPR/EABCN No. 17/2004

INFLATION SCARES AND FORECAST- BASED MONETARY POLICY

Athanasios Orphanides and John C Williams

INTERNATIONAL MACROECONOMICS

€ABCN

Euro Area Business Cycle Network

www.eabcn.org



Centre for Economic Policy Research

www.cepr.org

Available online at:

www.cepr.org/pubs/dps/DP4844.asp

INFLATION SCARES AND FORECAST-BASED MONETARY POLICY

Athanasios Orphanides, Board of Governors of the Federal Reserve System
John C Williams, Federal Reserve Bank of San Francisco

Discussion Paper No. 4844
January 2005

Centre for Economic Policy Research
90–98 Goswell Rd, London EC1V 7RR, UK
Tel: (44 20) 7878 2900, Fax: (44 20) 7878 2999
Email: cepr@cepr.org, Website: www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **INTERNATIONAL MACROECONOMICS**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as a private educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions. Institutional (core) finance for the Centre has been provided through major grants from the Economic and Social Research Council, under which an ESRC Resource Centre operates within CEPR; the Esmée Fairbairn Charitable Trust; and the Bank of England. These organizations do not give prior review to the Centre's publications, nor do they necessarily endorse the views expressed therein.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Athanasios Orphanides and John C Williams

ABSTRACT

Inflation Scares and Forecast-Based Monetary Policy*

Central bankers frequently emphasize the critical importance of anchoring private inflation expectations for successful monetary policy and macroeconomic stabilization. In most monetary policy models, however, expectations are already anchored through the assumption of rational expectations and perfect knowledge of the economy. In this Paper, we re-examine the role of inflation expectations by positing, instead, that agents have imperfect knowledge of the precise structure of the economy and policy-makers' preferences, and rely on a perpetual learning technology to form expectations. We find that with learning, disturbances can give rise to endogenous inflation scares, that is, significant and persistent deviations of inflation expectations from those implied by rational expectations, even at long horizons. The presence of learning increases the sensitivity of inflation expectations and the term structure of interest rates to economic shocks, in line with the empirical evidence. We also explore the role of private inflation expectations for the conduct of efficient monetary policy. Under rational expectations, inflation expectations equal a linear combination of macroeconomic variables and as such provide no additional information to the policy-maker. In contrast, under learning, private inflation expectations follow a time-varying process and provide useful information for the conduct of monetary policy.

JEL Classification: E52

Keywords: inflation forecasts, learning, policy rules and rational expectations

Athanasios Orphanides
Board of Governors of the
Federal Reserve
Washington DC 20551
USA
Tel: (1 202) 452 2654
Fax: (1 202) 452 2301
Email: aorphanides@frb.gov

John C Williams
Federal Reserve Bank of San
Francisco
101 Market Street
San Francisco, CA 94105
USA
Tel: (1 415) 974 2240
Email: john.c.williams@sf.frb.org

*The publication of this Paper is funded by the Euro Area Business Cycle Network (www.eabcn.org). This Network provides a forum for the better understanding of the euro area business cycle, linking academic researchers and researchers in central banks and other policy institutions involved in the empirical analysis of the euro area business cycle. We would like to thank George Evans, Martin Ellison, Ben Friedman, Peter Ireland, Allan Meltzer, Lars Svensson, an anonymous referee, and participants at presentations at the University of California, Berkeley, the Norges Bank, meetings of the Econometric Society, the American Economic Association, and the Society for Computational Economics, the Federal Reserve Bank of Atlanta Conference on Learning, March 21/22, 2003, and the Euro Area Business Cycle Network conference at the European Central Bank, December 5/6, 2003, for useful comments and discussions on earlier drafts. The opinions expressed are those of the authors and do not necessarily reflect the views of the Board of Governors of the Federal Reserve System or of management of the Federal Reserve Bank of San Francisco.

Submitted 16 December 2004

1 Introduction

Central bankers frequently emphasize the importance of anchoring inflation expectations for successful monetary policy. For example, as Federal Reserve Chairman Greenspan observed in May 2001:

“We have often pointed before to the essential role that low inflation expectations play in containing price pressures and promoting growth. Any evident tendency in financial markets or in household and business attitudes for such expectations to trend higher would need to factor importantly into our policy decisions.”

When private inflation expectations become unmoored from the central bank’s objectives—episodes that Goodfriend (1993) characterizes as “inflation scares”—macroeconomic stabilization can suffer. Such episodes are easily identified in the monetary history of the United States and other nations. Following the experience with high and volatile inflation in the 1970s, Federal Reserve Chairman Paul Volcker identified the problem of anchoring inflation expectations as crucial for policy success, noting that: “With all its built-in momentum and self-sustaining expectations, [the inflationary process] has come to have a life of its own” (Volcker, 1980). Given these concerns, central banks regularly monitor and analyze information regarding inflation expectations reflected in surveys or financial markets.¹

Relative to the attention that central bankers place on private inflation expectations, there has been comparatively little research that focuses on how these expectations could become unmoored from policymakers’ objectives and the types of monetary policies that might mitigate this problem. Two explanations of how private inflation expectations could become unmoored have received attention in the literature. In one, promoted by Clarida, Gali, and Gertler (2000) in their analysis of Federal Reserve policy in the 1970s, the central bank’s policy fails to satisfy the “Taylor principle” according to which the central bank raises real interest rates when inflation rises above target and vice versa. Under those conditions, inflation expectations are not anchored and may move independently of economic

¹In addition, central bank internal forecasts are at the center of policy deliberations at inflation-targeting central banks and have arguably been equally important for policy decisions in non-inflation-targeting central banks such as the Federal Reserve and the European Central Bank. There is a large literature evaluating the usefulness of internal forecasts for policy (see Levin, Wieland, and Williams, 2003, and references therein); our focus in this paper is on the formation of private inflation forecasts and their usefulness for policymaking.

fundamentals. In the second, the central bank's inflation target is assumed to change from time to time and private agents only gradually recognize these shifts (see Bomfim et al. 1997, Kozicki and Tinsley 2001a,b, and Erceg and Levin 2003). According to these explanations, anchoring inflation expectations should be straightforward in practice, and requires only that the central bank hold to a constant long-run inflation target and satisfy a basic stability condition. Once these conditions have been fulfilled, presumably central banks would no longer need to be so concerned with private inflation expectations.

One potential source of this apparent disconnect between the weight central bankers place on inflation expectations and the conclusion of policy evaluations conducted in the literature may be the rigid imposition of rational expectations in macroeconomic models with an assumed fixed and known structure. In standard linear models with fixed coefficients, once a linear monetary policy rule is specified, inflation expectations can be represented as a linear function of economic outcomes.² Under rational expectations, economic agents are assumed to know these functions and mechanically form expectations accordingly. But what if agents are, in fact, uncertain of the structure of the model, are concerned about possible structural change, or are simply uncertain of the values of model parameters? Once imperfect knowledge of this type is acknowledged, the tight mechanical link from economic outcomes to the process of expectations formation breaks down. As stressed by Friedman (1979) and Sargent (1993), the explicit learning process that economic agents are assumed to employ to form expectations should then be examined.

In this paper, we break the tight link between inflation expectations and observable macro variables implied by rational expectations by positing that agents do not know with certainty the parameters of the model but instead update their estimates based on the information available to them. We show that in a model with private agent learning, inflation expectations drift endogenously in response to macroeconomic disturbances even when policy satisfies standard stability conditions and the long-run inflation objective is a constant target. Importantly, this drift is not confined to short horizons but manifests itself in long-

²For the purposes of this discussion we assume the existence of a well behaved unique rational expectations solution. See Bernanke and Woodford (1997), Evans and Honkapohja (2001b) and Bullard and Mitra (2002) for comparisons of outcome- and forecast-based policies in terms of equilibrium stability and determinacy.

horizon expectations. We interpret these movements in long-horizon inflation expectations, which appear unrelated to fundamentals, as inflation scares. As we show, the prevalence and severity of these endogenous inflation scares is determined by the monetary policy in place, with policies that emphasize output stabilization being more prone to generating inflation scares.

We show that the movements in expectations resulting from the presence of perpetual learning in the economy are consistent with the appearance of excess sensitivity of long-horizon inflation expectations and long-term bond yields to transitory aggregate shocks. In particular, learning induces large positive correlations between long-run inflation expectations and transitory shocks that would not be present under rational expectations with perfect knowledge. We also demonstrate that these correlations depend on the policy regime in place, specifically, they are significantly smaller if the central bank is vigilant in responding to inflationary threats and the public can be certain of its long-run inflation inflation. These results provide an explanation for the responses of long-term nominal interest rates to news reported in the literature. (See Kuttner, 2001, Gürkaynak et al, 2003, and references therein.)

Finally, we show that in the presence of perpetual learning careful monitoring and responding to the public's inflation expectations may lead to significant improvements in economic stabilization performance that are not evident under rational expectations with perfect knowledge. In our model, forecast-based and outcome-based policies are isomorphic under rational expectations. When knowledge is imperfect, however, there is information in private forecasts beyond what is contained in the central bank's forecast. This information is useful for monetary policy: monitoring and responding to private inflation expectations, in addition to actual inflation, leads to improved policy outcomes. In our analysis we also differentiate between the public's expectations and the policymaker's inflation forecasts, and explore the marginal value of this additional information for policy design.

The remainder of the paper is organized as follows. Section 2 describes the model. In Section 3, optimal policy under perfect knowledge is derived and analyzed. Section 4 describes the process by which agents learn. Section 5 explores the incidence of inflation

scars in the model with learning. Section 6 examines the model’s predictions regarding the sensitivity of long-run expectations to economic conditions under perfect knowledge and learning. Section 7 computes optimized policy rules under learning and examines the usefulness of private inflation forecasts in setting monetary policy. Section 8 concludes.

2 The Model Economy

We adopt a simple two-equation macroeconomic model that gives rise to a nontrivial inflation-output variability tradeoff. The properties of this model are described in greater detail in Orphanides and Williams (2004).

The central bank’s objective is to design a policy rule that minimizes the loss, denoted by \mathcal{L} , equal to the weighted average of the asymptotic variances of the output gap, y , and of deviations of inflation, π , from the target rate, π^* ,

$$\mathcal{L} = (1 - \omega)Var(y) + \omega Var(\pi - \pi^*), \quad (1)$$

where $Var(z)$ denotes the unconditional variance of variable z , and $\omega \in (0, 1]$ is the relative weight placed on inflation stabilization.

We assume that the policymaker can set policy during period t so as to determine the intended level of the output gap for period $t + 1$, x_t , subject to a control error, u_{t+1} ,

$$y_{t+1} = x_t + u_{t+1}, \quad u \sim \text{iid}(0, \sigma_u^2). \quad (2)$$

Inflation is determined by a modified Lucas supply function that allows for some intrinsic inflation persistence,

$$\pi_{t+1} = \phi \pi_{t+1|t}^e + (1 - \phi)\pi_t + \alpha y_{t+1} + e_{t+1}, \quad e \sim \text{iid}(0, \sigma_e^2), \quad (3)$$

where π^e is the private agents’ expected inflation rate based on time t information, y is the output gap, $\phi \in (0, 1)$, $\alpha > 0$, and e is a serially uncorrelated innovation. In this setting, an interpretation of $1 - \phi$ is the fraction of agents who raise prices based on the latest observed inflation rate.³ For these agents, price-setting is invariant to the expectations

³This specification is also studied by Clark, Goodhart, and Huang (1999) and Lengwiler and Orphanides (2002). It is similar to specifications of price setting behavior where a portion of inflation expectations is indexed to past inflation, as in Galí and Gertler (1999) and Christiano, Eichenbaum, and Evans (2001).

formation mechanism. The fraction ϕ , then, serves as an index of the sensitivity of inflation movements to the expectations formation mechanism in this economy and becomes a crucial parameter in the model. If ϕ is small, expectations and their evolution are unimportant in this economy.

3 Optimal Policy under Perfect Knowledge

We begin by considering the benchmark case of “perfect knowledge,” where private agents know the structure of the economy and the central bank’s policy. In this case, expectations are rational in that they are consistent with the true data-generating process of the model economy. Later we turn to the case of imperfect knowledge, where agents do not know the structural parameters of the model, but instead must form expectations based on estimated forecasting models.

Under the assumption of perfect knowledge, the optimal policy is given by the Euler equation that relates the intended output gap to the inflation gap and one lead of the intended output gap:

$$x_t = E_{t-1} \left\{ x_{t+1} - \frac{\omega}{1-\omega} \frac{\alpha}{1-\phi} (\pi_{t+1} - \pi^*) \right\}. \quad (4)$$

This expression can be equivalently restated in a number of ways, two of which we consider here. In the first, the optimal policy relates the intended output gap to the inflation gap, the difference between the observed inflation rate and its target. We refer to such rules as “outcome-based” in that they respond to observed outcomes of inflation. In the second, the intended output gap is related to the difference between the expected rate of inflation and the target. We refer to these rules as “forecast-based” rules.

Specifying monetary policy in terms of an outcome-based rule, the intended output gap is given by:

$$x_t = -\theta_\pi (\pi_t - \pi^*), \quad (5)$$

where $\theta_\pi > 0$ measures the responsiveness of the intended output gap to the inflation gap. The optimal value of θ_π , denoted by θ_π^* is given by

$$\theta_\pi^* = \frac{\omega}{2(1-\omega)} \left(-\frac{\alpha}{1-\phi} + \sqrt{\left(\frac{\alpha}{1-\phi} \right)^2 + \frac{4(1-\omega)}{\omega}} \right) \quad \text{for } 0 < \omega < 1. \quad (6)$$

In the limit, when ω equals unity (that is, when the policymaker is not at all concerned with output stability), the policymaker sets the real interest rate so that inflation is expected to return to its target in the next period. The optimal policy in the case $\omega = 1$ is given by: $\theta_\pi^* = \frac{1-\phi}{\alpha}$. It is straightforward to show that the optimal value of θ_π is increasing with ω and the ratio $\frac{1-\phi}{\alpha}$.

Given a monetary policy rule of this form, inflation expectations are given by:

$$\pi_{t+1|t}^e = \frac{\alpha\theta_\pi}{1-\phi}\pi^* + \frac{1-\phi-\alpha\theta_\pi}{1-\phi}\pi_t. \quad (7)$$

Substituting this expression for expected inflation into equation (3) yields the rational expectations solution for inflation for a given monetary policy,

$$\pi_{t+1} = \frac{\alpha\theta_\pi}{1-\phi}\pi^* + \left(1 - \frac{\alpha\theta_\pi}{1-\phi}\right)\pi_t + e_{t+1} + \alpha u_{t+1}. \quad (8)$$

The autocorrelation of inflation is decreasing in ω , with a limiting value approaching unity when ω approaches zero and zero when ω equals one. That is, if the central bank cares only about output stabilization, the inflation rate becomes a random walk, while if the central bank cares only about inflation stabilization, the inflation rate displays no serial correlation.

As noted above, the optimal policy rule can be rewritten in terms of the expected inflation gap:

$$x_t = -\theta_{\pi^e}(\pi_{t+1|t}^e - \pi^*), \quad (9)$$

where $\theta_{\pi^e} > 0$ measures the responsiveness of the intended output gap to the expected inflation gap. The optimal value of θ_{π^e} is proportional to the optimal value of θ_π (the responsiveness to the actual output gap), with the factor of proportionality equal to the inverse of the autocorrelation of the inflation rate. Specifically,

$$\theta_{\pi^e}^* = \frac{1-\phi}{1-\phi-\alpha\theta_\pi^*}\theta_\pi^*, \quad (10)$$

for $\omega \in (0, 1)$. In the limiting case of $\omega \rightarrow 1$, the optimal value of θ_{π^e} becomes infinite and the equivalence between the optimal policies breaks down. We limit our analysis to values of $\omega \in (0, 1)$.

In the following, we consider two values of ϕ , a baseline value, 0.90, and a smaller value, 0.75.⁴ For smaller values of ϕ , the effect of learning on inflation dynamics is muted owing to the smaller role of expectations. To ease comparisons of policy and model properties for the two values of ϕ , we set α so that the optimal policy under perfect knowledge is identical in the two cases. Specifically, for $\phi = 0.75$ we set $\alpha = 0.25$ and for $\phi = 0.90$, we set $\alpha = 0.10$. In all cases, we assume $\sigma_e = \sigma_u = 1$.

Figure 1 shows the optimal values of θ_π and θ_{π^e} for values of ω between zero and one. Note that the optimal value of each parameter depends only on ω and the ratio $\alpha/(1 - \phi)$ thus is invariant to the two model parameterizations considered here. As seen in the figure, θ_{π^e} is much more sensitive to ω than is θ_π . This increased sensitivity to ω reflects the reduction in the autocorrelation of inflation as ω increases.

4 The Economy with Perpetual Learning

We now relax the assumption that private agents have perfect knowledge of all structural parameters and the policymaker's preferences. Instead, we posit that agents must infer the information necessary for forming expectations by observing historical data, in essence acting like econometricians who know the correct specification of the economy but are uncertain about the parameters of the model. In particular, we assume that private agents update the coefficients of their model for forecasting inflation using least squares learning with finite memory. Least squares learning possesses a number of desirable properties: it is straightforward to implement, it appears to correspond closely to the practice of real-world forecasters, and in model simulations it yields accurate forecasts that perform nearly as well as forecasts based on complete knowledge of the economy (see Orphanides and Williams 2004).

Estimation with finite memory reflects agents' concern for changes in the structural parameters of the economy. To focus our attention on the role of imperfections in the

⁴Note that the baseline value of $\phi = 0.90$ places very little weight on lagged inflation in the Phillips curve. In the limiting case of $\phi = 1$ lagged inflation does not directly affect inflation at all. In conjunction with our maintained assumption that shocks are i.i.d., this limiting case is of little interest in our model and we therefore do not consider it. In that case, the optimal policy is trivial under rational expectations and the model lacks dynamics. Gaspar and Smets (2002) have studied a similar model with learning with purely forward-looking inflation determination.

expectations formation process itself, we do not explicitly model the properties of structural change that would justify such concerns. That is, we do not include shocks to the structural parameters of the model in our simulations. Nor do we model the policymaker’s knowledge or learning, but instead focus on the implications of policy based on simple time-invariant rules that do not require explicit treatment of the policymaker’s learning problem.

As in Orphanides and Williams (2004), we model “perpetual learning” by assuming that agents employ a constant gain in their recursive least squares estimation problem. In essence, this assumes that agents place greater weight on more recent observations in estimation.⁵ This algorithm is equivalent to applying weighted least squares where the weights decline geometrically with the distance in time between the observation being weighted and the most recent observation. This approach is closely related to the use of fixed sample lengths or rolling-window regressions to estimate a forecasting model (Friedman 1979). In our model, this learning mechanism implies that a simple AR process with finite memory is used for forecasting. This approach can be conveniently generalized in more complicated models to an economy where agents use VARs for forecasting based on finite memory estimation.

As already noted, the reduced form of inflation under perfect knowledge in our model is given by an AR(1). Correspondingly, we assume that agents attempt to estimate the coefficients of the following equation:

$$\pi_i = c_{0,t} + c_{1,t}\pi_{i-1} + v_i. \tag{11}$$

To fix notation, let X_i and c_i be the 2×1 vectors, $X_i = (1, \pi_{i-1})'$, and $c_i = (c_{0,i}, c_{1,i})'$. Using data through period t , the least squares regression parameters for equation (11) can be written in recursive form:

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

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

where κ_t is the gain. With least squares learning and infinite memory, $\kappa_t = 1/t$, so as t increases, κ_t converges to zero. As a result, as the data accumulate, this mechanism

⁵Inflation expectations with learning based on such constant gain algorithms have been investigated in detail by Sargent (1999), Evans and Honkapohja (2001a), and Evans and Ramey (2001).

converges to the correct expectations function and the economy converges to the perfect knowledge benchmark solution. As noted above, to formalize perpetual learning we replace the decreasing gain in the infinite memory recursion with a small constant gain, $\kappa > 0$.

With imperfect knowledge, expectations are based on the perceived law of motion of the inflation process governed by the perpetual learning algorithm described above. The model under imperfect knowledge consists of the structural equation for inflation (3), the output gap equation (2), the monetary policy rule (5), and the one-step-ahead forecast for inflation, given by

$$\pi_{t+1|t}^e = c_{0,t} + c_{1,t}\pi_t, \quad (14)$$

where $c_{0,t}$ and $c_{1,t}$ are updated according to equations (12) and (13).

In the limit of perfect knowledge (that is, as $\kappa \rightarrow 0$), the expectations function above converges to rational expectations, and the stochastic coefficients for the intercept and slope collapse to:

$$c_0^P = \frac{\alpha\theta_\pi\pi^*}{1-\phi},$$

$$c_1^P = \frac{1-\phi-\alpha\theta_\pi}{1-\phi}.$$

As we deviate from this limiting case, for small positive κ , expectations are imperfectly rational in that agents need to estimate the reduced-form equations they use to form expectations. Nonetheless, as shown in Orphanides and Williams (2004), expectations are nearly rational in that the forecasts are close to being efficient, and the reduced form parameters of the process governing expectations, $c_{0,t}$ and $c_{1,t}$, remain close to what their values would be under perfect knowledge, c_0^P and c_1^P . In particular, as shown in Table 2 of that paper, agents' inflation expectations are nearly as accurate as those that would obtain under perfect knowledge. Moreover, assuming that agents form expectations using constant-gain least squares, agents' expectations are more accurate than alternative forecasts based on autoregressive models of inflation estimated using ordinary least squares, even when the OLS forecasting model includes multiple lags of inflation.⁶

⁶Of course, if an agent could know perfectly both the structure of the model and the forecast method of all other forecasters, she could construct a forecast model with time-varying parameters that would yield forecasts that would perform marginally better than those of other agents. We assume that is infeasible.

5 Learning and Inflation Scares

As noted in the introduction, inflation scares, i.e., increases in long-run inflation expectations—evidenced by shifts in the yield curve—that are unexplained by economic developments are a recurring feature of the U.S. economy (Goodfriend, 1993, Ireland, 1996). Although some instances of inflation scares may be associated with discrete events, others appear to develop endogenously through a confluence of economic developments. In this section, we examine the response of inflation, expected inflation, and output to shocks in our model economy. A related issue that has long puzzled researchers is the high correlation between movements in the entire yield curve and a wide variety of apparently transitory shocks. We take that issue up in the following section.

In calibrating the model for our illustrative simulations, we set $\kappa = 0.05$. (See Orphanides and Williams (2004) for a discussion of the sensitivity of results to κ .) We concentrate on the baseline parameterization $\phi = 0.9$ and $\alpha = 0.1$. To illustrate the effects of learning under different policies, we consider three pairs of alternative policies, corresponding to the optimal policies under perfect knowledge for policymakers with preferences with a relative weight on inflation, ω : 0.25, 0.50, and 0.75. For the forecast-based policy rule, we assume that the policymaker observes and responds to the private forecast. Note that this does not necessarily correspond to the policymaker's own forecast, which may incorporate other information.

5.1 The Response of the Economy to an Inflation Shock

We first consider the dynamic response of the model economy to a one-period 2 percentage point shock to inflation. In our model, the responses of inflation and inflation expectations to an output shock (or policy control error) are observationally equivalent to a shock to inflation (after appropriate scaling) so we do not report on it separately. Note that although the model is linear in the limiting case of perfect knowledge, under least squares learning the model responses depend nonlinearly on the initial values of the states c and R . In the following, we report the average response from 1000 simulations, each of which starts from initial conditions drawn from the relevant steady-state distribution.

Under perfect knowledge, the shock prompts a policy response starting in the following period, leading to a temporary decline in the output gap and a gradual disinflation. The solid lines in Figure 2 report the results under perfect knowledge for this experiment. As expected, the speed at which inflation is brought back to target depends on the monetary policy response, with the more aggressive policy yielding a sharper decline in output and a more rapid return of inflation to target. But in all three cases, output and inflation return to baseline within a few periods.

Imperfect knowledge with learning prolongs the dynamic response of inflation and output to the inflation shock. Consider first the case of the policymaker who responds to actual inflation, shown by the dashed lines in Figure 2. Especially when the central bank places significant weight on output stabilization (bottom panel), the economy stays away from the baseline much longer and the effects of the original shock decay quite slowly.

These differences can be traced to the evolution of the inflation expectations mechanism. As the economy evolves following a shock, agents' estimates of the intercept and the autocorrelation of inflation climb somewhat relative to their perfect knowledge benchmarks. This leads to a slight but persistent rise in inflation expectations, relative to what would be expected under rational expectations, slowing the return of the economy to the baseline. When the central bank places greater weight on inflation stabilization (top panel) the evolution of the economy deviates less from the perfect knowledge benchmark. Because the serial correlation of the inflation process is much smaller in this case, the inflation expectations process is better anchored and less influenced by the learning dynamics.

Relative to the policy based on observed inflation, the inflation forecast-based policy delivers a smaller and less persistent rise in inflation. The dash-dotted lines show the simulated responses of output and inflation when the policymaker follows the rule that responds to the public's inflation forecast with the policy parameter chosen based on perfect knowledge as before. Under this policy rule, the rise in inflation expectations beyond that implied by perfect knowledge elicits a more aggressive response than in the case of the policy that responds to observed inflation. The more substantial decrease in output helps stabilize inflation and inflation expectations.

5.2 Simulation of Serially Correlated Shocks

Next we consider the dynamic responses of the model economy to a set of serially correlated shocks. We examine the effect of such a serially correlated sequence of shocks for two reasons. First, such a sequence of shocks amplifies the effects of learning in the model and thus provides a useful test to explore the interaction of policy and learning. Importantly, since the model is non-linear under learning, the economy's response following a sequence of shocks cannot be inferred simply by scaling and adding up the responses to an individual shock discussed earlier. Second, such unanticipated and infrequent events (given our assumption of i.i.d. innovations) are of the kind that have posed the greatest challenge to policy and modeling historically, as evidenced, for instance, by the events of the 1970s. This experiment is also of interest as an illustration of the importance of initial conditions regarding the formation of inflation expectations for the economy's response to a shock. Recall that the response of inflation does not depend on the "source" of the shocks, that is, on whether we assume the shocks are due to policy errors or to other disturbances. The shock we examine is 2 percentage points in period one, and it declines in magnitude from periods two through eight; in periods nine and beyond there is no shock.

With perfect knowledge, the series of inflationary shocks causes a gradual rise in the inflation rate until the shocks dissipate and subsequently a decline, as shown by the solid lines in Figure 3. The rise in inflation prompts a policy response leading to a temporary decline in the output gap and subsequently a gradual rise towards the baseline. Since the model is linear in this limiting case, these responses are simply the sum of scaled responses to a single shock, as shown in Figure 2. Thus, as before, the speed at which inflation is brought back to target depends somewhat on the monetary policy response. However, in each case, output and inflation return to baseline well before the twentieth period.

Perpetual learning amplifies and prolongs the response of inflation and output to the sequence of shocks. For example, consider the case of the policymaker who responds to actual inflation, shown by the dashed lines in Figure 3 and compare that to the response to a single shock, shown in Figure 2. In Figure 3, the shocks cause inflation to rise above the target level and stay there, while, for the policy that emphasizes output stabilization,

inflation continues to rise even after the shocks to the system stop. As noted earlier, the persistence imparted by learning is inversely related to the strength of the policy response to observed inflation gaps. This is further amplified following a series of correlated shocks. As seen in the upper middle panel, with $\theta_\pi = 0.8$, the peak inflation response of a bit more than 2 percentage points is not appreciably larger than would occur under perfect knowledge. The return of inflation to target, however, is much more gradual. Inflation peaks about 3 percentage points above target when $\theta_\pi = 0.6$, and remains more than 2 percentage points above targets after 20 periods. The results are even more dramatic when $\theta_\pi = 0.4$. In that case, inflation plateaus at 4-1/2 percentage points above target. At the same time, the output gap is consistently minus one percent. The steady downward pressure of maintaining a small output gap in the first few periods is insufficient to overcome the effects of a stubborn buildup of high and persistent inflation expectations. The gradual disinflation prescription that would be optimal with perfect knowledge destabilizes the inflation expectations process in this case and yields stagflation—the simultaneous occurrence of persistently high inflation and low output.

The deterioration of the response of inflation under learning, relative to our perfect knowledge benchmark, is considerably smaller with a forecast-based policy (the dash-dotted lines in the figure). As noted earlier, under this policy rule, the rise in inflation expectations beyond that implied by perfect knowledge elicits a more aggressive response than in the case of the policy that responds to observed inflation. This is especially important when a sequence of shocks, as used in this illustration, threatens to temporarily destabilize the inflation expectations process. For the first two cases, corresponding to values of θ_{π^e} of 3.8 and 1.6, respectively, the peak response of inflation is only modestly above that that obtains under perfect knowledge, and the inflation gap closes reasonably quickly. Even with $\theta_{\pi^e} = 0.8$, the peak inflation response is only 3-1/2 percentage points and the inflation rate is 1-1/2 percentage points above target after 20 periods, 3 percentage points lower than in the case of the policy rule that responds to observed inflation.

As can be seen from these examples, although outcome- and forecast-based policies are isomorphic in the limit of perfect knowledge, with perpetual learning they differ importantly.

Policies responding to private agent’s forecasts of inflation, in particular appear better suited to control apparent instabilities in inflation, following unfavorable shocks.

6 The Term Structure of Inflation Expectations and Bond Yields

Economists have long been puzzled by the apparent excess sensitivity of yields on long-run government bonds to shocks. Shiller (1979) and Mankiw and Summers (1984) point out that long-term interest rates appear to move in the same direction following changes in short-term interest rates and “overreact” relative to what would be expected if the expectations hypothesis held and expectations were assumed to be rational. Changes in the federal funds rate appear to cause long-term interest rates to generally move considerably and in the same direction (Cook and Hahn, 1989, Roley and Sellon, 1995, Kuttner, 2001). Kozicki and Tinsley (2001a,b), Cogley (2002), and Gürkaynak et al (2003), suggest that this sensitivity could be attributed to movements in long-run inflation expectations that differ from those implied by standard linear rational expectations macro models with fixed and known parameters. Our results point to an important role for learning-induced inflation expectations dynamics in explaining this phenomenon and in this section we examine this mechanism in some additional detail.

6.1 The Response of Inflation Expectations to Shocks

We start by examining the responses of short- and long-run inflation expectations to transitory and persistent shocks. We are interested in examining the evolution of inflation expectations at the one-period ahead horizon, which determines the inflation and output dynamics in our model, as well as at longer horizons, which relate more closely to the historical narrative descriptions of inflation scares and the evolution of bond yields. The one-period inflation dynamics in our model are governed by the autoregressive process (14). Under rational expectations, this is a fixed parameter process that can be used to compute the rational k-step ahead forecast of inflation. The parameters of the process depend on policy and model structure, but given policy, they are fixed. Consider for example the case of a policy responding to inflation, θ_π . Then, given the reduced form parameters of the

inflation process, c_0 and c_1 , the law of iterated expectations can be easily applied to obtain forecasts at all horizons from the model.

With imperfect knowledge the translation of the forecasting model agents use to derive one-step ahead inflation expectations into longer-term expectations is not immediate. As a baseline case, we assume that agents use their reduced form estimates of the process governing the one-period ahead forecast, (11), as if it represents the correct model of the economy and use the law of iterated expectations with their latest estimates of that process, $c_{0,t}$ and $c_{1,t}$, as if these parameters were fixed. This is closer to the practice of employing a fixed parameter VAR estimated with the latest data and finite memory to obtain long-term horizon forecasts. (See e.g. Campbell and Shiller (1991) for an application to long-term bond yields and the term structure of interest rates and Orphanides and Williams (2002) for an application to inflation forecasting.)

Another alternative is to estimate a separate model for each desired long-term forecast horizon (with finite memory). Thus, instead of relying on equation (11), to forecast inflation at all horizons, agents may recursively estimate the reduced form process:

$$\pi_i = c_{0,k,t} + c_{1,k,t}\pi_{i-k} + v_i. \tag{15}$$

for each horizon, k , and use this horizon-specific forecasting model to form their expectations. This procedure is closer to a practice commonly employed for recursive estimation and out-of-sample forecasting in the presence of concerns about parameter instability of the forecasting model. (See e.g. Stock and Watson (1999) and Orphanides and van Norden (2003) for applications to simulated real-time inflation forecasting experiments.) We will refer to this as the “horizon-specific” forecasting model.

Note that in the limiting case of perfect knowledge (that is as $\kappa \rightarrow 0$), both the horizon-specific and baseline forecasting models produce identical forecasts. The slope coefficient in the horizon-specific model, in that case, simply equals the k -step ahead coefficient of the perfect knowledge benchmark economy. As with our one-period forecasting model, either of these two multi-period ahead forecasting technologies collapses to the standard rational expectations case in the perfect knowledge limit.

Figures 4 and 5 show the evolution of inflation expectations when the economy is sub-

jected to the shocks described in our previous experiments (shown in Figures 2 and 3, respectively). In each case, we present the evolution of inflation expectations at the one-period- and five-period-ahead horizons. For the longer horizon, these figures show expectations corresponding to our baseline forecasting model. (Expectations using the horizon-specific forecasting model for the five-period-ahead horizon are qualitatively similar.)

The solid lines in Figure 4 show the evolution of expectations under perfect knowledge following a one-period shock to inflation. As can be seen, for all three policies considered, the five-year ahead inflation expectations (right panels) are little affected by the shock, which mostly affects the evolution of the one-period ahead expectation (left panels). The initial response and speed of adjustment are influenced by the responsiveness of policy, as expected. but the one-period ahead expectation quickly reverts to baseline, after a few periods in each case.

Learning significantly prolongs the impact of the shock on the one-period-ahead inflation expectation and, unlike the perfect knowledge benchmark, also implies a significant response of longer-run expectations as well. This is most evident for the case of policy rules responding to lagged inflation, dashed lines. As can be seen, long-term and short-term expectations under learning co-move more closely than under rational expectations. Further, longer-term expectations under learning significantly “overreact” to the temporary shock relative to what would be expected with perfect knowledge.

Figure 5 reports the parallel experiment examining the evolution of the economy to a sequence of serially correlated shocks. This experiment illustrates how the *long-term* inflation expectations may become unhinged from the policymakers objective for a prolonged period, especially for a policy that places relatively little emphasis on price stability (bottom panels). The problem is evident for forecast-based policies as well, but is less severe under these policies.

6.2 Quantifying the Excess Sensitivity of Expectations to Shocks

We now quantify the sensitivity of inflation expectations at various horizons to economic developments in the model and explore how it is affected by the monetary policy regime. We first examine the unconditional correlation between inflation forecasts and actual infla-

tion in the model. We then analyze the correlation between changes in long-run inflation expectations and inflation surprises. The latter analysis corresponds closely to the empirical analysis of the response of long-term nominal interest rates to news noted earlier. (See also Beechey 2004, for additional analysis of the reaction of the term structure to news in the presence of perpetual learning.)

One simple way to summarize the sensitivity of inflation expectations at various horizons is by examining the regression-based slope coefficient of a regression of the k -step-ahead inflation forecast implied by the private agent's evolving forecasting model on the observed inflation rate and a constant:

$$\pi_{t+k|t}^e = a_{0,k} + a_{1,k}\pi_t + \epsilon_t.$$

This is determined by the policy pursued and the expectations formation process. For an outcome-based policy, under perfect knowledge, the k -step ahead slope coefficient, $a_{1,k}$, is given by $(\frac{1-\phi-\alpha\theta_\pi}{1-\phi})^k$. For policy rules corresponding to a policymaker who puts nontrivial weight on inflation stabilization, then, the slope coefficient becomes very small even for moderate values of k .

Under learning, inflation expectations are more persistent than under rational expectations with perfect knowledge. Table 1 reports the resulting slope coefficients from simulation experiments for the three alternative outcome-based policies examined above. We report the results for the one-, three-, five-, and ten-step-ahead forecasts. We compute results using our baseline forecasting model and the horizon-specific forecasting model. Relative to the case of rational expectations, under learning inflation expectations exhibit greater sensitivity to actual inflation. With the policy that responds relatively timidly to inflation (lower panel), and for the case when expectations are relatively more important determinants of actual inflation ($\phi = 0.9$) the expectations at all three forecast horizons shown exhibit behavior we would associate with a unit-root process in our baseline parameterization ($\kappa = 0.05$). Even with a policy that responds more aggressively to inflation (top panel) inflation forecasts at the three- and five-period-ahead horizons can be substantial whereas it is nearly zero under rational expectations. The sensitivity of inflation expectations to movements in actual inflation varies with the parameterization of the model and to illus-

trate this variation we report results for two alternative values for κ for each value of ϕ examined.

The analysis in Table 1 implicitly assumes that agents do not incorporate any explicit knowledge, say from pronouncements from policymakers, regarding the policymaker's ultimate inflation objective in forming expectations. If the central bank could communicate its numerical inflation target to the public, it would simplify the private agents' forecasting problem. Because the adoption and clear communication of such a target is a key part of the inflation targeting strategy that several central banks have adopted over the past decade or so, it is of interest to examine the sensitivity of inflation expectations to shocks in this case. To do so we perform a parallel set of simulations to those reported in Table 1 under the assumption that the public exactly knows the value of π^* and explicitly incorporates this information in forming inflation expectations.⁷ This also allows us to examine the extent to which the excess sensitivity of the term structure of inflation expectations to shocks should be seen as being determined by uncertainty regarding the dynamics of the economy or uncertainty regarding just the long-run inflation target.

As shown in Table 2, even with the assumption of a known inflation target inflation expectations can be substantially more sensitive to shocks than in the rational expectations benchmark. Evidently, even under the assumption that the expectations in the very long-run are tied-down with a fixed and known inflation target, learning regarding the dynamics of the inflation process can induce substantial deviations in longer-term expectations from the rational expectations benchmark. As with the case of an unknown target, these deviations are larger with policy that responds relatively timidly to inflation and for the case when expectations are relatively more important determinants of actual inflation.

Comparison of Table 2 with Table 1 confirms that inflation expectations under learning are generally much less sensitive to actual inflation when the inflation target is assumed to be known by the public. Indeed, the comparison indicates that the benefit of better anchored inflation expectations that is associated with successful communication of the central bank's

⁷To be sure, even in an explicit inflation targeting regime, the public may remain uncertain regarding the policymaker's inflation target, π^* , so that this assumption of a perfectly known inflation target may be seen as an illustrative limiting case. See Orphanides and Williams (2004) for further analysis and discussion of the effects of greater transparency of monetary policy in this model.

inflation target can be significant. As stressed by King (2002), this is consistent with the experience of the U.K. following the adoption of inflation targeting and the independence of the Bank of England. He notes that “inflation has been less persistent—in the sense that shocks die away more quickly—under inflation targeting than for most of the past century.”

Next we examine the model predictions regarding the sensitivity of long-run inflation expectations to inflation “surprises” which facilitates a more direct comparison with results regarding the response of long-term nominal interest rates to news. Table 3 reports the estimated slope coefficient of a regression of the change in the inflation rate forecast k periods ahead on the lagged inflation surprise, defined to be the actual inflation rate less last period’s forecast for inflation.

$$\pi_{t+k|t}^e - \pi_{t+k|t-1}^e = b_{0,k} + b_{1,k}(\pi_t - \pi_{t|t-1}^e) + \eta_t.$$

The basic methodology is the same as before and results are presented for the three-, five-, and ten-period horizons. The first column reports the assumed coefficient in the policy rule and relevant horizon. The second column reports the resulting slope coefficient, $b_{1,k}$, under rational expectations. This is close to zero for longer horizons under all policies, and, for any given horizon, is decreasing in θ_π . The next four columns report the estimated coefficients for the baseline model with learning for different values of κ and ϕ with an unknown inflation target. The final four columns report the corresponding results when the inflation target is assumed to be known by the public.

In the model under learning, when policy responds relatively weakly to inflation and the inflation target is not known by the public, the resulting correlation between changes in long-run inflation expectations and inflation surprises is considerably larger than when policy is vigilant against inflation and the long-run target is known by the public. This obtains for short as well as long horizons where these correlations would be virtually zero under rational expectations with perfect knowledge. For example, consider the case of $\kappa = 0.025$ and $\phi = 0.90$. Then, under the policy with $\theta_\pi = 0.43$ and assuming the long-run inflation target is not known by the public, the sensitivity coefficient at the ten-step-ahead horizon is 0.18. By comparison, under the policy with $\theta_\pi = 0.79$ and a known inflation target, the coefficient is 0.01.

These results provide an explanation of the findings of apparent overreaction of the response of long-term bond yields to current shocks, such as reported by Kuttner (2001) and Gürkaynak et al (2003), without resorting to the assumption that the central bank’s long-run inflation target is changing over time. In addition, they explain why this sensitivity would be significantly reduced if an institutional change solidifies a central bank’s inflation objective. For instance, Gürkaynak et al (2003) report that the sensitivity of U.K. forward rates to news regarding inflation was significantly reduced when the Bank of England was granted operational independence in 1997. This is consistent with the prediction of our model which would attribute this change the Bank’s enhanced ability to focus its policies towards its statutory inflation target following operational independence.

7 Imperfect Knowledge and the Design of Monetary Policy

The examples reported above illustrate how the behavior of the economy can differ significantly under outcome- and forecast-based policy rules that would be identical under perfect knowledge. We now consider the relative performance of optimized outcome- and forecast-based rules in terms of the unconditional variances of output and inflation assuming serially uncorrelated shocks.

7.1 Efficient Outcome- and Forecast-based Simple Rules

We start by examining the characteristics and performance of efficient simple one-parameter outcome- and forecast-based policy rules. The solid line in the upper panel of Figure 6 shows the best obtainable pairs of the standard deviations of inflation and the output gap under the assumption of perfect knowledge. Figure 7 shows the corresponding policy response parameters. The solid line in the upper panel of Figure 7 reports the corresponding optimal values of θ_π for an outcome-based rule; the solid line in the lower panel report the optimal values of θ_{π^e} for a forecast-based rule.

Within the class of one-parameter rules, policy should respond to expected inflation when inflation stabilization is weighted heavily in the objective, but should respond to observed inflation when output stabilization is relatively more important. The dashed line in Figure 6 shows the frontier for the one-parameter outcome-based rule; the dash-dotted

line shows the frontier for the one-parameter forecast-based rule. (As before, the central bank is assumed to respond to the private forecast of inflation in the case of the forecast-based rule.) As seen in the figure, neither class of rules dominates the other, and both do significantly worse than would result under perfect knowledge.⁸

The forecast-based one-parameter rule is more effective at stabilizing inflation than the outcome-based rule. The reason for this result is seen in the structural equation for inflation given by equation (3). In our calibration, inflation depends importantly on expected inflation; therefore, responding to expected inflation is an effective strategy to control inflation. More intriguing is the finding that responding to expected inflation is dominated when the policymaker is sufficiently concerned about output stabilization. Responding too strongly to expected inflation generates excessive variability of the output gap and the preferred policy responds instead to the actual inflation rate. Evidently, for the policymaker concerned primarily with output fluctuations and willing to downplay variability in inflation, expected inflation proves an excessively noisy measure of underlying inflation.

The efficient outcome-based rules respond more aggressively to deviations of inflation from target under learning than implied by perfect knowledge. As seen in the top panel of Figure 7, the efficient choice of θ_π is higher under imperfect knowledge than under perfect knowledge. This result holds across all values of ω . This finding is a manifestation of the need for greater vigilance against inflation when knowledge is imperfect, as discussed in detail in Orphanides and Williams (2004).

The efficient forecast-based rule is more aggressive under learning than under perfect knowledge only when the relative weight on inflation stabilization is relatively low. The reasoning for the more aggressive policy response is the same as in the case of outcome-based rules. Greater vigilance against inflation mitigates against inflation expectations from becoming uncoupled from the policy objective. As can be seen in the lower panel of Figure 7, however, for high values of ω the efficient response is more aggressive under perfect knowledge than learning. The optimal value of θ_{π^e} implied by perfect knowledge is very high

⁸Although not shown in the figure, the difference between the behavior of the economy under outcome- and forecast-based rules is greatest when expected inflation plays a dominant role in determining inflation: For values of ϕ below 0.9, the differences in the frontiers become smaller, and for larger values, the differences increase.

when the policymaker is primarily concerned with inflation stabilization. Under imperfect knowledge, inflation expectations become “noisy” in this economy. Responding aggressively to this noise is counterproductive; instead, the efficient simple rule is characterized by a muted response to inflation expectations.

7.2 Responding to both Actual and Forecasts of Inflation

We now examine the performance and characteristics of policy rules that respond to both observed inflation and the private forecast of inflation. The thin solid line in Figure 8 shows the outcomes under this efficient two-parameter rule. Responding efficiently to both expected and actual inflation outperforms rules responding to either only actual or only expected inflation.

The two-parameter rule uses information regarding the two determinants of inflation in this model: past actual inflation and the private forecast of inflation. To dissect the features of these rules, we compare their properties to rules that respond to the one-step-ahead forecast of inflation implied by the model, which we denote π^p , as opposed to the private forecast of inflation, π^e . Such a rule incorporates information about both observed inflation and the public’s forecast of inflation but constrains how this information is used relative to the two-parameter rule. In particular, the implied ratio of the response to expected inflation to that to observed inflation is given by $\phi/(1 - \phi)$. We use such rules as a benchmark to compare against the efficient two-parameter rules.

The thin dashed line in the figure shows the outcomes when policy responds to the policymaker’s one-period-ahead forecast of inflation, π^p , assuming that the policymaker knows the structural equation for inflation.⁹ This rule performs slightly better than the efficient simple forecast-based rule but does not dominate the simple outcome-based rule. It performs noticeably worse than the two-parameter rule. Evidently, the public’s forecast contains valuable information for the conduct of monetary policy beyond its direct effect on inflation. Examination of the coefficients of the two-parameter efficient rule, shown in Figure 9, indicates that the ratio of the response to expected inflation to observed inflation

⁹Note that this abstracts from the potential pitfalls associated with forecasts based on an incorrect model as in Levin, Wieland and Williams (2003) or on imperfect data as in Orphanides (2003).

is lower than implied by a rule that responds to the policymaker’s forecast (the ratio is the same for $\omega = 1$). That is, the efficient response in the two-parameter rule down-weighs the information contained in the public’s inflation forecast.

7.3 Optimal Policy with Imperfect Knowledge

Up to this point we have restricted ourselves to simple one- and two-parameter simple rules. With imperfect knowledge, optimal policy is described by a nonlinear function of all five states of the system, $\{\pi_t, c_{0,t}, c_{1,t}, R_{1,2,t}, R_{2,2,t}\}$, plus a constant. We have evaluated more complicated rules that respond linearly to all of these states and expected inflation and find that the additional terms yield trivial improvements in economic performance.

8 Conclusion

Central banks around the world pay close attention to inflation expectations, including surveys, market-based measures, and forecasts. One reason for this concern is the possible outbreak of inflation scares, i.e., unusual increases in inflation expectations, that appear to be a recurring phenomenon. In this paper, we explore the properties of endogenous fluctuations in the formation of expectations resulting from a process of perpetual learning and examine its implications for the design of forecast-based monetary policy. Under rational expectations and perfect knowledge, long-run inflation expectations are well anchored and do not budge in response to aggregate shocks. With learning, however, large shocks or a sequence of shocks can dislodge that anchor and an inflation scare may ensue. Inflation expectations can then move substantially away from the policymaker’s target. In this way, our model suggests an important role for learning-induced inflation expectations dynamics for explaining the appearance of “excess sensitivity” of long-term inflation expectations and nominal interest rates to aggregate shocks that is observed in the data.

We also find that under learning private inflation expectations contain potentially valuable information for the setting of monetary policy. In particular, policies that respond to both observed inflation and private inflation expectations yield significant improvements in macroeconomic performance over simple rules that respond to observed inflation.

References

- Beechey, Meredith. "Learning and the Term Structure of Interest Rates." University of California, Berkeley, mimeo, April 2004.
- Bernanke, Ben S. and Woodford, Michael. "Inflation Forecasts and Monetary Policy." *Journal of Money, Credit and Banking*, November 1997, 29(4) Part 2, pp. 653–84.
- Bomfim, Antulio, Robert Tetlow, Peter von zur Muehlen, and John C. Williams. "Expectations, Learning and the Cost of Disinflation," in: *Monetary Policy and the Inflation Process*, Basel, Switzerland: Bank for International Settlements, 1997.
- Bullard, James B. and Mitra, Kaushik. "Learning About Monetary Policy Rules." *Journal of Monetary Economics*, September 2002, 49(6), pp. 1105–1129.
- Campbell, John Y. and Robert J. Shiller, "Yield Spreads and Interest Rate Movements: A Bird's Eye View." *The Review of Economic Studies*, May 1991, 58(3), Special Issue: The Econometrics of Financial Markets, pp. 495–514.
- Christiano, Lawrence J.; Eichenbaum, Martin and Evans, Charles L. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." NBER Working Paper 8403, July 2001.
- Clarida, Richard, Jordi Gali, and Mark Gertler. "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory," *Quarterly Journal of Economics*, February 2000, pp. 147–180.
- Clark, Peter; Goodhart Charles A. E. and Huang, Haizhou. "Optimal Monetary Policy Rules in a Rational Expectations Model of the Phillips Curve." *Journal of Monetary Economics*, April 1999, 43(2), pp. 497–520.
- Cogley, Timothy. "An Exploration of Evolving Term Structure Relations." Mimeo, Arizona State, December 2002.
- Cook, T. and Hahn, T., "The Effect of Changes in the Federal Funds Rate Target on Market Interest Rates in the 1970s." *Journal of Monetary Economics*, 1989, 24, pp. 331–351.
- Erceg, Christopher J. and Andrew Levin. "Imperfect Credibility and Inflation Persistence," *Journal of Monetary Economics*. 50(4) May 2003, pp. 915–944.
- Evans, George and Ramey, Garey. "Adaptive Expectations, Underparameterization and the Lucas Critique." University of Oregon mimeo, May 2001.
- Evans, George and Honkapohja, Seppo. *Learning and Expectations in Macroeconomics*. Princeton: Princeton University Press, 2001a.
- Evans, George and Honkapohja, Seppo. "Expectations and the Stability Problem for Optimal Monetary Policy." Mimeo, 2001b.
- Friedman, Benjamin M. "Optimal Expectations and the Extreme Information Assumptions of 'Rational Expectations' Macromodels." *Journal of Monetary Economics*, January 1979, 5(1), pp. 23–41.
- Gali, Jordi and Gertler, Mark. "Inflation Dynamics: A Structural Economic Analysis." *Journal of Monetary Economics*, October 1999, 44(2), pp. 195–222.

- Gaspar, Vitor and Smets, Frank. "Monetary Policy, Price Stability and Output Gap Stabilisation." *International Finance*, Summer 2002, 5(2), pp. 193–211.
- Goodfriend, Marvin. "Interest Rate Policy and the Inflation Scare Problem: 1979–1992." *Economic Quarterly*, Federal Reserve Bank of Richmond, 79 (1), Winter 1993, pp. 1–23.
- Gürkaynak, Refet, Brian Sack and Eric Swanson, "The Excess Sensitivity of Long Term Interest Rates: Evidence and Implication for Macroeconomic Models." Board of Governors of the Federal Reserve System, FEDS Working Paper 2003-50, 2003.
- Greenspan, Alan. "Economic Developments." Remarks Before the Economic Club of New York, New York, May 24, 2001.
- Ireland, Peter. "Long-Term Interest Rates and Inflation: A Fisherian Approach." *Economic Quarterly*, Federal Reserve Bank of Richmond, 82 (1), Winter 1996, pp. 21–34.
- King, Mervyn. "The Inflation Target Ten Years On." Speech at the London School of Economics, November 19, 2002.
- Kuttner, Kenneth N., "Monetary policy surprises and interest rates: Evidence from the Fed funds futures market." *Journal of Monetary Economics*, June 2001, 47(3), pp. 523–544.
- Kozicki, Sharon and Tinsley, Peter A. "Term Structure Views of Monetary Policy Under Alternative Models of Agent Expectations." *Journal of Economic Dynamics and Control*, 2001a, 25, pp. 149–184.
- Kozicki, Sharon and Tinsley, Peter A. "What Do You Expect? Imperfect Policy Credibility and Tests of the Expectations Hypothesis." Federal Reserve Bank of Kansas City Working Paper 01-02, April 2001b.
- Lengwiler, Yvan and Orphanides, Athanasios. "Optimal Discretion." *Scandinavian Journal of Economics*, June 2002, 104(2), pp. 261–276.
- Levin, Andrew; Wieland, Volker and Williams, John. "The Performance of Forecast-Based Monetary Policy Rules under Model Uncertainty." *American Economic Review*, June 2003, 93(3), pp. 622–645.
- Mankiw, N. Gregory and Lawrence H. Summers, "Do Long-Term Interest Rates Overreact to Short-Term Interest Rates?" *Brookings Papers on Economic Activity*, 1984, 1, pp. 223–242.
- Orphanides, Athanasios. "Monetary Policy Evaluation with Noisy Information." *Journal of Monetary Economics*, 2003, 50(3), 605–631.
- Orphanides, Athanasios and van Norden, Simon. "The Reliability of Inflation Forecasts Based on Output Gap Estimates in Real Time." CIRANO Scientific Series 2003s-1, January 2003.
- Orphanides, Athanasios and Williams, John C. "Monetary Policy Rules with Unknown Natural Rates." *Brookings Papers on Economic Activity*, 2002, pp. 63–145.
- Orphanides, Athanasios and Williams, John C. "Imperfect Knowledge, Inflation Expectations, and Monetary Policy." in: *The Inflation Targeting Debate*, ed. by Ben S. Bernanke and Michael Woodford, Chicago: University of Chicago Press, 2004.

- Roley, V.V. and Sellon, G.H., "Monetary Policy Actions and Long Term Interest Rates." Federal Reserve Bank of Kansas City Economic Quarterly, 1995, 80, pp. 77-89.
- Sargent, Thomas J. *Bounded Rationality in Macroeconomics*. Oxford and New York: Oxford University Press, Clarendon Press, 1993.
- Sargent, Thomas J. *The Conquest of American Inflation*. Princeton: Princeton University Press, 1999.
- Shiller, Robert J., "The Volatility of Long-Term Interest Rates and Expectations Models of the Term Structure." *Journal of Political Economy* December 1979, 87(6), 1190-1219.
- Stock, James H. and Watson, Mark W. "Forecasting Inflation." *Journal of Monetary Economics*, October 1999, 44(2), pp. 293-335.
- Volcker, Paul. "A Rare Opportunity." Remarks before the 43rd annual dinner of the Tax Foundation, New York City, December 3, 1980.

Table 1: Sensitivity of Inflation Forecasts

Correlation Between Inflation Forecasts and Actual Inflation									
	Rational Expectations	Imperfect Knowledge							
		Baseline Expectations Model				Horizon-Specific Expectations Model			
		$\kappa = 0.025$		$\kappa = 0.050$		$\kappa = 0.025$		$\kappa = 0.050$	
		$\phi = 0.75$	$\phi = 0.90$	$\phi = 0.75$	$\phi = 0.90$	$\phi = 0.75$	$\phi = 0.90$	$\phi = 0.75$	$\phi = 0.90$
<i>Policy: $\theta_\pi = 0.79$</i>									
1-step-ahead forecast	0.21	0.29	0.40	0.36	0.58	0.29	0.40	0.36	0.58
3-step-ahead forecast	0.01	0.13	0.27	0.25	0.53	0.11	0.26	0.21	0.51
5-step-ahead forecast	0.00	0.10	0.24	0.22	0.52	0.09	0.22	0.18	0.48
10-step-ahead forecast	0.00	0.09	0.21	0.20	0.55	0.08	0.19	0.16	0.46
<i>Policy: $\theta_\pi = 0.62$</i>									
1-step-ahead forecast	0.38	0.48	0.62	0.57	0.82	0.48	0.62	0.57	0.82
3-step-ahead forecast	0.06	0.24	0.45	0.40	0.78	0.21	0.43	0.35	0.76
5-step-ahead forecast	0.01	0.17	0.39	0.35	0.78	0.14	0.35	0.28	0.74
10-step-ahead forecast	0.00	0.14	0.33	0.31	0.83	0.12	0.29	0.23	0.73
<i>Policy: $\theta_\pi = 0.43$</i>									
1-step-ahead forecast	0.57	0.69	0.84	0.78	0.97	0.69	0.83	0.78	0.97
3-step-ahead forecast	0.18	0.43	0.71	0.64	0.99	0.40	0.69	0.58	0.98
5-step-ahead forecast	0.06	0.33	0.65	0.58	1.02	0.28	0.62	0.49	1.01
10-step-ahead forecast	0.00	0.25	0.60	0.54	1.15	0.19	0.53	0.37	1.12

Notes: The table reports the slope coefficient from a regression of the k-step-ahead inflation forecast implied by the private estimated forecasting model on observed inflation.

Table 2: Sensitivity of Inflation Forecasts with Known Inflation Target

Correlation Between Inflation Forecasts and Actual Inflation									
	Rational Expectations	Imperfect Knowledge							
		Baseline Expectations Model				Horizon-Specific Expectations Model			
		$\kappa = 0.025$		$\kappa = 0.050$		$\kappa = 0.025$		$\kappa = 0.050$	
		$\phi = 0.75$	$\phi = 0.90$	$\phi = 0.75$	$\phi = 0.90$	$\phi = 0.75$	$\phi = 0.90$	$\phi = 0.75$	$\phi = 0.90$
<i>Policy: $\theta_\pi = 0.79$</i>									
1-step-ahead forecast	0.21	0.23	0.28	0.26	0.39	0.23	0.28	0.26	0.38
3-step-ahead forecast	0.01	0.04	0.11	0.09	0.28	0.04	0.10	0.07	0.26
5-step-ahead forecast	0.00	0.01	0.05	0.04	0.23	0.01	0.05	0.03	0.21
10-step-ahead forecast	0.00	0.00	0.02	0.01	0.23	0.01	0.02	0.01	0.17
<i>Policy: $\theta_\pi = 0.62$</i>									
1-step-ahead forecast	0.38	0.42	0.50	0.47	0.68	0.41	0.50	0.47	0.68
3-step-ahead forecast	0.06	0.13	0.25	0.21	0.57	0.11	0.24	0.18	0.55
5-step-ahead forecast	0.01	0.05	0.16	0.13	0.52	0.04	0.14	0.09	0.49
10-step-ahead forecast	0.00	0.01	0.07	0.06	0.52	0.01	0.06	0.03	0.44
<i>Policy: $\theta_\pi = 0.43$</i>									
1-step-ahead forecast	0.57	0.63	0.74	0.69	0.91	0.63	0.74	0.69	0.91
3-step-ahead forecast	0.18	0.30	0.52	0.44	0.87	0.28	0.51	0.40	0.86
5-step-ahead forecast	0.06	0.17	0.42	0.33	0.88	0.25	0.39	0.27	0.86
10-step-ahead forecast	0.00	0.06	0.29	0.22	0.98	0.04	0.25	0.12	0.94

Notes: The table reports the slope coefficient from a regression of the k-step-ahead inflation forecast implied by the private estimated forecasting model on observed inflation.


Table 3: Sensitivity of Forward Inflation Rate to Inflation Surprises

Correlation Between Changes in Inflation Forecasts and Inflation Surprises									
	Rational Expectations	Imperfect Knowledge—Baseline Expectations Model							
		Unknown Target				Known Inflation Target			
		$\kappa = 0.025$		$\kappa = 0.050$		$\kappa = 0.025$		$\kappa = 0.050$	
		$\phi = 0.75$	$\phi = 0.90$	$\phi = 0.75$	$\phi = 0.90$	$\phi = 0.75$	$\phi = 0.90$	$\phi = 0.75$	$\phi = 0.90$
<i>Policy: $\theta_\pi = 0.79$</i>									
3-step-ahead forecast	0.01	0.06	0.10	0.11	0.18	0.04	0.07	0.06	0.12
5-step-ahead forecast	0.00	0.04	0.07	0.09	0.15	0.01	0.03	0.02	0.08
10-step-ahead forecast	0.00	0.03	0.05	0.08	0.15	0.00	0.01	0.01	0.06
<i>Policy: $\theta_\pi = 0.62$</i>									
3-step-ahead forecast	0.05	0.13	0.19	0.20	0.31	0.10	0.16	0.14	0.24
5-step-ahead forecast	0.01	0.07	0.13	0.15	0.27	0.03	0.09	0.07	0.18
10-step-ahead forecast	0.00	0.05	0.08	0.11	0.25	0.00	0.03	0.02	0.12
<i>Policy: $\theta_\pi = 0.43$</i>									
3-step-ahead forecast	0.19	0.27	0.36	0.36	0.50	0.24	0.32	0.30	0.42
5-step-ahead forecast	0.06	0.17	0.27	0.28	0.48	0.12	0.22	0.20	0.35
10-step-ahead forecast	0.00	0.09	0.18	0.23	0.51	0.03	0.11	0.10	0.31


Notes: The table reports the slope coefficient from a regression of the change in the forecast of inflation k periods in the future on the lagged inflation surprise, equal to the difference between the actual inflation rate and the private forecast of inflation.

Figure 1

Optimal Response to Observed Inflation Gap under Perfect Knowledge



Optimal Response to Expected Inflation Gap under Perfect Knowledge



Notes: The top panel shows the optimal response to the observed inflation gap corresponding to the alternative weights ω ; the bottom panel shows the optimal response to the expected output gap inflation gap.



Figure 2

Evolution of Economy Following an Inflation Shock
($\phi = 0.9, \alpha = 0.1$)



Output

Inflation

Biased towards inflation control: $\theta_\pi = 0.8$ or $\theta_{\pi^e} = 3.8$



Balanced Preferences: $\theta_\pi = 0.6$ or $\theta_{\pi^e} = 1.6$



Biased toward Output control: $\theta_\pi = 0.4$ or $\theta_{\pi^e} = 0.8$







Figure 3

Evolution of Economy Following a Series of Inflation Shocks
($\phi = 0.90, \alpha = 0.10$)



Output

Inflation

Biased towards inflation control: $\theta_\pi = 0.8$ or $\theta_{\pi^e} = 3.8$



Balanced Preferences: $\theta_\pi = 0.6$ or $\theta_{\pi^e} = 1.6$



Biased toward Output control: $\theta_\pi = 0.4$ or $\theta_{\pi^e} = 0.8$







Figure 4

Evolution of Inflation Expectations Following an Inflation Shock



One-step-ahead expectations

Five-step-ahead expectations

Biased towards inflation control: $\theta_\pi = 0.8$ or $\theta_{\pi^e} = 3.8$



Balanced Preferences: $\theta_\pi = 0.6$ or $\theta_{\pi^e} = 1.6$



Biased toward Output control: $\theta_\pi = 0.4$ or $\theta_{\pi^e} = 0.8$







Figure 5

Evolution of Inflation Expectations Following a Series of Inflation Shocks



One-step-ahead expectations

Five-step-ahead expectations

Biased towards inflation control: $\theta_\pi = 0.8$ or $\theta_{\pi^e} = 3.8$



Balanced Preferences: $\theta_\pi = 0.6$ or $\theta_{\pi^e} = 1.6$



Biased toward Output control: $\theta_\pi = 0.4$ or $\theta_{\pi^e} = 0.8$





Figure 6

Performance of Optimized One-parameter Policy Rules
($\phi = 0.9, \alpha = 0.1$)





Figure 7
 Optimized Response to Observed Inflation in One-parameter Rule
 ($\phi = 0.9, \alpha = 0.1$)



Optimized Response to Expected Inflation in One-parameter Rule




Figure 8

Performance of Optimized One- and Two-parameter Policy Rules
($\phi = 0.9, \alpha = 0.1$)




Figure 9

Optimized Coefficients of Two-Parameter Policy Rule
($\phi = 0.9, \alpha = 0.1$)

