Inflation Scares and Forecast-Based Monetary Policy

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

April 2003

Abstract

Central banks pay close attention to inflation expectations. In standard models, however, inflation expectations are tied down by the assumption of rational expectations and should be of little independent interest to policy makers. In this paper, we relax the assumption of rational expectations with perfect knowledge and reexamine the role of inflation expectations in the economy and in the conduct of monetary policy. Agents are assumed to have imperfect knowledge of the precise structure of the economy and the policymakers' preferences. Expectations are governed by a perpetual learning technology. 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. The presence of learning raises 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.

KEYWORDS: Inflation forecasts, policy rules, rational expectations, learning.

JEL Classification System: E52

Correspondence: Orphanides: Federal Reserve Board, Washington, D.C. 20551, Tel.: (202) 452-2654, e-mail: Athanasios.Orphanides@frb.gov. Williams: Federal Reserve Bank of San Francisco, 101 Market Street, San Francisco, CA 94105, Tel.: (415) 974-2240, e-mail: John.C.Williams@sf.frb.org.

* We would like to thank George Evans, Ben Friedman, Peter Ireland, Lars Svensson, and participants at the Federal Reserve Bank of Atlanta Conference on Learning, March 21-22, 2003, and the Januray 2003 Meeting of the American Economic Association for useful comments and discussions on earlier drafts. The opinions expressed are those of the authors and do not necessarily reflect views of the Board of Governors of the Federal Reserve System or the Federal Reserve Bank of San Francisco.

1 Introduction

What is the role of inflation expectations in the monetary policy process? Policymakers in the United States and many other nations have long recognized that for monetary policy to be successful, inflation expectations must be well anchored. Indeed, successful policymaking, by its very nature, is forward-looking in its orientation and policymakers have often stressed the importance of preemptive policy action.[1] Monitoring information regarding inflation expectations, as reflected in surveys or financial markets, for instance, has been part of the policy process at many central banks.[2] Forecasts of inflation are at the center of policy deliberations at inflation-targeting central banks, and have arguably been as important for policy decisions in non-inflation-targeting central banks such as the Federal Reserve and the European Central Bank.

Why do inflation expectations receive so much attention at central banks? One reason often cited is the lagged effect of monetary policy actions on output and inflation.[3] Given the policy lag, it makes sense that policy decisions should be based on expected future conditions when the policy action will first take hold. However, a number of studies, including ones conducted using models developed at central banks that incorporate an important role for expectations formation and policy lags, find that simple policy rules prescribing that the central bank set interest rates in response to recent outcomes of inflation and economic activity similar to the Taylor (1993) rule provide remarkably good performance.[4] Indeed,

[1] Recent examples of such policymaker views at the Federal Reserve can be found in Greenspan (2001), Meyer (2002) and Bernanke (2003). Views of policymakers from other central banks are reflected in King (2000), Issing (2000), Gjedrem (2001), Bollard (2002), and elsewhere. At the Federal Reserve, recognition of the value of pre-emptive policies can be traced virtually to the founding of the System (Orphanides, 2003).

[2] In the United States, the Federal Reserve Bank of Philadelphia publishes a quarterly report from a Survey of Professional Forecasters. The European Central Bank, Bank of England, Reserve Bank of New Zealand, Reserve Bank of Australia and Sveriges Riksbank regularly report on similar survey forecasts. Information regarding inflation expectations derived from comparisons of prices on inflation-indexed and nominal Government securities is also regularly presented by several central banks, including the European Central Bank, Reserve Bank of Australia and Sveriges Riksbank. Such information is also available at the Federal Reserve (Greenspan, 2000).

[3] In addition to policy lags, Batini and Haldane (1999) identify two other arguments in favor of forecast-based rules: forecasts encompass a wide variety of information and responding to inflation forecasts as opposed to current inflation reduces the variability of output. Giannoni and Woodford (2002) and Svensson and Woodford (2003) show that optimal policies can be formulated in terms of forecasts of the target variables and argue that such policies are robust to certain types of model uncertainty.

[4] See Bryant, Hooper, Mann (1993) and Taylor (1999a) for collections of policy evaluations studies, and Taylor (1999b), Orphanides and Williams (2002), Levin and Williams (2003), and references therein for more recent studies. Sims (2002) and Pagan (2003) provide recent reviews of models currently used at central

Levin, and Wieland, Williams (2003) find that forecast-based policy rules provide only trivial gains over rules that instead respond to current inflation and output in several rational expectations macro models.[5] These findings call into question the traditional reasons given why policymakers should be concerned with inflation expectations.

A potential source of this apparent disconnect between policy practice and policy modeling may be identified in the rigid imposition of rational expectations in macroeconometric models with an assumed fixed and known structure. The policy evaluations described above generally assume a fixed and perfectly known structure of the economy and specify that expectations are model consistent. In linear fixed-parameter models of this nature, for example, once the rule specifying monetary policy is specified, inflation expectations can be represented as a fixed linear function of economic outcomes.[6] Economic agents are then assumed to form expectations mechanically based on these simple linear functions of economic outcomes that are assumed to be perfectly known. These assumptions greatly simplify the analysis from a modeling perspective. But what if agents are, in fact, less than perfectly certain of the structure of the model, its time invariance, or simply the values of the model parameters? Once imperfect knowledge is acknowledged, the tight mechanical link from economic outcomes to the expectations formation process is broken. 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 instead.

Concern for misspecification of the expectations formation process is not merely a theoretical curiosity. Episodes when expectations appeared to have become unmoored from the policymakers' objectives can be easily identified in the monetary history of the United States and other nations. An example of such an episode was seen in the United States at the very end of the 1970s inflationary experience. Reflecting on the evolution of inflation

banks.

[5]This result is not surprising in small-scale models for which the "optimal" policy can be approximated or even fully described by a few variables capturing recent economic outcomes; see Svennson (1997), Ball (1999), Rudebusch and Svensson (1999), Orphanides and Wieland (2000), Orphanides (2003) and Woodford (2002). But, Levin, Wieland, and Williams' analysis includes medium- and large-scale macro models in which optimal policy involves responses to hundreds of variables.

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

expectations in December, 1980, Chairman Volcker noted: "With all its built-in momentum and self-sustaining expectations, [the inflationary process] has come to have a life of its own." Fears of inflation or deflation, whether entirely justified from the policymakers' perspective or not, seemed to have influenced actual decision making and economic behavior at times, presenting real complications for policy decisions. This is the essence of the "inflation scare" problem described by Goodfriend (1993) as complicating monetary policy decisions in practice.

In this paper we break the tight link between inflation expectations and observable macro variables by positing that agents do not know with certainty the parameters of the model but instead constantly update their estimates based on available information. We explore two related issues. First, we examine the occurrence and properties of inflation scares, defined to be deviations of inflation expectations from those implied by rational expectations, under learning. Under rational expectations, long-run inflation expectations are well anchored and are therefore insensitive to shocks. In contrast, under perpetual learning, we find that long-run inflation expectations drift endogenously in response to macroeconomic disturbances in a pattern supported by the evidence on the excess sensitivity of yields on long-term bonds to aggregate shocks. The prevalence and severity of endogenous inflation scares is affected by the monetary policy in place, with policies that emphasize output stabilization more prone to inflation scares. Second, we compare the performance of policies based on observed outcomes to those based on inflation expectations. In our model, forecast-based and outcome-based policies are isomorphic under the assumption of rational expectations. As we show, however, forecast- and outcome-based policies are no longer identical when knowledge is imperfect and inflation expectations cannot be summarized as a simple function of inflation outcomes. Under these circumstances, monitoring and responding to the public's inflation expectations, in addition to monitoring the evolution of actual inflation leads to improved policy outcomes. In our analysis we also differentiate between the public's expectations and the policymaker's inflation forecasts under the assumption that the policymaker knows the structure of the economy and explore the marginal value of reliance on additional information about the economy for policy design.

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 (2003).[7]

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^*,

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

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

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

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

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

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

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.[8] For this fraction, price setting is invariant to the expectations 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.

[7] See also Clark, Goodhart, and Huang (1999) and Lengwiler and Orphanides (2002).

[8] This specification, where a portion of inflation expectations is indexed to past inflation, is similar to those of Gali and Gertler (2000) and Christiano, Eichenbaum, and Evans (2001).

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, in which case 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 rate and one lead of the intended output gap:

$$x_t = E_{t-1}\left\{x_{t+1} - \frac{\omega}{1-\omega}\frac{\alpha}{1-\phi}\pi_{t+1}\right\}. \tag{4}$$

This expression can can be equivalently restated in a number of equivalent 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 the 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:

$$x_t = -\theta_\pi(\pi_t - \pi^*), \tag{5}$$

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

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

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:

$$\pi^e_{t+1} = \frac{\alpha\theta}{1-\phi}\pi^* + \frac{1-\phi-\alpha\theta}{1-\phi}\pi_t. \tag{7}$$

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

$$\pi_{t+1} = \frac{\alpha\theta}{1-\phi}\pi^* + (1 - \frac{\alpha\theta}{1-\phi})\pi_t + e_{t+1} + \alpha u_{t+1}. \tag{8}$$

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:

$$x_t = -\theta_{\pi^e}(\pi^e_{t+1} - \pi^*), \tag{9}$$

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,

$$\theta^*_{\pi^e} = \frac{1-\phi}{1-\phi-\alpha\theta^*_\pi}\theta^*_\pi, \tag{10}$$

for $\omega \in (0,1)$. In the limiting case of $\omega \to 1$, 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: 0.75 and 0.90. For smaller values of ϕ, the affect 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 for values of ω between zero and one. Note that the optimal value of each parameter depends only on ω and the ratio $\alpha/(1-\phi)$ are thus is invariant to the two model parameterizations considered here. As seen in the chart, θ_{π^e} is much more sensitive to ω than 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. We focus on least squares learning because of its desirable properties, straightforward implementation, and close correspondence to the practice of real-world forecasters. 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 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 (2003), 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.[9] 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

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

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 employ 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:

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

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

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 converges to the correct expectations functions 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

$$\pi_{t+1}^e = c_{0,t} + c_{1,t}\pi_t, \tag{14}$$

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 $\kappa \to 0$), 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 employ to form expectations. Nonetheless, as shown in Orphanides and Williams (2003), 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.

5 Learning and Inflation Scares

As noted in the introduction, inflation scares, increases in long-run inflation expectations—evidenced by shifts in the yield curve—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 in response to a wide variety of innovations. 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 (2003) for a discussion of the sensitivity of results to κ.) 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 response to an output shock (or policy control error) is observationally equivalent (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. The shock occurs in the first period and is equal to 2 percentage points.

With 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. This can be seen in Figure 2, which reports this experiment for the case $\alpha = 0.10$ and $\phi = 0.90$. (We consider the second calibration below.) 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 (red) 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 to inflation, 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 (green) 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.

The differences between the outcome- and forecast-based rules are less pronounced when the role of expectations in the determination of inflation is smaller. Figure 3 repeats the same experiments as described above, but assuming that $\phi = 0.75$ and $\alpha = 0.25$. Qualitatively, the results are the same as before, but the quantitative effects of learning are more muted and the differences between the simulated outcomes from the two rules are likewise smaller.

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 affects of learning in the model and thus provides a useful test bed to explore the interaction of policy and learning. Importantly, since the model is non-linear under learning, the response of the economy 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, as evidenced for instance by the events of the 1970s. This is also of interest as an illustration of the importance of initial conditions regarding the formation of inflation expectations for the response of the economy to a shock. 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 4. (As before, we discuss the case $\alpha = 0.10$ and $\phi = 0.90$. Figure 5 shows the parallel results for the alternative calibration $\alpha = 0.25$ and $\phi = 0.75$.) 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, 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. Examine first the case of the policymaker who responds to actual inflation, shown by the dashed lines in Figure 4 and compare that to the response to a single shock, shown in Figure 2. The shocks cause inflation to persistently rise above the target level and, for the policy placing greater emphasis on output stabilization, to continue 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 in this case, 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 (green) 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.

5.3 The Response of Inflation Expectations to Temporary 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 horizon, 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 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, $c_{0,t}$ $c_{1,t}$ 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 $\kappa \to 0$), 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.

In Figures 6 and 7 we show the evolution of inflation expectations when the economy

is subjected to the shocks described in our previous experiments (Figures 2–5). 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 were qualitatively similar in this case.)

The solid (blue) lines in Figure 6 show the evolution of expectations under perfect knowledge following a one-period shock to inflation. In each panel the thicker line indicates the response of the one-period ahead inflation expectation and the thinner line the response of the five-period ahead inflation expectation. As can be seen, for all three policies considered, the five-year ahead inflation expectations are little affected by the shock, which mostly affected the evolution of the one-period ahead expectation. The initial response and speed of adjustment are influenced by the responsiveness of policy, as would expected. but the one-period ahead expectation quickly reverts to normal, 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, it is transmitted to longer-run expectations as well. This is most evident for the case of policy rules responding to lagged inflation, dashed (red) 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 appear to significantly "overreact" to the temporary shock relative to what would be expected with perfect knowledge.

Figure 7 reports the parallel experiment examining the evolution of the economy to a sequence of serially correlated shocks, as shown in Figures 4 and 5. 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 panel). The problem is evident for forecast-based policies as well, but is less severe under these policies.

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 and 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 Gurkaynak, Sack and Swanson (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.

One 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:

$$\pi^e_{t+k} = a_{0,k} + a_{1,k}\pi_t + u_t.$$

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 the 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, on the other hand, the implied persistence of inflation expectations can be considerably higher. Table 1 reports the resulting slope coefficients from simulation experiments for the three alternative outcome-based policies also examined before. We report the results for the one-, three- and five-step-ahead forecasts, as well as for the forecast

of average inflation over the next five periods. In each case, we compare results under rational expectations, with our baseline forecasting model as well as the horizon-specific forecasting model under learning. As can be seen, relative to the sensitivity under rational expectations, under learning inflation expectations exhibit greater sensitivity to 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 illustrate 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 explicitly knowledge regarding the policymaker's ultimate inflation objective in forming expectations. To the extent a central bank could communicate successfully a numerical inflation target to the public, however, the private agents' forecasting problem would become considerably simpler. Since 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.[10] This also allows us to examine the extend 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. Table 2 reports the parallel results for the case of a known target. As expected, inflation expectations under learning are less

[10]To 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.

sensitive to inflation when the inflation target is assumed to be known. However, even with that assumption, inflation expectations can be substantially more sensitive to shocks relative to 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.

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. The particular example of the serially correlated shocks is not, however, representative of the "typical" pattern of serially uncorrelated shocks that we assumed in deriving the optimal policies. We now consider the relative performance of outcome- and forecast-based rules with 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. For comparison, the solid line in the upper panels of Figure 8 show the best obtainable pairs of the standard deviations of inflation and the output gap under the assumption of perfect knowledge. The solid lines in the middle panel report the corresponding optimal values of θ_π for an outcome-based rule; the solid lines of the lower panel report the optimal values of θ_{π^e} for a forecast-based rule. In each case, the left-hand panel reports results for the $\phi = 0.90$ specification; the right-hand panel for the $\phi = 0.75$ specification.

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 lines

in the upper panels show the frontier for the one-parameter outcome-based rule; the dash-dotted curves show the frontiers 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 rule dominates the other, and both do significantly worse than would result under perfect knowledge. Consistent with the results of our simulation evidence, 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.

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 middle panels, 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 (2003).

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. For high values of ω, however, 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 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 lines in the upper panel of Figure 9 show the outcomes under this efficient two-parameter rule. These rule significantly outperform the one-parameter rules in the case of $\phi = 0.90$; the differences are smaller when $\phi = 0.75$. The response coefficients of the two-parameter rules are shown in the lower panels.

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. 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 upper panel show the outcomes when policy responds to the policymaker's one-period-ahead forecast of inflation, denoted by π^p, assuming that the policymaker knows the structural equation for inflation. As seen in the figure, this rule performs slightly better than the efficient simple forecast-based rule does not dominate the simple outcome-based rule. Not surprisingly, it performs 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 indicates that the ration of the response to expected inflation to

observed inflation 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.

As before, the benefits of responding to forecasts of inflation are the greatest when inflation expectations have a large weight in determining inflation. The greater the extent that inflation is determined by a *fixed* relationship to past inflation, that is, the smaller is ϕ, then the less important expectations are for policy in this model.

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 economic forecasts. Inflation scares, unusual increases in inflation expectations, appear to be a recurring phenomenon of concern to policymakers. But model-based monetary policy evaluations suggest that outcome-based monetary policy rules similar to the Taylor Rule, that respond to observed output and inflation, do nearly as well at achieving policy goals as rules based on forecasts. Existing research has provided little insight into why central banks are so fixated on inflation expectations.

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 "excess sensitivity" of long-term inflation expectations to aggregate shocks than 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

Ball, Laurence. "Policy Rules for Open Economies." In Taylor, John B. ed., *Monetary Policy Rules.* Chicago: University of Chicago, 1999.

Batini, Nicoletta and Haldane, Andrew. "Forward-Looking Rules for Monetary Policy," in John B. Taylor, ed., *Monetary Policy Rules.* Chicago: University of Chicago Press, 1999, pp. 157–92.

Bernanke Ben S. "'Constrained discretion' and monetary policy." Remarks Before the Money Marketeers of New York University, New York, New York, February 3, 2003.

Bernanke, Ben S. and Mishkin, Frederic S. "Inflation Targeting: A New Framework for Monetary Policy?" *Journal of Economic Perspectives*, Spring 1997, 11(2), pp. 97–116.

Bernanke, Ben S. and Woodford, Michael. "Inflation Forecasts and Monetary Policy. *Journal of Money, Credit and Banking*, 29(4), 653-84, Part 2, November 1997.

Bollard, Alan. "The Evolution of Monetary Policy." A speech to the Rotary Club of Wellington, November 25, 2002.

Bullard, James B. and Mitra, Kaushik. "Learning About Monetary Policy Rules." *Journal of Monetary Economics*, September 2002, 49(6), 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, 495-514.

Carey, Kevin and Faruquee, Hamid. "Credibility Acquisition without Recession: Lessons from the 1994 Inflation Scare." mimeo, International Monetary Fund, January 2003.

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.

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, 331-351.

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, 2001.

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.

Giannoni, Marc P. and Woodford, Michael. "Optimal Interest-Rate Rules: I. General Theory." Princeton University manuscript, August 2002.

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.

Gurkaynak, Refet, Brian Sack and Eric Swanson, "The excess Sensitivity of Long Term Interest Rates: Evidence and Implication for Macroeconomic Models." Mimeo, Federal Reserve Board, 2003.

Greenspan, Alan. "Technology and the economy." Remarks Before the Economic Club of New York, New York, New York January 13, 2000.

Greenspan, Alan. " Economic developments." Remarks Before the Economic Club of New York, New York, May 24, 2001.

Gjedrem, Svein. "Monetary policy, forecasts and market communication." Remarks Before The Norwegian School of Management's Centre for Monetary Economics, June 7, 2001.

Ireland, Peter. "Long-Term Interest Rates and inflation: A Fisherian Approcah." *Economic Quarterly*, Federal Reserve Bank of Richmond, 82 (1), Winter 1996, pp. 21–34.

Issing, Otmar. "Monetary policy in a new environment" Speech at Bundesbank-BIS Conference, Frankfurt am Main, September 29, 2000.

King, Mervyn. "Monetary Policy: Theory and Practice." Address to the joint luncheon of the American Economic Association and the American Finance Association, Boston, January 7, 2000.

Kuttner, Kenneth N., "Monetary policy surprises and interest rates: Evidence from the Fed funds futures market." *Journal of Monetary Economics*, June 2001, 47(3), 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, 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 T. and Williams, John C. "Robust Monetary Policy with Competing Reference Models." *Journal of Monetary Economics*, 2003, Forthcoming.

Levin, Andrew; Wieland, Volker and Williams, John. "The Performance of Forecast-Based Monetary Policy Rules under Model Uncertainty." *American Economic Review*, 2003, forthcoming.

Mankiw, N. Gregory and Lawrence H. Summers, "Do Long-Term Interest Rates Overreact to Short-Term Interest Rates?" *Brookings Papers on Economic Activity*, 1984, 1:1984, 223-242.

Meyer, Laurence H., 2002. *Rules and Discretion*, Remarks At the Owen Graduate School of Management, Vanderbilt University, Nashville, Tennessee, January 16.

Orphanides, Athanasios. "Historical Monetary Policy Analysis and the Taylor Rule." March 2003.

Orphanides, Athanasios. "Monetary Policy Evaluation with Noisy Information." *Journal of Monetary Economics*, 2003b, forthcoming.

Orphanides, Athanasios and Wieland, Volker. "Inflation Zone Targeting." *European Economic Review*, June 2000, 44(7), pp. 1351–1387.

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 Imperfect Knowledge." in: *Inflation Targeting*, ed. by Michael Woodford, Chicago: University of Chicago Press, 2003, forthcoming.

Pagan, Adrian. "Report on Modelling and Forecasting at the Bank of England." London, Bank of England, January 2003.

Phelps, Edmund S. and Taylor, John B. "Stabilizing Powers of Monetary Policy under Rational Expectations." *Journal of Political Economy*, 1977, 85(1), pp. 163–190.

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, 77-89.

Rudebusch, Glenn D. and Svensson, Lars E. O. "Policy Rules for Inflation Targeting," in John B. Taylor, ed., *Monetary Policy Rules*, Chicago: University of Chicago Press, 1999, pp. 203–53.

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.

Sims, Christopher. "The Role of Models and Probabilities in the Monetary Policy Process." *Brookings Papers on Economic Activity*, 2002 (2), pp. 1–62.

Stock, James H. and Watson, Mark W. "Forecasting Inflation." *Journal of Monetary Economics*, October 1999, 44(2), pp. 293–335.

Svensson, Lars E. O. "Inflation Forecast Targeting: Implementing and Monitoring Inflation Targets," *European Economic Review*, 41(6), 1111-1146, June 1997.

Svensson, Lars E. O. and Woodford, Michael. "Implementing Optimal Policy through Inflation-Forecast Targeting." in: *Inflation Targeting*, ed. by Michael Woodford, Chicago: University of Chicago Press, 2003, forthcoming.

Taylor, John B. "Discretion versus Policy Rules in Practice." *Carnegie Rochester Conference Series on Public Policy*, 1993, 39, pp. 195–214.

Taylor, John B. ed., *Monetary Policy Rules.* Chicago: University of Chicago, 1999a.

Taylor, John B. "The Robustness and Efficiency of Monetary Policy Rules as Guidelines for Interest Rate Setting by the European Central Bank." *Journal of Monetary Economics*, June 1999b, 43(3), pp. 655–79.

Volcker, Paul. "A Rare Opportunity." Remarks before the 43rd annual dinner of the Tax Foundation, New York City, December 3, 1980.

Williams, John C. "Simple Rules for Monetary Policy." *Federal Reserve Bank of San Francisco Economic Review*, 2003, forthcoming.

Woodford, Michael. *Interest and Prices: Foundations of a Theory of Monetary Policy.* Manuscript in preparation for Princeton University, 2002

Table 1: Sensitivity of Inflation Forecasts

	Correlation Between Inflation Forecasts and Actual Inflation								
	Rational Expectations	Imperfect Knowledge							
		Baseline Expectations Model				Horizon-Specific Expectations Model			
		$\kappa = 0.025$		$\kappa = 0.050$		$\kappa = 0.025$		$\kappa = 0.050$	
		$\phi = 0.75$	$\phi = 0.90$	$\phi = 0.75$	$\phi = 0.90$	$\phi = 0.75$	$\phi = 0.90$	$\phi = 0.75$	$\phi = 0.90$
Policy: $\theta_\pi = 0.79$									
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
Policy: $\theta_\pi = 0.62$									
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
Policy: $\theta_\pi = 0.43$									
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: 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. For the 5-period average forecast, the slope coefficients from the one- through five-step-ahead forecast regressions are averaged.

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			
		$\kappa = 0.025$		$\kappa = 0.050$		$\kappa = 0.025$		$\kappa = 0.050$	
		$\phi = 0.75$	$\phi = 0.90$	$\phi = 0.75$	$\phi = 0.90$	$\phi = 0.75$	$\phi = 0.90$	$\phi = 0.75$	$\phi = 0.90$
Policy: $\theta_\pi = 0.79$									
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
Policy: $\theta_\pi = 0.62$									
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
Policy: $\theta_\pi = 0.43$									
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: 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. For the 5-period average forecast, the slope coefficients from the one- through five-step-ahead forecast regressions are averaged.

Figure 1

Optimal Response to Observed Inflation Gap under Perfect Knowledge

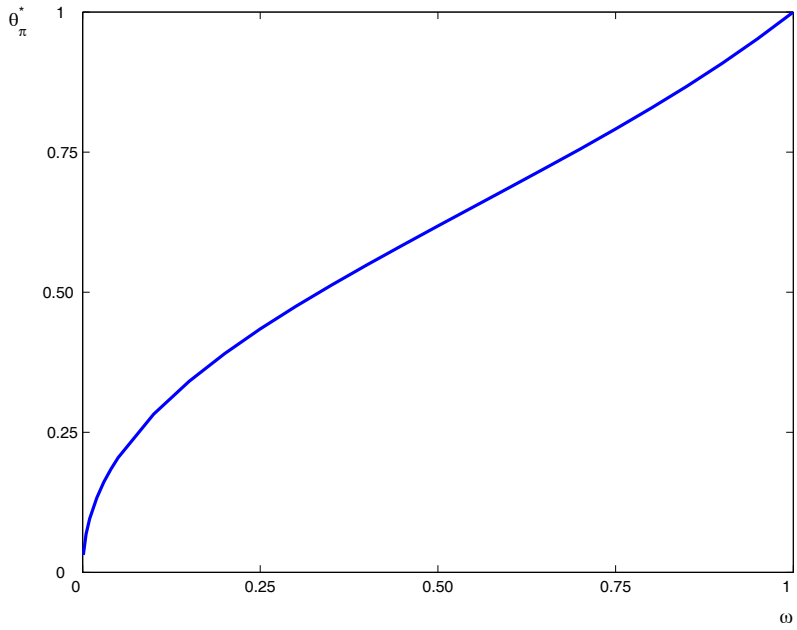

Optimal Response to Expected Inflation Gap under Perfect Knowledge

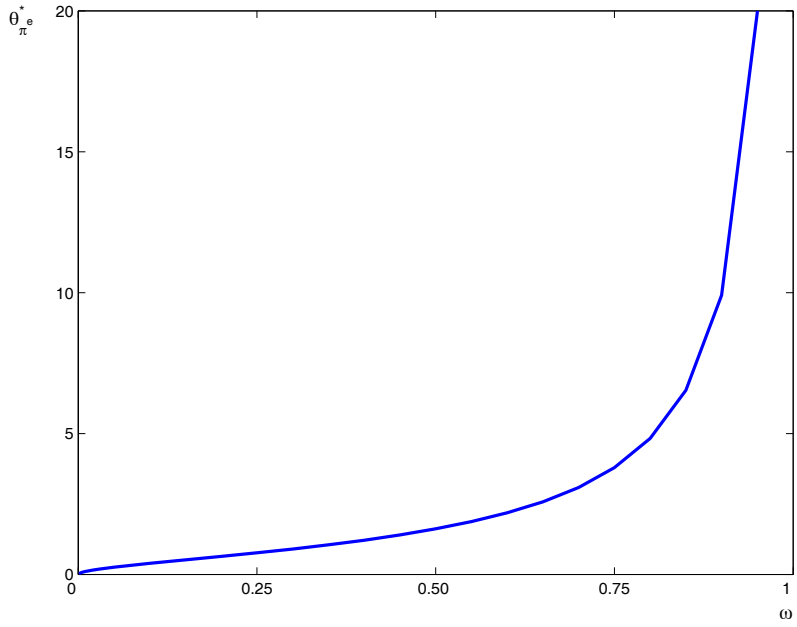

Notes: The top panel shows the optimal response to the observed inflation gap corresponding to the alternative weights ω; the bottom panel shows the optimal response to the expected output gap inflation gap.

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* $\theta_{\pi^e} = 3.8$

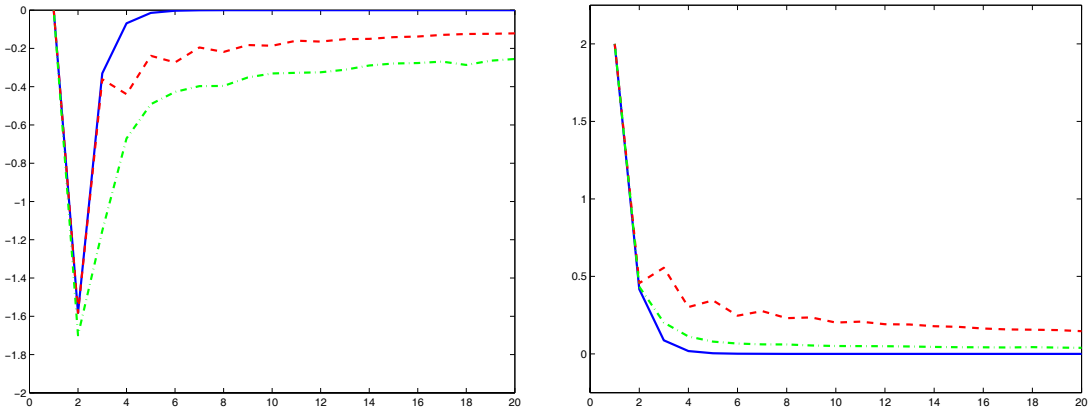

Balanced Preferences: $\theta_\pi = 0.6$ *or* $\theta_{\pi^e} = 1.6$

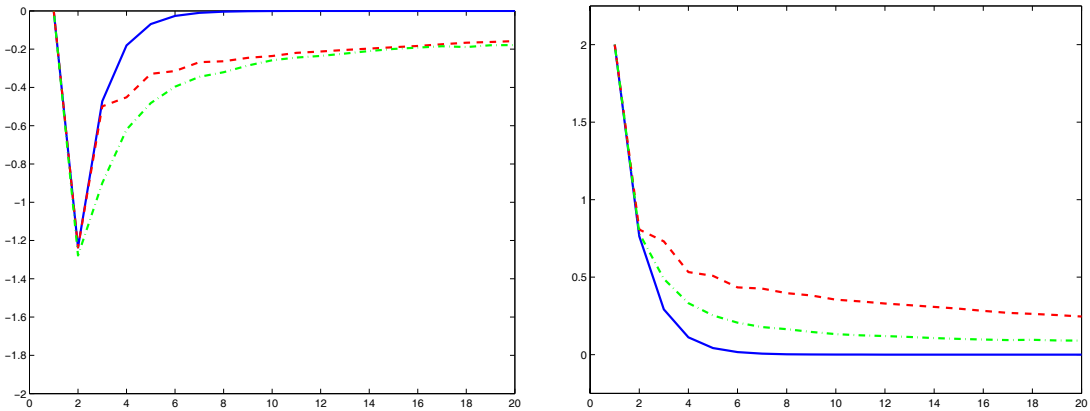

Biased toward Output control: $\theta_\pi = 0.4$ *or* $\theta_{\pi^e} = 0.8$

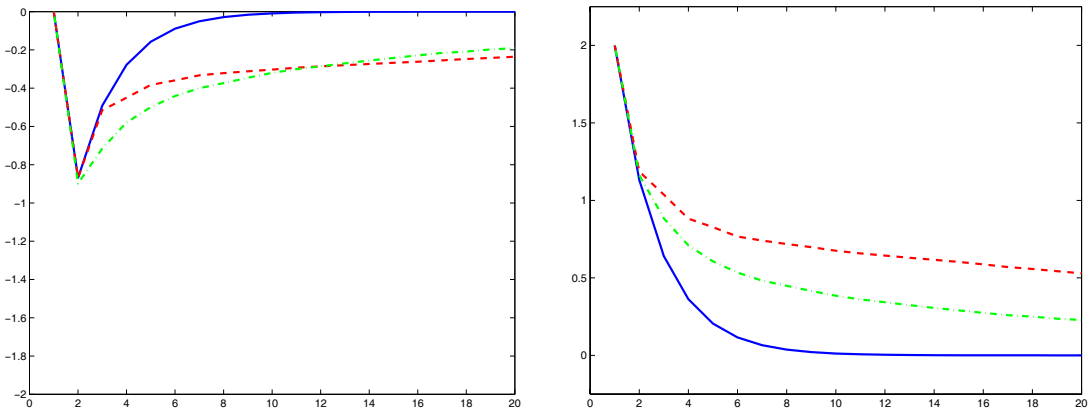

Figure 3

Evolution of Economy Following an Inflation Shock
$(\phi = 0.75, \alpha = 0.25)$

Output **Inflation**

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

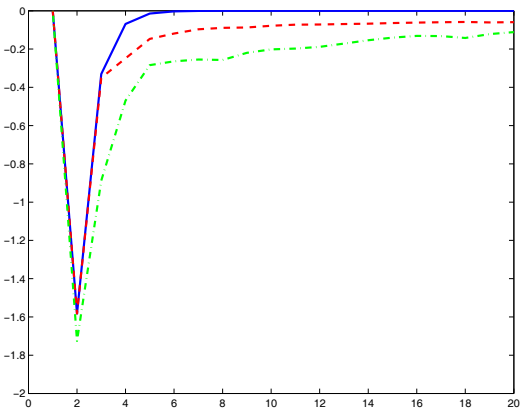
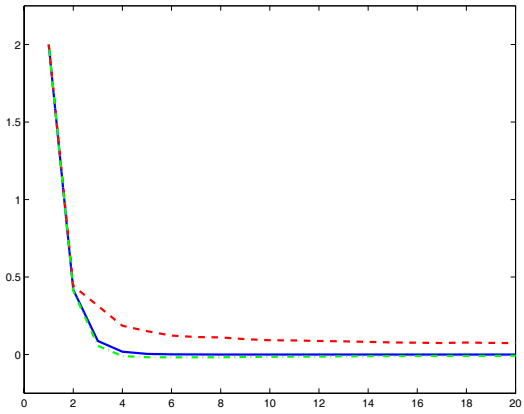

Balanced Preferences: $\theta_\pi = 0.6$ *or* $\theta_{\pi^e} = 1.6$

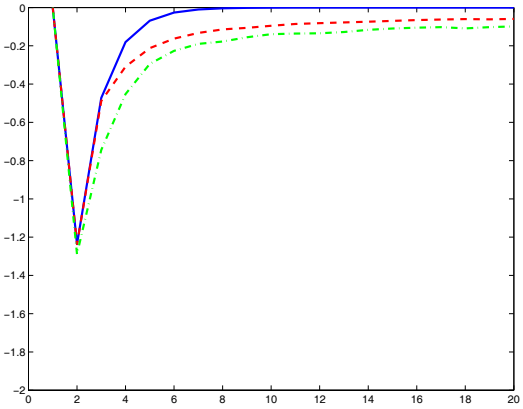
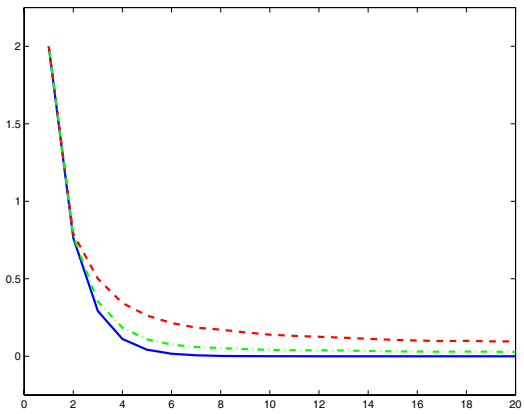

Biased toward Output control: $\theta_\pi = 0.4$ *or* $\theta_{\pi^e} = 0.8$

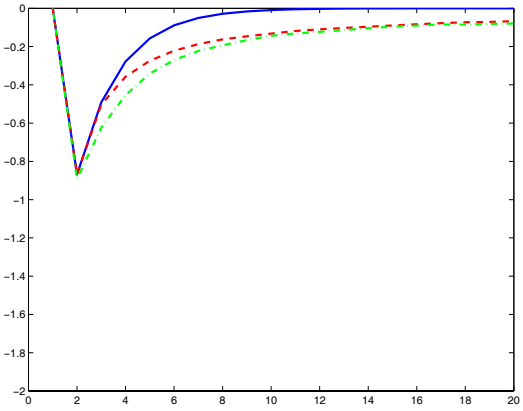
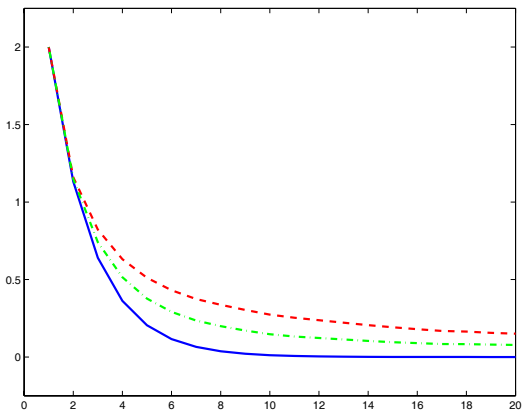

Figure 4

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

Output **Inflation**

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

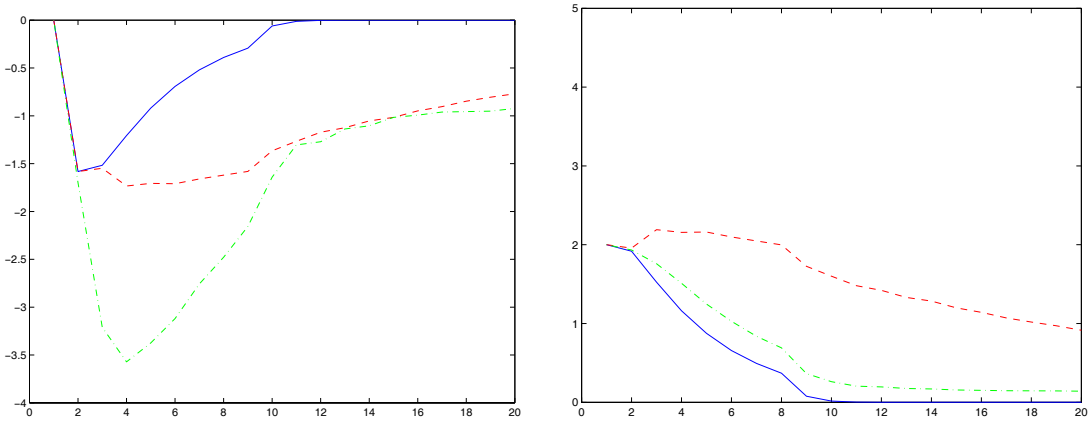

Balanced Preferences: $\theta_\pi = 0.6$ *or* $\theta_{\pi^e} = 1.6$

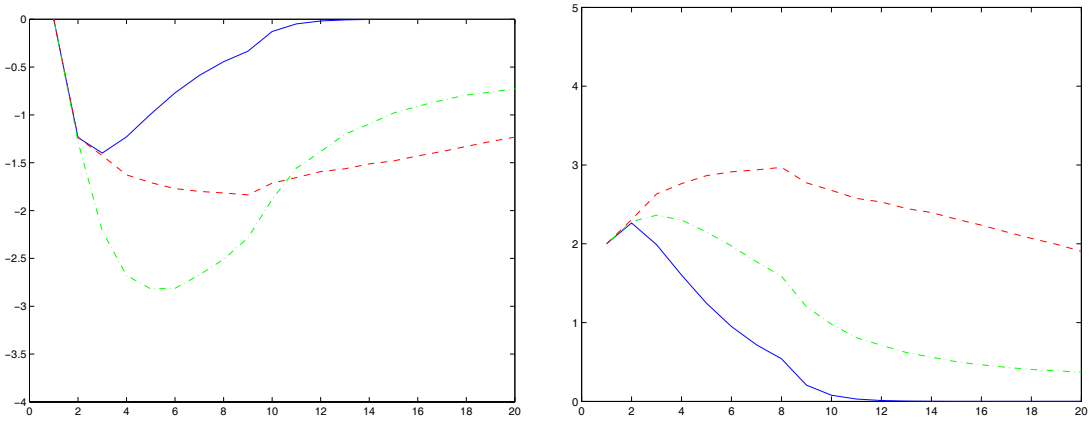

Biased toward Output control: $\theta_\pi = 0.4$ *or* $\theta_{\pi^e} = 0.8$

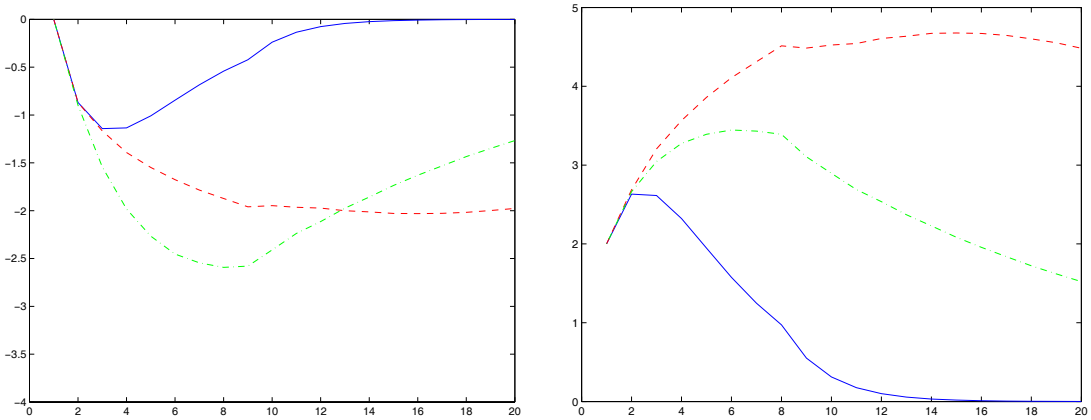

Figure 5

Evolution of Economy Following a Series of Inflation Shocks
$(\phi = 0.75, \alpha = 0.25)$

Output **Inflation**

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

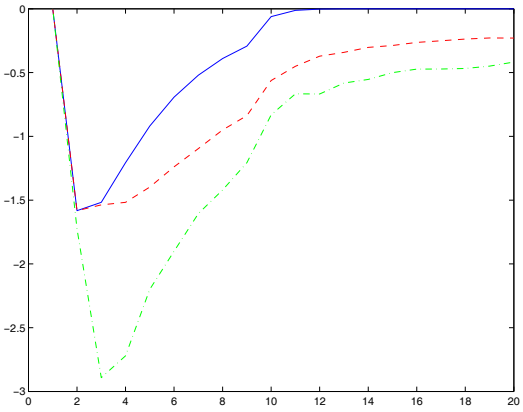
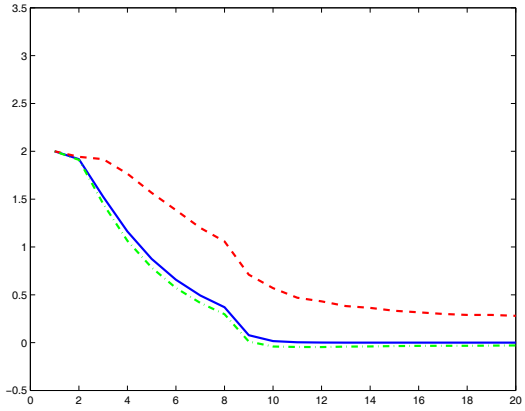

Balanced Preferences: $\theta_\pi = 0.6$ *or* $\theta_{\pi^e} = 1.6$

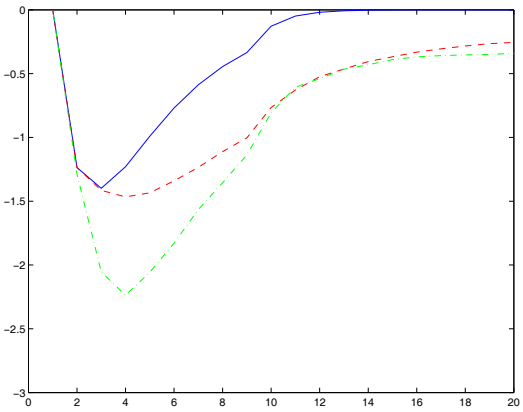
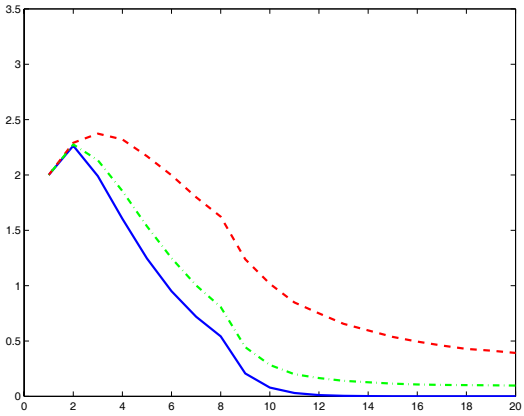

Biased toward Output control: $\theta_\pi = 0.4$ *or* $\theta_{\pi^e} = 0.8$

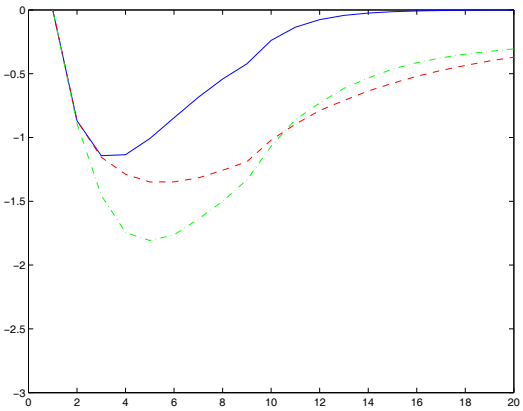
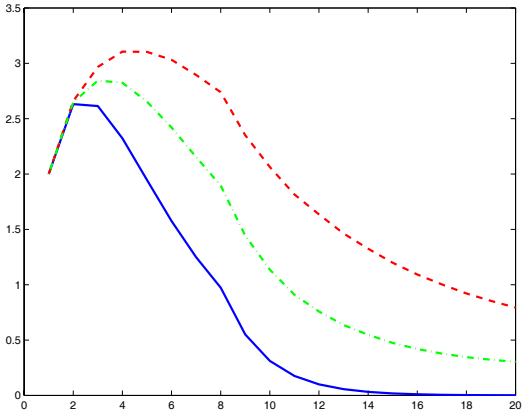

Figure 6

Evolution of Inflation Expectations Following an Inflation Shock

$(\phi = 0.75, \alpha = 0.25)$ $(\phi = 0.9, \alpha = 0.1)$

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

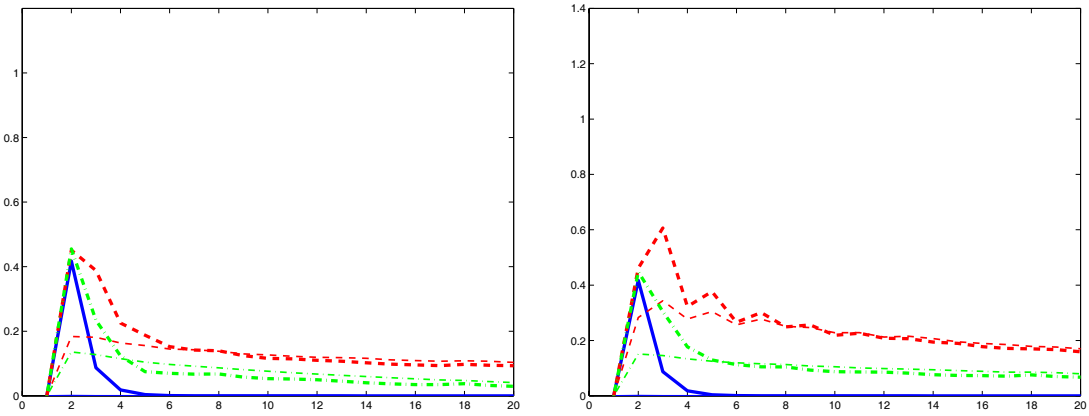

Balanced Preferences: $\theta_\pi = 0.6$ *or* $\theta_{\pi^e} = 1.6$

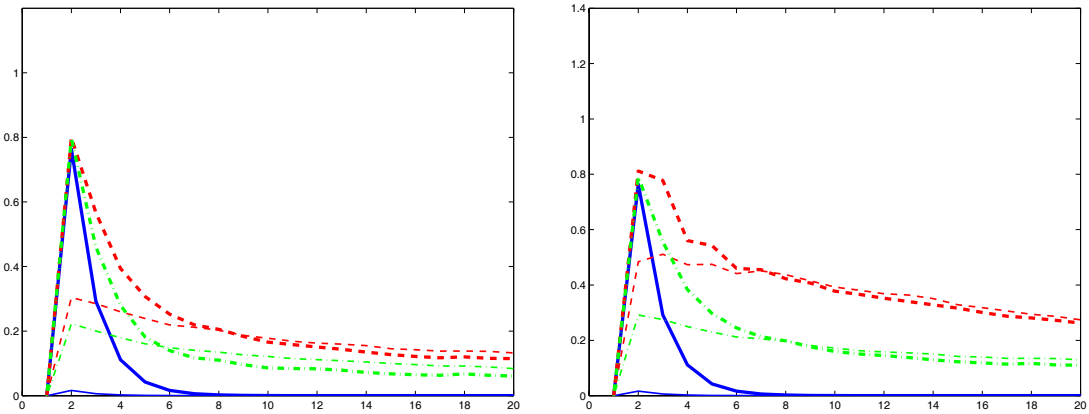

Biased toward Output control: $\theta_\pi = 0.4$ *or* $\theta_{\pi^e} = 0.8$

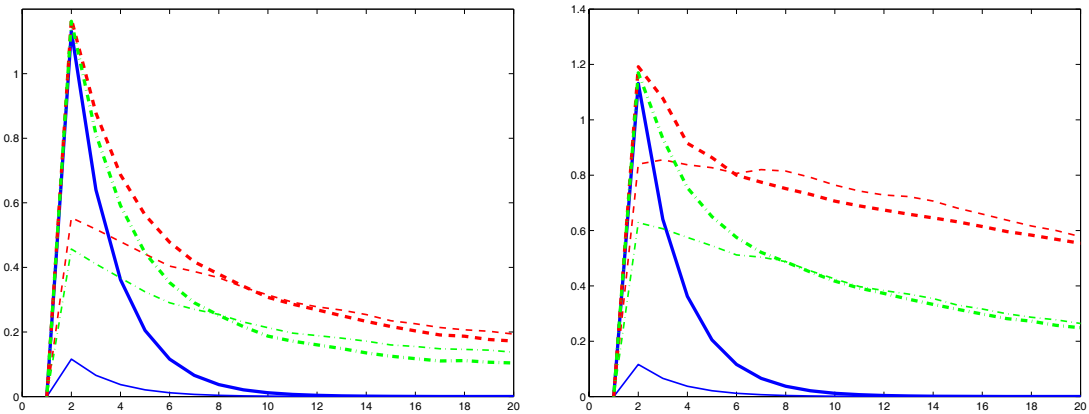

Figure 7

Evolution of Inflation Expectations Following a Series of Inflation Shocks

$(\phi = 0.75, \alpha = 0.25)$ $(\phi = 0.9, \alpha = 0.1)$

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

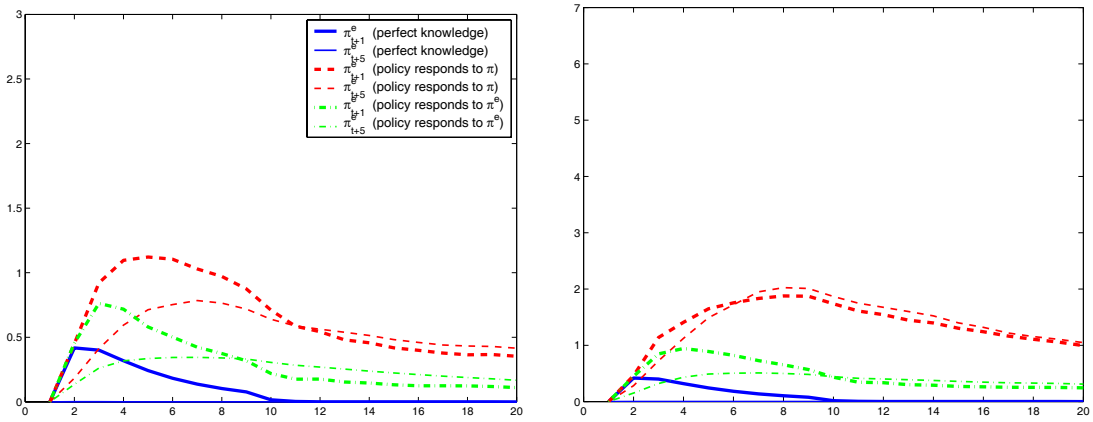

Balanced Preferences: $\theta_\pi = 0.6$ *or* $\theta_{\pi^e} = 1.6$

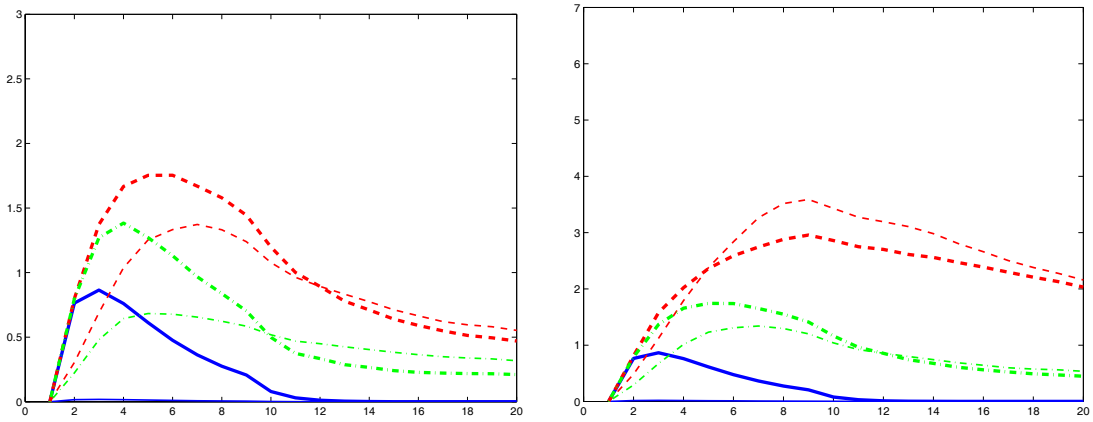

Biased toward Output control: $\theta_\pi = 0.4$ *or* $\theta_{\pi^e} = 0.8$

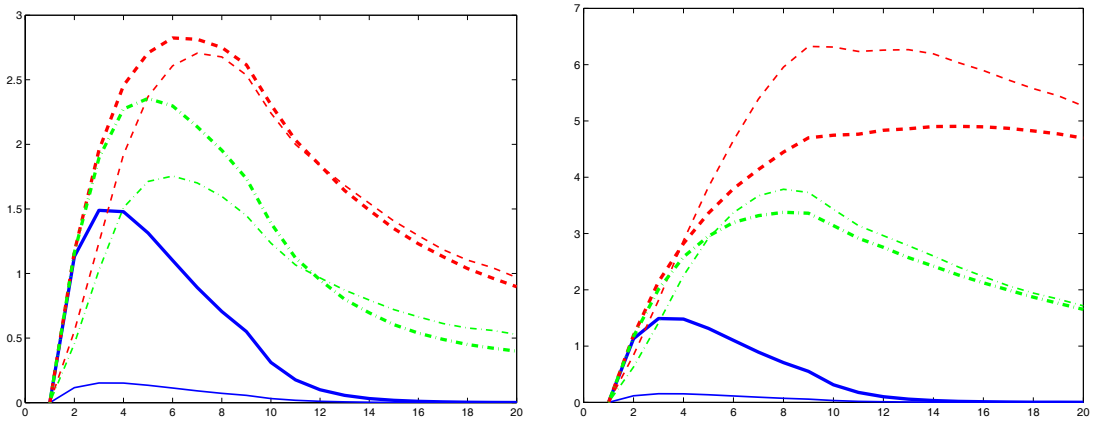

Figure 8

Performance of Optimized One-parameter Policy Rules

$\phi = 0.90,\ \alpha = 0.10$ $\phi = 0.75,\ \alpha = 0.25$

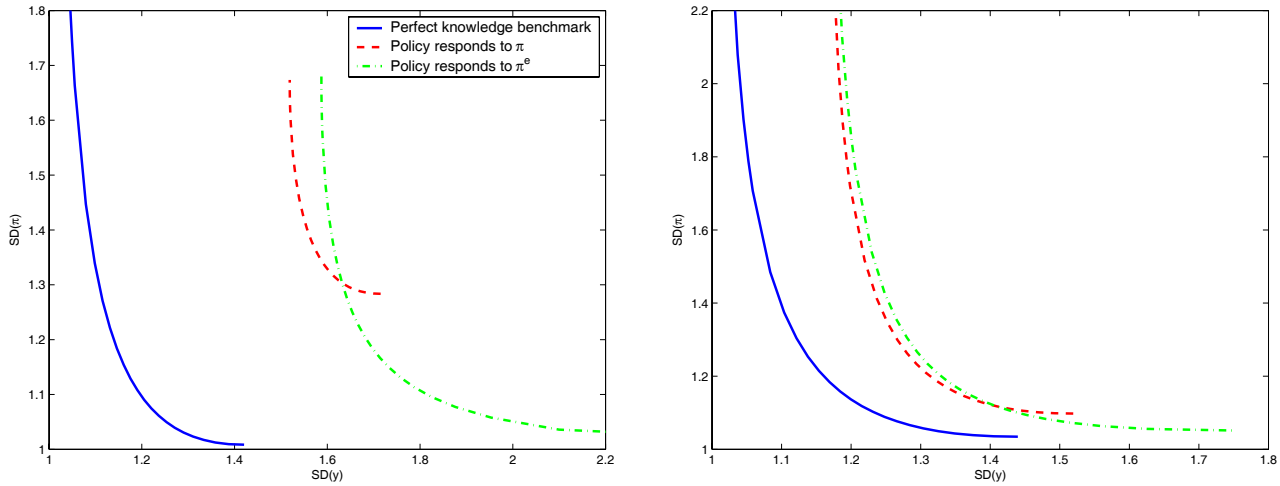

Optimized Response to Observed Inflation in One-parameter Rule

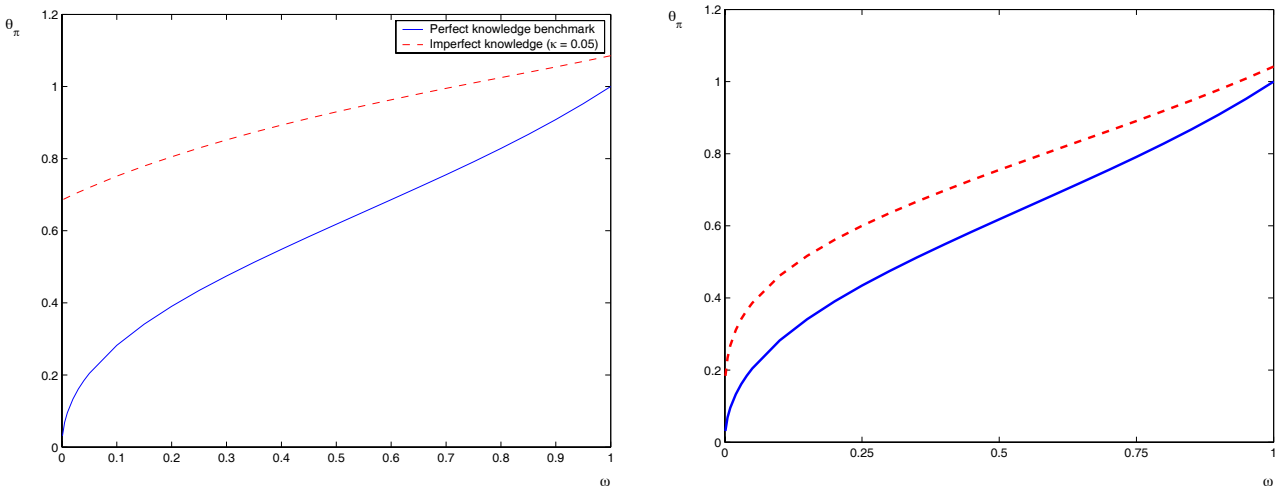

Optimized Response to Expected Inflation in One-parameter Rule

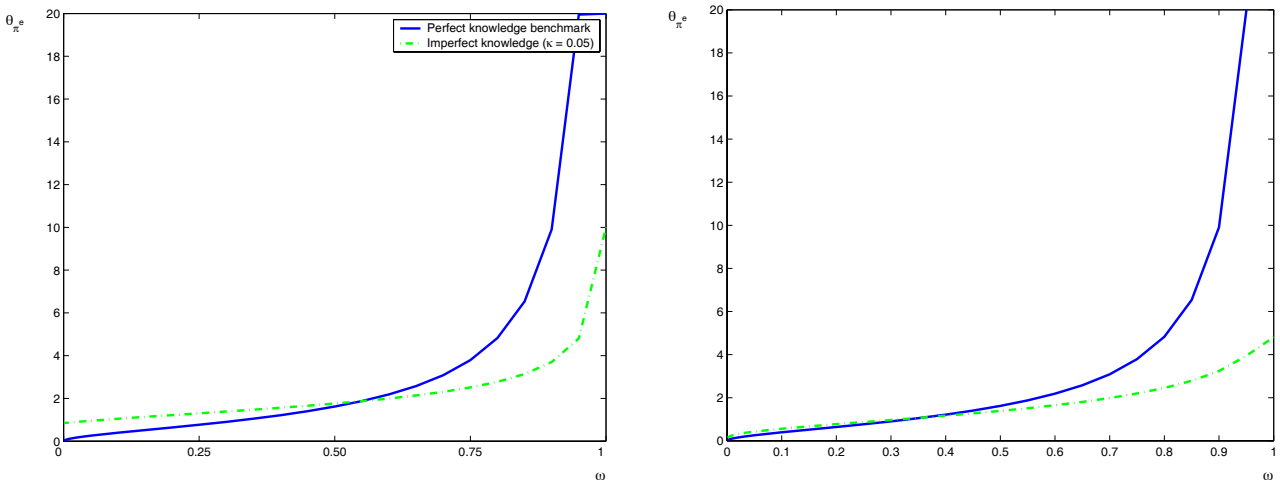

Figure 9

Performance of Optimized One- and Two-parameter Policy Rules

$\phi = 0.90, \alpha = 0.10$ $\phi = 0.75, \alpha = 0.25$

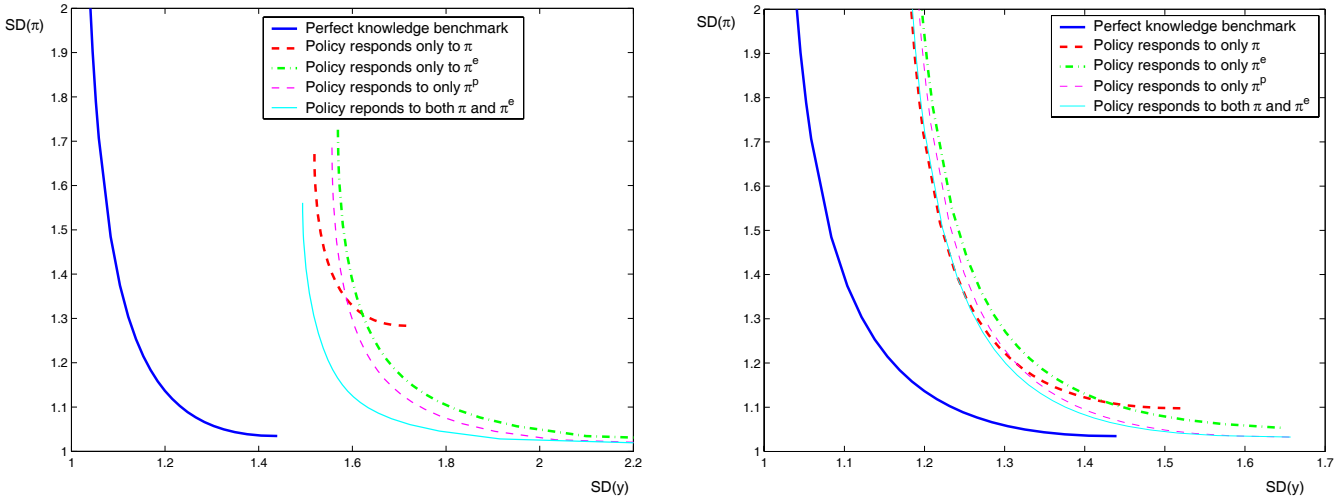

Optimized Coefficients of Two-Parameter Policy Rule