

# Are Inflation Expectations Rational?\*

David Andolfatto  
*Simon Fraser University*

Scott Hendry  
*Bank of Canada*

Kevin Moran  
*Université Laval*

January 4, 2005

## Abstract

Simple econometric tests reported in the literature consistently report what appears to be a bias in inflation expectations. These results are commonly interpreted as constituting evidence overturning the hypothesis of rational expectations. In this paper, we investigate the validity of such an interpretation.

The main tool utilized in our investigation is a computational dynamic general equilibrium model capable of generating aggregate behavior similar to the data along a number of dimensions. By construction, the model embedded the assumption of rational expectations. Standard regressions run on equilibrium realizations of inflation and inflation expectations nevertheless reveal an apparent bias in inflation expectations. In these simulations, the null hypothesis of rational expectations is incorrectly rejected in a large percentage of cases; a result that casts some doubt on conventional interpretations of the evidence.

## 1 Introduction

By various measures, inflation expectations appear to evolve sluggishly relative to actual inflation. In particular, expectations appear to underestimate inflation during periods of rising inflation, and overestimate inflation dur-

---

\*We thank Bob Amano, Paul Gomme, Ken Kasa, Andrew Levin, Césaire Meh, and seminar participants at the Bank of Canada, the 2002 CEA meeting, the 2002 North American Summer Meeting of the Econometric Society and the 2003 Annual Conference of De Nederlandsche Bank for useful comments and suggestions. Andolfatto acknowledges the financial support of SSHRCC. The views expressed in this paper are those of the authors. No responsibility for them should be attributed to the Bank of Canada.

ing periods of falling inflation.<sup>1</sup> Related to this phenomenon is a stylized fact documented by a sizable empirical literature that measured inflation expectations appear to be empirically inconsistent with unbiasedness and efficiency.<sup>2</sup> These results have sometimes been interpreted as evidence against the rational expectations hypothesis.

In this paper, we question whether conventional econometric tests that reject the unbiasedness of inflation expectations necessarily constitutes evidence that rejects the rationality of inflation expectations. Our skepticism in this matter was initially motivated by two facts. First, it seems that every econometric test of rationality has been based on a relatively small sample of data. Thus, it may be worthwhile to explore whether conventional tests of rationality applied to small samples have any reasonable power. Second, we know from the work of several authors that ‘sluggish’ (adaptive) expectation formation may in fact be optimal (rational) in environments that feature some type of ‘information friction;’ see Muth (1960), Brunner, *et. al.* (1980), Andolfatto and Gomme (2003), Erceg and Levin (2003), and Andolfatto, *et. al.* (2004). To the extent that information frictions are a feature of reality, there is good reason to believe that even rational inflation expectations may (in small samples) fail conventional tests of unbiasedness and efficiency.

Our methodological approach to this issue is to run the conventional econometric tests of unbiasedness on the artificial data generated by a conventional business cycle model that embeds the rational expectations hypothesis, and then report the frequency with which these tests reject the null hypothesis of rational expectations (unbiased expectations). The strategy followed here is similar to the one employed by Kozicki and Tinsley (2001), who argue that frequent empirical rejections of the expectation hypothesis of the term structure could be the result of individuals learning only gradually about shifts in the objectives of the Federal Reserve. In earlier contributions, Lewis (1988, 1989) uses a similar idea to examine whether gradual learning might generate the ‘forward discount’ puzzle observed in

---

<sup>1</sup>For example, Dotsey and DeVaro (1995) find that during the disinflation of 1980.1–1983.3, expected inflation exceeded actual inflation in all but three periods using eight-month-ahead forecasts, and in every period using one-year-ahead forecasts. Likewise, Delong (1997) reports that during the U.S. inflationary episode of the 1970s, a consensus private sector inflation forecast underestimated the actual inflation rate in every year and that, remarkably, in each and every year inflation was actually expected to *fall* (Figure 6.9, pg. 267).

<sup>2</sup>See, for example, Thomas (1999), Roberts (1997), Croushore (1997), Mehra (2002), and Mankiw, Reis, and Wolfers (2003).

foreign exchange data.

The benchmark model we employ is the limited participation model developed by Christiano and Gust (1999). We extend their environment to include a Taylor rule that is subject to relatively infrequent ‘regime shifts,’ which we model as occasional persistent shifts in the monetary authority’s underlying long-run inflation target. As is standard, we also assume that the interest rate is subject to transitory shocks, which are interpreted as instances when the monetary authority wishes to deviate from its rule for exogenous reasons (e.g., to react to the financial market developments). In the spirit of Andolfatto and Gomme (2003), we assume that our model agents cannot separately observe these two shocks. Instead, market participants must make rational inferences based on the limited information at their disposal. This signal extraction problem gives rise to a learning rule that shares some features with adaptive expectations processes.

Our prior for this exercise was that conventional econometric tests of rationality would incorrectly reject the null hypothesis of rationality on a frequent basis and that the reason for this high rejection rate would lie in the nature of the information friction concerning the conduct of monetary policy. As it turns out, we were only partly right.

As expected, we find that standard econometric tests of unbiasedness using simulated data frequently and incorrectly reject the null hypothesis. Furthermore, the frequency of rejection falls steadily as the sample size of each simulation is increased. These results confirm our conjecture that conventional econometric tests of rationality are not very powerful in small samples.

When we repeated the experiment in the context of an environment with no signal extraction problem, the number of rejections fell significantly. But surprisingly, the rejection rate did not fall dramatically as expected. Evidently, it appears that the high rate of rejection may be attributable to several factors, one of which may be the information friction alluded to above. But an even more important factor may simply be the relatively short sample periods utilized by econometricians, in conjunction with the relatively persistent behavior of inflation and inflation expectations. Extending the sample period of simulations and altering the monetary policy rule to induce less persistence generated far fewer rejections of the null in our experiments.

Our paper is organized as follows. Section 2 presents a summary of the type of evidence that is used to show that measured inflation expectations

fail simple unbiasedness tests. Section 3 develops the model used in our simulations. Section 4 describes the model’s calibration and Section 5 provides a brief evaluation of the model’s dynamics. Section 6 presents our Monte Carlo simulations along with our interpretations. Section 7 offers a brief summary and conclusion.

## 2 Empirical Evidence on Inflation Expectations

Survey data represent one of the tools commonly used to identify economic agents’ inflation expectations.<sup>3</sup> Figure 1 illustrates a typical path for such data. It depicts the (mean) forecast for one-year-ahead inflation, as measured by the Livingston Survey, as well as the inflation rate that eventually prevailed.<sup>4</sup> The ‘sluggishness’ of inflation expectations alluded to above is clearly visible in the figure. That is, note how inflation expectations appear to lag (and underestimate) the actual inflation rate in the early part of the sample, which is characterized by generally rising inflation. Likewise, inflation expectations appear to lag (and overestimate) the actual inflation rate over the latter part of the sample, which is characterized by generally falling inflation.

Several studies (e.g., Mankiw *et. al.* 2003, Mehra 2002, Thomas 1999, Roberts 1997, and Croushore 1997) assess whether economic agents make systematic errors when forecasting inflation, i.e. whether measured inflation expectations are rational. These studies commonly emphasize one key characteristic of rationality; namely, if inflation expectations are rational, they should be unbiased.<sup>5</sup> Testing for unbiasedness is typically implemented by running the following regression:

---

<sup>3</sup>Other methods make use of futures market data (Dotsey and DeVaro, 1995) or the yield differences between non-indexed and indexed Treasury bills (Shen and Corning, 2001).


<sup>4</sup>The Livingston survey is maintained by the Federal Reserve Bank of Philadelphia; its history and current structure are described in Croushore (1997). Other survey data on U.S. inflation expectations include those in the survey of households conducted by the Institute for Social Research at the University of Michigan and the more recently established Survey of Professional Forecasters, also maintained by the Federal Reserve Bank of Philadelphia.

<sup>5</sup>Some studies also attempt to verify whether inflation expectations are efficient by testing whether any information publicly available at the time the forecasts were made could reduce expectational errors.

$$\pi_t = a_0 + a_1 E_{t-k} [\pi_t] + e_t; \tag{1}$$

where  $\pi_t$  is the net annualized rate of inflation and  $E_{t-k} [\pi_t]$  is a measure of the inflation rate expected for date  $t$ , with the expectation formed at date  $t - k$ . Under the null hypothesis of rational expectations,  $H_0 : a_0 = 0$  and  $a_1 = 1$ .

FIGURE 1  
Inflation and Expected Inflation  
(Livingston Survey)



Actual inflation: Y/Y increase in all-items CPI. Source: Federal Reserve Bank of St. Louis. Expected inflation: One-Year Ahead (Mean) Estimate from the Livingston Survey. Source: Federal Reserve Bank of Philadelphia.

Typically (though not always), the hypothesis of unbiasedness tends to be rejected, particularly in small samples; in larger samples, the hypothesis is

rejected less often. For example, Table 1 reports Thomas' (1999) regression results using the Livingston data. These results have been updated and confirmed by Mehra (2002). Table 1 shows that unbiasedness is rejected decisively in various subsamples, but not rejected in the full sample.

TABLE 1				
Testing for Unbiasedness				
	$\hat{a}_0$	$\hat{a}_1$	$\chi^2(2)$	Significance
1960.1–1980.2	0.62	1.29	18.11	0.000
1980.3–1997.4	0.13	0.88	9.37	0.009
1960.1–1997.4	0.13	1.12	2.30	0.317

This pattern of frequent rejections of the null in small samples but less frequent rejections in larger samples is consistent with the thesis of the present paper: incomplete information and learning about a few significant shifts in monetary policy (among other factors) can lead the expectations of rational agents to appear biased.<sup>6</sup>

### 3 Model

This section presents the model used to perform our experiments. We describe the optimization problem of households, firms, and financial intermediaries, as well as the policy rule followed by the monetary authority. In addition, we provide a detailed description of the shifts that affect this rule and the learning of private agents about these shifts. The key assumption is that changes in the inflation target of monetary authorities cannot be observed directly by private agents and must instead be inferred using observed inflation outcomes.

The model belongs to the class of limited participation models, as introduced by Lucas (1990). This class of model assumes that households are relatively slow to adjust their nominal savings in response to monetary policy shocks. As a result, monetary injections affect the loanable funds of banks and thus the supply of credit to firms, with consequent effects on economic activity. Christiano and Gust (1999) and Christiano, Eichenbaum and Evans (1997, 1998) demonstrate that this class of models can replicate

---

<sup>6</sup>This possibility is also discussed in Thomas (1999) and Mehra (2002).

the main dynamic effects of monetary policy shocks. In particular, even though prices are not sticky by assumption, smooth movements in prices can arise as an equilibrium phenomenon.

### 3.1 Household Sector

The economy is populated by a continuum of identical, infinitely-lived households with preferences given by:

$$E_0 \sum_{t=0}^{\infty} \beta^t U(c_t, n_t + h_t), \quad (2)$$

where  $c_t$  denotes consumption,  $n_t$  denotes employment,  $h_t$  represents the time-cost associated with portfolio adjustment, and  $0 < \beta < 1$  denotes a subjective discount factor.

At the beginning of period  $t$ , a household possesses  $M_t$  dollars which is divided between (liquid) cash  $M_t^c$  and (illiquid) interest-bearing deposits  $M_t^d$ ; i.e.,  $M_t = M_t^c + M_t^d$ . The allocation of money across cash and deposits is ‘sticky’ in the sense of being fixed during period  $t$  (i.e., having been determined during period  $t - 1$ ). At the beginning of the period, the household earns  $W_t n_t$  dollars in the labor market and so has  $M_t^c + W_t n_t$  dollars available to make purchases of consumption  $c_t$  and new capital goods  $x_t$ ; i.e., assume the following cash-in-advance constraint:

$$P_t (c_t + x_t) \leq M_t^c + W_t n_t, \quad (3)$$

where  $P_t$  denotes the period  $t$  price level.

Households own the economy’s capital stock  $k_t$ , which they rent to firms at the real rental rate  $r_t$ . Assume that capital depreciates at rate  $0 < \delta < 1$ , so that  $k_{t+1} = (1 - \delta)k_t + x_t$ . Let  $R_t$  denote the gross nominal interest rate earned on cash deposits  $M_t^d$  and let  $D_t$  denote the nominal dividend income that a household earns through its ownership of the financial sector. At the end of the period, households receive their capital rental and dividend income, and their return on deposits. These revenues, combined with any liquid assets remaining subsequent to their goods purchases sum up to their end-of-period financial wealth, which is then divided between next period’s liquid and illiquid assets; i.e.,

$$M_{t+1}^c + M_{t+1}^d = P_t r_t k_t + R_t M_t^d + D_t + [M_t^c + W_t n_t - P_t c_t - P_t x_t]. \quad (4)$$

Observe that the term in the square brackets above will equal zero whenever the cash-in-advance constraint (3) binds.

Following Christiano *et. al.* (1998), Christiano and Gust (1999) and Aiyagari and Braun (1998), we assume that adjusting the household wealth portfolio is costly, but not infinitely so. The purpose of this assumption is to generate persistence in the liquidity effect following a shock to monetary policy. To this end, we assume that:

$$h_t = H\left(\frac{M_{t+1}^c}{M_t^c}\right), \quad (5)$$

where  $H(\cdot)$  is an increasing and convex function.

Thus, the household's choice problem involves choosing a stochastic process  $\{c_t, n_t, h_t, x_t, k_{t+1}, M_{t+1}^c, M_{t+1}^d\}_{t=0}^{\infty}$  to maximize expected utility (2) subject to (3), (4), (5) and initial conditions  $k_0, M_0^c, M_0^d$ . Note that the expectations operator in (2) represents an expectation that is based on all available information to the household. As we will discuss in further detail below, the specification of  $E_0$  will vary depending on the information structure we assume.

### 3.2 Business Sector

The business sector consists of goods-producing firms and financial intermediaries. Firms produce output  $y_t$  using capital  $k_t$  and labor  $n_t$  according to a standard neoclassical production function  $y_t = a_t F(k_t, n_t)$ . The technology parameter  $a_t$  evolves according to:

$$\ln a_t = (1 - \rho) \ln a + \rho \ln a_{t-1} + v_t,$$

with  $v_t \sim N(0, \sigma_v^2)$  and  $0 < \rho < 1$ . Assume that  $v_t$  is realized at the beginning of the period.

Firms rent capital and labor from households. Capital is paid in the usual way; i.e., with claims to the output produced within the period. However, following the literature, we assume that households demand to be paid up-front in government cash. Because firms begin each period with no cash, they must borrow the necessary funds from financial intermediaries. Assuming that firms borrow just enough cash to finance the period wage bill, the amount borrowed will equal  $W_t n_t$  dollars. If the loan is payable at the end of the period, then the total owing at that time is equal to  $R_t W_t n_t$ , where

$R_t$  denotes the gross nominal interest rate on cash loans. In this case, the firm's period profit function is given by:

$$P_t a_t F(k_t, n_t) - P_t k_t - R_t W_t n_t. \quad (6)$$

Note that because workers must be paid in borrowed cash, the nominal interest rate here serves as a type of tax on production. In this way, an unexpected increase in the nominal interest rate has the potential to depress economic activity.

Financial intermediaries accept money deposits  $M_t^d$  from households and extend money loans  $W_t n_t$  to firms. If intermediation is costless, then competition dictates that the interest rate charged on loans and earned on deposits will be equated to some common value  $R_t$ . We assume that  $R_t$  is determined in a competitive market for money loans, but that this interest rate can be influenced (and if desired, entirely determined) by the central bank through injections (or withdrawals) of new cash  $X_t$  accruing to intermediaries.

Thus, intermediaries generate end-of-period revenues equal to  $R_t(M_t^d + X_t)$  and end-of-period expenses (deposit repayments) equal to  $R_t M_t^d$ , for a net profit equal to  $D_t = R_t X_t$ . These profits are remitted to households at the end of the period as a dividend payment.

### 3.3 Monetary Policy

The monetary authority sets the net nominal interest rate  $i_t \equiv R_t - 1$  according to the following Taylor rule:

$$i_t = (1 - \theta) [r^* + \pi_t^* + \lambda(\pi_t - \pi_t^*) + \gamma(y_t - y^*)] + \theta i_{t-1} + u_t, \quad (7)$$

where  $r^*$  denotes the steady-state real (net) interest rate,  $\pi_t^*$  denotes an inflation target,  $y^*$  denotes the steady-state level of GDP, and  $u_t$  denotes an exogenous transitory monetary policy shock. The parameter  $0 \leq \theta < 1$  indexes the degree to which the monetary authority wishes to smooth out interest rate movements. As mentioned above, the monetary authority achieves any particular  $i_t$  with an appropriate lump-sum injection/withdrawal of cash  $X_t$  to/from intermediaries.

We assume that the monetary policy rule is subject to two types of disturbances. The first disturbance consists of the monetary policy shocks  $u_t$  defined above. We interpret these shocks as the reaction of monetary authorities to economic factors, such as financial stability concerns, not articulated by the rule (7). Alternatively, the shocks could be understood

as stemming from the imperfect control exercised by monetary authorities over the growth rate of money supply. Under either interpretation, we view these shocks as possessing little persistence. Accordingly, we assume that their evolution is governed by the following process:

$$u_t = \phi_1 u_{t-1} + e_t, \tag{8}$$

with  $0 \leq |\phi_1| \ll 1$  and  $e_t \sim N(0, \sigma_e^2)$ .

The second disturbance to monetary policy is as follows. We assume that the inflation target of the monetary authorities  $\pi_t^*$ , while remaining constant for extended periods of time, is nevertheless subject to occasional, persistent shifts. We see two possible interpretations for these shifts. First, they could correspond to changes in economic thinking that lead monetary authorities to modify their views about the proper rate of inflation to pursue. Delong (1997), for example, argues that the Great Inflation of the 1970s, and its eventual termination by the Federal Reserve at the beginning of the 1980s, was a result of shifting views about the shape of the Phillips curve and, more generally, about the nature of the constraints under which monetary policy is conducted. Alternatively, a change in the inflation target could reflect the appointment of a new central bank chair, whose preferences over inflation outcomes differ from their predecessor's. Under either interpretation, we view these shifts as exhibiting significant duration, in the order of five to ten years.

We express these shifts with the variable  $z_t \equiv \pi_t^* - \pi^*$ , so that  $z_t$  constitutes the deviation of the current target of authorities  $\pi_t^*$  from its long term (time-invariant) mean  $\pi^*$ . We assume that the following process, a mixture of a Bernoulli trial and a normal random variable, expresses how  $z_t$  evolves over time:

$$z_t = \begin{cases} z_{t-1} & \text{w.p. } \phi_2; \\ s_t & \text{w.p. } 1 - \phi_2, \quad s_t \sim N(0, \sigma_s^2); \end{cases} \tag{9}$$

with  $0 \ll \phi_2 < 1$ .

Notice that the process for  $z_t$  shares some similarities with a random walk. Specifically, the conditional expectation of  $z_t$  is close to  $z_{t-1}$  when  $\phi_2$  is close to one. In contrast with a random walk however, the process is not affected by innovations every period and is stationary. The process also differs from a standard autoregressive process in that the decay of a given impulse will be sudden and complete, rather than smooth and gradual. We

believe that this characterization of the regime shifts accords well with recent episodes of monetary history and with our suggested interpretations of these shifts.

### 3.4 Information Structure and Learning

In our experiments below, we consider two information structures, distinguished by whether agents can directly observe the inflation target  $\pi_t^*$  or not. We refer to the case in which individuals do observe  $\pi_t^*$  as *complete information*; and the case in which they do not as *incomplete information*. In both cases, agents are assumed to understand the structure of monetary policy in the sense they know the underlying parameters  $(\theta, \lambda, \gamma, \phi_1, \phi_2, \sigma_e, \sigma_s)$  governing the policy rule (7).

The assumption of complete information supposes that the monetary authority is willing and able to communicate changes in its inflation target. However, credibly communicating such shifts might be difficult for several reasons. For example, although a new central bank head may make a strong aversion for inflation known in public announcements, the lack of precision of these announcements may leave private agents uncertain as what they imply quantitatively for the inflation target. Agents might, as a result, modify their beliefs about the inflation target of monetary authorities only once several periods of lower inflation have been observed. More explicit announcements of changes in the inflation target might also suffer, at least initially, from similar credibility problems.

When agents have incomplete information, they only observe a mixture of the two shocks  $z_t$  and  $u_t$ , so that  $\pi_t^*$  is not directly observable. To illustrate the nature of the signal extraction problem in this case, consider the following sequence of events. Suppose that initially, the inflation target is set to its long-term mean; i.e.,  $\pi_0^* = \pi^*$ . At any date  $t > 0$  then, the interest rate is given by:

$$i_t = (1 - \theta) [r^* + \pi^* + \lambda(\pi_t - \pi^*) + \gamma(y_t - y^*)] + \theta i_{t-1} + u_t. \quad (10)$$

Now, suppose that at some time  $t > 0$ , the monetary authority changes its inflation target to some other value  $\pi_t^* \neq \pi^*$ , so that the interest rate now satisfies:

$$i_t = (1 - \theta) [r^* + \pi_t^* + \lambda(\pi_t - \pi_t^*) + \gamma(y_t - y^*)] + \theta i_{t-1} + u_t.$$

This latter equation can be expressed as:

$$i_t = (1 - \theta) [r^* + \pi^* + \lambda(\pi_t - \pi^*) + \gamma(y_t - y^*)] + \theta i_{t-1} + \epsilon_t, \quad (11)$$

where  $\epsilon_t \equiv (1 - \theta)(1 - \lambda)z_t + u_t$  and  $z_t \equiv (\pi_t^* - \pi^*)$ .

Comparing equations (10) and (11) shows that from the perspective of an agent whose initial belief about the inflation target was  $\pi^*$ , the observed shock to the policy rule  $\epsilon_t$  is a combination of the regime shift  $(1 - \theta)(1 - \lambda)z_t$  and the transitory shock  $u_t$ . The signal extraction that agents face therefore entails separating  $\epsilon_t$  into its persistent and transitory components. Having solved this problem, agents can then back out an estimate of the shift in the inflation target  $z_t = (\pi_t^* - \pi^*)$ .

In the case of incomplete information, we assume that agents use a Kalman filter to solve the signal extraction problem. Note that the disturbances to monetary policy  $z_t$  and  $u_t$ , along with the observed shock to the monetary policy rule  $\epsilon_t$ , can be expressed as follows:

$$\begin{aligned} \begin{bmatrix} z_{t+1} \\ u_{t+1} \end{bmatrix} &= \begin{bmatrix} \phi_2 & 0 \\ 0 & \phi_1 \end{bmatrix} \begin{bmatrix} z_t \\ u_t \end{bmatrix} + \begin{bmatrix} \kappa_{t+1} \\ e_{t+1} \end{bmatrix}; \\ \epsilon_t &= \begin{bmatrix} (1 - \theta)(1 - \lambda) & 1 \end{bmatrix} \begin{bmatrix} z_t \\ u_t \end{bmatrix}; \end{aligned} \quad (12)$$

where

$$\kappa_{t+1} = \begin{cases} (1 - \phi_2) & \text{w.p. } \phi_2; \\ s_{t+1} - \phi_2 z_t & \text{w.p. } 1 - \phi_2. \end{cases}$$

Observe that under this definition,  $E_t \kappa_{t+1} = 0$ .

The system defined in (12) along with the specification of  $e_t$  and  $\kappa_t$  define a state-space system (Hamilton, 1994, Chapter 13), where the first equation in (12) is the state equation and the second equation in (12) is the observation equation. When applied to such a system, the Kalman filter delivers estimates of the two unobserved states ( $z_t$  and  $u_t$ ) conditional on all observed values of  $\epsilon_t$  up to and including period  $t$ . These estimates are updated sequentially as new information becomes available according to:

$$\begin{bmatrix} E_t z_t \\ E_t u_t \end{bmatrix} = \begin{bmatrix} E_{t-1} z_t \\ E_{t-1} u_t \end{bmatrix} + K_t (\epsilon_t - E_{t-1} \epsilon_t),$$

where  $K_t$  represents the gain matrix of the filter and governs the extent to which unexpected movements in the observed shock  $\epsilon_t$  are attributed to the persistent shift  $z_t$  or to the transitory shock  $u_t$ . The system in (12) can then be used to compute expected future shocks:

$$\begin{aligned} \begin{bmatrix} E_t z_{t+1} \\ E_t u_{t+1} \end{bmatrix} &= \begin{bmatrix} \phi_2 & 0 \\ 0 & \phi_1 \end{bmatrix} \begin{bmatrix} E_t z_t \\ E_t u_t \end{bmatrix}; \\ E_t \epsilon_{t+1} &= \begin{bmatrix} (1-\theta)(1-\lambda) & 1 \end{bmatrix} \begin{bmatrix} E_t z_{t+1} \\ E_t u_{t+1} \end{bmatrix}. \end{aligned}$$

The projections arising from the Kalman filter represent the best linear forecasts of the unobserved variables based on available information. However, since  $z_t$  is not normally distributed (conditional on information up to date  $t$ ) but is instead a mixed Bernoulli-normal variate, the Kalman filter is not optimal.<sup>7</sup> This implies that agents using this filter are not strictly rational and could potentially improve on their forecasts of future monetary policy disturbances by using another (significantly more complicated) non-linear filter. We performed a number of tests to check whether this source of ‘non-rationality’ was likely to be quantitatively important in the experiments performed below. It turns out that in repeated simulations of the model, the instances where agents could have actually improved their forecasts in a statistically significant manner are not frequent.<sup>8</sup>

### 3.5 Equilibrium

An equilibrium for this economy consists of an allocation  $\{c_t, n_t, h_t, x_t, k_{t+1}, M_{t+1}^c, M_{t+1}^d\}_{t=0}^\infty$ , dividend payments  $\{D_t\}_{t=0}^\infty$ , a price system  $\{P_t, W_t, R_t, r_t\}_{t=0}^\infty$  and monetary injections  $\{X_t\}_{t=0}^\infty$  such that, given the stochastic process  $\{a_t, z_t, u_t\}_{t=0}^\infty$  and initial condition  $(k_0, M_0^c, M_0^d)$ :

1. The allocation maximizes (2) subject to (3), (4), (5);
2. The allocation maximizes (6) and  $D_t = R_t X_t$ ;
3. Monetary injections  $\{X_t\}_{t=0}^\infty$  are consistent with (7);
4. Markets clear at each date; i.e.,  $c_t + k_{t+1} = a_t F(k_t, n_t) + (1-\delta)k_t$  and  $M_t^d + X_t = W_t n_t$ ; and

---

<sup>7</sup>When one considers that the only source of variation in  $z_t$  arises from a normal variable, it must be that in an unconditional sense,  $z_t$  is distributed normally. Considering however the high values of  $\phi_2$  used in our calibration, this normal behavior will only appear after a large number of data have been observed.

<sup>8</sup>As our main results do not depend critically on the learning mechanism, for the sake of brevity, we omit describing our investigation of this issue. For the interested reader, details of our tests are available on request.

5. Expectations are consistent with the stochastic process generated by  $\{a_t, z_t, u_t\}_{t=0}^{\infty}$  and optimal behavior, given available information.

## 4 Calibration

We solve the model by linearly approximating policy and pricing functions around the nonstochastic steady state; see King and Watson (1998). Nominal variables are transformed in the usual way to render them stationary.

### 4.1 Preferences and Technology

The model is calibrated to quarterly data and assign standard values for the discount factor  $\beta = 0.99$  and depreciation rate  $\delta = 0.025$ . The production function is Cobb-Douglas  $F(k, n) = k^\alpha n^{1-\alpha}$ , with  $\alpha = 0.36$ . The scale parameter for the technology shock is set to  $a = 1$  and the persistence parameter is set to  $\rho = 0.95$ . The standard deviation of the innovations to technology is set to  $\sigma_v = 0.0035$ , so that the variance of simulated GDP corresponds roughly with observation.

Following Christiano and Gust, preferences are parameterized according to:

$$U(c, n + h) = \log \left[ c_t - \psi_0 \frac{(n + h)^{1+\psi_1}}{1 + \psi_1} \right].$$

As is well-known, this specification of preferences implies that time allocation is independent of wealth; see also Greenwood, *et. al.* (1988). We set  $\psi_1 = 0.4$ , which implies a (steady-state) wage elasticity of labor supply of 2.5; the value used by Christiano and Gust (1999). The scale parameter  $\psi_0$  is set to 2.15, so that employment is normalized to unity in the steady-state.

Following Aiyagari and Braun (1998), the portfolio adjustment cost function is assumed to be quadratic:

$$H(M_{t+1}^c/M_t^c) = \frac{\tau}{2} \left( \frac{M_{t+1}^c}{M_t^c} - \pi^* \right)^2.$$

In a steady-state, the cash balances grow at the rate of inflation so that  $h^* = 0$ . We set the adjustment cost parameter to  $\tau = 15$ , which corresponds roughly to the specification in Christiano *et. al.* (1998) and Christiano and Gust (1999). In an economic sense, adjustment costs turn out to be small in the sense that in our simulations, they never exceed 0.1% of steady-state GDP or hours.

## 4.2 Monetary Policy

The long-run inflation rate is set to  $\pi^* = 0.0123$ , which implies an annual inflation rate of 5% in the steady-state. According to the monetary policy rule (7), the nominal interest rate reacts to deviations of inflation from its current target ( $\lambda$ ), to deviations of output from its trend ( $\gamma$ ), and to its own lagged value ( $\theta$ ). Calibrating these values turns out to be not so straightforward since the sizable literature on this subject reports estimates of Taylor rules using a wide variety of specifications and methods.<sup>9</sup> Furthermore, some parameter combinations lead to non-uniqueness (or nonexistence) of stable equilibria; see Christiano and Gust (1999). Nevertheless, the literature can be used to restrict the parameter space in a plausible way.

Beginning with Taylor (1993), it has been argued that the interest rate must react aggressively to inflation to avoid the possibility of ‘expectation traps’ in which high inflation becomes a self-fulfilling prophesy. Accordingly, we set  $\lambda = 2$  in our benchmark calibration. This value is only slightly larger than the value used by Taylor and is in line with recent empirical estimates; see Erceg and Levin (2003), English *et. al.* (2003), and Schorfheide (2003). We analyze the sensitivity of our results for  $\lambda \in [2, 4]$ .

Similarly, we set the benchmark smoothing parameter to  $\theta = 0.50$ , which is in the range of recent empirical estimates; again, see Erceg and Levin (2003), Kozicki and Tinsley (2003), and Schorfheide (2003). Note that these empirical exercises allow for the possibility of within-sample regime shifts in monetary policy and so are consistent with our environment.

The parameter  $\gamma$  governs the response of the interest rate to changes in the output gap. For our benchmark calibration, we set  $\gamma = 0.25$ , which is lower than most other empirical estimates, but consistent with the value obtained in Erceg and Levin (2003). We explore the sensitivity of our results for  $\gamma \in [0, 0.25]$ . Higher values of  $\gamma$  are avoided since they increase the likelihood of multiple equilibria; see Christiano and Gust (1999).

The parameters  $\phi_2$  and  $\sigma_s$  govern the dynamics of the inflation target  $z_t$ . The parameter  $\phi_2$  governs the mean duration of any given regime. In our benchmark calibration, we set  $\phi_2 = 0.975$ , which implies a mean duration of about 10 years. This value is consistent with our earlier interpretation

---

<sup>9</sup>Empirical estimates of Taylor rules differ on whether the estimated rule is forward looking (Clarida *et. al.*, 2000) or is based solely on contemporaneous variables (Taylor, 1993); and whether parameters are obtained by single-equation estimation (Erceg and Levin 2003, English, *et. al.* 2003), or as part of a system-wide estimation (Schorfheide 2003, Kozicki and Tinsley, 2003).

of regimes corresponding to particular central bank heads or the life of a particular economic paradigm that dominates FOMC committee decisions. It is also in line with the estimates reported in Erceg and Levin (2003) and Schorfheide (2003).

The parameter  $\sigma_s$  represents the standard deviation of the distribution from which a regime shift is drawn when it occurs. We set  $\sigma_s = 0.01$ , which implies that a one standard deviation policy shift changes the inflation target by 4% on an annual basis. Again, this value is consistent with the empirical results reported in Erceg and Levin (2003), Kozicki and Tinsley (2003), and Schorfheide (2003).

The parameters governing the transitory monetary policy shock are set to  $(\phi_1, \sigma_e) = (0.10, 0.005)$ . The low persistence of these shocks is consistent with our interpretation that they represent temporary deviations from the policy rule in place. The standard deviation is at the high end of estimates reported elsewhere in the literature. Sensitivity analysis reveals that our main conclusions are robust to different values for these parameters within a reasonable range.

## 5 Model Evaluation

In this section, we provide a brief evaluation of the model by reporting a set of conventional (RBC) second moments and some impulse response functions. Table 1 reports three statistics for a set of variables  $x$ . The first column records the percent standard deviation of variable  $x$ ; the second column records the standard deviation of variable  $x$  relative to GDP; and the third column records the first-order autocorrelation of variable  $x$ . These measurements are based on HP-filtered data simulated from the model under its benchmark calibration.

**Table 1****Benchmark Model***Selected Variables and Moments*

Variable ( $x$ )	$\sigma(x)$	$\sigma(x)/\sigma(y)$	$\rho(x)$
GDP	1.68	1.00	0.63
Consumption	1.30	0.78	0.57
Investment	2.98	1.77	0.69
Employment	1.59	0.95	0.58
Real Wage	0.63	0.38	0.58
Inflation	0.61	0.37	0.14
Expected Inflation	0.28	0.17	0.66
Technology	0.70	0.42	0.69

The main thing to take away from Table 1 is that the model's dynamic properties are more or less consistent with the data and standard monetary business cycle models; e.g., see Cooley and Hansen (1995). As with most monetary models, monetary shocks contribute very little to the volatility of real variables. Note that the model appears to embed a quantitatively important propagation mechanism in the sense that the volatility of the technology shock is significantly less than the volatility in GDP. As well, note that expected inflation is considerably more stable and persistent than actual inflation.

Figure 2 records the economy's response to a standard deviation decline in the inflation target. This type of regime change initially generates a substantial and persistent increase in the interest rate, which is consistent with what we know of such episodes (Erceg and Levin 2003). In our model, such a shock also generates a mild recession. Note that inflation falls gradually, with inflation expectations consistently overestimating actual inflation for several periods. In the long-run, however, such a policy change results in lower interest rates and higher GDP (recall that the interest rate here acts like a tax on economic activity).

FIGURE 2  
 Dynamics Following a Regime Change  
 (Lower Inflation Target)

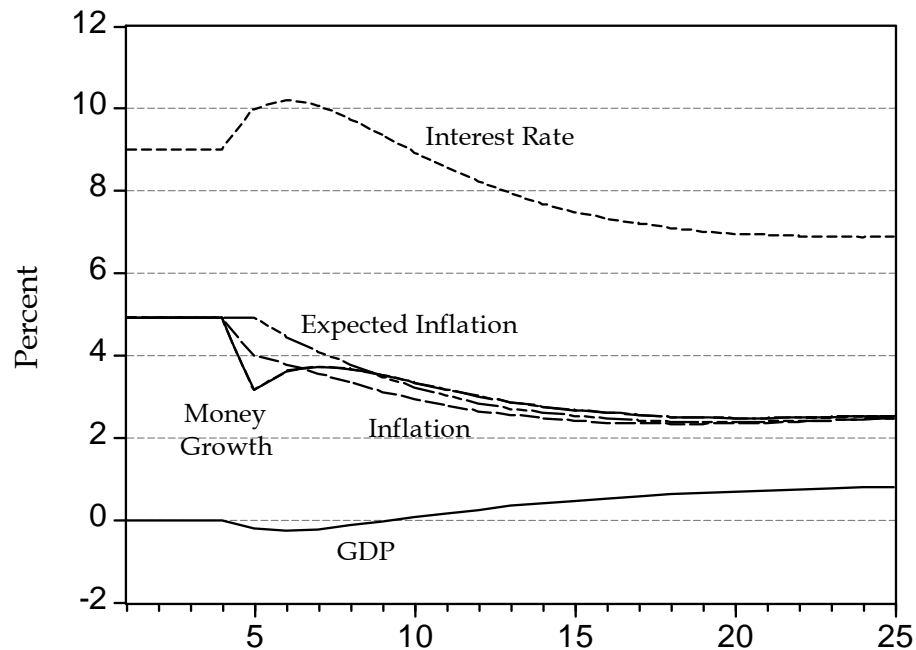
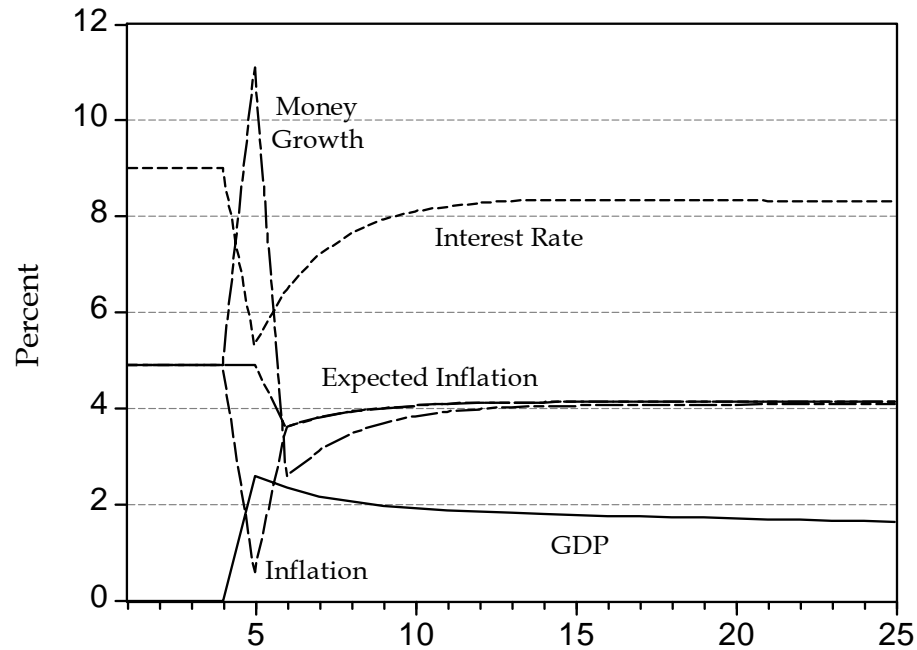


Figure 3 records the economy's response to a standard deviation increase in the technology shock. As is standard in these environments, the sudden surge in output is disinflationary. While the monetary policy rule calls for a higher interest rate as output rises above normal levels, the drop in inflation initially requires a significant infusion of liquidity (i.e., a decline in the interest rate) to stabilize inflation around its target level. Evidently, this latter effect dominates the former. Note that expected inflation deviates from actual inflation for only one period following the shock. The subsequent transition dynamics are drawn out over a hundred periods.

FIGURE 3  
Dynamics Following a Positive  
Technology Shock



Our basic conclusion is that the model's dynamics are not unreasonable. To the extent that this is true, the model can serve as a useful laboratory for the experiment that follows.

## 6 Testing the Unbiasedness Hypothesis

### 6.1 The Experiment

We treat our benchmark model as the data generating process (DGP) and ask what an econometrician would conclude about the unbiasedness of inflation expectations if presented with data from our DGP.

To this end, we simulate the model 1000 times, generating realizations for inflation and expected inflation 80 periods (20 years) in length for each simulation. Because measured inflation expectations are often based on one-year ahead forecasts, we record two sets of results based on one-quarter ahead and year-over-year forecasts. In addition, we report results for both the complete and incomplete information environments in order to evaluate the impact of the information frictions stressed above.

For each simulation, we estimate equation (1) and record the parameter estimates  $\hat{a}_0, \hat{a}_1$  and test the null hypothesis  $H_0 : (\hat{a}_0, \hat{a}_1) = (0, 1)$ . For one-quarter ahead expectations, the residuals in equation (1) are not serially correlated under the null and the standard  $\chi^2$  test statistic is valid. However, with year-over-year expectations, the sample observations overlap because the frequency of the data is quarterly while the forecast horizon is one year. To correct for the serial correlation induced by this overlap, an HAC-robust variance-covariance matrix with a uniform window of three lags is constructed for the residuals. This procedure follows Diebold and Mariano (1995) and Hansen and Hodrick (1980) as well as Thomas (1999).

Finally, we also record the fraction of the simulations for which the test statistic rejects  $H_0$  at a significance level lower than 5%. When the null hypothesis is true and the test is correctly specified, this fraction should be close to 5%, the theoretical size of the test. We interpret results where this fraction is significantly higher than 5% as evidence suggesting that simple tests of unbiasedness constitute a poor test of the rational expectations hypothesis.

## 6.2 Results

Our results for the benchmark calibration are reported in Table 2. The first two columns record the median (over 1000 simulations) parameter estimates for the regression equation (1), while the last column records the fraction of times the null hypothesis is rejected over 1000 simulations.

**Table 2**  
**Parameter Estimates and Hypothesis Tests**

	$\hat{a}_0$ (Median)	$\hat{a}_1$ (Median)	Rejection Rate for $H_0$
One-quarter Forecast			
<i>Complete Info</i>	0.0008	0.8924	12.5%
<i>Incomplete Info</i>	0.0014	0.8369	17.1%
Four-quarter Forecast			
<i>Complete Info</i>	0.0017	0.8091	34.0%
<i>Incomplete Info</i>	0.0021	0.7585	38.1%

As we expected, the rejection rate for  $H_0$  in the incomplete information environment is considerably higher than 5% for both the one-quarter ahead and four-quarter ahead forecasts. In the latter case, the hypothesis of unbiased expectations is rejected over 38% of the time, despite the fact that the model embeds within it the hypothesis of rational expectations.

We are also not surprised to discover that when the information friction is de-activated, the number of rejections decreases. However, we are surprised to learn that this information friction does not appear to be the main reason behind the high rejection rate. In particular, the rejection rates in the complete information environment, while somewhat lower, are nevertheless considerably larger than 5% as well. Evidently, there is more to the puzzle than just a Peso problem here. Below, we try to isolate other possible reasons for why rejections of the null occur so frequently.

### 6.3 Further Analysis

Upon further investigation we found that the presence of monetary shocks plays a relatively small role in explaining the high frequency of rejections. Thus, in what follows, we abstract completely from money shocks (and hence learning) so that the only source of disturbance is the technology shock. The results reported below would essentially remain unchanged if we were to include money shocks and learning.

The most obvious avenue to explore is whether the high rejection rate is simply a small-sample issue. To investigate this hypothesis, we repeat

the experiment described above by alternatively considering sample periods of 80, 160, 250 and 500 periods (20, 40, 62.5 and 125 years, respectively). Table 3 reports the rejection rates we obtain for 1000 simulations with each sample period.

**Table 3**  
**Rejection Rates for Different Sample Periods**

Periods	One-quarter Forecast	Four-quarter Forecast
80	17.7 %	41.2 %
160	14.4 %	29.7 %
250	13.1 %	20.7 %
500	12.3 %	15.2 %

The results in Table 3 reveal that sample size certainly does appear to play a role in influencing the probability of rejection. However, sample size does not appear to be the only force at work here. In particular, note that the frequency of rejections remains significantly larger than 5% for very large samples.

At this stage, we conjecture that the high rejection rate may be somehow related to what appears to the near nonstationarity of the simulated data. Under our benchmark calibration, the model is very near  $I(1)$ , with the highest eigenvalue equal to 0.9937. This near nonstationarity appears to arise for two reasons. First, the technology shock exhibits a high degree of persistence (autocorrelation coefficient equal to 0.95). Second, the relatively strong response of monetary policy to the output gap appears to contribute to the ‘instability’ of the model’s dynamics (see Christiano and Gust, 1999).

Table 4 demonstrates how the rejection rate falls as the persistence of the technology shock ( $\rho$ ) is progressively reduced (1000 simulations, 80 periods in length). Note that as the technology shock becomes more transitory in nature, the rejection rate for the one-quarter ahead forecasts falls close to the theoretical size of the test. The rejection rate for the year-over-year forecasts also falls significantly, but remains considerably greater than 5%.

**Table 4**  
**Rejection Rates for Different Values of  $\rho$**

$\rho$	One-quarter Forecast	Four-quarter Forecast
0.95	17.7 %	41.2 %
0.90	13.1 %	36.3 %
0.50	5.7 %	17.6 %
0.01	5.6 %	15.8 %

We next evaluate to what extent the monetary policy rule itself contributes to the high rejection rate. Table 5 records how the rejection rate responds to changes in  $\gamma$ , the parameter that governs the interest rate response to changes in the output gap. Once again, the experiment consists of 1000 simulations, 80 periods in length.

**Table 5**  
**Rejection Rates for Different Values of  $\gamma$**

$\gamma$	One-quarter Forecast	Four-quarter Forecast
0.25	17.7 %	41.2 %
0.15	15.3 %	39.0 %
0.10	14.2 %	38.3 %
0.00	6.2 %	17.9 %

For  $\gamma = 0$ , the model's largest eigenvalue drops from its benchmark value of 0.9937 to 0.9875. As a result, we see from Table 5 that the rejection rate is not much different than 5% for the one-quarter ahead forecast. On the other hand, the rejection rate remains fairly high for the year-over-year forecast.

In the last experiment, we examine how the rejection rate varies for different values of  $\lambda$ , the parameter that governs the interest rate response to changes in the inflation gap. Table 5 records the rejection rate values of  $\lambda$  ranging from 2 to 5. Again, the results are based on 1000 simulations, 80 periods in length.

**Table 6**  
**Rejection Rates for Different Values of  $\lambda$**

$\lambda$	One-quarter Forecast	Four-quarter Forecast
2	17.7 %	41.2 %
3	13.5 %	37.0 %
4	12.0 %	34.8 %
5	11.6 %	34.3 %

From Table 6, we see that a more aggressive reaction to the inflation gap results in fewer rejections, presumably because of an induced increase in the mean-reverting behavior of inflation (and inflation expectations). The rejection rate, however, does not appear to be as sensitive to  $\lambda$  as it is to  $\gamma$  (increasing  $\lambda$  leads to only very small decreases in the system's largest eigenvalue). Even for high values of  $\lambda$ , the rejection rate remains considerably above 5%.

## 7 Summary and Conclusion

Simple econometric tests reported in the literature consistently report that inflation expectations appear to be biased. This finding is commonly interpreted as constituting evidence overturning the hypothesis of rational expectations. The analysis conducted above casts some doubt on the validity of such an interpretation.

The main tool utilized in our investigation was a computational dynamic general equilibrium model capable of generating aggregate behavior similar to the data along a number of dimensions. By construction, the model embedded the assumption of rational expectations. Standard regressions run on equilibrium realizations of inflation and inflation expectations nevertheless revealed an apparent bias in inflation expectations. In these simulations, the null hypothesis of rational expectations was incorrectly rejected in a large percentage of cases.

We conjectured that the primary factor behind this high rejection rate may have something to do with how agents must cope with processing information when shifts in monetary policy regimes are not directly observ-

able. Our conjecture proved to be correct in the sense that the rejection rate for our simulated data was clearly higher when our model agents were confronted with a signal extraction problem defined over the conduct of monetary policy. However, even in the absence of any information friction, the rejection rate remained surprisingly high.

Upon further investigation, it seems that the high rate of rejection may be attributable to several factors, one of which may be the information friction alluded to above. But an even more important factor may simply be the relatively short sample periods utilized by econometricians, in conjunction with the relatively persistent behavior of inflation and inflation expectations. Extending the sample period of simulations and altering the monetary policy rule to induce less persistence generated far fewer rejections of the null in the experiments conducted above.

## References

1. S. R. Aiyagari and R. A. Braun (1998). “Some Models to Guide Policymakers,” *Carnegie-Rochester Conference Series on Public Policy*, 48(1): 1–42.
2. D. Andolfatto and P. Gomme (2003). “Monetary Policy Regimes and Beliefs,” *International Economic Review*, 44: 1–30.
3. D. Andolfatto, S. Hendry and K. Moran (2004). “Labor Markets, Liquidity, and Monetary Policy Regimes,” *Canadian Journal of Economics*, 37(2): 392–420.
4. K. Brunner, A. Cukierman and A. Meltzer (1980). “Stagflation, Persistent Unemployment and the Permanence of Economic Shocks,” *Journal of Monetary Economics*, 6: 467–492.
5. L. J. Christiano, M. Eichenbaum and C. L. Evans (1997). “Sticky Price and Limited Participation Models of Money: A Comparison,” *European Economic Review*, 41: 1201–1249.
6. L. J. Christiano and C. J. Gust (1999). “Taylor Rules in a Limited Participation Model,” Working Paper 7017, National Bureau of Economic Research.
7. R. Clarida, J. Gali and M. Gertler (2000). “Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory,” *Quarterly Journal of Economics*, 115: 147–180.
8. T. F. Cooley and G. D. Hansen (1995). “Money and the Business Cycle,” Chapter 7 in *Frontiers of Business Cycle Research*, edited by T. F. Cooley, Princeton University Press, Princeton, New Jersey.
9. D. Croushore (1997). “The Livingston Survey: Still Useful After All These Years,” *Federal Reserve Bank of Philadelphia Business Review*, March–April: 1–12.
10. J. B. DeLong (1997). “America’s Peacetime Inflation: The 1970s,” in *Reducing Inflation: Motivation and Strategy*, edited by C. D. Romer and D. H. Romer, The University of Chicago Press, Chicago.
11. F. X. Diebold and R. S. Mariano (1995). “Comparing Predictive Accuracy,” *Journal of Business and Economic Statistics*, 13: 253–263.

12. M. Dotsey and J. L. DeVaro (1995). "Was the Disinflation of the Early 1980s Anticipated?" *Federal Reserve Bank of Richmond Economic Quarterly*, 81(4): 41–59.
13. W. B. English, W. R. Nelson and B. P. Sack (2003). "Interpreting the Significance of the Lagged Interest Rate in Estimated Monetary Policy Rules," *Contributions to Macroeconomics*, 3(1), Article 5.
14. C. J. Erceg and A. T. Levin (2003). "Imperfect Credibility and Inflation Persistence," *Journal of Monetary Economics*, 50: 915–944.
15. J. Greenwood, Z. Hercowitz and G. Huffman (1988). "Investment, Capacity Utilization, and the Real Business Cycle," *American Economic Review*, 78: 402–417.
16. J. D. Hamilton (1994). *Time Series Analysis*, Princeton University Press, Princeton, New Jersey.
17. L. P. Hansen and R. J. Hodrick (1980). "Forward Exchange Rates as Optimal Predictors of Future Spot Rates: An Econometric Analysis," *Journal of Political Economy*, 88: 105–121.
18. R. G. King and M. W. Watson (1998). "The Solution to Linear Difference Systems Under Rational Expectations," *International Economic Review*, 39: 1015–1026.
19. S. Kozicki and P. A. Tinsley (2001). "What Do You Expect? Imperfect Policy Credibility and Tests of the Expectations Hypothesis," Working Paper 01–02, Federal Reserve Bank of Kansas City.
20. S. Kozicki and P. A. Tinsley (2003). "Permanent and Transitory Policy Shocks in an Empirical Macro Model with Asymmetric Information," Working Paper 03–09, Federal Reserve Bank of Kansas City.
21. K. K. Lewis (1988). "The Persistence of the 'Peso Problem' When Policy is Noisy," *Journal of International Money and Finance*, 7: 5–21.
22. K. K. Lewis (1989). "Changing Beliefs and Systematic Rational Forecast Errors with Evidence from Foreign Exchange," *American Economic Review*, 79: 621–636.
23. R. E. Lucas (1990). "Liquidity and Interest Rates," *Journal of Economic Theory*, 50: 237–264.

24. N. G. Mankiw, R. Reis and J. Wolfers (2003). "Disagreement About Inflation Expectations," forthcoming *NBER Macroeconomic Annual*.
25. Y. P. Mehra (2002). "Survey Measures of Expected Inflation: Revisiting the Issues of Predictive Content and Rationality," *Federal Reserve Bank of Richmond Quarterly Review*, 88:17–36.
26. J. F. Muth (1960). "Optimal Properties of Exponentially Weighted Forecasts," *Journal of the American Statistical Association*, 55: 299–306.
27. J. M. Roberts (1997). "Is Inflation Sticky?" *Journal of Monetary Economics*, 39: 173–196.
28. F. Ruge-Murcia (2003). "Inflation Targeting Under Asymmetric Preferences," *Journal of Money, Credit, and Banking*, 35: 763–785.
29. F. Schorfheide (2003). "Learning and Monetary Policy Shifts," Manuscript: University of Pennsylvania.
30. P. Shen and J. Corning (2001). "Can TIPS Help Identify Long-Term Inflation Expectations?" *Federal Reserve Bank of Kansas City Economic Review*, 4th Quarter: 61–87.
31. J. B. Taylor (1993). "Discretion versus Policy Rules in Practice," *Carnegie-Rochester Conference Series on Public Policy*, 39: 195–214.
32. L. B. Thomas (1999). "Survey Measures of Expected U.S. Inflation," *Journal of Economic Perspectives*, 13: 125–144.