WP 27-07

DAVID ANDOLFATTO
Simon Fraser University, Canada
and
The Rimini Centre for Economic Analysis, Italy

SCOTT HENDRY
Bank of Canada

KEVIN MORAN
Université Laval, Canada

.

"ARE INFLATION EXPECTATIONS RATIONAL?"

Are Inflation Expectations Rational?

David Andolfatto[a], Scott Hendry[b], Kevin Moran[c,*]

[a]Simon Fraser University and Rimini Centre for Economic Analysis; [b]Bank of Canada; [c]Université Laval

Abstract

Several recent papers report evidence of an apparent statistical bias in inflation expectations and interpret these findings as overturning the rational expectations hypothesis. In this paper, we investigate the validity of such an interpretation. We present a computational dynamic general equilibrium model capable of generating aggregate behavior similar to the data along several dimensions. By construction, model agents form "rational" expectations. We run a standard regression on equilibrium realizations of inflation and inflation expectations over sample periods corresponding to those tests performed on actual data and find evidence of an apparent bias in inflation expectations. Our experiments suggest that this incorrect inference is largely the product of a small sample problem, exacerbated by short-run learning dynamics in response to infrequent shifts in monetary policy regimes.

Keywords: Regime changes; Learning dynamics; Monte Carlo experiments; Sample size.

JEL classification: E47, E52, E58

1.   Introduction

It is well-known that measured inflation expectations appear to evolve "sluggishly" relative to the time-path of actual inflation. In particular, expectations tend to underestimate inflation during periods of rising inflation, and overestimate inflation during periods of falling inflation.[1] This pattern is clearly evident in Figure 1, which plots U.S. inflation and the Livingston Survey measure of expected inflation over the sample period 1960-2005.[2]

Insert Figure 1 here

Assuming that expectations are measured accurately, the evidence in Figure 1 suggests that individuals are prone to making persistent forecast errors. In other words, inflation expectations appear to be biased; a property that is inconsistent with the commonly maintained hypothesis of "rational expectations."[3]

Several studies report evidence against the hypothesis of "rational" expectations; see, for example, Mankiw *et. al.* (2003), Mehra (2002), Thomas (1999), Roberts (1997), and Croushore (1997). The test of this hypothesis is typically formalized as follows. Let π_t denote the inflation rate and let $E_{t-k}[\pi_t]$ represent a measure of the inflation rate expected at date t, conditional on information available at date $t-k$. Using this data, run the following regression:

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

Under the null hypothesis of rational (unbiased) expectations, we have $H_0 : (a_0, a_1) = (0, 1)$.

---

[1] This stylized fact appears robust to various measures of inflation expectations. See, for example, Dotsey and DeVaro (1995) and DeLong (1997).

[2] The Livingston survey is maintained by the Federal Reserve Bank of Philadelphia; its history and structure are described in Croushore (1997).

[3] For the purpose of this paper, we shall take "rational expectations" to mean unbiased expectations, or expectations that display serially uncorrelated forecast errors.

Typically, 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.[4] Table 1 shows that unbiasedness is rejected decisively in various subsamples, but not in the full sample.[5]

Insert Table 1 here

In this paper, we ask a simple question: Should rejections of unbiasedness, based on simple econometric tests such as the one described above, be construed as evidence against rational expectations?

We believe that there are good reasons to interpret such rejections with a grain of salt. First, as the evidence in Table 1 suggests, limitations on sample size are probably important here. Second, we know from the work of several authors that "sluggish" (adaptive) expectation formation may in fact be optimal (rational) in environments where agents cannot directly disentangle the effects of persistent and transitory shocks; see Muth (1960), Brunner, *et. al.* (1980), Andolfatto and Gomme (2003), and Erceg and Levin (2003). We think this is important because much of recent U.S. monetary history can plausibly be interpreted as displaying relatively infrequent "regime shifts" subject to extended "learning" periods. Rational expectations in such an environment still implies serially uncorrelated forecast errors–but only in population. In small samples, this type of learning friction is very likely to manifest itself as serially correlated forecast errors.[6]

The methodological approach we adopt here is to run the commonly-employed econometric tests of unbiasedness on the artificial data generated by a conventional monetary business cycle model that embeds the hypothesis of rational expectations, and then report the fre-

---

[4]These results have been updated and confirmed by Mehra (2002).

[5]Mankiw *et. al.* (2003) reject the null even in their full sample (Table 4).

[6]This possibility has also been raised by Thomas (1999) and Mehra (2002).

quency with which these tests reject the null hypothesis of rational expectations (unbiased expectations).[7]

The benchmark model we employ is a more or less conventional New-Keynesian model. We extend this 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, interpreted as instances when the monetary authority wishes to deviate from its rule for exogenous reasons. 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.

We find that standard econometric tests of unbiasedness using simulated data frequently reject the null hypothesis. As expected, the frequency of these rejections fall as sample size is increased. When we consider a version of the model that "shuts down" the information friction, the frequency of rejections is reduced significantly at all sample sizes. Together, these results suggest that the combination of optimal learning over relatively infrequent regime changes can lead to expectation dynamics that will, in the short-run (i.e., over short samples), display apparently biased expectations and serially correlated errors. Evidence of such behavior cannot, however, be construed as evidence against the rational expectations hypothesis.

Our paper is organized as follows. Section 2 presents a simple empirical model that conveniently captures the estimation/inference problems that arise when learning dynamics are present in small samples. Section 3 develops the economic model used in our simulations.

---

[7]This strategy is similar to that of Lewis (1988, 1989), who studies foreign exchange data, and Kozicki and Tinsley (2005a) who study the term structure.

Section 4 describes the model's calibration and Section 5 provides a brief evaluation of the model's properties. Section 6 presents our Monte Carlo simulations along with our interpretations. Section 7 offers a brief summary and conclusion.

2. An Empirical Model

In this section, we develop a simple empirical model of a stochastic inflation process that incorporates learning. We then simulate this model to test whether the rational expectations generated by the Kalman filter are unbiased. This will provide a quick illustration of the challenges involved with tests of unbiasedness that will be expanded upon with the development of a complete economic model in the next sections.

Assume that inflation π_t follows an exogenous stochastic process:

$$\pi_t = \mu + z_t + e_t; \tag{2}$$

with:

$$z_t = \begin{cases} z_{t-1} & \text{w.p.} & \phi; \\ s_t & \text{w.p.} & 1 - \phi; \end{cases} \tag{3}$$

where s_t is $N(0, \sigma_s^2)$ and e_t is $N(0, \sigma_e^2)$. Hence, the inflation rate has a long-run (unconditional) mean μ and is subject to two types of shocks: a persistent component z_t (think of ϕ as close to unity); and a purely transitory component e_t. Here, it is natural to interpret z_t as the current inflation "regime."

We want to consider two environments distinguished by the nature of information available to agents. In the "complete information" environment, agents can observe (among other things) the regime-shock $z^t \equiv \{z_t, z_{t-1}, ..., z_0\}$ directly. Clearly, a rational one-period-ahead inflation forecast in this case is:

$$E\left[\pi_{t+1} \mid z^t\right] = \mu + \phi z_t, \tag{4}$$

where by construction, forecast errors are serially uncorrelated.

In the "incomplete information" environment, agents cannot observe the regime-shock directly; they only observe the entire history of inflation-shocks $\pi^t \equiv \{\pi_t, \pi_{t-1}, ..., \pi_0\}$. Hence, agents face a signal-extraction problem. In what follows, we assume that agents employ a Kalman filter to update their expectations.[8] In particular, let $b_t \equiv E\left[z_t \mid \pi^t\right]$ denote the expected value of the "hidden" state variable z_t conditional on history π^t. Then the sequence $\{b_t\}$ can be computed recursively according to:

$$b_{t+1} = (1 - \alpha)b_t + \alpha(\pi_t - \mu); \tag{5}$$

where $0 < \alpha < 1$ denotes the (stationary) Kalman gain (itself a function of underlying parameters). In this case, a "rational" one-period-ahead inflation forecast is given by:

$$E\left[\pi_{t+1} \mid b_t\right] = \mu + \phi b_t. \tag{6}$$

Note that the Kalman filter also generates serially uncorrelated forecast errors (in population).[9]

In comparing (6) and (5) with (4), it is apparent that the assumed structure of information matters very much for how rational agents update their inflation forecasts in response to new information. For example, under complete information, the inflation forecast in (4) depends

---

[8]Strictly speaking, the Kalman filter is not "fully rational" in that it restricts agents to use a linear forecasting rule. We think our conclusions would be unchanged under a (significantly more complicated) non-linear filter because first, our results are consistent with those of Andolfatto and Gomme (2003), who do use an optimal non-linear filter (albeit, in a much simpler two-regime environment); second, the Kalman filter generates serially uncorrelated forecast errors (in population); and third, the analysis in Appendix 1 shows that the instances where the filter forecasts could have been improved (in a statistically significant manner) are infrequent.

[9]This is true whether or not we use the steady-state Kalman filter.

solely on the persistent "regime" parameter z_t and not at all on the "transitory" component e_t. If the probability of regime-change is small ($\phi \to 1$), then expectations in this environment are almost always stable; shifting occasionally when a regime change occurs.

Under incomplete information, however, agents cannot immediately disentangle the reasons underlying a sudden exogenous change in contemporaneous inflation. Based on what they know of the underlying stochastic process, they "hedge their bets" and assign some probability to each possible source in accordance with (5). One implication of this hedging behavior is that even transitory shocks will have persistent effects on the dynamics of inflation forecasts. A second and related implication is that inflation forecasts do not fully adjust contemporaneously with a regime shock (as is the case in the complete information environment).

Note that the "learning dynamics" just described do not matter "in population." That is, empirically, one should find that inflation expectations are unbiased and forecast errors are serially uncorrelated in a long enough sample. But it seems clear enough that over "short" sample periods, these learning dynamics could matter when estimating equations like (1). For example, consider simulating the process in (2)-(3) under incomplete information, generating a sample of length T. If the probability of regime change is sufficiently small, then the number of such changes expected to be observed in a given sample T diminishes quickly as T becomes smaller. Imagine the extreme case in which no regime change is actually observed. This implies that the entire variation in inflation over the sample period is driven by the serially uncorrelated shock e_t. However, observe that the rational inflation forecast (5)-(4) would in this case necessarily generate serially correlated forecast errors, as well as reduce the tight correlation that might otherwise obtain between expected and realized inflation. In general, the presence of only a few (say one or two) regime changes in the small samples used to study (1) could in this way drive OLS estimates away from (0,1) and affect the inference drawn from the estimates, leading to frequent rejection of the unbiasedness hypothesis.

We can evaluate the potential quantitative relevance of the intuition developed above by considering an explicit numerical example. Here, we choose the following parameter values: $\mu = 0.0123$, $\phi = 0.975$, $\sigma_s = 0.010$, and $\sigma_e = 0.007$. These parameters imply a (stationary) Kalman gain $\alpha = 0.26$. Assuming a quarterly time-period, these parameters imply a long-run (annualized) inflation rate of 5% and an average regime duration of 40 quarters.

Now, consider the following experiment. Given the parameters above, fix a sample length T and simulate the stochastic process for inflation. Based on this simulated data, construct a measure of expected inflation for both the complete information environment (4) and the incomplete information environment (6). Run the regression (1) on this simulated data (for $k = 1$). Test the null hypothesis $H_0 : (a_0, a_1) = (0, 1)$. Repeat this 1000 times, record the median of the estimates $(\hat{a}_0, \hat{a}_1)$, and report the rejection rate for the null hypothesis. Redo all of this for a sequence of increasing sample sizes $T \in \{80, 160, 250, 1000\}$.

Table 2 records the results of this experiment in the complete information environment. As expected, the median estimate $(\hat{a}_0, \hat{a}_1)$ in this case is close to $(0, 1)$ even in small samples; furthermore, there is evidence of $(\hat{a}_0, \hat{a}_1) \rightarrow (0, 1)$ as $T \rightarrow \infty$. Note that the null hypothesis is rejected (incorrectly) about 12% of the time in these simulations at sample size $T = 80$; with the rejection rate declining steadily as T is increased.[10]

Insert Table 2 here

The next table reports the analogous results under the incomplete information environment.

Insert Table 3 here

---

[10]In Tables 2 and 3, standard deviations of estimates are computed using the HAC-consistent covariance matrix advocated by Newey and West, 1994. This method uses a flexible, data-dependent, lag structure. Results are obtained using the Parzen kernel, but are robust to other methods as well (uniform, Bartlett, etc.).

Table 3 confirms our suspicion that the combination of "optimal learning" over relatively infrequent regime changes can lead to expectation dynamics that will, in the short-run, display serially correlated errors. In a small sample, there is a high likelihood of (incorrectly) rejecting the null hypothesis. As in the complete information environment, observe that $(\hat{a}_0, \hat{a}_1) \rightarrow (0, 1)$ as $T \rightarrow \infty$.

All-in-all, these results point to the possibility that recent studies purporting to find evidence overturning "rational expectations" are making incorrect inferences based on inappropriately small samples.[11] How seriously one takes this small sample issue is likely to depend on one's view about the existence and frequency of "regime shocks" in the inflation process; and whether one believes that agents in the economy can easily disentangle persistent and transitory components in this process.

Of course, in reality, inflation is not an exogenous process. Inflation is likely determined in part by the conduct of monetary policy and in part by natural economic forces. If monetary policy is reasonably approximated by a Taylor rule, then it may make sense to interpret a monetary policy regime as a specific parameterization of this rule. Among other things, a Taylor rule specifies a "long-run" inflation target as a parameter. Based on our reading of the historical evidence, we do not view it far-fetched to think of "regime-shifts" in monetary policy as manifesting themselves in the form of periodic changes in this inflation target parameter.

To enhance the plausibility of our hypothesis, we re-cast it in the context of a suitably calibrated model that broadly mimics the quantitative behavior of the economy along a number of important dimensions.

---

[11] The qualitative nature of these results is not altered when we consider year-over-year ($k = 4$) forecasts. The interested reader can refer to Appendix 2, where such a case is considered.

3. An Economic Model

We pursue our investigation by embedding the regimes and learning framework within a variant of the now-standard New Keynesian model. As the model as been analyzed in detail elsewhere, our description here will be brief.[12]

3.1. Household Sector

There exists a continuum of identical, infinitely-lived households with preferences given by:

$$E_0 \sum_{t=0}^{\infty} \beta^t u(c_{1t}, c_{2t}, n_t), \tag{7}$$

where c_{1t} denotes consumption of a "cash" good, c_{2t} denotes the consumption of a "credit" good, n_t denotes employment, and $0 < \beta < 1$ is the households' subjective discount factor. The expectation operator E_0 will be discussed in detail below, as its form will depend on the assumed information structure.

At the beginning of period t, a household is in possession of M_t dollars. It then receives a beginning-of-period cash transfer T_t from the monetary authority. This transfer is related to the monetary authority's management of the short-term interest rate through its policy rule (to be described below). The household's cash-in-advance constraint is given by:

$$P_t c_{1t} \leq M_t + T_t, \tag{8}$$

where P_t is the dollar price of the cash good.

Households own the economy's capital k_t, which can be employed at intensity h_t, so that

---

[12]See for example, Erceg, Henderson and Levin (2000), Ireland (2001), and Christiano, Eichenbaum and Evans (2005). We experimented with several different model specifications, including the limited-participation model of Christiano and Gust (1999). While different models fit the data better or worse along various dimensions, our conclusions about empirical tests of rationality in inflation expectations are not sensitive to model choice.

available capital services are given by $h_t k_t$. These services are rented at the real rental rate r_t^k. Utilizing the stock of capital at rate h_t entails a per-unit real resource cost of $\tau(h_t)$.[13] The ownership of capital therefore generates a (nominal) income flow $P_t \left[r_t^k h_t - \tau(h_t) \right] k_t$ for the household.

Labor is paid a (competitive) nominal wage W_t and the household receives a nominal dividend D_t based on its ownership of a (monopolistically competitive) intermediate-goods sector. Hence, the household's nominal income at date t is given by:

$$W_t n_t + P_t \left[r_t^k h_t - \tau(h_t) \right] k_t + D_t.$$

Out of this income, the household purchases the credit good c_{2t} and the investment good x_t (both of which sell for the same nominal price). Any unused income is carried forward in the form of cash M_{t+1}. Augmenting this carry-forward is any unspent cash from (8) in the event that the cash-in-advance constraint does not bind. Hence, we have:

$$M_{t+1} + P_t \left[c_{2t} + x_t \right] = W_t n_t + P_t \left[r_t^k h_t - \tau(h_t) \right] k_t + D_t + \left[M_t + T_t - P_t c_{1t} \right]. \qquad (9)$$

Investment x_t augments the capital stock according to:

$$k_{t+1} = (1 - \delta)k_t + [1 - S(x_t/x_{t-1})]x_t, \qquad (10)$$

where $\delta \in (0, 1)$ is a constant capital depreciation rate, and $[1 - S(x_t/x_{t-1})]x_t$ summarizes the process by which current and past levels of investment increase the available stock of capital.[14] The function S is such that $S(1) = S'(1) = 0$ and $S''(1) \equiv \kappa > 0$.

---

[13] See Christiano, Eichenbaum and Evans (2005).

[14] Again, see Christiano, Eichenbaum and Evans (2005).

The household's decision involves choosing a stochastic process

$$\{c_{1t}, c_{2t}, n_t, x_t, h_t, k_{t+1}, M_{t+1}\}_{t=0}^{\infty}$$

to maximize expected lifetime utility (7) subject to the cash-in-advance constraint (8), the budget constraint (9), the capital accumulation equation (10) and initial conditions k_0, x_{-1}, M_0. As we will discuss below, the specification of the expectation operator E_0 under which this maximization is conducted will depend on the information structure we assume.

3.2. Business Sector

The final-good, y_t, is produced by assembling a continuum of intermediate goods $y_{j,t}$ for $j \in (0,1)$ that are imperfect substitutes with a constant elasticity of substitution ϑ. The production function is defined as:

$$y_t \leq \left(\int_0^1 y_{j,t}^{\frac{\vartheta-1}{\vartheta}} \, dj \right)^{\frac{\vartheta}{\vartheta-1}}, \qquad \vartheta > 1. \tag{11}$$

Final-good producing firms behave competitively, maximizing profits and taking the market price of the final-good P_t as well as the price of each intermediate-good prices $P_{j,t}$, $j \in (0,1)$ as given.

Profit maximizing by the final-good producing firms implies the following input demand function for the intermediate good j

$$y_{jt} = \left(\frac{P_{j,t}}{P_t} \right)^{-\vartheta} y_t, \tag{12}$$

which represents the economy-wide demand for intermediate good j as a function of its relative price and of the economy's total output of final-good y_t.

Finally, imposing zero profits in the sector leads to the following description of the final-

good price index P_t:

$$P_t = \left(\int_0^1 P_{j,t}^{1-\vartheta} \, dj \right)^{\frac{1}{1-\vartheta}}. \tag{13}$$

Intermediate-good producing firm j uses capital and labor services $k_{j,t}$ and $n_{j,t}$ to produce $y_{j,t}$ units of its differentiated good, according to the following constant-returns-to-scale technology:

$$y_{j,t} \leq a_t k_{j,t}^{\alpha} (n_{j,t} - \eta)^{1-\alpha}, \quad \alpha \in (0,1), \tag{14}$$

where $\eta \geq 0$ represents a fixed production cost and a_t is a standard aggregate technology shock realized at the beginning of the period.[15]

Following a variant of the Calvo (1983) mechanism, assume that with probability ξ, a firm must charge for its good last period's price indexed by the steady-state (gross) rate of inflation Π^*; whereas with probability $1 - \xi$, the firm is free to re-optimize and fix a new price.[16]

If firm j receives the signal to re-optimize at date t, it chooses a price $\hat{P}_{j,t}$, as well as contingency plans for $n_{j,t+q}$, $k_{j,t+q}$ for all $q \geq 0$ that maximize its discounted, expected (real) profit flows for the period where it will not be able to re-optimize again. The profit maximization problem is the following:

$$\max_{\{k_{j,t}, n_{j,t}, \hat{P}_{j,t}\}} E_t \left[\sum_{q=0}^{\infty} (\beta\xi)^q \lambda_{t+q} D_{j,t+q} / P_{t+q} \right], \tag{15}$$

---

[15]The fixed cost η is set to a value that implies zero economic profits in steady state and thus validates the assumption of no entry int the sector.

[16]This specification of the Calvo mechanism follows Yun (1996). To introduce inertial movements in inflation, authors often use the alternative assumption that when the re-optimization signal is not received, the price is indexed to the preceding period's rate of inflation (Christiano, Eichenbaum and Evans, 2005). Since the learning behaviour discussed below also imparts inertia to inflation, we use the indexation to average inflation in our benchmark model. Our regressions results are not sensitive to this assumption.

with the profit flow $D_{j,t+q}$ given by:

$$D_{j,t+q} = (\hat{P}_{j,t}\Pi^*)^q y_{j,t+q} - P_{t+q} r^k_{t+q} k_{j,t+q} - W_{t+q} n_{j,t+q}, \tag{16}$$

where ξ^q expresses the probability that $\hat{P}_{j,t}$ remains in effect (including of indexation) at time $t+q$ and $(\Pi^*)^q$ is a cumulative indexation factor. Profit maximization is subject to the demand for good j (12) and to the production function (14). Again, note that the nature of the expectation operator E_t will depend on the assumed structure of information.

The symmetry in the demand for each intermediate good implies that all firms allowed to re-optimize choose the same price $\hat{P}_{j,t}$, which we denote $\hat{P}_t$. Considering the definition of the price index in (13) and the fact that at the economy's level, a fraction $1 - \xi$ of intermediate-good producing firms re-optimize, the aggregate price index P_t evolves according to:

$$P_t^{1-\vartheta} = \xi(\Pi^*)^{1-\vartheta} P_{t-1}^{1-\vartheta} + (1 - \xi)(\hat{P}_t)^{1-\vartheta}. \tag{17}$$

3.3. Monetary Policy

Let π_t denote the net inflation rate and let i_t denote the net nominal interest rate. Further, let (r^*, y^*) denote the net real rate of interest and level of real GDP that would prevail in this economy in a deterministic steady-state. We assume a Taylor rule of the following form:

$$i_t = (1 - \theta)\left[r^* + \pi_t^* + \lambda(\pi_t - \pi_t^*) + \gamma(y_t - y^*)\right] + \theta i_{t-1} + u_t, \tag{18}$$

where π_t^* denotes the monetary authority's date t inflation target 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 alluded to earlier, the monetary authority achieves any particular i_t with an appropriate lump-sum

injection/withdrawal of cash T_t to/from the household sector. Accordingly, the aggregate stock of money evolves as follows:

$$\overline{M}_{t+1} = \overline{M}_t + T_t. \tag{19}$$

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 (18). 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:

$$u_t = \phi_1 u_{t-1} + e_t, \tag{20}$$

with $0 \leq |\phi_1| << 1$ 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 π_t^*, while remaining fixed 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 π_t^* from its long term (time-invariant) mean π^*. We assume that the following process, a mixture of a Bernoulli trial and a normal random variable, expresses how z_t evolves over time:

$$
z_t = \begin{cases} z_{t-1} & \text{w.p.} \quad \phi_2; \\ s_t & \text{w.p.} \quad 1 - \phi_2, \quad s_t \sim N(0, \sigma_s^2); \end{cases} \tag{21}
$$

with $0 << \phi_2 < 1$.

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 ϕ_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.

Note that, by combining the definition $z_t \equiv \pi_t^* - \pi^*$ with the Taylor rule (18), we can write:

$$
i_t = (1 - \theta)\left[r^* + \pi^* + \lambda(\pi_t - \pi^*) + \gamma(y_t - y^*)\right] + \theta i_{t-1} + \epsilon_t \tag{22}
$$

where $\epsilon_t \equiv (1 - \theta)(1 - \lambda)z_t + u_t$. Hence, under complete information, agents are assumed to observe the inflation target π_t^* directly (or, equivalently, they can observe z_t, u_t separately). Under incomplete information, agents must make rational inferences over the inflation target (as they can only observe ϵ_t and not the composition of ϵ_t across z_t and u_t).

3.4. The Kalman Filter

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 ϵ_t can be expressed as follows:

$$
\begin{bmatrix} z_{t+1} \\ u_{t+1} \end{bmatrix} = \begin{bmatrix} \phi_2 & 0 \\ 0 & \phi_1 \end{bmatrix} \begin{bmatrix} z_t \\ u_t \end{bmatrix} + \begin{bmatrix} \kappa_{t+1} \\ e_{t+1} \end{bmatrix} ; \ \epsilon_t = \begin{bmatrix} (1-\theta)(1-\lambda) & 1 \end{bmatrix} \begin{bmatrix} z_t \\ u_t \end{bmatrix} ; \quad (23)
$$

where

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

Observe that $E_t \kappa_{t+1} = 0$.

The system defined in (23) along with the specification of e_t and κ_t define a state-space system (Hamilton, 1994, Chapter 13), where the first equation in (23) is the state equation and the second equation in (23) 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 ϵ_t up to and including period t. These estimates are updated sequentially as new information becomes available according to:

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

where K_t represents the gain matrix of the filter and governs the extent to which unexpected movements in the observed shock ϵ_t are attributed to the persistent shift z_t or to the transitory shock u_t. The system in (23) can then be used to compute expected future shocks:

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

As for the empirical model studied earlier, 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 Bernouilli-normal variate, the Kalman filter is not–strictly speaking–fully optimal. This implies that agents using this filter are not fully 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; see Appendix 1.

3.5. Equilibrium

An equilibrium for this model economy consists of: an allocation

$$\{c_{1t}, c_{2t}, n_t, x_t, h_t, y_t, k_{t+1}, M_{t+1}\}_{t=0}^{\infty}, \{y_{j,t}, k_{j,t}, n_{j,t}, D_t\}_{t=0}^{\infty} \, \forall j \in (0,1);$$

a price-system $\{P_t, W_t, r_t, i_t\}_{t=0}^{\infty}$, $\{\hat{P}_{j,t}\}_{t=0}^{\infty} \, \forall \, j \in (0,1)$; a monetary policy rule, including injections $\{T_t\}_{t=0}^{\infty}$; and expectations $\{E_t\}_{t=0}^{\infty}$ such that, given a known stochastic process over $\{a_t, z_t, u_t\}_{t=0}^{\infty}$, the following conditions hold:

1. Given expectations, the allocation maximizes (7) subject to (8), (9), (10);

2. Given expectations, the allocation and price-setting behavior $\{\hat{P}_{j,t}\}_{t=0}^{\infty} \; \forall \; j \in (0,1)$ maximizes (15) subject to (12) and (14), generating a nominal dividend D_t in accordance with (16);

3. $\hat{P}_{j,t} = \hat{P}_t \; \forall \; j \in (0,1)$ with the aggregate price level P_t satisfying (17);

4. Monetary injections $\{T_t\}_{t=0}^{\infty}$ are chosen to be consistent with the desired nominal interest rate as determined by (22);

5. The allocation and prices are consistent with individual and aggregate behavior; i.e.

$$
\begin{aligned}
h_t k_t &= \int_0^1 k_{j,t} \, dj; \; n_t = \int_0^1 n_{j,t} \, dj; \; M_t = \overline{M_t}; \\
y_t &= c_{1t} + c_{2t} + x_t + \tau(h_t) k_t = a_t F(h_t k_t, n_t - \eta).
\end{aligned}
$$

6. Expectations $\{E_t\}_{t=0}^{\infty}$ are formed in a manner consistent with the stochastic properties induced by [1]-[5], given available information.

4. Solution Method and Calibration

As nominal variables are non-stationary (with the exception of the nominal interest rate), we transform them in the usual manner (e.g. by dividing each nominal variable either by the price-level or the money supply) to render a stationary solution. After calibrating the model's parameters, the deterministic steady-state of the system is computed and a first-order approximation of the equilibrium system is formed around the steady-state, with all variables expressed as percent deviations from their steady-state values. The algorithm delivers a first-order system in which all such deviations are expressed as functions of predetermined state variables and exogenous shocks; see King and Watson (2002) for details.[17]

---

[17] Finding solutions for the incomplete information case requires only a slight modification to the King and Watson (2002) algorithm, with forecasts of the monetary policy shocks based on the Kalman filter described

4.1. Preferences and Technology

The model is calibrated to a quarterly frequency. Preferences are given by:

$$u(c_{1t}, c_{2t}, n_t) = \zeta \ln(c_{1t}) + (1 - \zeta) \ln(c_{2t}) - Bn_t. \tag{24}$$

We set $\zeta = 0.84$, which implies an interest-rate elasticity of the velocity of money equal to 0.20, as in Cooley and Hansen (1995). B is set to 2.6 so that steady-state hours represent 0.3 of total available time.

The parameter ϑ, the price elasticity of demand for each intermediate good, is set to 5. This implies that the steady-state markup of price over marginal costs is 25 percent, a standard value in the literature. The parameter ξ, expressing the probability a given intermediate-good producer does not re-optimize, is set to 0.6, implying that on average, prices are re-optimized every $1/(1-0.6) = 2.5$ quarters. This value is similar to those arrived at when estimating similar monetary DSGE models (e.g. Ireland, 2001). The parameter $\kappa = S''(1)$, governing the severity of the adjustment costs in capital accumulation, is set to 1, the estimate arrived at in Christiano, Eichenbaum and Evans (2005) in their model specification that is closest to the present one.

The production function of intermediate-good producing firms is Cobb-Douglas in capital and (net) labor services, i.e. $F(k_t, n_t) = a_t k_t^\alpha (n_t - \eta)^{1-\alpha}$, and aggregate technology shocks evolve according to

$$\ln a_t = (1 - \rho) \ln a + \rho \ln a_{t-1} + v_t, \ v_t \sim N(0, \sigma_v^2).$$

The scale parameter for the technology shock is set to $a = 1$. Further, we set the autocorrelation (ρ) and standard deviation (σ_v) parameters to the standard values of 0.95 and 0.005,

earlier. Details and programs are available from the authors.

respectively.

The monopolistic competitive structure of the sector producing intermediate goods implies that these producers enjoy positive profits. We set the fixed labor cost η to $1/3$, which renders these steady-state profits zero. The parameter α is set to 0.5, which implies a steady-state labour income share of 0.6.[18]

Following Christiano, Eichenbaum and Evans (2005), we calibrate the capital utilization function $\tau(h_t)$ by requiring that h equals 1 in steady state and by setting $\tau(1) = 0$. This leaves the steady-state computations unaffected by the capital utilization decision. We then set the elasticity $\sigma_\tau = \tau''(1)/\tau'(1)$ to 0.3, which allows our model to match the peak response of capital utilization following monetary policy shocks documented by Christiano, Eichenbaum and Evans (2005).

Finally, following standard arguments, β is set to 0.988 and δ to 0.019, implying a steady-state (quarterly) capital-output of 13, an investment-output ratio of 0.25, and a (quarterly) net real interest rate $r^* = 0.012$.

4.2. Monetary Policy

The long-run (quarterly) inflation rate is set to $\pi^* = 0.0123$, which implies steady-state inflation of 5% at an annualized rate. According to the monetary policy rule (18), the nominal interest rate reacts to deviations of inflation from its current target (λ), to deviations of output from its trend (γ), and to its own lagged value (θ). We use results from the sizeable literature estimating these rules to choose plausible parameter values.[19]

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

---

[18] The steady state GDP share of labour is $1 - \alpha$ divided by the markup 1.25.

[19] Empirical estimates of Taylor rules are arrived at using various methodologies: rules can be forward looking (Clarida *et. al.*, 2000) or based solely on contemporaneous variables (Taylor, 1993); as well, 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, 2005b).

becomes a self-fulling prophesy. Accordingly, we set $\lambda = 2$ in our benchmark calibration. This value is only slightly larger that the original value used by Taylor and is in line with recent empirical estimates; see Erceg and Levin (2003), English *et. al.* (2003), and Schorfheide (2005). We set the smoothing parameter to $\theta = 0.50$, which is in the range of recent empirical estimates; again, see Erceg and Levin (2003), Kozicki and Tinsley (2005b), and Schorfheide (2005) (all three empirical exercises explicitly allow for regime shifts in monetary policy and so are consistent with our environment). The parameter γ, the response of the interest rate to changes in the output gap is set to $\gamma = 0.0$. This follows results where fully-developed DSGE models, estimated using maximum likelihood, find estimates of γ that are not statistically significant (e.g. Ireland, 2001).[20]

The parameters ϕ_2 and σ_s govern the dynamics of the inflation target z_t. In particular, ϕ_2 governs the mean duration of any given regime and σ_s represents the standard deviation of the distribution from which a regime shift is drawn when it occurs. We set $\phi_2 = 0.975$, which implies a mean duration of shifts of 10 years (40 model periods). This value is consistent with our earlier interpretation of regimes corresponding to the tenure of 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 (2005). 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 (2005b), and Schorfheide (2005).

Finally, parameters related to transitory monetary policy shocks are $(\phi_1, \sigma_e) = (0, 0.007)$. The persistence parameter here is consistent with the estimates provided in Schorfheide (2005), but the variance of the innovation is slightly above his estimated range. We chose σ_e to match the model's predicted GDP volatility with that of the data; i.e., see Table 4.

---

[20]Further, Erceg and Levin (2003) find only a small value (0.25) for this parameter using a different estimation technology.

5.   Properties of the Benchmark Model

In this section, we report some of the properties of our benchmark model (incomplete information over monetary policy regimes) to show the reader that the model is not strongly at variance with actual economies. To this end, we begin with Table 4, which reports a set of moments describing the model's covariance structure in relation to the U.S. economy. In this table, $\sigma(x)$ measures the percent standard deviation of variable x; $\rho(x, GDP)$ measures the contemporaneous correlation of variable x with real GDP; and $\rho(x, x_{-1})$ measures the first-order autocorrelation coefficient for variable x.[21]

Insert Table 4 here

Table 4 reveals that the model's cyclical properties are not grossly inconsistent with the data. In fact, the model performs about as well as any other quantitative general equilibrium model along these dimensions. Of particular interest is the behavior of inflation and inflation expectations, since these two variables constitute the basis for the regression equation (1). While the model does not mimic the behavior of these two variables exactly, it seems reasonably "close" for the purpose at hand.

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.[22]

---

[21]Source: GDP, expenditure components, labor input (establishment survey) and inflation (CPI) variables are from the Federal Reserve Bank of St. Louis *Fred* database, 1954Q1–2005Q3. TFP is computed as $\ln(GDP) - 0.64 \log(Hours)$. Expected inflation is the expected increase in one-quarter-ahead CPI according to the *SPF* survey of the Federal Reserve Bank of Philadelphia (sample period 1981Q3–2005Q2).

[22]Some readers may wonder whether the apparent biasedness in inflation expectations is the product of

Insert Figure 2 here

The dynamic response of inflation and expected inflation in Figure 2 demonstrates that under incomplete information, even rational expectations may–in the "short-run"–exhibit serially correlated forecast errors. We suspect that this type of force may be one reason why conventional tests of rational expectations are frequently (and incorrectly) rejected in small samples.

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, although this effect is somewhat muted here owing to sluggish price adjustment. Nevertheless, even this relatively muted disinflation is penalized heavily by the monetary policy rule, which is what accounts for the brief acceleration in money growth (and slight decline in the nominal interest rate).

Insert Figure 3 here

Of course, the purpose of our paper is not to provide an extensive evaluation of the strengths and shortcomings of this particular business cycle model. The key observation here is that the model's dynamic properties are not unreasonable and that the model does embed the hypothesis of (near) rational expectations. The model therefore provides a reasonably good laboratory with which we can perform our experiments.

---

the suboptimality of our linear filter. As Andolfatto and Gomme (2003) demonstrate, the same dynamics are evident when expectations are formed optimally according to Bayes' rule.

6. Testing the Unbiasedness Hypothesis

Here, we repeat the experiment performed earlier using our simple empirical model. In particular, 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.

Table 5 records the results for the complete information environment.

Insert Table 5 here

Note that the parameter estimates and rejection rates in Table 5 are not that different from those reported in Table 2.

To quantify the role of information frictions, Table 6 reports the corresponding results for the incomplete information environment. Observe that here, as in Table 3, the median parameter estimates $(\hat{a}_0, \hat{a}_1)$ converge quickly to $(0, 1)$ as sample size is increased. The basic result here is similar to that found in the empirical model (although not as dramatic). For the sample sizes considered here, the effect of the information friction is to increase the rejection rate by roughly a factor of two. For sample sizes in the range commonly used in the empirical literature $(T = 80)$, the probability of incorrectly rejecting the null hypothesis is almost 30%.[23]

Insert Table 6 here

---

[23]This behaviour is also present when four-quarter-ahead expectations are considered. Appendix 2 shows that in that case, incomplete information leads OLS to again reject the null hypothesis of unbiasedness roughly twice as often than under complete information.

7. Conclusion

By various measures, inflation expectations appear to evolve "sluggishly" relative to actual inflation. Several recent papers report econometric evidence of an apparent bias in inflation expectations. These findings are commonly interpreted as constituting evidence that overturns the rational expectations hypothesis. Our analysis casts some doubt on the validity of such an interpretation.

Our main message can be summarized as follows. To the extent that U.S. monetary policy in the recent past has been characterized by "regime shifts" subject to "adaptive" (but rational) learning, the probability of incorrectly rejecting the null of unbiasedness is not insignificant (roughly in the order of 30%), given the sample sizes that are commonly employed in these tests.

Of course, to believe that our argument is empirically relevant, one must adopt the view that monetary policy is characterized by regime shifts in a world where the monetary authority is either unwilling or unable to signal its new objective in a credible and timely manner to market participants. If one was instead to adopt the view that sudden changes in policy are, for the most part, readily communicated to private agents, then our findings are weakened considerably. In particular, the probability of incorrectly rejecting the null in our complete information model is relatively small; i.e., in the order of 10%-15% for the sample sizes commonly employed.

Given the ink that has been spilled on the subject, we do not view it implausible to think of monetary policy as often operating in a less than transparent manner. Of course, even if this is true, our results do not prove that individuals do, in fact, form their inflation expectations rationally. But we do think that our findings provide an important caveat for those who are perhaps a little too eager to interpret visual and statistical evidence as contrary to the rational expectations hypothesis.

Appendix 1
On the Optimality of the Kalman Filter

The Kalman filter produces the best forecast of the future monetary policy disturbances ϵ_t out of the class of linear filters. In principle, better forecasts may be achievable through the use of a nonlinear filter. This fact raises the possibility that the rejections of H_0 reported in the text (in the incomplete information case only) may to some extent be the product of our restriction on the form of the optimal learning rule. Here, we try to evaluate the extent to which our restriction to a linear filter constitutes a reasonably good approximation.

The question we ask here is whether agents in our model, in forming their expectations with the Kalman filter, could have systematically improved on their forecasts. To this end, define v_t as the error that agents make when forecasting the monetary policy disturbance ϵ_t, where

$$v_t = \epsilon_t - E_{t-1}\left[\epsilon_t\right] = \epsilon_t - \left[\begin{array}{cc} (1-\theta)(1-\lambda) & 1 \end{array}\right] \left[\begin{array}{c} \hat{z}_{t|t-1} \\ \hat{u}_{t|t-1} \end{array}\right].$$

Next, estimate the following additional regression on each generated sample of data:

$$v_{t+1} = \theta_1(L)\epsilon_t + \theta_2(L)(\epsilon_t)^2 + \theta_3(L)(\epsilon_t)^3 + \theta_4(L)\epsilon_t\epsilon_{t-1} + e_t,$$

where $\theta_i(L)$, for $i = 1, 2, 3, 4$ are polynomials in the lag operator. We then test the null hypothesis:

$$\hat{H}_0 : \theta_1(L) = \theta_2(L) = \theta_3(L) = \theta_4(L) = 0.$$

The appropriate lag length for each simulation is the Schwarz (Bayesian) criterion. If this null hypothesis is rejected very often, or if it is always rejected when H_0 itself is rejected, this would suggest that our linear filter is at the source of our reported rejections of unbiasedness (H_0); rather than incomplete information and rational learning about the monetary policy

shifts.

The following table demonstrates that our restriction to a linear filter is not the main driver behind our results. The table reports, for our benchmark calibration (both complete and incomplete information environments), the frequency at which the null $\hat{H}_0$ was rejected (first column) and the frequency at which both H_0 and $\hat{H}_0$ were jointly rejected (second and third columns). Recall that even if there are no possible gain in forecasting, correcly-sized tests would still reject $\hat{H}_0$ 5% of the cases.

Insert Table 7 here

The table above shows that $\hat{H}_0$ was rejected in 5.8% of the simulations under complete information, with $H_0, \hat{H}_0$ jointly rejected in only 1.2% of the cases, whether the forecast horizon is $k = 1$ or $k = 4$. Under incomplete information, $\hat{H}_0$ is rejected slightly more frequently (9.8%); however, $H_0, \hat{H}_0$ is jointly rejected much less often, in less than 5% of the cases. These results show that only a small fraction of the rejections of the unbiasedness hypothesis we report can be attributed to the "quasi-rationality" of the Kalman filter.

Appendix 2
Year-over-Year Forecast Horizon

In the main body of the paper, we restricted attention to testing the null using a regression based on one-quarter-ahead expectations (i.e., $k = 1$ in equation 1). However, tests in the literature are often performed using a year-over-year horizon. As we demonstrate here, the choice of the forecast horizon can further increase the probability of incorrectly rejecting the null (in small samples); overall, however, the results discussed in the paper continue to hold.

Tables 7 and 8 report the results of our Monte Carlo experiment for the economic model under complete and incomplete information, for the four-quarter-ahead expectations ($k = 4$) case.[24]

Insert Table 8 here

Insert Table 9 here

---

[24]With year-over-year expectations ($k = 4$), the sample observations overlap because the frequency of the data is quarterly while the forecast horizon is one year. This overlap induces serial correlation of up to three lags, even under complete information. We continue to correct for serial correlation in a flexible manner by using the Newey and West, 1994, method. Results are robust to using alternative kernel definitions (uniform, Bartlett) and maximal lag lengths.

8.   References

1. D. Andolfatto and P. Gomme (2003). "Monetary Policy Regimes and Beliefs," *International Economic Review*, 44: 1–30.

2. K. Brunner, A. Cukierman and A. H. Meltzer (1980). "Stagflation, Persistent Unemployment and the Permanence of Economic Shocks," *Journal of Monetary Economics*, 6: 467–492.

3. G. A. Calvo (1983). "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics*, 12: 383-398.

4. L. J. Christiano, M. Eichenbaum and C. L. Evans (2005). "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 113: 1–45.

5. L. J. Christiano and C. J. Gust (1999). "Taylor Rules in a Limited Participation Model," Working Paper 7017, National Bureau of Economic Research.

6. 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.

7. 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.

8. D. Croushore (1997). "The Livingston Survey: Still Useful After All These Years," *Federal Reserve Bank of Philadelphia Business Review*, March-April: 1–12.

9. 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.

10. 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.

11. 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.

12. C. J. Erceg, D. W. Henderson, A. T. Levin (2000). "Optimal Monetary Policy with Staggered Wage and Price Contracts," *Journal of Monetary Economics*, 46: 281–313.

13. C. J. Erceg and A. T. Levin (2003). "Imperfect Credibility and Inflation Persistence," *Journal of Monetary Economics*, 50: 915–944.

14. J. D. Hamilton (1994). *Time Series Analysis*, Princeton University Press, Princeton, New Jersey.

15. P. N. Ireland (2001). "Sticky-Price Models of the Business Cycle: Specification and Stability," *Journal of Monetary Economics*, 47: 3–18.

16. R. G. King and M. W. Watson (2002). "System Reduction and Solution Algorithms for Singular Linear Difference Systems Under Rational Expectations," *Computational Economics*, 20: 57–86.

17. S. Kozicki and P. A. Tinsley (2005a). "What Do You Expect? Imperfect Policy Credibility and Tests of the Expectations Hypothesis," *Journal of Monetary Economics*, 52: 421–447.

18. S. Kozicki and P. A. Tinsley (2005b). "Permanent and Transitory Policy Shocks in an Empirical Macro Model with Asymmetric Information," *Journal of Economic Dynamics and Control*, 29: 1985–2015.

19. K. K. Lewis (1988). "The Persistence of the 'Peso Problem' When Policy is Noisy," *Journal of International Money and Finance*, 7: 5–21.

20. K. K. Lewis (1989). "Changing Beliefs and Systematic Rational Forecast Errors with Evidence from Foreign Exchange," *American Economic Review*, 79: 621–636.

21. N. G. Mankiw, R. Reis and J. Wolfers (2003). "Disagreement About Inflation Expectations," *NBER Macroeconomics Annual*, 18(1): 209–248.

22. 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.

23. J. F. Muth (1960). "Optimal Properties of Exponentially Weighted Forecasts," *Journal of the American Statistical Association*, 55: 299–306.

24. W. K. Newey and K. D. West (1994). "Automatic Lag Selection in Covariance Matrix Estimation," *The Review of Economic Studies*, 61: 631–653.

25. J. M. Roberts (1997). "Is Inflation Sticky?" *Journal of Monetary Economics*, 39: 173–196.

26. F. Schorfheide (2005). "Learning and Monetary Policy Shifts," *Review of Economic Dynamics*, 8(2): 392–419.

27. J. B. Taylor (1993). "Discretion versus Policy Rules in Practice," *Carnegie-Rochester Conference Series on Public Policy*, 39: 195–214.

28. L. B. Thomas (1999). "Survey Measures of Expected U.S. Inflation," *Journal of Economic Perspectives*, 13: 125–144.

29. T. Yun (1996). "Nominal Price Rigidity, Money Supply Endogeneity, and Business Cycles," *Journal of Monetary Economics*, 37: 345–370.

Table 1

Testing for Unbiasedness (Thomas, 1999)

	$\hat{a}_0$	$\hat{a}_1$	$\chi^2(2)$	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

Table 2

Tests under Complete Information (*Empirical Model*)

Sample size T	80	160	250	1000
$\hat{a}_0$	0.0004	0.0003	0.0001	0.0001
$\hat{a}_1$	0.9663	0.9772	0.9894	0.9982
% times H_0 is rejected	12.3%	10.5%	8.5%	8.1%

Table 3

Tests under Incomplete Information (*Empirical Model*)

Sample size T	80	160	250	1000
$\hat{a}_0$	0.0035	0.0015	0.0008	0.0003
$\hat{a}_1$	0.7004	0.8767	0.9331	0.9868
% times H_0 is rejected	40.8%	28.0%	21.6%	13.7%

Table 4: Business Cycle Statistics

U.S. Economy vs. Benchmark Model

Variable (x)	$\sigma(x)$		$\rho(x, GDP)$		$\rho(x, x_{-1})$	
	Data	Model	Data	Model	Data	Model
GDP	1.60	1.58	1.00	1.00	0.85	0.52
Consumption	1.24	0.90	0.88	0.91	0.85	0.32
Investment	5.00	2.03	0.91	0.73	0.90	0.90
Hours	1.86	0.93	0.89	0.97	0.90	0.48
Inflation	0.49	0.38	0.33	0.66	0.31	0.44
Expected Inflation	0.12	0.28	0.49	0.59	0.55	0.56
TFP	0.74	0.64	0.68	0.56	0.71	0.69

Table 5

Tests Under Complete Information (*Economic Model*)

Sample size T	80	160	250	1000
$\hat{a}_0$	0.0004	0.0002	0.0001	0.0000
$\hat{a}_1$	0.9609	0.9797	0.9900	0.9975
% times H_0 is rejected	14.1%	10.6%	8.9%	6.7%

Table 6				
Tests Under Incomplete Information (*Economic Model*)				
Sample size T	80	160	250	1000
$\hat{a}_0$	0.0013	0.0006	0.0004	0.0001
$\hat{a}_1$	0.8778	0.9436	0.9688	0.9934
% times H_0 is rejected	27.6%	24.2%	19.4%	11.9%

Table 7		
Frequency of Rejections of $\hat{H}_0$		
	% times $\hat{H}_0$ is rejected	% times $\hat{H}_0$ and H_0 are rejected
		$(k=1)$ $\qquad$ $(k=4)$
Complete Information	5.8%	1.2% $\qquad$ 1.2%
Incomplete Information	9.8%	4.4% $\qquad$ 4.7%

Table 8				
Tests under Complete Information (*Economic Model*, $k = 4$)				
Sample size T	80	160	250	1000
$\hat{a}_0$	0.0009	0.0005	0.0003	0.0001
$\hat{a}_1$	0.9138	0.9525	0.9765	0.9937
% times H_0 is rejected	15.6%	9.5%	6.2%	3.3%

Table 9				
Tests under Incomplete Information (*Economic Model*, $k = 4$)				
Sample size T	80	160	250	1000
$\hat{a}_0$	0.0026	0.0013	0.0008	0.0002
$\hat{a}_1$	0.7513	0.8834	0.9376	0.9859
% times H_0 is rejected	31.6%	19.3%	12.0%	3.3%

C:/swp50/temp/graphics/figure1__1.pdf

C:/swp50/temp/graphics/figure2__2.pdf

C:/swp50/temp/graphics/figure3__3.pdf