DNB STAFF REPORTS

David Andolfatto, Scott Hendry and Kevin Moran

Inflation Expectations and Learning about Monetary Policy

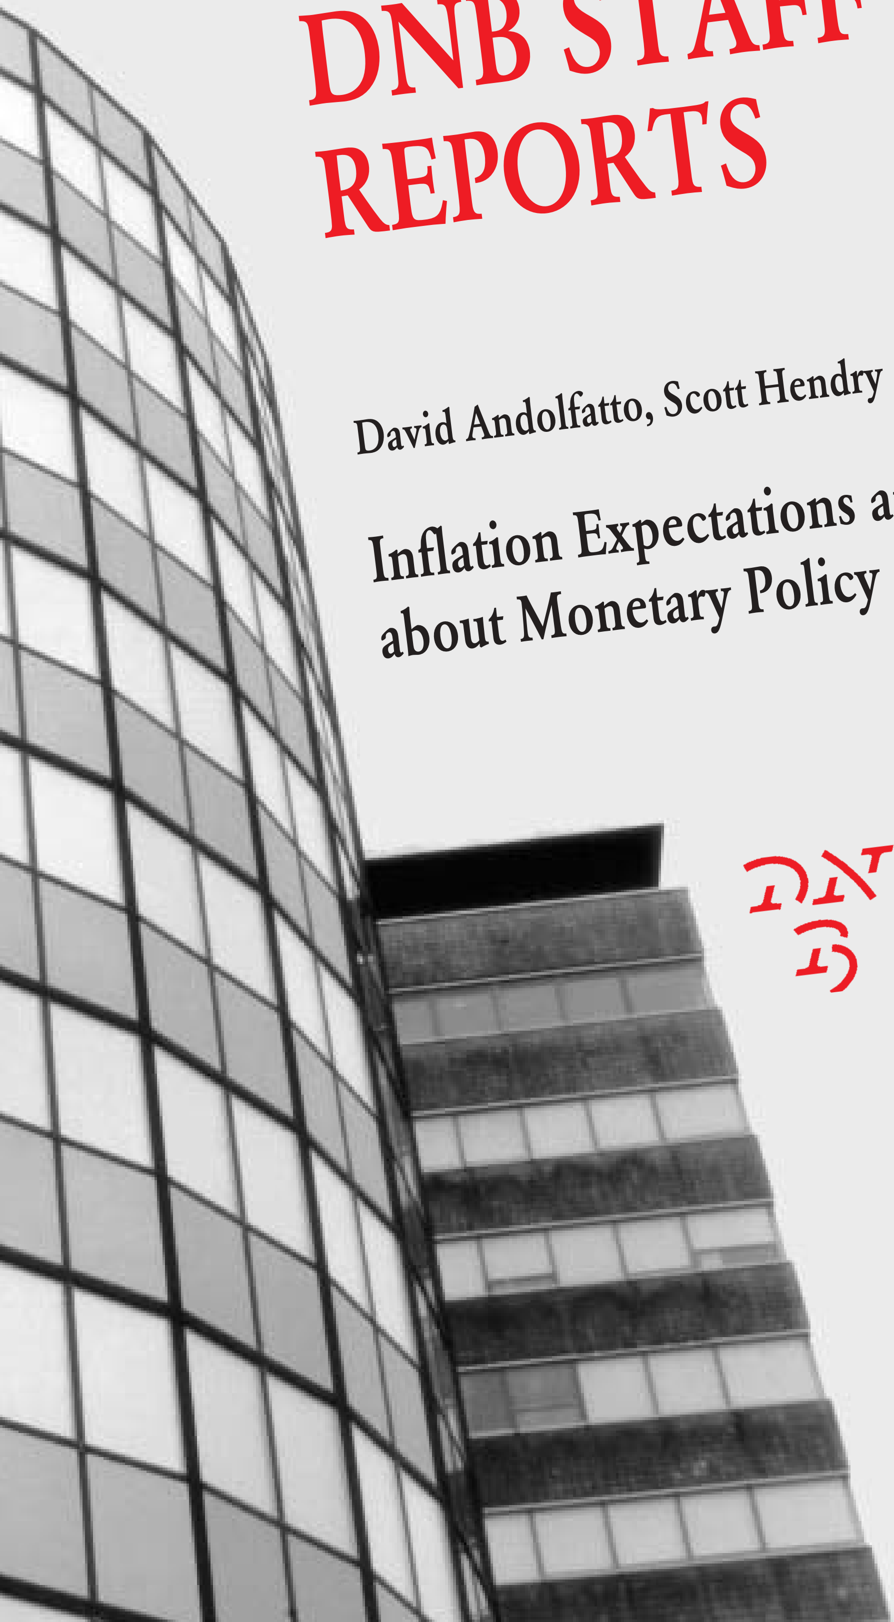

No. 121/2004

Inflation Expectations and Learning about Monetary Policy

David Andolfatto, Scott Hendry and Kevin Moran

DNB Staff Reports 2004, No. 121
De Nederlandsche Bank

Corresponding author: Kevin Moran
e-mail: kmoran@bank-banque-canada.ca

Subscription orders for DNB Staff Reports and requests for specimen copies
should be sent to:

De Nederlandsche Bank NV
Communications Department
Westeinde 1
P.O. Box 98
1000 AB Amsterdam
The Netherlands

Internet: *www.dnb.nl*

Inflation Expectations and Learning about Monetary Policy *

David Andolfatto, *Simon Fraser University*
Scott Hendry, *Bank of Canada*
Kevin Moran, *Bank of Canada*

First Version, February 2002. This Version, October 2003.

Abstract

By various measures, inflation expectations appear to evolve sluggishly relative to actual inflation. Further, they often fail conventional tests of unbiasedness. These observations are sometimes interpreted as evidence against rational expectations.

We embed, within a standard monetary DSGE model, an information friction and a learning mechanism regarding the interest rate targeting rule followed by monetary policy authorities. The learning mechanism enables optimizing economic agents to distinguish between transitory shocks to the policy rule and occasional shifts in the inflation target of monetary authorities.

We show that the model's simulated data is consistent with the empirical evidence. When the information friction is activated, simulated inflation expectations fail conventional unbiasedness tests much more frequently than in the complete information case when this friction is shut down. We interpret these results as suggesting that an important size distortion may be present when conventional tests of unbiasedness are applied to relatively small samples dominated by a few significant shifts in monetary policy and sluggish learning about these shifts.

JEL classification: E47, E52, E58

*The authors wish to acknowledge the expert research assistance of Veronika Dolar. We thank Paul Gomme, Andrew Levin, Césaire Meh and seminar participants at the Bank of Canada, the 2002 CEA meeting, and the 2002 North American Summer Meeting of the Econometric Society for useful comments and discussions. The views expressed in this paper are those of the authors. No responsibility for them should be attributed to the Bank of Canada.

1. Introduction

By various measures, inflation expectations appear to evolve sluggishly relative to actual inflation. Expectations tend to under-predict in periods of rising inflation and over-predict during periods of diminishing inflation.[1] Related to this sluggishness phenomenon is the stylized fact, documented by a sizable empirical literature, that measured inflation expectations often reject the hypotheses of unbiasedness and efficiency[2]. These results have sometimes been interpreted as evidence against rational behavior on the part of economic agents.

This paper assesses whether an information friction over, and a learning mechanism about, the interest rate targeting rule followed by monetary policy, once embedded in a standard monetary Dynamic Stochastic General Equilibrium (DSGE) model, can lead simulated data to replicate quantitatively the empirical evidence against the unbiasedness of inflation expectations.

The information friction we introduce is the following. We assume that the interest rate targeting rule followed by monetary authorities is affected by transitory shocks but also, occasionally, by persistent shifts in the inflation target anchoring the rule. We interpret the transitory shocks in the standard way, as instances when monetary authorities wish to deviate from their rule for a short period, in order to react to financial shocks for example. We envision the occasional shifts in the inflation target as reflecting changes in economic thinking about the optimal inflation rate, or the appointment of a new central bank head with different preferences over inflation outcomes. Importantly, we also assume that these transitory shocks and persistent shifts cannot be separately observed (nor credibly revealed). Consequently, market participants must solve a signal extraction problem to distinguish between the two components, giving rise to a learning rule that shares some features with adaptive expectations processes.[3]

Next, we calibrate the parameters of this signal extraction problem and embed it within the limited participation environment of Christiano and Gust [1999]. We then repeatedly simulate the model and carry out, on the artificial data, unbiasedness tests equivalent to those performed on measured inflation expectations.

Our simulations identify substantial differences in the outcomes of the unbiasedness tests when the information friction is active, relative to the complete information case when it is shut off. Specifically, the fraction of rejections when complete information is assumed never deviates significantly from the level suggested by the size of the tests. In contrast, when the information friction is activated, the tests reject the null hypothesis of unbiasedness much more frequently –between two to fives times– although our model

[1]For example, Dotsey and DeVaro [1995] uncover economic agents' expectations about U.S. inflation –using commodity futures data– over the disinflationary episode of 1980:1–1983:3 and find that expected inflation exceeded actual inflation in all but three periods for the eight-month forecasts and in every period for the one-year forecasts. 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).

[2]See Thomas [1999], Roberts [1997], and Croushore [1997], as well as the references cited in these papers.

[3]Muth [1960] demonstrates that the optimal learning rule in such a signal extraction problem resembles adaptive expectation processes.

embodies the 'rational expectations' solution concept by construction. Interestingly, these differences are much attenuated when the sample size of each simulation is increased significantly.

These results lead us to propose the following interpretation of the empirical rejections of the unbiasedness hypothesis. The process by which economic agents form inflation expectations may be fundamentally sound, but the presence of a few significant shifts in monetary policy, coupled with relatively sluggish learning about these shifts, may lead to significant size distortions of the tests. Further, while this effect may be sufficient to trigger excessive rejections of the null in small samples, it should tend to disappear as the sample size grows.

Environments with information frictions and learning effects similar to the one presented here have been used previously, notably to rationalize the persistent responses of real variables following monetary policy shocks.[4] The present paper makes a two-fold contribution to this literature.

First, we locate the signal extraction problem within an interest rate targeting rule, rather than a monetary growth process. This feature, which our model shares with Erceg and Levin [2003], allows the learning literature to connect with the now-standard view about the proper modelling of monetary policy.

Second, we evaluate the incomplete information and learning framework not by its capacity to generate persistence in the dynamics of real variables, as Erceg and Levin [2003] do, but by its ability to replicate the dynamic relationship that exists between realized and expected inflation. More precisely, we specify parameter values for the underlying components of the monetary policy process and verify whether the rejection of the unbiasedness hypothesis emerges as an implication of these parameter values. By contrast, Erceg and Levin [2003] choose parameter values in order to match the relationship between realized and expected inflation and concentrate their analysis on the implication of their chosen specification for real variables. As such, the analysis presented here complements theirs and broadens the scrutiny over the empirical relevance of incomplete information and learning effects.

The strategy followed in this paper is similar to that employed by Kozicki and Tinsley [2001a], who argue that the frequent empirical rejections of the expectation hypothesis of the term structure could be the result of economic agents learning only gradually about shifts in the objectives of the Federal Reserve. The authors embed a learning mechanism similar to ours in a simple macroeconomic environment and assess whether the expectation hypothesis is rejected by the simulated data. In earlier contributions, Lewis [1988, 1989] uses similar intuition to verify whether sluggish learning can generate the 'forward discount' puzzle observed in foreign exchange market data.

The rest of this paper is structured as follows. Section 2 describes the stylized fact that measured inflation expectations fail simple unbiasedness tests. Section 3 presents the

[4] Recent contributions include Moran [1999], Andolfatto and Gomme [2003], Andolfatto et al. [2003], and Erceg and Levin [2003], who analyze closed economies, and Sill and Wrase [1999] who study an open-economy environment. In an early contribution using a different modelling technology but appealing to very similar ideas, Brunner et al. [1980] analyze the properties of a stochastic IS–LM model in which agents cannot distinguish between permanent and transitory shocks to real and nominal variables.

model used in our simulation, essentially the one of Christiano and Gust [1999]. Section 4 makes precise our view of monetary policy as an interest rate targeting rule affected by two types of disturbances: transitory shocks to the rule and occasional, persistent shifts in the inflation target of monetary authorities. This section also describes the mechanics of the Kalman filter used by economic agents to solve the signal extraction problem and distinguish one component of monetary policy disturbances from the other. Section 5 explains the calibration strategy we utilize. Section 6 describes our Monte Carlo simulations and presents our results. Section 7 discusses our results and concludes.

2. Empirical Evidence on Inflation Expectations

One of the tools commonly used to identify economic agents' inflation expectations is survey data.[5] To illustrate a typical path for such data, Figure 1 depicts the (mean) forecast for one-year-ahead inflation (as measured by the Livingston Survey) as well as the inflation rate that eventually prevailed.[6] The sluggishness discussed above is clear: in times of generally rising inflation, like the 1970s, expected inflation tends to under-predict realized inflation. In contrast, in times of falling inflation like the 1980s and 1990s, the forecasts appear to over-predict inflation.

Several studies examine the statistical properties of such inflation expectations, with the objective of testing for departures from rationality. Such departures, usually identified as rejections of unbiasedness and efficiency, appear to be a common conclusion of this literature.[7] The unbiasedness tests are typically conducted by testing $H_0 : a_0 = 0;\ a_1 = 1$ using the following simple regression equation:

$$\pi_t = a_0 + a_1 \pi_t^e + \epsilon_t, \tag{1}$$

where π_t is the net, annualized rate of inflation from period t-k to period t and π_t^e is the expectation of π_t formed at time t-k.

We identify the rejection of unbiasedness, defined using (1), as the stylized fact that the model should replicate. For illustrative purposes, we reproduce below one of Thomas [1999]'s regressions. Run with data from the Livingston Survey, it shows the following estimated equation:

$$\begin{aligned} \pi_t = \underset{(0.41)}{0.134} + \underset{(0.08)}{0.88} E_{t-2}[\pi_t], \end{aligned} \tag{2}$$

[5] Other methods include uncovering inflation expectations from futures market data [Dotsey and DeVaro, 1995] or comparing yields on inflation-indexed and non-indexed Treasuries [Shen and Corning, 2001].

[6] The Livingston survey was started by J. A. Livingston, a business journalist in the Philadelphia area and is now maintained by the Federal Reserve Bank of Philadelphia. Croushore [1997] describes the history of the survey and its current structure. Other survey data on inflation expectations include those arising from 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. Thomas [1999] discusses the three surveys. Note finally that the Conference Board of Canada also has produced, since 1988, survey data on (Canadian) inflation expectations.

[7] Thomas [1999] conducts unbiasedness and efficiency tests on the three sources of survey data. Croushore [1997] reviews the tests conducted over the years on the Livingston data.

Figure 1: Realized Inflation versus Measured Expectations (Livingston Survey)

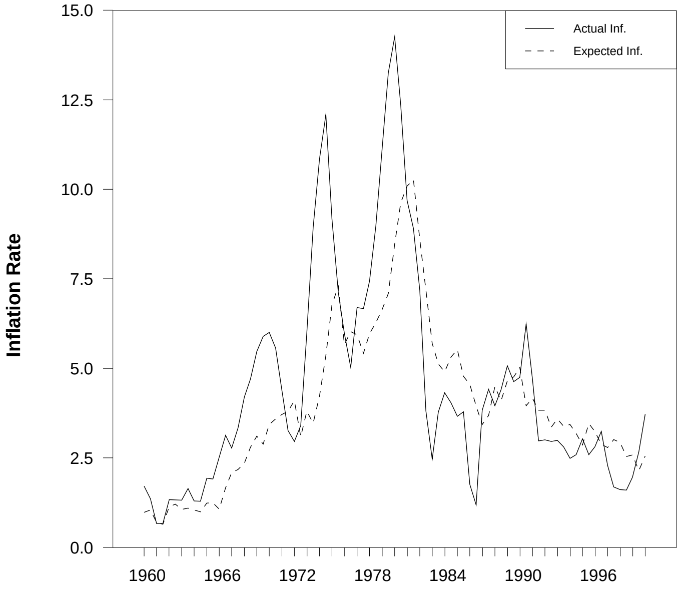

where the sample used is 1980:3 to 1997:4. The data are of semi-annual frequency and the expectations have a one-year-ahead horizon (two periods).[8] These estimation results lead to a rejection of H_0.[9]

Note that the rejection of H_0 can be overturned in some large samples, where the positive forecasting errors of the 1970s appear to cancel the mainly negative ones of the 1980s. This suggests that the rejections of the hypothesis of unbiasedness could simply be due to a small sample problem. As mentioned in the introduction and discussed below in Section 6, this is precisely what our results imply.

Interestingly, once defined with a quasi-difference specification, the unbiasedness hypothesis continues to be rejected in large samples.[10] Future research might investigate

[8] This frequency is not standard across all sources of inflation expectations data. In the model developed below, a period corresponds to one quarter and we report simulation results obtained with one-quarter-ahead and four-quarter-ahead expectations.

[9] Note that a correction for serial correlation in the residuals must be introduced when constructing the test statistic. Thomas [1999] reports the results of estimating (1) on other samples and with alternative measures of inflation expectations. On balance, the evidence points to rejections of the null hypothesis, especially when the sample considered is small. A similar regression run with the Canadian data on inflation expectations, over the sample running from 1988Q1 to 2001Q1, yields the following estimate:

$$\pi_t = \underset{(0.29)}{0.29} + \underset{(0.11)}{0.77} E_{t-4}[\pi_t]$$

with, again, an easy rejection of H_0.

[10] Consider estimating the following regression:

$$\pi_{t+k} - \pi_t = a_0 + a_1(E_t[\pi_{t+k}] - \pi_t) + u_t,$$

and testing $H_0:\ a_0 = 0,\ a_1 = 1$. Observe that under the null, both this regression and (1) are identical.

whether this facet of the relationship between realized and expected inflation could be replicated by our incomplete information and learning environment.

3. The Model

The model used in this paper is very similar to the one in Christiano and Gust [1999]. We therefore provide only an overview and refer interested readers to the original paper. Note that the main nominal rigidity appearing in the model is the assumption of *limited participation*, one of the standard ways of introducing monetary non-neutralities in a DSGE model. In contrast, Erceg and Levin [2003] use nominal price and wage stickiness to achieve this non-neutrality. Note that we could redo our analysis with these nominal rigidities but the robustness of the results we obtain, documented in Section 6.4 below, suggests that using nominal price or wage rigidity would not alter our main conclusions.

3.1 Households

The model economy comprises a continuum of identical, infinitely-lived households. At the start of every period, a household's wealth consists of k_t units of capital, M_t^c units of liquid financial assets, and M_t^d units of illiquid assets (deposited at a financial intermediary).[11] During the course of the period, households rent their capital to firms, allocate their time between work and leisure, choose desired levels of consumption and investment and, finally, choose how to allocate their financial assets into the cash and deposits they will carry over to the next period.

The purchase of consumption and investment goods must be carried out with liquid assets. Available liquid assets consist of beginning-of-period balances (M_t^c) and wage payments. This assumption leads to the following liquidity constraint:

$$P_t c_t + P_t(k_{t+1} - (1-\delta)\, k_t) \leq M_t^c + W_t n_t, \tag{3}$$

where c_t is per-household consumption, $(k_{t+1} - (1-\delta)k_t)$ is investment, P_t is the nominal price of goods, W_t is the nominal wage, and n_t is labor supply.

At the end of the period, households receive their capital rental income and return on deposits. These revenues, combined with any liquid assets remaining from their goods

Nevertheless, Dolar and Moran [2002] report that the evidence against the null is much more robust using this alternative regression. Further, the estimates of a_1 arising from this regression are almost always between zero and one.

Note that this specification is very similar to the one often used to document the forward discount puzzle:

$$e_{t+1} - e_t = b_0 + b_1(f_t - e_t) + u_t,$$

with e_t the spot exchange rate at time t and f_t the forward exchange rate. Researchers have often proposed learning effects as one potential explanation of the very frequent empirical rejections of $H_0: \ b_0 = 0, \ b_1 = 1$ one obtains from this regression. See Froot and Thaler [1990] and Taylor [1995] for a discussion.

[11] In what follows, lower case variables k_t and n_t express the levels of capital and work supplied by the households; upper case variables K_t and N_t represent the quantities of these variables demanded by firms. Moreover, M_t^d and M_t^c express households' holdings of liquid assets while $\overline{M_t}$ denotes the total supply of such assets.

purchases, sum up to their end-of-period financial wealth, which is allocated between next period's liquid and illiquid assets. The following budget constraint arises:

$$M_{t+1}^d + M_{t+1}^c \leq r_{kt}k_t + R_t^d M_t^d + (M_t^c + W_t n_t - P_t c_t - P_t(k_{t+1} - (1-\delta)k_t)), \qquad (4)$$

where r_{kt} is the rental rate on capital and R_t^d the return on illiquid assets.

Households choose a plan for consumption, investment, labor supply, and financial asset allocation in order to maximize their lifetime utility. Hence, they solve the following problem:

$$\max_{[c_{t+k}, n_{t+k}, k_{t+k+1}, M_{t+k+1}^c, M_{t+k+1}^d]|_{k=0}^{\infty}} E_t \sum_{k=0}^{\infty} \eta^{t+k} U(c_{t+k}, n_{t+k} + AC_{t+k}) \qquad (5)$$

where $U(.,.)$ is the period utility function, η the time discount, and where the maximization is done with respect to (3),(4) and initial levels k_t, M_t^c, and M_t^d.

Note that the term $n_{t+k} + AC_{t+k}$ represents the time costs of market activities, in terms of leisure foregone. The term AC_t represents the costs households must incur to adjust their liquid asset portfolios. The functional form selected for these costs is as follows:[12]

$$AC_t = \tau(M_{t+1}^c / M_t^c - \mu)^2, \qquad (6)$$

with μ the steady-state growth rate of the total supply of liquid assets.

3.2 Firms

Firms combine labor and capital inputs to produce the economy's output. They have access to the following constant returns to scale production function:

$$Y_t = A_t K_t^{\theta} N_t^{1-\theta},$$

where A_t denotes a transitory productivity shock that evolves according to the following process:

$$A_t = (1 - \rho_A)A + \rho_A \cdot A_{t-1} + \nu_t^A, \;\; \nu_t^A \sim N(0, \sigma_A^2). \qquad (7)$$

Firms rent capital and hire labor to maximize per-period profits. Since firms pay their capital rental expenditures directly from revenues, the first-order condition for the choice of capital is the familiar one:

$$r_{kt} = \theta A_t (K_t/N_t)^{\theta-1}. \qquad (8)$$

In contrast, it is assumed that a given fraction (denoted $1 - J_t$) of the firms' wage costs must be paid in advance. In order to do so, firms must borrow the necessary funds from financial intermediaries at the rate R_t^l. This assumption leads to the following first-order condition for labor demand:[13]

$$((1 - J_t)R_t^l + J_t)W_t/P_t = (1 - \theta)A_t(K_t/N_t)^{\theta}. \qquad (9)$$

[12] Expressing these costs in terms of leisure rather than goods is not important for the results.

[13] Note that because the production function features constant returns to scale, these efficiency conditions also hold for the aggregate values of capital and labor demand in the economy and, hereafter, K_t and N_t represent those aggregate quantities.

The evolution of J_t (which we refer to as a money demand shock) is exogenous and obeys the following:

$$J_t = \rho_J J_{t-1} + \nu_t^J, \;\; \nu_t^J \sim N(0, \sigma_J^2). \tag{10}$$

3.3 Financial Intermediaries

Financial intermediaries accept deposits from households and lend the receipts to firms. Further, they are the recipients of any injection of liquid assets that the central bank engineers to support its monetary policy rule.[14] The revenues of the intermediaries are thus the total amount lent multiplied by the lending rate while their expenses are the total deposits received multiplied by the deposit rate. Profits are thus:

$$R_t^l B_t - R_t^d (M_t^d + X_t),$$

where B_t is total lending and X_t represents injections of liquid assets from the central bank. The assumption of perfect competition in the financial sector ensures that, in equilibrium, $R_t^l = R_t^d \equiv R_t$.

3.4 Equilibrium

An equilibrium for this artificial economy is defined as follows. It consists of a vector of allocations $(c_{t+k}, n_{t+k}, k_{t+k+1}, M_{t+k+1}^c, M_{t+k+1}^d, N_{t+k}, K_{t+k}, B_{t+k})|_{k=0}^{\infty}$, of prices $(P_{t+k}, W_{t+k}, R_{t+k}, rk_{t+}$ of exogenous variables $(A_{t+k}, J_{t+k}|_{k=0}^{\infty})$, and of starting values (k_t, M_t^c, M_t^d). These allocations, prices, exogenous variables, and starting values are such that households maximize lifetime utility, as stated by (5), firms and financial intermediaries maximize profits, and, finally, the following market-clearing equilibrium conditions are met:

$$c_t + K_{t+1} - (1-\delta)K_t = Y_t;$$

$$M_t^c + M_t^d = \overline{M_t};$$

$$M_t^d + X_t = B_t = W_t N_t;$$

$$N_t = n_t;$$

$$K_t = k_t.$$

4. Monetary Policy

4.1 The Monetary Policy Rule

Monetary authorities target the nominal interest rate. This targeting is made precise by assuming that the desired nominal rate, denoted i_t^*, is the following function of macroeconomic conditions:

$$i_t^* = r^{ss} + \pi_t^T + \alpha(\pi_t - \pi_t^T) + \beta y_t, \tag{11}$$

[14] They cannot profit from these injections, however, because it is assumed these injections are deposited in the households' accounts.

where r^{ss} is the steady-state value of the real interest rate, π_t^T is the inflation target of monetary authorities at time t, and y_t represents the output gap. Note that $r^{ss} + \pi_t^T$ represents the steady-state value of the nominal rate; (11) thus signifies that monetary authorities will increase rates relative to steady-state when price pressures threaten to push inflation over the current target or when the output gap is positive.

It is often conjectured that, instead of rapidly moving the nominal rate to reach the targeted level, monetary authorities implement gradual changes in rates that only eventually converge to that level. Such a smoothing motive can be represented mathematically by assuming that the actual rate implemented by the central bank will be the following weighted sum of the targeted rate and the preceding period's rate;

$$i_t = (1-\rho)i_t^* + \rho i_{t-1}, \tag{12}$$

where the coefficient ρ governs the extent of smoothing exercised by the monetary authorities.

Finally, monetary authorities regularly deviate from their rule. These deviations (discussed further below) are referred to as monetary policy shocks and are denoted by the variable u_t. Combining equations (11) and (12), as well as introducing the u_t shocks, leads us to the following characterization of monetary policy:

$$i_t = (1-\rho)[r^{ss} + \pi_t^T + \alpha(\pi_t - \pi_t^T) + \beta y_t] + \rho i_{t-1} + u_t. \tag{13}$$

It should be emphasized that the instrument by which monetary authorities implement the rule (13) remains the growth rate of money supply. The significance of this rule is that the central bank manipulates this growth rate ($\mu_t = \frac{\overline{M_{t+1}}}{\overline{M_t}}$), in a manner such that the observed relationship between nominal rates, inflation, and output that emerges obeys (13).

4.2 Monetary Policy Shocks and Monetary Policy Shifts

We envision that monetary policy, as expressed by the interest rate rule in (13), is subject to two types of disturbances. The first disturbance consists of the monetary policy shocks mentioned above (the variable u_t). We interpret these disturbances as the reaction of monetary authorities to economic factors, such as financial stability concerns, not articulated in the rule (13). Alternatively, these shocks could be understood as errors stemming from the imperfect control exercised by monetary authorities over the growth rate of money supply (μ_t). Under either interpretation however, we envision these shocks as possessing little persistence. Accordingly, we assume that their evolution is governed by the following process:

$$u_{t+1} = \phi_1 u_t + e_{t+1}, \tag{14}$$

with $0 \leq \phi_1 << 1$ and $e_{t+1} \sim N(0, \sigma_e^2)$.

The second disturbance to monetary policy is as follows. We assume that, while remaining constant for extended periods of time, the inflation target of the monetary authorities π_t^T 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 Philips 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 head, whose preferences over inflation outcomes differ from their predecessor's. Under either interpretation, we envision these shifts to exhibit significant duration, in the order of five to ten years, say.

Mathematically, we express these shifts in the inflation target by the variable $z_t \equiv \pi_t^T - \pi_0^T$, so that z_t constitutes the deviation of the current target of authorities (π_t^T) from its long term unconditional mean (π_0^T). We assume that the following process, a mixture of a Bernoulli trial and a normal random variable, expresses how z_t evolves over time:

$$z_{t+1} = \begin{cases} z_t & \text{with probability } \phi_2; \\ g_{t+1} & \text{with probability } 1 - \phi_2, \quad g_{t+1} \sim N(0, \sigma_g^2). \end{cases} \tag{15}$$

with $0 << \phi_2 < 1$. Notice that, in some ways, the process for z_t shares some similarities with a random walk specification. Specifically, with a high value of ϕ_2 the conditional expectation of z_{t+1} is close to z_t. In contrast with a random walk however, the process is not affected by innovations every period and is, ultimately, stationary.

On the other hand, the process differs from a standard autoregressive process in that the decay of a given impulse will be sudden and complete, rather than the gradual decline characteristic of autoregressive processes. 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.[15]

Below, we refer to some model simulations as representing complete information. By this we shall mean that economic agents can observe the exact decomposition of monetary policy disturbances between their z_t and u_t parts. In such a case, although uncertainty remains (it arises from the innovations e_{t+1} and g_{t+1}), agents have sufficient information to compute the correct conditional expectations concerning future monetary policy.

4.3 Incomplete Information and Learning

Credibly communicating shifts in the inflation target might be difficult for monetary authorities. For example, although a new central bank head with a strong aversion for inflation might indicate this aversion in public announcements, economic agents may be uncertain as what these announcements mean for the quantitative inflation targets. As a result, they might treat these announcements with scepticism and only modify their beliefs about the inflation target of monetary authorities once several periods of lower inflation have been observed. Announcements of explicit, quantitative changes in the inflation

[15] It is a question left for further research to determine how much difference it would make, in practice, to model the shifts as arising from a random walk process with very low innovation variance.

target might suffer, at least initially, from similar credibility problems.[16] Alternatively, central banks sometimes do not make explicit announcements about their inflation target but let economic agents decipher as best they can announcements of a more general nature.

To capture the spirit of this information problem, we assume that the z_t shifts are unobservable to economic agents. They only observe a mixture of the z_t shifts and the u_t shocks. Agents thus face a signal extraction problem that is solved using the Kalman filter.

Recalling (13), assume that at time t, the long-run inflation target is changed from its unconditional mean of π_0^T to π_t^T. Assume also, for notational purposes, that the response to the output gap –the coefficient β– is zero. The rule is thus:

$$i_t = (1-\rho)[r^{ss} + \pi_t^T + \alpha(\pi_t - \pi_t^T)] + \rho i_{t-1} + u_t. \tag{16}$$

Now, rewrite (16) by adding and subtracting π_0^T two times:

$$i_t = (1-\rho)[r^{ss} + \pi_t^T + (\pi_0^T - \pi_0^T) + \alpha(\pi_t - \pi_t^T + \pi_0^T - \pi_0^T)] + \rho i_{t-1} + u_t,$$

or, rearranging terms,

$$i_t = (1-\rho)[r^{ss} + \pi_0^T + \alpha(\pi_t - \pi_0^T)] + \rho i_{t-1} + \underbrace{(1-\rho)(1-\alpha)(\pi_t^T - \pi_0^T)) + u_t}_{\epsilon_t^*}. \tag{18}$$

Equation (18) illustrates the following point. From the point of view of an economic agent whose initial belief about the inflation target of monetary authorities was π_0^T, the observed shock to the policy rule (ϵ_t^*) is a combination of a persistent shift $(1-\rho)(1-\alpha)(\pi_t^T - \pi_0^T)$ and of the transitory disturbance to the rule u_t. The signal extraction problem that economic agents face thus entails separating ϵ_t^* into its persistent and transitory components. Next, given their knowledge of the rule and its parameters (α and ρ), agents can back-out an estimate of $\pi_t^T - \pi_0^T$, the shift in the inflation target.[17]

The signal extraction problem is carried out using the Kalman filter. Notice that the evolution of ϵ_t^*, the observed shock to the monetary policy rule in (18), can be expressed within the following system:

$$\begin{bmatrix} z_{t+1} \\ u_{t+1} \end{bmatrix} = \begin{bmatrix} \phi_2 & 0 \\ 0 & \phi_1 \end{bmatrix} \cdot \begin{bmatrix} z_t \\ u_t \end{bmatrix} + \begin{bmatrix} N_{t+1} \\ e_{t+1} \end{bmatrix}; \tag{19}$$

$$\epsilon_t^* = \begin{bmatrix} (1-\rho)(1-\alpha) & 1 \end{bmatrix} \cdot \begin{bmatrix} z_t \\ u_t \end{bmatrix}; \tag{20}$$

[16] Even after such announcements are made and credibility is largely established, substantial uncertainty over the weight attributed by the central bank to inflation outcomes within a targeted range might still remain. Ruge-Murcia [2003], for example, argues that contrary to stated weights, the inflation outcomes of the 1990s in Canada are consistent with asymmetric preferences of the Bank of Canada over its official target range.

[17] Note that a different type of learning could also be modelled. Agents could be thought of having imperfect knowledge about the *coefficients* of the rule (α, β, and ρ) and learn about these shifts by repeated observations of the interest rate changes engineered by monetary authorities. Empirical estimation of Taylor-type monetary policy rules have identified structural shifts in the parameters of such rules around 1980. See Clarida et al. [2000], for example. We plan to pursue the implications of this type of imperfect information in future work.

where N_{t+1} is defined as follows:

$$N_{t+1} = \begin{cases} (1-\phi_2)\, z_t, & \text{with probability } \phi_2; \\ g_{t+1} - \phi_2 z_t, & \text{with probability } 1-\phi_2. \end{cases}$$

Notice that, under the definition of N_t, $E_t[N_{t+1}] = 0$. The fact that $E_t[e_{t+1}] = 0$ was already assumed in equation (14).

Equations (19) and (20) define a state-space system [e.g. Hamilton, 1994, ch. 13], where (19) is the state equation and (20) the observation equation. When applied to such a system, the Kalman filter delivers forecasts of the two unobserved states (z_t and u_t) conditional on all observed values of ϵ_t^*. We assume that economic agents know the value of all the parameters of the problem so that the forecasts arising from the filter are feasible.

The projections underlying the Kalman filter are updated sequentially and represent the best linear forecasts of the unobserved variables based on available information. Further, if the variables in the dynamic system are normal, the forecasts arising from the filter are optimal.[18] Denote the forecasts of the two unobserved variables, given the information available at time t, as $\widehat{z}_{t|t}$ and $\widehat{u}_{t|t}$. Equations (19) and (20) can then be used to compute expected, future deviations of the interest rate from the benchmark rule, as follows:

$$\begin{bmatrix} \widehat{z}_{t+1}|_t \\ \widehat{u}_{t+1}|_t \end{bmatrix} = \begin{bmatrix} \phi_2 & 0 \\ 0 & \phi_1 \end{bmatrix} \cdot \begin{bmatrix} \widehat{z}_t|_t \\ \widehat{u}_t|_t \end{bmatrix}; \tag{22}$$

$$E_t\left[\epsilon_{t+1}^*\right] = \begin{bmatrix} (1-\rho)(1-\alpha) & 1 \end{bmatrix} \cdot \begin{bmatrix} \widehat{z}_{t+1}|_t \\ \widehat{u}_{t+1}|_t \end{bmatrix}. \tag{23}$$

Additional details on our implementation of the Kalman filter, which simply requires us to establish a formal correspondence between our notation and that in Hamilton [1994], are provided in Appendix 1.

Note that the information friction we assume is stronger, in some sense, than others often used in the literature, notably by Andolfatto and Gomme [2003]. These authors assume that the 'regime' part of monetary policy can take only a finite number of values (usually two). Such a restriction simplifies the learning problem of economic agents and usually produces quick transition of the beliefs following regime shifts. However, we consider that these 'two-point' learning problems understate the severity of the information friction over monetary policy faced by real-world economic agents.[19]

5. Calibration and Solution of the Model

There are three distinct areas of the model that require calibration: the model itself (preferences, technology, etc.), the parameters of the interest rate rule (13) and, finally,

[18] Examination of (15) shows that, conditional on the value of z_t, z_{t+1} is not normally distributed. But 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 the high values of ϕ_2 used in our calibration, however, this unconditional, normal behavior will only appear after a very large number of data have been observed.

[19] See Kozicki and Tinsley [2001b], page 165, for a similar argument.

the parameters governing the evolution of the shocks and shifts in the rule. Note that the model period corresponds to one quarter.

5.1 Preferences and Technology

The first part of the calibration exercise is straightforward as we adopt most of the choices made in Christiano and Gust [1999]. For example, the utility function is specialized to be:

$$U(c_t, n_t + AC_t) = \log[c_t - \psi_0 \frac{(n_t + AC_t)^{1+\psi_1}}{1+\psi_1}].$$

Note that under this specification of utility, no intertemporal smoothing motive is present in labor supply; the only factor affecting the decision of households is the real wage, with an elasticity of $1/\psi_1$.[20] We choose ψ_1 so that the elasticity is 2.5. The parameter ψ_0 is mainly a scale parameter and we fix its value to 2.15, which implies a steady-state value of around 1.0 for employment.

The parameter τ expresses the severity of the portfolio adjustment costs. We fix its value to 10.0, which, in a version of the model with monetary policy expressed as an exogenous process for money growth, generates sizable persistence following a monetary policy shock.

Other parameters governing preferences and technology appear in most models and standard values for their calibration are established: we thus fix η to 0.99, θ to 0.36, and δ to 0.025.

Below, we conduct two types of experiments, with regard to our assumption about the technology (A_t) and money demand (J_t) shocks. We first envision a world where disturbances to the monetary policy rule (u_t and z_t) are the only source of volatility. In such a world, we fix technology and money demand at their long run mean so that $A_t \equiv 1$ and $J_t \equiv 0$, $\forall t$.

In a second type of experiment, we want to add these two extra sources of volatility into the model. We thus reintroduce the technology shocks by using the familiar values of 0.95 for ρ_A and 0.005 for σ_A. Since no similar values are established for the money demand shock, we follow Christiano and Gust [1999] and apply the technology shock values to the process for J_t: we thus have $\rho_J = 0.95$ and $\sigma_J = 0.005$.

The model is solved using the first order approximation method and algorithms presented in King and Watson [1998]. Details on the solution method are available from the authors upon request.

5.2 Parameters of the Interest Rate Targeting Rule

According to the rule in (13), current interest rates are determined by the deviation of inflation from its current target (with a coefficient α), by the output gap (β), and by its own lagged values (ρ).

The evidence about the correct values for these coefficients is not precise, particularly because empirical studies of interest rate targeting rules [Taylor, 1993, Clarida et al., 2000,

[20]See Greenwood et al. [1988] for further details.

Nelson, 2000] often use specifications of (13) that, although similar in spirit to the one used here, differ in the details of the timing assumptions and definitions used. Further, some values of the triple (α, β, ρ) lead to non-uniqueness (or non existence) of stable equilibria in the model.[21]

We thus use such empirical evidence to suggest a range of reasonable values for the parameters and conduct a sensitivity analysis of our results to different values within that range. For example, many empirical studies report evidence that the behavior of monetary authorities is consistent with significant smoothing of interest rate changes. We thus use a range of [0, 0.5] for the parameter ρ. To ensure uniqueness of equilibria, we must fix the coefficient describing the response to inflation, α, to a relatively high value. We thus explore values in the range [2.0, 4.0] for that parameter. Finally, the same requirement of uniqueness suggests relatively low values for the response to the output gap β. We thus use a range of [0, 0.25]. Our benchmark specification sets $\alpha = 2.0$, $\rho = 0.25$, and $\beta = 0.25$.[22]

5.3 Shifts and Shocks to Monetary Policy

We now discuss the calibration of the processes governing the evolution of the shocks (the u_t variables) and the shifts (the z_t variables) in monetary policy.

Recall that ϕ_2 and σ_g govern the dynamics of the z_t variable. These parameters respectively express the expected duration of a particular regime and the standard deviation of the distribution from which the value of a regime shift, when one occurs, is drawn. ϕ_1 and σ_e denote the autocorrelation and innovation variance of the u_t shocks.

The interpretations suggested above for the shifts in the variable z_t, changes in economic thinking or appointments of new central bank heads, suggest that these shifts occur only infrequently, perhaps once every five or ten years. Transposed to the quarterly frequency we use, this corresponds to one shift, on average, every 20 to 40 periods. Such an average duration between shifts corresponds to values of ϕ_2 between 0.95 and 0.975. We use the slightly wider range of [0.95, 0.99] for ϕ_2, with 0.975 as the benchmark value.

Calibrating the standard deviation of the innovation in regime shifts, σ_g, is less straightforward. In our benchmark specification, we set it to 0.005, which implies that when a one standard deviation shift does occur, it corresponds to a change of 2%, on an annualized basis, in the inflation target of monetary authorities. We also explore the consequences of lower (0.0025) and higher (0.01) values for this parameter.

Note that one interpretation of the Romer and Romer [1989, 1994] dates is that they represent changes in the inflation target of the Federal Reserve, and as such, occurrences of z_t shifts.[23] Since there are seven such dates identified over a 40 year sample, this would

[21] See Christiano and Gust [1999] for a detailed examination of the ranges of values for which non-uniqueness obtains.

[22] Note that we define the output gap as the deviation of current output from its steady-state value. This is an incorrect definition, particularly in the presence of technology shocks that modify potential output significantly. A better measure of the output gap results when potential output is defined as the level of output that would obtain in a version of the model where all nominal frictions have been removed.

[23] In these papers, the authors analyze the minutes of FOMC deliberations, and identify dates at which the Federal Reserve Board decided to cause a recession in order to stop inflationary pressures. See Hoover

correspond to an expected duration of five to six years (or 20 to 25 quarters) for these shifts, placing the duration parameter within the range we use.[24]

Table 1. Parameter Calibration

Parameter	Symbol	Benchmark Value	Range Examined
		Preferences and Technology	
Elasticity of Labor Supply	ψ_1	0.4	-
Scaling of Labor Supply	ψ_0	2.15	-
Portfolio Adjustment Costs	τ	15	-
Discount Factor	η	0.99	-
Capital Share in Production	θ	0.36	-
Capital Depreciation Rate	δ	0.025	-
		Interest Rate Targeting Rule	
Response to Inflation	α	2.0	[2.0, 4.0]
Response to Output Gap	β	0.25	[0, 0.25]
Smoothing of Interest Rates	ρ	0.25	[0, 0.5]
		Shifts and Shocks to Monetary Policy	
Duration of Shifts	ϕ_2	0.975	[0.95, 0.99]
Standard Deviation of Shifts	σ_g	0.005	[0.0025 0.01]
Persistence in Shocks	ϕ_1	0.1	[0, 0.2]
Standard Deviation of Shocks	σ_e	0.005	[0.0025, 0.005]

Turning now to the calibration of the transitory shocks u_t, we simply use a range of [0, 0.2] for the autocorrelation parameter ϕ_1, with 0.1 as the benchmark value. We set the benchmark value of the variance of the innovations to these shocks, σ_e, to 0.005, in a symmetric way with the variance of the regime shifts, and experiment with lower values. Because u_t is equivalent to the interest rate shock in the monetary policy rule, a one-standard deviation value of 0.005 corresponds to a 2% innovation in the (annualized) rate. Considering that central banks usually change interest rates by much lower increments, a value of 0.005 for σ_e is probably an upper bound. The calibration values we use are summarized in Table 1.

Figure 2 illustrates the impact of the information friction (in our benchmark calibration) following a negative, one-standard deviation shift in z_t. Again, this shift corresponds to a decrease in the inflation target from 5% to 3%. The true z_t, along with agents' best estimates of that variable, appear in Panel A of the Figure.

Following the shift, economic agents assign some weight to the possibility that the observed disturbance to monetary policy was a regime shift and thus Panel A of the Figure shows that the agents' best estimate of z_t starts to decline towards the true value. However, agents also assign some weight to the possibility that the observed disturbance was a transitory shock u_t. Thus the middle panel of the graph shows that agents' best

and Perez [1994] and Leeper [1997] for a discussion of the methodology and results of Romer and Romer.

[24] Other results with which we could match our calibration of the z_t shifts are those in Owyang and Ramey [2001], where the authors identify shifts in the preferences of monetary authorities over inflation, within an empirical expression of the classic Barro and Gordon [1983] model.

Figure 2: The Learning Mechanism: The Benchmark Calibration

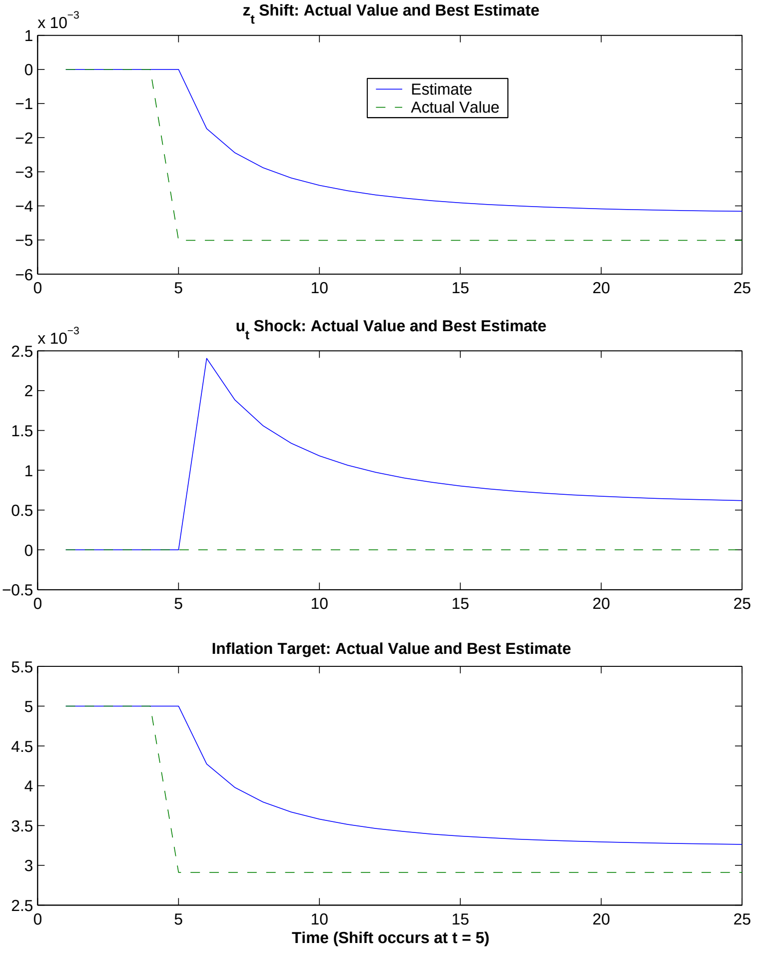

estimate of u_t rises after the initial period. Eventually, because the shift is persistent, agents doubt more and more that it might have come from the largely transitory u_t and become convinced it must have come from a z_t shift. Accordingly, the agents' best estimate of z_t and u_t respectively converge towards the true value and to zero.

Finally, the bottom panel of the Figure shows the progress of the agents' estimate of the annualized inflation target of the monetary authorities (implied by their estimates of z_t: recall the definition of ϵ_t^* in equation (18)). The panel shows that beliefs smoothly converge towards the true value of 3%.

6. Monte Carlo Simulation of the Model

6.1 Impulse Responses Following a Regime Shift

To develop intuition about the Monte Carlo results that follow, Figure 3 presents the impulse responses of the artificial economy following a shift in the inflation target of the monetary authorities. The shift is identical to the one illustrated in Figure 2: the inflation target is lowered, at time $t = 5$, from 5% on an annualized basis to 3%.

In the first eight panels of the graph, the solid lines express the case where agents have complete information about the shift. In contrast, the dashed lines relate the case when the information friction is active and the learning mechanism discussed above governs the formation of expectations.

The solid lines demonstrate that, following the implementation of the monetary policy shift, a very short downturn affects the economy: consumption, output and employment shrink for only one or two periods. Very rapidly, the positive, long-term effects of the decrease in inflation begin to take hold and all real variables increase, passing their initial levels and converging towards a higher steady-state. The dashed lines indicate that in the incomplete information case, this process takes several periods to firmly establish itself, a period during which all real aggregates are lower than they were in the full information case. This arises because economic agents assign a positive probability to the interest rate shock being a transitory hike in interest rates, with straightforward negative effects on the real economy.

The bottom two panels of Figure 3 illustrate the situation from a slightly different angle. The solid lines now depict realized inflation, and the dashed lines, expected inflation. The bottom-left panel refers to the complete information case. The graph shows that, apart from the initial surprise in the first period of the shock, economic agents have the correct inflation expectations. Now turn to the bottom-right graph, which depicts the incomplete information case. It shows inflation expectations lagging actual inflation for several periods, before converging. Notice that this feature replicates, in a qualitative fashion, the behavior of inflation expectations during the 1980s, as shown in Figure 1. In that figure, during a period of generally decreasing inflation, expectations –as measured by the Livingston Survey– over-predicted actual numbers for several quarters.

The gap between realized and expected inflation in the bottom-right graph of Figure 3 is a direct result of the learning behavior discussed in the preceding section. Initially, agents assign some weight to the possibility that the observed monetary disturbance was a transitory disturbance to the rule. They therefore do not expect it to last and think inflation might return to the initial, higher level of 5%. Over time, agents become convinced that a shift has indeed occurred and their expectations converge to values closer to the actual ones.

In the simulations carried out with the model, transitory shocks occur simultaneously with the shifts in the inflation target; therefore, a picture of the artificial economy's responses will not be very informative. However, the main idea of the graphs in Figure 3 remains: when a shift occurs, agents are likely to underestimate them for some time, and inflation expectations are likely to erroneously predict actual inflation for some time. It

Figure 3: Complete and Incomplete Information Responses to a Shift in the Inflation Target

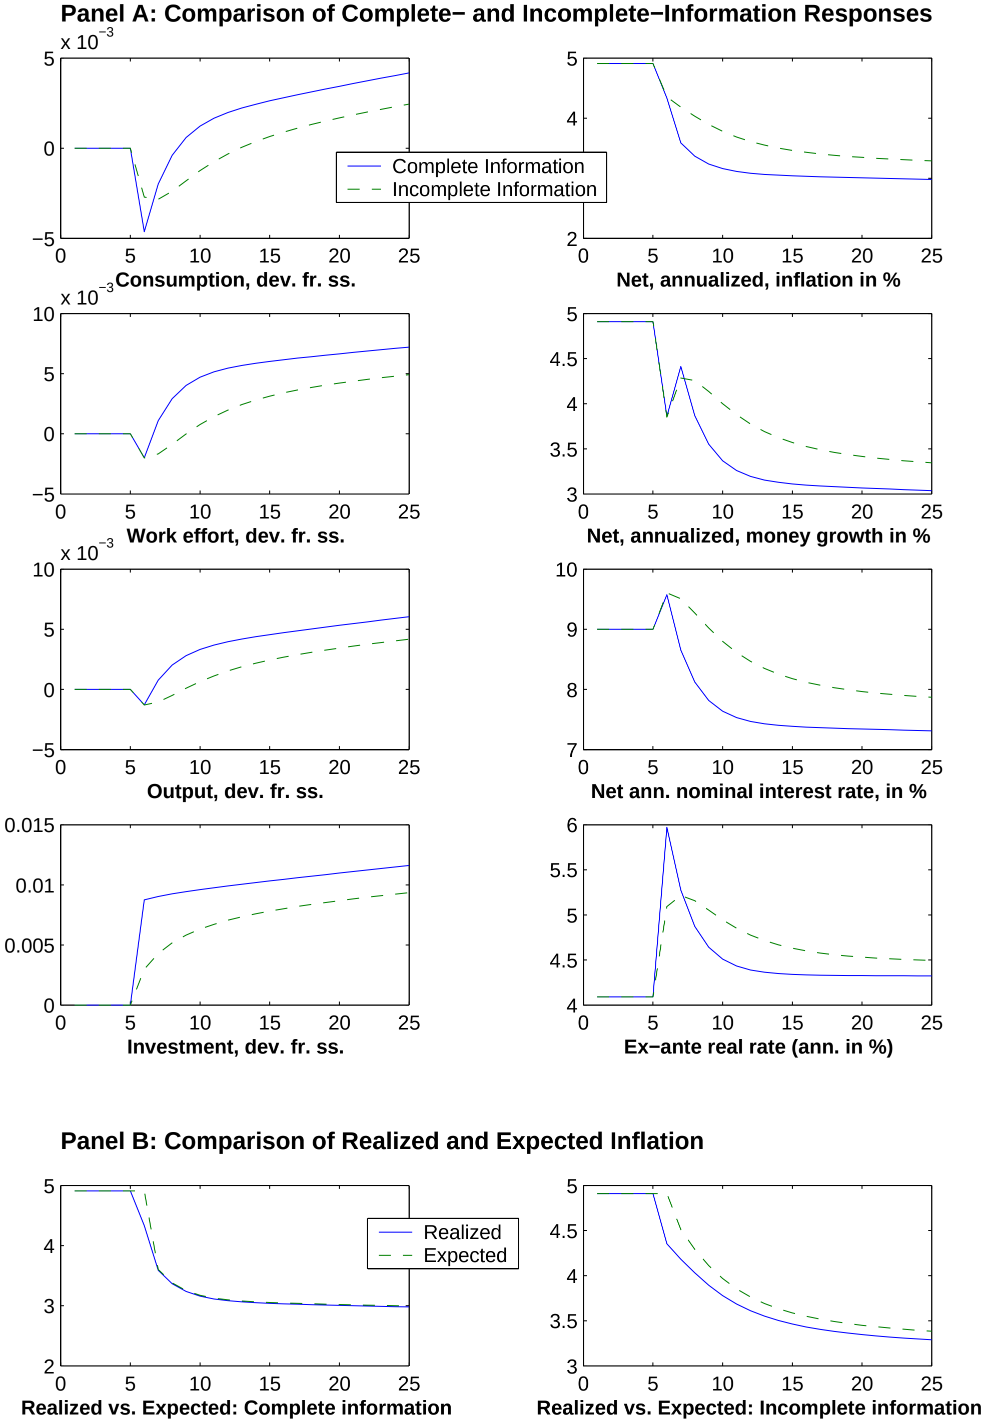

remains to be seen whether this effect is strong enough to generate empirical rejections of the unbiasedness hypothesis.

6.2 The Experiment

We treat our model as if it was the true Data Generating Process (DGP) of economic variables and assess what an empirical researcher, given outcomes from this DGP, would conclude about the unbiasedness of inflation expectations. To this end, a series of simulations of the model economy are performed.

In each of these simulations, a random realization of 80 periods is generated for both unobserved disturbances to the interest rate rule (the u_t and z_t variables)[25]. Economic agents' estimates of these shocks are computed and the model is solved according to this information. Two alternative measures of inflation, one-quarter-ahead inflation ($\pi_t \equiv P_{t+1}/P_t$) and four-quarter-ahead inflation ($\pi_t \equiv P_{t+4}/P_t$), are stored, along with the corresponding expectations of these quantities ($\pi_t^e \equiv E_t[P_{t+1}/P_t]$ and $\pi_t^e \equiv E_t[P_{t+4}/P_t]$).

Next, for each of these simulations, we carry out the test of unbiasedness discussed in Section 2. Recall that this involves estimating the regression

$$\pi_t = a_0 + a_1 \pi_t^e + \varepsilon_t; \tag{24}$$

and testing the null hypothesis $H_0 : a_0 = 0;\ a_1 = 1$. For each simulation, we record the estimates $\widehat{a_0}$ and $\widehat{a_1}$, as well as the appropriate test statistic about H_0.[26]

We repeat this simulation 1000 times. We present the output of these experiments as follows. In each of the figures below, the top graph is a histogram depicting the estimates of a_0 across the 1000 replications. It also depicts the median of the estimates. The middle graph depicts the estimates of a_1, again identifying the median. Finally, the third graph illustrates the results for the 1000 tests of H_0, showing a histogram of the test statistic along with its median and the 5% and 1% critical values associated with the test. We also indicate the fraction of the simulations for which the test statistic rejects H_0 at a significance level better than 5%. When the null hypothesis is true and the test is correctly specified, this fraction should be close to 5%, the size of the test. On the other hand, one can interpret results where this fraction is significantly higher than 5% as suggesting that the learning effects reduce the capability of the test to properly identify unbiasedness.

6.3 Results for the Benchmark Case

Figures 4 to 7 present the results of our Monte Carlo experiments using the benchmark calibration. Figures 4 and 5 illustrate the cases of one-quarter-ahead and four-quarter-ahead expectations, respectively, when information is complete. Figures 6 and 7 illustrate

[25] The empirical rejections of the unbiasedness hypothesis described in Section 2 are typically obtained with data samples of limited length.

[26] Under the null hypothesis and for one-quarter-ahead expectations, the expectation errors (the residuals in (24)) should not be serially correlated and we therefore use a simple F-statistic to test H_0. In the case of four-quarter-ahead expectations, the expectation errors would be correlated up to three lags even under H_0, because of the overlap between the horizon of the expectations and the frequency of the data. We thus use the Newey-West procedure to correct for serial correlation when computing the standard errors of the estimates. The test statistic is distributed as a χ^2.

the cases of one-quarter-ahead and four-quarter-ahead expectations, for the incomplete information case in which the information friction emphasized in the present paper is activated.

The top panel of Figure 4 is a histogram depicting the estimates of a_0 across the 1000 replications. It shows that the median estimate of a_0 is very close to zero. Similarly, the middle panel of the figure, depicting the estimates of a_1, shows a median very close to the hypothesized value of 1. Finally, the bottom panel confirms that these deviations from the respective values of 0 for α_0 and 1 for α_1 were not often significant: the test statistic for the hypothesis has a median around 0.70, when the 5% rejection region starts above 3. In fact, only 5.3% of the simulations lead the test statistic to reject the null at better than the 5% significance level. It appears that in the case of one-quarter-ahead expectations with complete information, the unbiasedness test performs just as it should.

Figure 5 presents the case when inflation expectations are measured as four-quarter-ahead expectations, with the information still complete. While the estimates of a_0 remain close to zero, the middle panel of the figure shows that the median estimate of a_1 is now around 0.96. However, the correction for serial correlation makes rejections harder to achieve, so that the test statistic rejects the null hypothesis in only about 9% of the cases, not drastically away from the 5% size of the test.

Over all, the complete information results in Figure 4 and 5 suggest that in such an economy, the simple tests of unbiasedness often performed in the empirical literature behave much as they should.

Let us now turn to the cases for which the information friction is activated. Figure 6 shows that for the one-quarter-ahead expectations, while the estimate of the constant parameter is once again not drastically different from zero, the distribution of the slope estimates is significantly skewed away from the hypothesized value of 1, yielding a median value of 0.82. The bottom panel of the figure illustrates that this skewness is reflected in the number of times H_0 is rejected: more than 20% of the cases feature a rejection of the null hypothesis, even though, by construction, our solution embodies the 'rational expectations' hypothesis.

Figure 7 shows that for the case of four-quarter-ahead expectations with incomplete information, similar results to those in Figure 6 emerge: the slope estimates are distributed significantly away from the hypothesized value of 1 and imply rejections of H_0 that are about five times more frequent than the normal rate of 5%.

These benchmark results suggest that the joint hypothesis of the model, the learning mechanism, and the calibration of the problem introduce significant size distortions in tests of unbiasedness of inflation expectations. These distortions arise because the relatively small samples with which these tests are performed are dominated by a few significant shifts in monetary policy that surprise agents and lead them to be, at least initially, confused about the true intentions of monetary authorities. The next section explores the extent to which the qualitative nature of the results expressed in Figures 4 to 7 are sensitive to the calibration of the model.

Figure 4: Results from Monte Carlo Simulations: One-Quarter-Ahead Expectations, Complete Information

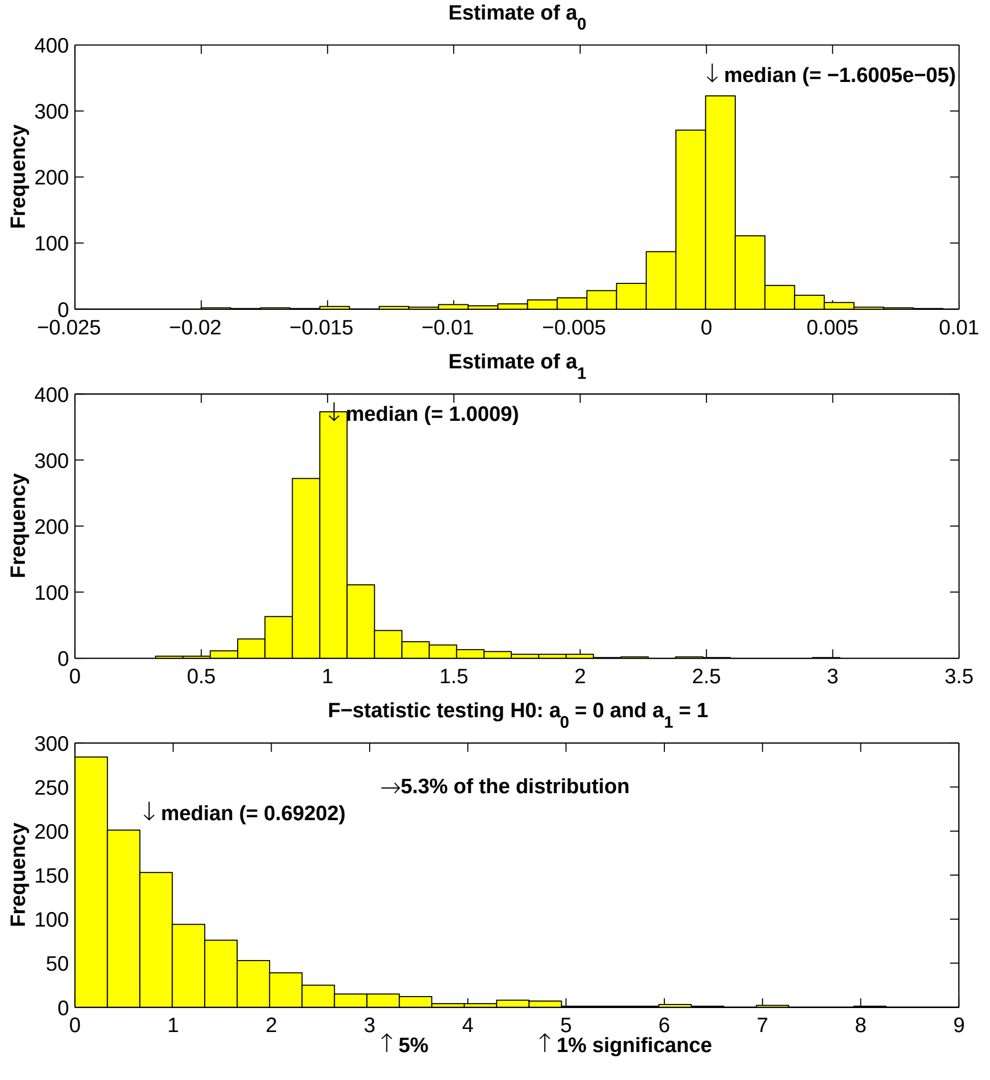

Figure 5: Results from Monte Carlo Simulations: Four-Quarter-Ahead Expectations, Complete Information

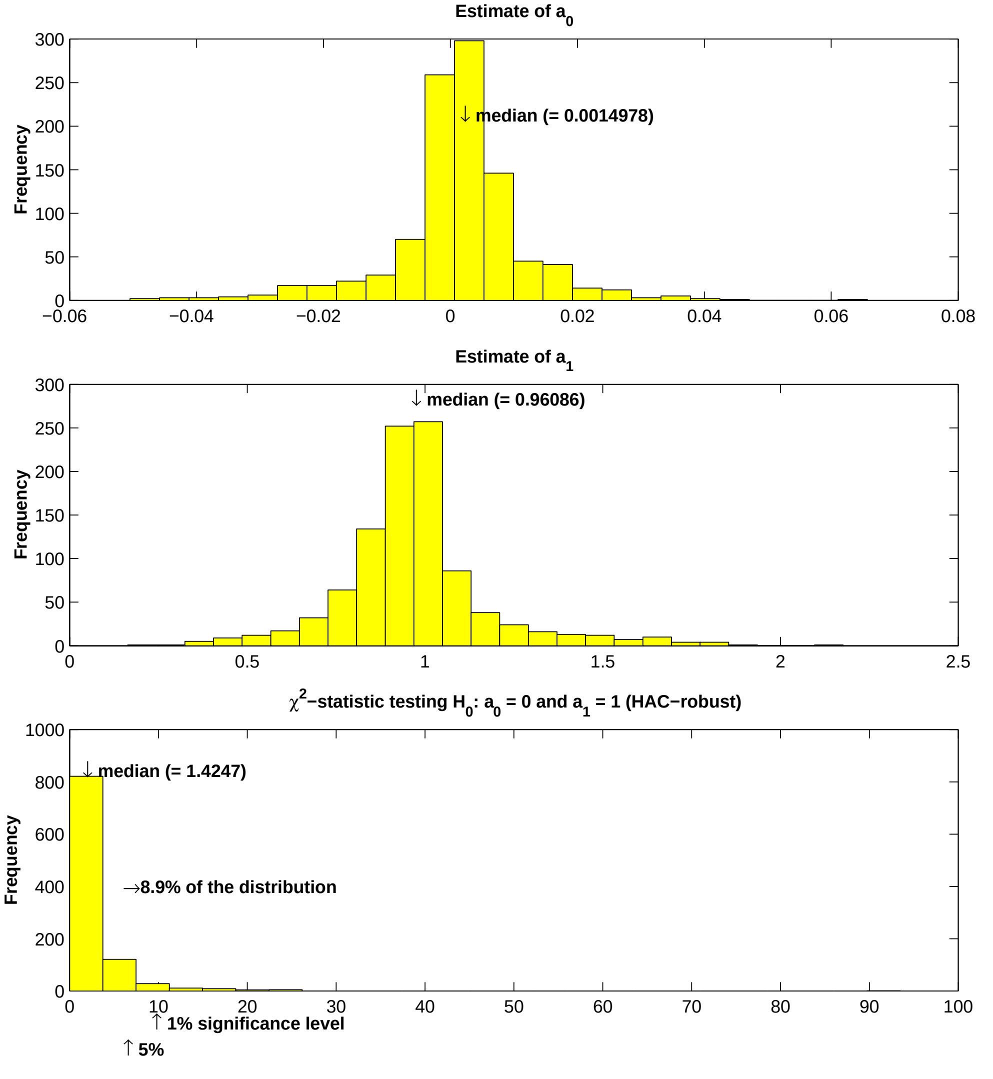

Figure 6: Results from Monte Carlo Simulations: One-Quarter-Ahead Expectations, Incomplete Information

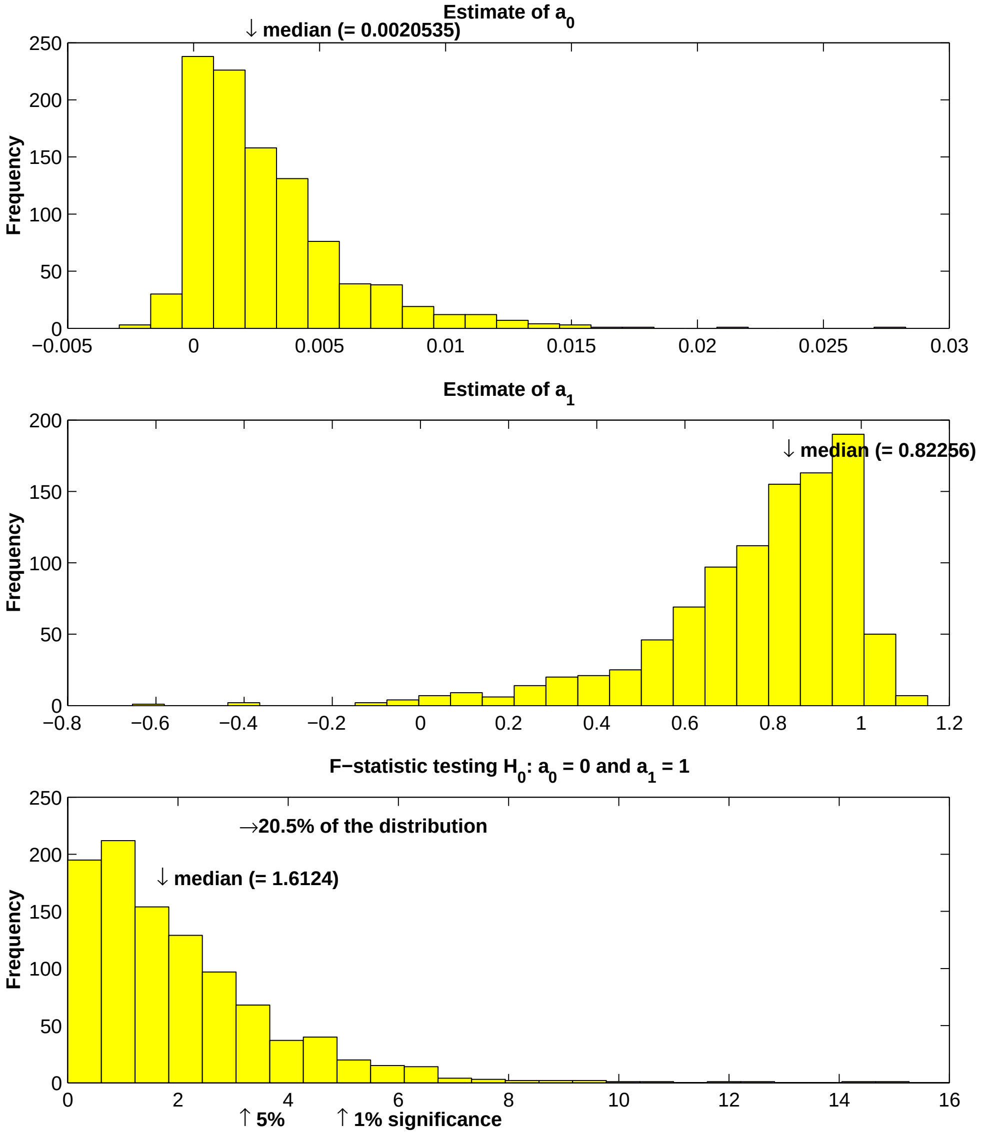

Figure 7: Results from Monte Carlo Simulations: Four-Quarter-Ahead Expectations, Incomplete Information

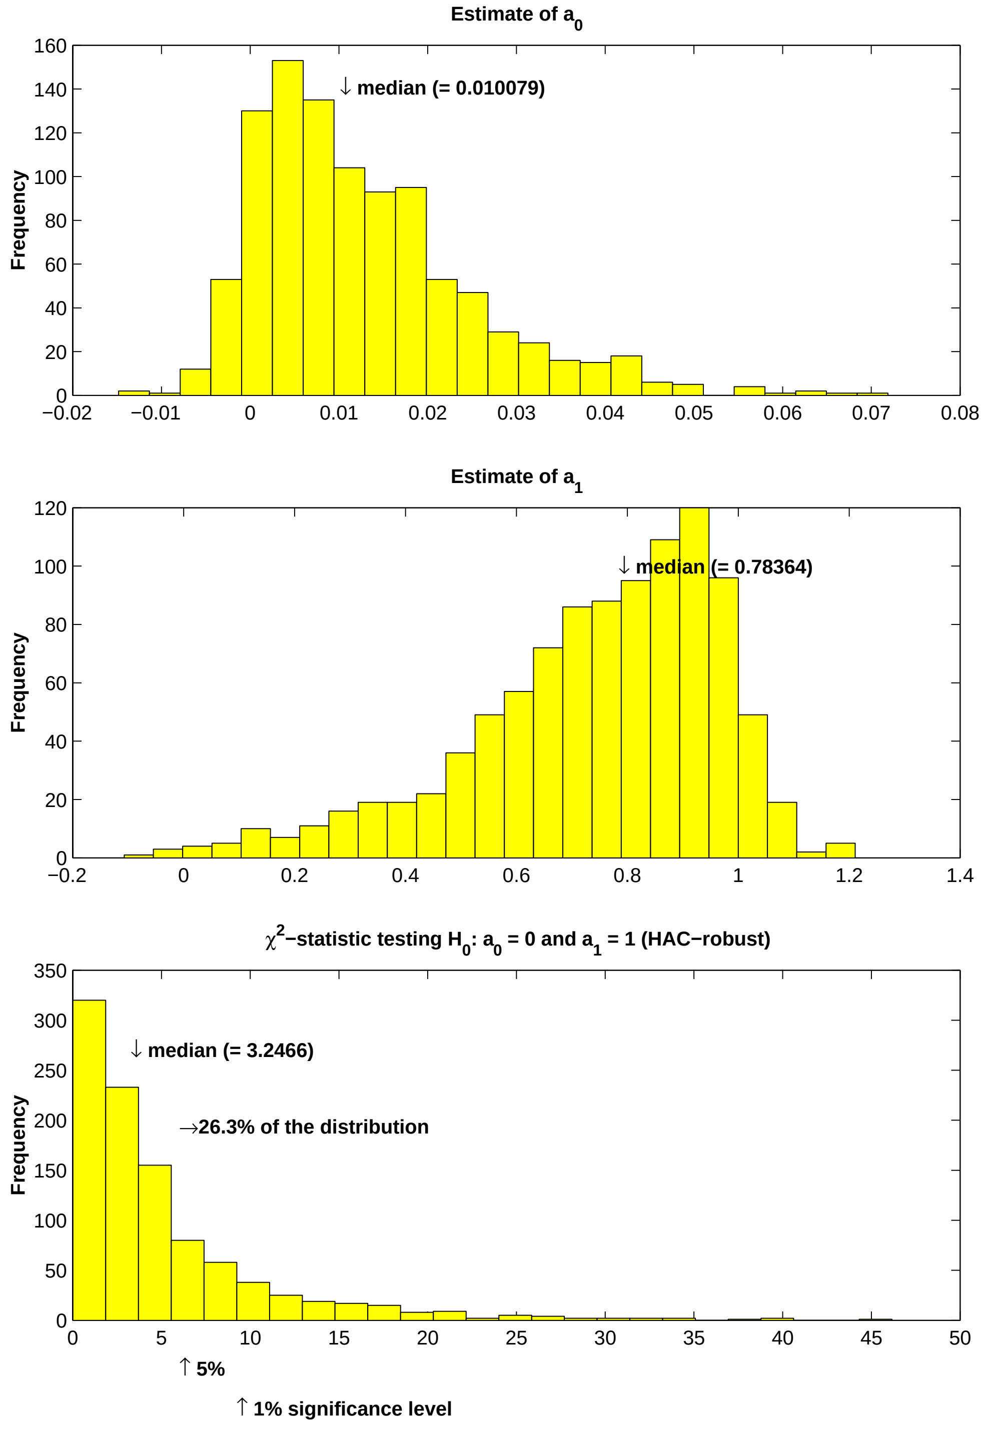

6.4 Sensitivity Analysis

To explore the sensitivity of the results to modifications in the calibration, we redo the analysis discussed above for several alternative specifications. The results are reported in Table 2. In that table, the first column indicates what kind of departure from the benchmark calibration is under study. The next two columns indicate the frequency at which the unbiasedness hypothesis is rejected when one-quarter-ahead expectations are used, in the complete information and incomplete information cases, respectively. The last two columns report corresponding results when the four-quarter-ahead expectations are utilized. To facilitate the comparison, the results from the benchmark cases are repeated at the beginning of the table.

The general impression one gets Table 2 is that the complete information cases generate rejections of H_0 as often, roughly, as the size of the tests implies. Particularly in the first column, the fraction of rejections seldom departs significantly from the level (5%) suggested by the size of the test. Although the numbers in the third column do depart more significantly from 5%, the departures are never excessive. In contrast, the incomplete information cases feature rejections of H_0 that are far more frequent. Although the precise numbers change from one case to the next, we observe rejections of H_0 two to five times more often when the information friction is activated. Interestingly, the fraction of rejections do not seem to depend on whether one-quarter-ahead or four-quarter-ahead expectations are used.

Turning now to specific cases, we see, in the first three departures from the benchmark case, that eliminating the response of monetary authorities to the output gap or modifying the extent to which interest rate changes are smoothed-in does not change the results markedly. However, increasing the aggressiveness in the response of monetary authorities to deviations of inflation from target (an increase of α from 2.0 to either 3.0 or 4.0) does modify the results substantially. The unbiasedness hypothesis is now rejected around 35% of the time when the information friction is active, while the corresponding numbers for the complete information case increase only slightly. These increases in the frequency of rejections is brought about because a high value for the coefficient α acts like a multiplier on the monetary policy shift. This is seen best by recalling the definition of ϵ_t^* in (18): a high value of α implies, for a given shift $(\pi_t - \pi_0^T)$, a stronger increase of interest rates.[27]

The second type of changes we experiment with are modifications in the processes governing the evolution of the two components to monetary policy. First, we modify the expected duration of a given shift in z_t, first from 0.975 to 0.99, then back to 0.95. Increasing the duration opens the gap between the complete and incomplete information cases somewhat, compared to the benchmark case. Decreasing the duration closes that gap. Modifications to the persistence of the transitory shocks, the next departures from benchmark we consider, does not modify the results markedly. Next, we experiment with changes in the variances of the shocks and shifts. The gist of these experiments is that increasing the variance in the shifts z_t, relative to the variance of the shocks u_t, increases

[27] While it may seem that stronger shifts would make learning easier, the high duration of a given shift implies that economic agents will not, at first, identify even sharp spikes in interest rates as arising from shifts in the inflation target. This intuition is also at play below when the departure from benchmark analyzed is an increase in the standard deviation of the shifts themselves.

the gap between the complete and incomplete information cases. Considering our previous assessment that the benchmark value for the variance in the shocks u_t (the parameter σ_e) is an upper bound, this result points to significant and continued differences between the complete and incomplete information cases.

Table 2. Sensitivity Analysis of the Results: Frequency of Rejections

ecification Examined	**One-Quarter-Ahead Expectations**		**Four-Quarter-Ahead Expect**	
	Complete Info.	**Incomplete Info.**	**Complete Info.**	**Incomplet**
nchmark case	5.3%	20.5%	8.9%	26.3%
nel A: Modifications to the Monetary Policy Rule				
response to output gap ($\beta = 0.0$)	5.5%	19.9%	8.5%	26.2%
smoothing of interest rates ($\rho = 0.0$)	5.5%	22.2%	9.4%	29.5%
reased smoothing of interest rates ($\rho = 0.5$)	5.1%	18.0%	7.8%	24.3%
re aggressive response to inflation ($\alpha = 3.0$)	5.7%	34.7%	10.6%	34.7%
st aggressive response to inflation ($\alpha = 4.0$)	7.2%	37.7%	12.7%	34.6%
nel B: Modifications to the Calibration of Monetary Shocks and Shifts				
ry high duration of regime shifts ($\phi_2 = 0.99$)	5.2%	22.4%	6.7%	35.3%
wer duration of regime shifts ($\phi_2 = 0.95$)	5.6%	15.4%	10.9%	20.6%
persistence in transitory shocks ($\phi_1 = 0.0$)	4.7%	19.8%	8.7%	27.8%
gher persistence of transitory shocks ($\phi_1 = 0.2$)	5.8%	19.4%	8.4%	25.5%
gher variance of regime shifts ($\sigma_g = 0.01$)	5.6%	37.9%	15.1%	34.6%
wer variance of regime shifts ($\sigma_g = 0.0025$)	4.8%	8.1%	7.3%	19.4%
wer variance of transitory shocks ($\sigma_e = 0.0025$)	5.6%	33.9%	15.2%	34.8%
nel C: Other Modifications				
reased number of repetitions[a]	5.8%	21.8%	9.6%	28.6%
gher portfolio adjustment costs ($\tau = 15.0$)	5.3%	19.6%	9.0%	26.5%
lusion of money demand shocks	5.4%	20.1%	8.8%	26.1%
lusion of technology shocks	5.4%	6.9%	9.2%	10.1%
reased length of simulated time series[b]	5.0%	10.8%	2.7%	4.9%

2500 repetitions using a 80 period sample and benchmark calibration
1000 repetitions using a 500 period sample

The last series of modifications to the benchmark calibration, in Panel 3, bring interesting observations. First, increasing the number of repetitions for the benchmark case (from 1000 to 2500) brings only small changes to the results. This experiment shows that 1000 repetitions is enough to get a good sense of the population distribution of the test statistics. Next, the panel shows that modifying the level of the costs of adjustments in the portfolio of economic agents or including money demand shocks as specified in (9) and (10) does not change the results significantly.

On the other hand, adding technology shocks, as specified in (7), attenuates the difference between the complete and incomplete information cases. Such a result actually validates our approach; because the technology shocks are perfectly observed by both economic agents and monetary authorities, one does not expect that the introduction of these shocks would imply a relationship between realized and expected inflation that is distorted. The rejections of the unbiasedness hypothesis identified in this paper thus truly arise from the limited information about monetary policy shocks.[28]

Finally, the last experiment shows that using 1000 repetitions of much longer samples (with 500 periods in each sample), results in a drastic reduction of the rejections, particularly for the incomplete information case. This last experiment strongly suggests that the observed rejections of the unbiasedness hypothesis may be due to size distortions brought about by the presence of a few significant shifts in monetary policy in the small samples typically used to study inflation expectations.

7. Discussion and Conclusion

Consider Figure 1 once again, which graphs the realized and expected (CPI) inflation (as measured by the Livingston Survey data). From this figure, one gets the sense that a few significant shifts in monetary policy during the 1970s and at the beginning of the 1980s surprised economic agents and, for a while, left them unsure of the true intentions of monetary authorities. Apart from the periods immediately following the shifts, the expectations of economic agents do not look completely out of line with realized inflation.

This paper shows that these empirical features can be reproduced by modelling these shifts in a standard monetary DSGE model, and assuming that agents must learn about the shifts over time, replicates this intuition. In the case of complete information about the shifts, the unbiasedness hypothesis is rejected with very low frequency, in keeping with the size of the tests. In contrast, when the information friction is active, the unbiasedness hypothesis is rejected much more often –between two and five times– than the size of the tests would imply, even though our model embeds the rational expectations solution concept by construction. Further, the likelihood of a rejection tends to be eliminated when the sample size increases, even though the information friction remains.

One needs to acknowledge that the Kalman filter used in the present paper may not be the optimal filter in small samples, for which the (asymptotic) normal behavior of the two unobserved components of monetary policy has not established itself. It would be interesting to verify the extent to which economic agents could improve on the Kalman filter estimates by using non-linear filters. Further, we must caution that this paper identifies a single inflation expectation as the average of survey participants' responses. The dispersion of the survey's participants around that average is neither analyzed nor modelled.

Overall, the results presented in this paper support the view that learning effects with regard to monetary policy, in addition to creating persistence in the responses of most macro-aggregates following monetary policy shocks, imply dynamics in the expectations of agents that replicate well some of the empirical evidence about measured inflation

[28] Of course, the belief that macroeconomic volatility in the last forty years was solely the result of technology shocks would imply, in the environment presented in the present paper, that the empirical rejections of the unbiasedness hypothesis could not have come from learning about monetary policy. We believe, however, that ample evidence exists of very significant monetary policy shocks having affected the macroeconomic outcomes of all major economies in the last four decades. Note also that in an environment with learning about the *parameters* of the monetary policy rule, technology shocks could potentially affect learning about monetary policy to a greater degree.

expectations.

Reinforcing this view, and thus the importance of incomplete information for modelling macroeconomic activity, could proceed by verifying that the incomplete information and learning framework replicates other facets of the relationship between realized and (measured) expected inflation. Notably, could the learning effects replicate the additional results against unbiasedness discussed in Section 2, or the evidence that inflation expectations are not efficient predictors of realized inflation? Alternatively, could they lead simulated data to match the dynamic correlation patterns linking realized and expected inflation? These questions could be easily addressed by the framework developed here.

Further, the evidence about shifts in the parameters of the interest rate rule [Clarida et al., 2000] opens interesting avenues for future research. For example, could such shifts, and least-square learning about these shifts on the part of economic agents, be responsible for the evidence against unbiasedness and efficiency in measured inflation expectations? Once again, straightforward modifications to the framework used here could be employed to address such questions.

References

D. Andolfatto and P. Gomme. Monetary policy regimes and beliefs. *International Economic Review*, 44:1–30, 2003.

D. Andolfatto, S. Hendry, and K. Moran. Labour markets, liquidity, and monetary policy regimes. *Canadian Journal of Economics*, 2003. forthcoming.

R. J. Barro and D. B. Gordon. A positive theory of monetary policy in a natural rate model. *Journal of Political Economy*, 91:589–610, 1983.

K. Brunner, A. Cukierman, and A. Meltzer. Stagflation, persistent unemployment and the permanence of economic shocks. *Journal of Monetary Economics*, 6:467–492, 1980.

L. J. Christiano and C. J. Gust. Taylor rules in a limited participation model. Working Paper 7017, National Bureau of Economic Research, 1999.

R. Clarida, J. Gali, and M. Gertler. Monetary policy rules and macroeconomic stability: Evidence and some theory. *Quarterly Journal of Economics*, 115:147–180, 2000.

D. Croushore. The Livingston survey: Still useful after all these years. *Federal Reserve Bank of Philadelphia Business Review*, pages 1–12, March-April 1997.

J. B. DeLong. America's peacetime inflation: The 1970s. In C. D. Romer and D. H. Romer, editors, *Reducing Inflation: Motivation and Strategy*. The University of Chicago Press, Chicago, 1997.

V. Dolar and K. Moran. Are survey-based inflation expectations biased? a reconsideration in the context of the forward discount puzzle. Bank of Canada, Manuscript, 2002.

M. Dotsey and J. L. DeVaro. Was the disinflation of the early 1980s anticipated? *Federal Reserve Bank of Richmond Economic Quarterly*, 81(4):41–59, 1995.

C. Erceg and A. Levin. Imperfect credibility and inflation persistence. *Journal of Monetary Economics*, 50:915–944, 2003.

K. Froot and R. Thaler. Anomalies: Foreign exchange. *Journal of Economic Perspectives*, 4:179–192, 1990.

J. Greenwood, Z. Hercowitz, and G. Huffman. Investment, capacity utilization, and the real business cycle. *The American Economic Review*, 78:402–417, 1988.

J. D. Hamilton. *Time Series Analysis*. Princeton University Press, Princeton, 1994.

K. D. Hoover and S. J. Perez. Post hoc ergo propter once more: An evaluation of 'does monetary policy matter' in the spirit of James Tobin. *Journal of Monetary Economics*, 34:47–73, 1994.

R. G. King and M. W. Watson. The solution of linear difference systems under rational expectations. *International Economic Review*, 39:1015–1026, 1998.

S. Kozicki and P. A. Tinsley. What do you expect? Imperfect policy credibility and tests of the expectations hypothesis. Working Paper 01-02, Federal Reserve Bank of Kansas City, 2001a.

S. Kozicki and P. A. Tinsley. Term structure views of monetary policy under alternative models of agent expectations. *Journal of Economic Dynamics and Control*, 25:149–184, 2001b.

E. M. Leeper. Narrative and VAR approaches to monetary policy: Common identification problems. *Journal of Monetary Economics*, 40:641–657, 1997.

K. K. Lewis. The persistence of the 'peso problem' when policy is noisy. *Journal of International Money and Finance*, 7:5–21, 1988.

K. K. Lewis. Changing beliefs and systematic rational forecast errors with evidence from foreign exchange. *The American Economic Review*, 79:621–636, 1989.

K. Moran. The welfare benefits of lower inflation: Accounting for the transition. Bank of Canada, Manuscript, 1999.

J. F. Muth. Optimal properties of exponentially weighted forecasts. *Journal of the American Statistical Association*, 55:299–306, 1960.

E. Nelson. UK monetary policy, 1972-97: A guide using Taylor rules. Working Paper 120, Bank of England, 2000.

M. Owyang and G. Ramey. Regime switching and monetary policy measurement. Working Paper 022A, Federal Reserve Bank of St Louis, 2001.

J. M. Roberts. Is inflation sticky? *Journal of Monetary Economics*, 39:173–196, 1997.

C. D. Romer and D. H. Romer. Does monetary policy matter? A new test in the spirit of Friedman and Schwartz. In B. S. Bernanke and K. Rogoff, editors, *NBER Macroeconomics Annual*, pages 121–170. The MIT Press, Cambridge, 1989.

C. D. Romer and D. H. Romer. Monetary policy matters. *Journal of Monetary Economics*, 34:75–88, 1994.

F. Ruge-Murcia. Inflation targeting under asymmetric preferences. *Journal of Money, Credit, and Banking*, 35:763–785, 2003.

P. Shen and J. Corning. Can TIPS help identify long-term inflation expectations? *Federal Reserve Bank of Kansas City Economic Review*, pages 61–87, Fourth Quarter 2001.

K. Sill and J. Wrase. Exchange rates and monetary policy regimes in Canada and the U.S. Working Paper 99-13, Federal Reserve Bank of Philadelphia, 1999.

J. B. Taylor. Discretion versus policy rules in practice. *Carnegie-Rochester Conference Series on Public Policy*, 39:195–214, 1993.

M. P. Taylor. The economics of exchange rates. *Journal of Economic Literature*, 33:13–47, 1995.

L. B. Thomas. Survey measures of expected U.S. inflation. *Journal of Economic Perspectives*, 13:125–144, 1999.

A Appendix

Recall equations (19), (20), describing the evolution of the observed monetary policy deviations from the benchmark rule:

$$\begin{bmatrix} z_{t+1} \\ u_{t+1} \end{bmatrix} = \begin{bmatrix} \phi_2 & 0 \\ 0 & \phi_1 \end{bmatrix} \cdot \begin{bmatrix} z_t \\ u_t \end{bmatrix} + \begin{bmatrix} N_{t+1} \\ e_{t+1} \end{bmatrix}; \quad \text{(A.1)}$$

$$\epsilon_t^* = \begin{bmatrix} (1-\rho)(1-\alpha) & 1 \end{bmatrix} \cdot \begin{bmatrix} z_t \\ u_t \end{bmatrix}; \quad \text{(A.2)}$$

with N_{t+1} defined as follows:

$$N_{t+1} = \begin{cases} (1-\phi_2)\, z_t & \text{with probability } \phi_2; \\ g_{t+1} - \phi_2 z_t & \text{with probability } 1-\phi_2, \end{cases}$$

and where, again, it was assumed that $e_{t+1} \sim N(0, \sigma_e^2)$ and $g_{t+1} \sim N(0, \sigma_g^2)$.

Compare this system to the one described in Hamilton (1994, chapter 13)'s discussion of state space models and the Kalman filter:

$$\begin{aligned} y_t &= \mathbf{A}' \cdot \mathbf{x}_t + \mathbf{H}' \cdot \xi_t + \mathbf{w}_t; \\ \xi_{t+1} &= \mathbf{F} \cdot \xi_t + \mathbf{v}_{t+1}; \\ E(\mathbf{v}_t \mathbf{v}_t') &= \mathbf{Q}; \\ E(\mathbf{w}_t \mathbf{w}_t') &= \mathbf{R}. \end{aligned}$$

The equivalence between the two systems is established by defining $y_t = \epsilon_t^*$, $x_t = 0$, $\xi_t = [z_t\ u_t]'$, $\mathbf{w}_t = 0$, $\mathbf{v}_t = [N_t\ e_t]'$ as well as the following matrices:

$$\mathbf{A} = 0;\ \mathbf{H} = \begin{bmatrix} (1-\rho)(1-\alpha) \\ 1 \end{bmatrix};\ \mathbf{F} = \begin{bmatrix} \phi_2 & 0 \\ 0 & \phi_1 \end{bmatrix};\ \mathbf{Q} = \begin{bmatrix} \sigma_N^2 & 0 \\ 0 & \sigma_e^2 \end{bmatrix};\mathbf{R} = 0.$$

Note that one can show that $\sigma_N^2 = (1-\phi_2)(1+\phi_2)\sigma_g^2$.

Denote the MSE (mean-squared-error) of the one-step-ahead forecasts of the unobserved states, conditional on time-t information, as $\mathbf{P}_{t+1}|_t$.[29] Conditional on starting values $\hat{\xi}_1|_0$ and $P_1|_0$, the following recursion for $\hat{\xi}_{t+1}|_t$ and $P_{t+1}|_t$ emerges:

$$\begin{aligned} \mathbf{K}_t &= \mathbf{F}\mathbf{P}_t|_{t-1}\mathbf{H}(\mathbf{H}'\mathbf{P}_t|_{t-1}\mathbf{H})^{-1}; && \text{(A.4)} \\ \widehat{\xi}_{t+1}|_t &= \mathbf{F}\hat{\xi}_t|_{t-1} + \mathbf{K}_t(y_t - \mathbf{H}'\hat{\xi}_t|_{t-1}); && \text{(A.5)} \\ \mathbf{P}_{t+1}|_t &= (\mathbf{F} - \mathbf{K}_t\mathbf{H}')\mathbf{P}_t|_{t-1}(\mathbf{F}' - \mathbf{H}\mathbf{K}_t') + \mathbf{Q}; && \text{(A.6)} \end{aligned}$$

The intuition behind this updating sequence is that, at each step, agents will use their *observed* forecasting errors $(y_t - \mathbf{H}' \cdot \widehat{\xi}_t|_{t-1})$ and their knowledge of the parametric form of the system to update their best estimates of the unobserved states ξ_t. The mechanics of this updating takes the form of linear projection and is detailed in Hamilton.

Under the assumption of normality of both processes (for z_t and u_t), the sequence $(\widehat{\xi}_{t+1}|_t)|_{t=1}^T$ represents the optimal one-step-ahead forecasts of the unobserved states.[30]

[29] So that $\mathbf{P}_{t+1}|_t = E_t\left[(\xi_{t+1} - \widehat{\xi}_{t+1}|_t)(\xi_{t+1} - \widehat{\xi}_{t+1}|_t)'\right]$.

[30] Even without the assumption of normality, the sequence of filtered estimates remains the best linear forecasts of the unobserved states conditional on time-t information.

Earlier publications in this series as from January 2003

No. 93/2003 Is Financial Market Volatility Informative to Predict Recessions?
N. Valckx, M.J.K de Ceuster and J. Annaert

Overview of papers presented at the 5th Annual Conference 'Global Linkages and Economic Performance' held at De Nederlandsche Bank, Amsterdam, 13 – 15 November 2002 (Nos. 94 – 101)

No. 94/2003 The Trilemma in History: Tradeoffs among Exchange Rates, Monetary Policies, and Capital Mobility
Maurice Obstfeld, Jay C. Shambaugh and Alan M. Taylor

No. 95/2003 Financial Globalization and Monetary Policy
Helmut Wagner and Wolfram Berger

No. 96/2003 Fancy a stay at the "Hotel California"?
Foreign Direct Investment, Taxation and Exit Costs
Holger Görg

No. 97/2003 Global vs. Local Competition
Patrick Legros and Konrad Stahl

No. 98/2003 Long-Term Global Market Correlations
William N. Goetzmann, Lingfeng Li and K. Geert Rouwenhorst

No. 99/2003 The Importance of Multinational Companies for Global Economic Linkages
W. Jos Jansen and Ad C.J. Stokman

No. 100/2003 Globalisation and Market Structure
J.Peter Neary

No. 101/2003 Is Wealth Increasingly Driving Consumption?
Tamim Bayoumi and Hali Edison

No. 102/2003 On the influence of capital requirements on competition and risk taking in banking
Peter J.G. Vlaar

No. 103/2003 Robust versus Optimal Rules in Monetary Policy: A note
Maria Demertzis and Alexander F. Tieman

No. 104/2003 The Impact of the Single Market on the Effectiveness of ECB Monetary Policy
Paul Cavelaars

No. 105/2003 Central Bank Transparency in Theory and Practice
Maria Demertzis and Andrew Hughes Hallett

No. 106/2003 The (A)symmetry of shocks in the EMU
Bastiaan A. Verhoef

No. 107/2003 Forecasting inflation: An art as well as a science!
Ard H.J. den Reijer and Peter J.G. Vlaar

No. 108/2003	On Wage Formation, Wage Development and Flexibility: a comparison between European countries and the United States *Marga Peeters and Ard den Reijer*
No. 109/2003	Foreign Banks and Credit Stability in Central and Eastern Europe A Panel Data Analysis *Ralph de Haas and Iman van Lelyveld*
No. 110/2003	Corporate Investment and Financing Constraints: Connections with Cash Management *Allard Bruinshoofd*
No. 111/2003	Bank provisioning behaviour and procyclicality *Jacob A. Bikker and Paul A.J. Metzemakers*
No. 112/2004	Optimal Investment Policies for Defined Benefit Pension Funds *Arjen Siegmann*
No. 113/2004	The Dynamic Properties of Inflation Targeting Under Uncertainty *Maria Demertzis and Nicola Viegi*
No. 114/2004	A Taylor Rule for the Euro Area Based on Quasi-Real Time Data *Yvonne Adema*

Overview of papers presented at the 6th Annual Conference 'Information and communication in central banking' held at De Nederlandsche Bank, Amsterdam, 12 – 14 November 2003 (Nos. 115 – 122)

No. 115/2004	Expectations and the Effects of Money Illusion *Ernst Fehr and Jean-Robert Tyran*
No. 116/2004	Optimal Monetary Policy with Imperfect Common Knowledge *Klaus Adam*
No. 117/2004	Public and Private Information in Monetary Policy Models *Jeffery D. Amato and Hyun Song Shin*
No. 118/2004	Does inflation targeting matter? *Laurence Ball and Niamh Sheridan*
No. 119/2004	Coordination of Expectations in Asset Pricing Experiments *Cars Hommes, Joep Sonnemans, Jan Tuinstra and Henk van de Velden*
No. 120/2004	Information variables for monetary policy in a small structural model of the euro area *Francesco Lippi and Stefano Neri*
No. 121/2004	Inflation Expectations and Learning about Monetary Policy *David Andolfatto, Scott Hendry and Kevin Moran*
No. 122/2004	Central Bank Communication and Output Stabilization *Marco Hoeberichts, Mewael Tesfaselassie and Sylvester Eijffinger*

De Nederlandsche Bank