

Bank of Canada



Banque du Canada

Working Paper 2002-30 / Document de travail 2002-30

A large, stylized white graphic of a classical building facade with a pediment and columns, set against a light gray background. The title and author information are centered within this graphic.

Inflation Expectations and Learning about Monetary Policy

by

David Andolfatto, Scott Hendry, and Kevin Moran

ISSN 1192-5434

Printed in Canada on recycled paper

Bank of Canada Working Paper 2002-30

October 2002

Inflation Expectations and Learning about Monetary Policy

by

David Andolfatto,¹ Scott Hendry,² and Kevin Moran²

¹Department of Economics
Simon Fraser University
Burnaby, British Columbia, Canada V5A 1S6
dandolfa@sfu.ca

²Monetary and Financial Analysis Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
shendry@bankofcanada.ca
kmoran@bankofcanada.ca

The views expressed in this paper are those of the authors.
No responsibility for them should be attributed to the Bank of Canada.

Contents

Acknowledgements.....	iv
Abstract/Résumé.....	v
1. Introduction.....	1
2. Empirical Evidence of Inflation Expectations.....	3
3. The Model.....	5
3.1 Households.....	5
3.2 Firms.....	6
3.3 Financial intermediaries.....	7
3.4 Equilibrium.....	7
4. Monetary Policy.....	7
4.1 The monetary policy rule.....	7
4.2 Monetary policy shocks and monetary policy shifts.....	8
4.3 Incomplete information and learning.....	9
5. Calibration and Solution of the Model.....	12
5.1 Preferences and technology.....	12
5.2 Parameters of the interest-rate-targeting rule.....	13
5.3 Shifts and shocks to monetary policy.....	13
6. Monte Carlo Simulation of the Model.....	15
6.1 Impulse responses following a regime shift.....	15
6.2 The experiment.....	16
6.3 Results for the benchmark case.....	17
6.4 Sensitivity analysis.....	18
7. Conclusion.....	20
References.....	22
Tables.....	25
Figures.....	27
Appendix A: Kalman Filter.....	34

Acknowledgements

The authors wish to acknowledge the expert research assistance of Veronika Dolar. We thank Walter Engert, 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.

Abstract

Various measures indicate that inflation expectations evolve sluggishly relative to actual inflation. In addition, they often fail conventional tests of unbiasedness. These observations are sometimes interpreted as evidence against rational expectations.

The authors embed, within a standard monetary dynamic stochastic general-equilibrium model, an information friction and a learning mechanism regarding the interest-rate-targeting rule that monetary policy authorities follow. 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 policy authorities.

The model's simulated data are 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. These results suggest that an important size distortion may occur when conventional tests of unbiasedness are applied to relatively small samples dominated by a few significant shifts in monetary policy and sluggish learning about those shifts.

JEL classification: E47, E52, E58

Bank classification: Business fluctuations and cycles; Economic models

Résumé

Divers indicateurs donnent à penser que les attentes d'inflation évoluent plus lentement que l'inflation observée. En outre, dans bien des cas, ils se révèlent entachés de biais lorsqu'on les soumet aux tests usuels d'absence de biais. Le rejet de ces tests est parfois interprété comme une réfutation de l'hypothèse de rationalité des attentes.

Les auteurs intègrent, dans un modèle monétaire stochastique dynamique d'équilibre général standard, un élément de friction relatif à l'acquisition de l'information et un mécanisme d'apprentissage concernant la règle de taux d'intérêt qu'appliquent les autorités monétaires. Le mécanisme d'apprentissage permet aux agents économiques ayant un comportement d'optimisation de distinguer les chocs temporaires que subit la règle de politique monétaire des modifications apportées à l'occasion à la cible d'inflation poursuivie.

Les données simulées dans le modèle sont conformes aux observations empiriques. Lorsque l'élément de friction est pris en compte, les attentes d'inflation simulées comportent un biais

d'après les tests standard, et ce, beaucoup plus souvent que dans le cas d'une information complète, sans friction. Ces résultats indiquent qu'une forte distorsion de niveau est possible lorsque les tests standard d'absence de biais sont appliqués à des échantillons relativement restreints caractérisés par seulement quelques changements d'orientation importants de la politique monétaire et que les agents mettent beaucoup de temps à reconnaître ces derniers.

Classification JEL : E47, E52, E58

Classification de la Banque : Cycles et fluctuations économiques; Modèles économiques

1. Introduction

Various measures indicate that inflation expectations evolve sluggishly relative to actual inflation. Expectations tend to underpredict inflation during periods of rising inflation and overpredict it during periods of diminishing inflation.¹ Related to this sluggishness phenomenon is the stylized fact, documented by an extensive 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 behaviour 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 authorities, once embedded in a standard monetary dynamic stochastic general-equilibrium (DSGE) model, can lead simulated data to quantitatively replicate the empirical evidence against the unbiasedness of inflation expectations.

The following information friction is introduced. We assume that the interest-rate-targeting rule followed by monetary policy authorities is affected by transitory shocks but also, occasionally, by persistent shifts in the inflation target that anchors the rule. We interpret the transitory shocks in the standard way, as instances of monetary policy authorities wishing to deviate from their rule for a short period; for example, to react to financial shocks. We assume that the occasional shifts in the inflation target reflect changes in economic thinking about the optimal inflation rate, or the appointment of a new central bank head with different preferences for 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 developed by Christiano and Gust (1999). We then repeatedly simulate the model and perform unbiasedness tests on the artificial data equivalent to those performed on measured inflation expectations.

Our simulations identify substantially different outcomes in the unbiasedness tests when the information friction is active compared with the complete-information case when it is shut down. Specifically, the fraction of rejections when complete information is assumed never deviates significantly from the level suggested by the

¹For example, Dotsey and DeVaro (1995) uncover economic agents' expectations about U.S. inflation—using commodity futures data—over the disinflationary episode of 1980Q1–1983Q3, and find that expected inflation exceeded actual inflation in all but three periods for the eight-month forecasts and in each 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, 267).

²See Thomas (1999), Roberts (1997), and Croushore (1997), and the references they cite.

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

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 five times—although our model embodies the “rational expectations” solution concept by construction. Interestingly, these differences are much attenuated when the sample size of each simulation is increased significantly.

Given these results, we 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 a few significant shifts in monetary policy, coupled with relatively sluggish learning about those shifts, can lead to significant size distortions of the tests. Furthermore, while this effect may be sufficient to trigger excessive rejections of the null hypothesis in small samples, it should disappear as the sample size grows.

Environments with information frictions and learning effects similar to the one we describe have been used previously, notably to rationalize the persistent responses of real variables following monetary policy shocks.⁴ Our paper makes a twofold 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 that of Erceg and Levin (2001), allows the learning literature to connect with the now-standard view of 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 (2001) 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. Conversely, Erceg and Levin (2001) choose parameter values to match the relationship between realized and expected inflation and concentrate on the implication of their chosen specification for real variables. Our analysis complements theirs and broadens the scrutiny of the empirical relevance of incomplete information and learning effects.

Our strategy is similar to that of Kozicki and Tinsley (2001a), who argue that the frequent empirical rejections of the term structure’s expectation hypothesis could be the result of economic agents learning only gradually about shifts in the Federal Reserve’s objectives. Kozicki and Tinsley 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

⁴Recent contributions include Andolfatto and Gomme (1999), Moran (1999), Andolfatto, Hendry, and Moran (2000), Erceg and Levin (2001), 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, Cukierman, and Meltzer (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.

“forward discount” puzzle observed in foreign exchange market data.

This paper is organized as follows. Section 2 describes the stylized fact that measured inflation expectations fail simple unbiasedness tests. Section 3 describes the model used in our simulation, essentially the one developed by Christiano and Gust (1999). Section 4 details 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 policy authorities. Section 4 also describes the mechanics of the Kalman filter, used by economic agents to solve the signal-extraction problem and to distinguish one component of monetary policy disturbances from another. Section 5 explains the calibration strategy we utilize. Section 6 describes our Monte Carlo simulations and our results. Section 7 concludes.

2. Empirical Evidence of Inflation Expectations

Survey data are one of the tools commonly used to identify economic agents’ inflation expectations.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 described earlier is clear: in times of generally rising inflation, such as the 1970s, expected inflation tends to underpredict realized inflation. In contrast, in times of falling inflation, such as the 1980s and 1990s, the forecasts appear to overpredict 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:

$$\pi_t = a_0 + a_1 \pi_t^e + \epsilon_t, \quad (1)$$

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

⁵Other methods include uncovering inflation expectations from futures market data (as in Dotsey and DeVaro 1995) or comparing yields on inflation-indexed and non-indexed treasuries (see Shen and Corning 2001).

⁶The 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 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) describes the three surveys. The Conference Board of Canada also has produced, since 1988, survey data on (Canadian) inflation expectations.

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

of Thomas' (1999) regressions. Run with data from the Livingston survey, it shows the following estimated equation:

$$\pi_t = 0.134 + 0.88E_{t-2}[\pi_t], \quad (2)$$

(0.41) (0.08)

where the sample used is 1980Q3 to 1997Q4. 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

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 unbiasedness hypothesis could simply be owing to a small sample problem. As stated in section 1 and described in section 6, this is precisely what our results imply.

Interestingly, once it is defined with a quasi-difference specification, the unbiasedness hypothesis continues to be rejected in large samples.¹⁰ Future research might investigate whether this facet of the relationship between realized and expected inflation could be replicated by our incomplete information and learning environment.

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

⁹A 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 being 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:

$$\pi_t = 0.29 + 0.77E_{t-4}[\pi_t],$$

(0.29) (0.11)

with, again, an easy rejection of H_0 .

¹⁰Consider estimating the following regression:

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

and testing $H_0 : a_0 = 0, a_1 = 1$. Under the null hypothesis, both this regression and (1) are identical. Nevertheless, Dolar and Moran (2002) report that the evidence against the null hypothesis is much more robust using this regression. Furthermore, the estimates of a_1 arising from the regression are almost always between zero and one. 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 for the frequent empirical rejections of $H_0 : b_0 = 0, b_1 = 1$ obtained from this regression. See Froot and Thaler (1990) and Taylor (1995) for a discussion.

3. The Model

The model we used is very similar to the one developed by Christiano and Gust (1999). We therefore provide only an overview and refer interested readers to the original paper. 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 (2001) use nominal price and wage stickiness to achieve this non-neutrality. We could redo our analysis with these nominal rigidities, but the robustness of our results, described in section 6.4, 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 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:

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

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 labour 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 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:

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

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, labour supply, and financial asset allocation to maximize their lifetime utility. Hence, they solve the

¹¹In 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' liquid asset holdings, while \overline{M}_t denotes the total supply of such assets.

following problem:

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

where $U(., .)$ is the period utility function and 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 .

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

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

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

3.2 Firms

Firms combine labour 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:

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

Firms rent capital and hire labour 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:

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

In contrast, it is assumed that a given fraction (denoted $1 - J_t$) of the firms' wage costs must be paid in advance. 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 labour demand:¹³

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

The evolution of J_t (which we call a money-demand shock) is exogenous and obeys the following:

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

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

¹³Because the production function features constant returns to scale, these efficiency conditions also hold for the aggregate values of capital and labour demand in the economy. Hereafter, K_t and N_t represent those aggregate quantities.

3.3 Financial intermediaries

Financial intermediaries accept deposits from households and lend the receipts to firms. Furthermore, they are the recipients of any liquid assets that the central bank injects into the economy to support its monetary policy rule.¹⁴ The revenues of the intermediaries are therefore 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 into the economy of liquid assets by 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 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 $(P_{t+k}, W_{t+k}, R_{t+k}, r k_{t+k})|_{k=0}^\infty$, 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 the following market-clearing equilibrium conditions are met:

$$\begin{aligned} c_t + K_{t+1} - (1 - \delta)K_t &= Y_t; \\ M_t^c + M_t^d &= \overline{M}_t; \\ M_t^d + X_t &= B_t = W_t N_t; \\ N_t &= n_t; \\ K_t &= k_t. \end{aligned}$$

4. Monetary Policy

4.1 The monetary policy rule

Monetary policy 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:

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

where r^{ss} is the steady-state value of the real interest rate, π_t^T is the inflation target of the monetary policy authorities at time t , and y_t represents the output gap.

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

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 policy 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:

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

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

Monetary policy authorities regularly deviate from their rule. These deviations (described in section 4.2) are called 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 to the following characterization of monetary policy:

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

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, ($\mu_t = \frac{M_{t+1}}{M_t}$), 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 assume that monetary policy, as expressed by the interest rate rule in (13), is subject to two types of disturbances. The first consists of the monetary policy shocks referred to in section 4.1 (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 to be errors stemming from the imperfect control exercised by monetary policy authorities over the growth rate of the money supply (μ_t). Under either interpretation, however, we envision that these shocks have little persistence. Accordingly, we assume that their evolution is governed by the following process:

$$u_{t+1} = \phi_1 u_t + e_{t+1}, \quad (14)$$

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

The second disturbance that monetary policy is subject to is as follows. We assume that, while remaining constant for extended periods of time, the monetary policy authorities' inflation target, π_t^T , is nevertheless subject to occasional, persistent shifts. We see two possible interpretations of these shifts. First, they could

correspond with changes in economic thinking that lead monetary policy 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 for inflation outcomes differ from their predecessor's. Under either interpretation, we envision that these shifts will have a significant duration, in the order of, say, five to ten years.

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:

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

with $0 \ll \phi_2 < 1$. 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 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.¹⁵

In section 4.3, we refer to some model simulations as representing complete information. By this we 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

To credibly communicate shifts in the inflation target might be difficult for monetary policy authorities. For example, although a new central bank head with a strong aversion to inflation might indicate this aversion in public announcements, economic agents may be uncertain as to what these announcements mean for the quantitative inflation targets. As a result, they might treat the announcements with skepticism and modify their beliefs about the monetary policy authorities' inflation target only after observing several periods of lower inflation. Announcements of

¹⁵It is left for future 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.

explicit, quantitative changes in the inflation 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 observe only 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:

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

Rewrite (16) by adding and subtracting π_0^T two times:

$$i_t = (1 - \rho)[r^{ss} + \pi_t^T + (\pi_0^T - \pi_0^T) + \alpha(\pi_t - \pi_t^T + \pi_0^T - \pi_0^T)] + \rho i_{t-1} + u_t, \quad (17)$$

or, rearranging terms,

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

Equation (18) illustrates that, from the viewpoint of an economic agent whose initial belief about the monetary policy authorities' inflation target 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 the transitory disturbance to rule u_t . The signal-extraction problem that economic agents face thus entails separating ϵ_t^* into its persistent and transitory components. Then, 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

As stated earlier, the signal-extraction problem is solved using the Kalman filter. The evolution of ϵ_t^* , the observed shock to the monetary policy rule in (18), can be

¹⁶Even 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 (2001), 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.

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

expressed within the following system:

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

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

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

$$N_{t+1} = \begin{cases} (1 - \phi_2) z_t, & \text{with probability } \phi_2; \\ g_{t+1} - \phi_2 z_t, & \text{with probability } 1 - \phi_2. \end{cases} \quad (21)$$

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, chapter 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. Furthermore, if the variables in the dynamic system are normal, the forecasts arising from the filter are optimal.¹⁸ We denote the forecasts of the two unobserved variables, given the information available at time t , as $\hat{z}_{t|t}$ and $\hat{u}_{t|t}$. Equations (19) and (20) can then be used to compute expected future deviations of the interest rate from the benchmark rule, as follows:

$$\begin{bmatrix} \hat{z}_{t+1|t} \\ \hat{u}_{t+1|t} \end{bmatrix} = \begin{bmatrix} \phi_2 & 0 \\ 0 & \phi_1 \end{bmatrix} \cdot \begin{bmatrix} \hat{z}_{t|t} \\ \hat{u}_{t|t} \end{bmatrix}; \quad (22)$$

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

Additional details on our implementation of the Kalman filter (which requires nothing more than to establish a formal correspondence between our notation and that by Hamilton) are provided in Appendix A.

The information friction that we assume is stronger, in some sense, than others often used in the literature—notably by Andolfatto and Gomme (1999), who 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

¹⁸Examination 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 behaviour will appear only after a very large number of data have been observed.

and usually produces quick transition of the beliefs following regime shifts. We consider, however, 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

Three distinct areas of the model require calibration: the model itself (preferences, technology, etc.), the parameters of the interest rate rule (13), and the parameters governing the evolution of the shocks and shifts in the rule. 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 specified to be:

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

Under this specification of utility, no intertemporal smoothing motive is present in the labour 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 that expresses monetary policy 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.

We conduct two types of experiments regarding 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 addition, 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 . Because no similar values are established for the money-demand shock, we follow Christiano and Gust and apply the technology shock values to the process for J_t : we thus have $\rho_J = 0.95$ and $\sigma_J = 0.005$.

¹⁹See Kozicki and Tinsley (2001b, 165) for a similar argument.

²⁰See Greenwood, Hercowitz, and Huffman (1988) for further details.

The model is solved using the first-order approximation method and algorithms given in King and Watson (1998). Details of 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, Gali, and Gertler 2000, 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. Furthermore, 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 behaviour of monetary policy 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 the 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. 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 describe 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

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

²²We 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.

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 per cent, on an annualized basis, in the inflation target of monetary policy authorities. We also explore the consequences of lower (0.0025) and higher (0.01) values for this parameter.

One interpretation of the Romer and Romer (1989, 1994) dates is that they represent changes in the inflation target of the Federal Reserve, and therefore occurrences of z_t shifts.²³ Because seven such dates are identified over a 40-year sample, this would 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

To calibrate 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 per cent 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. Table 1 summarizes the calibration values that we use.

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 per cent to 3 per cent. The true z_t , along with agents' best estimate of that variable, appears in the top panel 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 the top panel of Figure 2 shows that the agents' best estimate of z_t starts to decline towards the true value. Agents also assign some weight, however, to the possibility that the observed disturbance was a transitory shock, u_t . Thus the middle panel of the graph shows that agents' best 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

²³In their papers, Romer and Romer analyze the minutes of FOMC deliberations, and identify dates at which the Federal Reserve Board decided to cause a recession to stop inflationary pressures. See Hoover and Perez (1994) and Leeper (1997) for a discussion of Romer and Romer's methodology and results.

²⁴Other 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 policy authorities over inflation, within an empirical expression of the classic Barro and Gordon (1983) model.

from the largely transitory u_t , and they become convinced that it must have come from a z_t shift. Accordingly, the agents' best estimate of z_t and u_t , respectively, converges towards the true value and to zero.

The bottom panel of Figure 2 shows the progress of the agents' estimate of the monetary policy authorities' annualized inflation target (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 per cent.

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 shows the impulse responses of the artificial economy following a shift in the monetary policy authorities' inflation target. The shift is identical to the one illustrated in Figure 2: the inflation target is lowered, at time $t = 5$, from 5 per cent on an annualized basis to 3 per cent.

In Panel A of Figure 3, the solid lines represent the case where agents have complete information about the shift. In contrast, the dashed lines represent the case where information friction is active and the learning mechanism governs the formation of expectations.

The solid lines indicate 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, during which all real aggregates are lower than they were in the full-information case. This occurs 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.

Panel B of Figure 3 illustrates the situation from a slightly different angle. The solid lines depict realized inflation, and the dashed lines expected inflation. The left graph in Panel B shows the complete-information case: apart from the initial surprise in the first period of the shock, economic agents have the correct inflation expectations. The right graph in Panel B depicts the incomplete-information case: inflation expectations lag actual inflation for several periods before converging. This feature replicates, in a qualitative fashion, the behaviour 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—overpredicted actual numbers for several quarters.

The gap between realized and expected inflation in the right graph of Panel B, Figure 3, is a direct result of the learning behaviour described in section 4. 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 they think that inflation might return to the initial, higher level of 5 per cent. Over time, agents become convinced that a shift has indeed occurred and their expectations converge to values that are closer to the actual ones.

In the simulations performed using 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. The main picture given by the graphs in Figure 3, however, 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 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, Monte Carlo simulations of the model economy are performed 1000 times.

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-quarters-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 perform the unbiasedness test described in section 2. Recall that this involves estimating the regression

$$\pi_t = a_0 + a_1 \pi_t^e + \varepsilon_t, \quad (24)$$

and testing the null hypothesis $H_0 : a_0 = 0; a_1 = 1$. For each simulation, we record the estimates \hat{a}_0 and \hat{a}_1 , as well as the appropriate test statistic about H_0 .²⁶

Figures 4 to 7 show the results of these simulations using the benchmark calibration. In each of those figures, the top panel is a histogram that depicts the estimates of a_0 across the 1000 replications. It also depicts the median of the estimates. The middle panel depicts the estimates of a_1 , again identifying the median. The bottom

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

²⁶Under 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-quarters-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 .

panel illustrates the results of the 1000 tests of H_0 , showing a histogram of the test statistic along with its median and the 5 per cent and 1 per cent 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 per cent. When the null hypothesis is true and the test is correctly specified, this fraction should be close to 5 per cent, the size of the test. On the other hand, one can interpret results where this fraction is significantly higher than 5 per cent to suggest that the learning effects reduce the capability of the test to properly identify unbiasedness.

6.3 Results for the benchmark case

Figures 4 and 5 illustrate the cases of one-quarter-ahead and four-quarters-ahead expectations, respectively, when information is complete. Figures 6 and 7 illustrate the cases of one-quarter-ahead and four-quarters-ahead expectations for the incomplete-information case, in which the information friction that we emphasized is activated.

The top panel of Figure 4 is a histogram that depicts 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, which depicts the estimates of a_1 , shows a median very close to the hypothesized value of 1. 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 per cent rejection region starts above 3. In fact, only 5.3 per cent of the simulations lead the test statistic to reject the null at better than the 5 per cent 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 depicts the case when inflation expectations are measured as four-quarters-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. The correction for serial correlation, however, makes rejections harder to achieve, so that the test statistic rejects the null hypothesis in only about 9 per cent of the cases, not drastically away from the 5 per cent size of the test.

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

Let us now examine 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 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 Figure 6 illustrates that this skewness is reflected in the number of times H_0 is rejected: more than 20 per cent 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-quarters-ahead expectations with incomplete information, results are similar to those in Figure 6: 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 per cent.

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 unbiasedness tests 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, at least initially, to be confused about the true intentions of monetary policy authorities. Section 6.4 analyzes the extent to which the qualitative nature of the results expressed in Figures 4 to 7 are sensitive to the calibration of the model.

6.4 Sensitivity analysis

To analyze the sensitivity of the results to modifications in the calibration, we redo the above analysis for several alternative specifications. Table 2 reports the results. Column one indicates the kind of departure from the benchmark calibration that is under study. Columns two and three indicate the frequency with which the unbiasedness hypothesis is rejected when one-quarter-ahead expectations are used, in the complete-information and incomplete-information cases, respectively. Columns four and five report the corresponding results when the four-quarters-ahead expectations are utilized. To facilitate the comparison, the results from the benchmark cases are repeated at the beginning of the table.

Table 2 gives the general impression that the complete-information cases generate rejections of H_0 as often, roughly, as the size of the tests implies. Particularly in column one, the fraction of rejections seldom departs significantly from the level (5 per cent) suggested by the size of the test. Although the numbers in column three do depart more significantly from 5 per cent, 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 does not seem to depend upon whether one-quarter-ahead or four-quarters-ahead expectations are used.

For specific cases, in the first three departures from the benchmark case, eliminating the response of monetary policy 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 of the monetary policy authorities' response to deviations of inflation from the target (an increase of α from 2.0 to either 3.0 or 4.0) does modify the results substantially. The unbiasedness hypothesis is then rejected around 35 per cent of the time when the information friction is active, while the corresponding numbers for the complete-information case increase only slightly. The frequency of rejections increases because a high value for the coefficient α acts like a multiplier on the monetary policy shift. This is best illustrated

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 in interest rates.²⁷

We also experiment with modifications to the processes governing the evolution of the two components to monetary policy. 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 with the benchmark case. Decreasing the duration closes that gap. Modifications to the persistence of the transitory shocks, the next departures from benchmark that we consider, do not modify the results markedly. Next, we experiment with changes in the variances of the shocks and shifts. These experiments show that increasing the variance in the shifts z_t , relative to the variance of the shocks u_t , increases 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.

The last series of modifications to the benchmark calibration, in Panel C of Table 2, bring interesting observations. First, increasing the number of repetitions for the benchmark case (to 2500 from 1000) 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. Second, Panel C 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 policy authorities, one does not expect that the introduction of those shocks would imply a distorted relationship between realized and expected inflation. The rejections of the unbiasedness hypothesis that we identify thus truly arise from the limited information about monetary policy shocks.²⁸

The last experiment shows that using 1000 repetitions of much longer samples (with 500 periods in each sample) drastically reduces the rejections, particularly for the incomplete-information case. This last experiment strongly suggests that

²⁷While 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 when the departure from benchmark analyzed is an increase in the standard deviation of the shifts themselves.

²⁸Of course, the belief that macroeconomic volatility in the last 40 years was solely the result of technology shocks would imply, in the environment we describe, 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 40 years. 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.

the rejections of the unbiasedness hypothesis may be the result of size distortions caused by a few significant shifts in monetary policy in the small samples typically used to study inflation expectations.

7. Conclusion

Figure 1, which graphs the realized and expected consumer price index inflation (as measured by the Livingston survey data), suggests 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 policy authorities. Apart from the periods immediately following the shifts, the expectations of economic agents do not appear to be completely out of line with realized inflation.

We have shown that these empirical features can be represented by modelling the shifts in a standard monetary DSGE model, and by assuming that agents must learn about the shifts over time. 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. Furthermore, the likelihood of rejection tends to be eliminated when the sample size increases, even though the information friction remains.

We acknowledge that the Kalman filter may not be optimal in small samples, for which the (asymptotic) normal behaviour 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. Furthermore, we must caution that we identify 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, our results 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 expectations.

This view, and thus the importance of incomplete information for modelling macroeconomic activity, could be reinforced by verifying that the incomplete information and learning framework replicates other facets of the relationship between realized and (measured) expected inflation. Notably, our framework could be used to determine whether 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, the framework could be used to determine whether the learning effects could lead simulated data to match

the dynamic correlation patterns linking realized and expected inflation.

The evidence of shifts in the parameters of the interest rate rule (see Clarida, Galí, and Gertler 2000) opens interesting avenues for future research. Straightforward modifications to our framework could be used to determine whether such shifts, and least-square learning about the shifts on the part of economic agents, are responsible for the evidence against unbiasedness and efficiency in measured inflation expectations.

References

Andolfatto, D. and P. Gomme. 1999. "Monetary Policy Regimes and Beliefs." Federal Reserve Bank of Cleveland Working Paper No. 9905R.

Andolfatto, D., S. Hendry, and K. Moran. 2000. "Inflation Expectations, Unemployment and Persistent Liquidity Effects." Bank of Canada. Manuscript.

Barro, R.J. and D.B. Gordon. 1983. "A Positive Theory of Monetary Policy in a Natural Rate Mode." *Journal of Political Economy* 91: 589–610.

Brunner, K., A. Cukierman, and A. Meltzer. 1980. "Stagflation, Persistent Unemployment and the Permanence of Economic Shocks." *Journal of Monetary Economics* 6: 467–92.

Christiano, L. and C. Gust. 1999. "Taylor Rules in a Limited Participation Model." National Bureau of Economic Research Working Paper No. 7017.

Clarida, R., J. Galí, and M. Gertler. 2000. "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory." *Quarterly Journal of Economics* 115(1): 147–80.

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

DeLong, J.B. 1997. "America's Peacetime Inflation: The 1970s." In *Reducing Inflation*, edited by C.D. Romer and D.H. Romer. The University of Chicago Press.

Dolar, V. and K. Moran. 2002. "Are Survey-Based, Inflation Expectations Biased? A Reconsideration of the Evidence in the Light of the Forward Discount Puzzle." Bank of Canada. Manuscript.

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

Erceg, C. and A. Levin. 2001. "Imperfect Credibility and Inflation Persistence." Finance and Economics Discussion Series 2001-45, Federal Reserve Board.

Froot, K. and R. Thaler. 1990. "Anomalies: Foreign Exchange." *Journal of Economic Perspectives* 4(3): 179–92.

Greenwood, J., Z. Hercowitz, and G. Huffman. 1988. "Investment, Capacity Utilization, and The Real Business Cycle." *American Economic Review* 78(3): 402–17.

-
- Hamilton, J.D. 1994. *Time Series Analysis*. New Jersey: Princeton University Press.
- Hoover, K.D. and S.J. Perez. 1994. "Post Hoc Ergo Propter Once More: An Evaluation of 'Does Monetary Policy Matters' in the Spirit of James Tobin." *Journal of Monetary Economics* 34: 47–73.
- King, R.G. and M.W. Watson. 1998. "The Solution of Linear Difference Systems Under Rational Expectations." *International Economic Review* 39(4): 1015–26.
- Kozicki, S. and P.A. Tinsley. 2001a. "What Do You Expect? Imperfect Policy Credibility and Tests of the Expectations Hypothesis." Federal Reserve Bank of Kansas City Working Paper 01-02.
- _____. 2001b. "Term Structure Views of Monetary Policy Under Alternative Models of Agent Expectations." *Journal of Economic Dynamics and Control* 25: 149–84.
- Leeper, E.M. 1997. "Narrative and VAR Approaches to Monetary Policy: Common Identification Problems." *Journal of Monetary Economics* 40: 641–57.
- Lewis, K.K. 1988. "The Persistence of the 'Peso Problem' when Policy is Noisy." *Journal of International Money and Finance* 7: 5-21.
- _____. 1989. "Changing Beliefs and Systematic Rational Forecast Errors with Evidence from Foreign Exchange." *American Economic Review* 79(4): 621–36.
- Moran, K. 1999. "Welfare Benefits of Lower Inflation: Accounting for the Transition." Bank of Canada. Manuscript.
- Muth, J.F. 1960. "Optimal Properties of Exponentially Weighted Forecasts." *Journal of the American Statistical Association* 55: 299–306.
- Nelson, E. 2000. "UK Monetary Policy, 1972-97: A Guide Using Taylor Rules." Bank of England Working Paper 2001-120.
- Owyang, M. and G. Ramey. 2001. "Regime Switching and Monetary Policy Measurement." Federal Reserve Bank of St. Louis. Manuscript.
- Roberts, J.M. 1997. "Is Inflation Sticky?" *Journal of Monetary Economics* 39(2): 173–96.
- Romer, C.D. and D.H. Romer. 1989. "Does Monetary Policy Matter? A New Test

in the Spirit of Friedman and Schwartz.” *NBER Macroeconomics Annual* 121–170. Cambridge: MIT Press.

Romer, C.D. and D.H. Romer. 1994. “Monetary Policy Matters.” *Journal of Monetary Economics* 34: 75–88.

Ruge-Murcia, F. 2001. “Inflation Targeting under Asymmetric Preferences.” Université de Montréal Working Paper No. 2001-04.

Shen, P. and J. Corning. 2001. “Can TIPS Help Identify Long-Term Inflation Expectations?” *Reserve Bank of Kansas City Economic Review*, Fourth Quarter.

Sill, K. and J. Wrase. 1999. “Exchange Rates and Monetary Policy Regimes in Canada and the U.S.” Federal Reserve Bank of Philadelphia Working Paper 99-13.

Taylor, J.B. 1993. “Discretion Versus Policy Rules in Practice.” *Carnegie-Rochester Conference Series on Public Policy* 39: 195–214.

Taylor, M.P. 1995. “The Economics of Exchange Rates.” *Journal of Economic Literature* 33: 13–47.

Thomas Jr., L.B. 1999. “Survey Measures of Expected U.S. Inflation.” *Journal of Economic Perspectives* 13(4): 125–44.

Table 1. Parameter Calibration

Parameter	Symbol	Benchmark Value	Range Examined
Preferences and Technology			
Elasticity of labour supply	ψ_1	0.4	-
Scaling of labour 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]

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

Specification Examined	One Quarter Ahead		Four Quarters Ahead	
	Complete	Incomplete	Complete	Incomplete
Benchmark case	5.3%	20.5%	8.9%	26.3%
<i>Panel A: Modifications to the Monetary Policy Rule</i>				
No response to output gap ^a	5.5%	19.9%	8.5%	26.2%
No smoothing of interest rates ^b	5.5%	22.2%	9.4%	29.5%
Increased smoothing of interest rates ^c	5.1%	18.0%	7.8%	24.3%
More aggressive response to inflation ^d	5.7%	34.7%	10.6%	34.7%
Most aggressive response to inflation ^e	7.2%	37.7%	12.7%	34.6%
<i>Panel B: Modifications to the Calibration of Monetary Shocks and Shifts</i>				
Very high duration of regime shifts ^f	5.2%	22.4%	6.7%	35.3%
Lower duration of regime shifts ^g	5.6%	15.4%	10.9%	20.6%
No persistence in transitory shocks ^h	4.7%	19.8%	8.7%	27.8%
Higher persistence of transitory shocks ⁱ	5.8%	19.4%	8.4%	25.5%
Higher variance of regime shifts ^j	5.6%	37.9%	15.1%	34.6%
Lower variance of regime shifts ^k	4.8%	8.1%	7.3%	19.4%
Lower variance of transitory shocks ^l	5.6%	33.9%	15.2%	34.8%
<i>Panel C: Other Modifications</i>				
Increased number of repetitions ^m	5.8%	21.8%	9.6%	28.6%
Higher portfolio adjustment costs ⁿ	5.3%	19.6%	9.0%	26.5%
Inclusion of money-demand shocks	5.4%	20.1%	8.8%	26.1%
Inclusion of technology shocks	5.4%	6.9%	9.2%	10.1%
Increased length of simulated time series ^o	5.0%	10.8%	2.7%	4.9%

^a $\beta = 0.0$

^b $\rho = 0.0$

^c $\rho = 0.5$

^d $\alpha = 3.0$

^e $\alpha = 4.0$

^f $\phi_2 = 0.99$

^g $\phi_2 = 0.95$

^h $\phi_1 = 0.0$

ⁱ $\phi_1 = 0.2$

^j $\sigma_g = 0.01$

^k $\sigma_g = 0.0025$

^l $\sigma_e = 0.0025$

^m2500 repetitions using an 80-period sample and benchmark calibration

ⁿ $\tau = 15.0$

^o1000 repetitions using a 500-period sample

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




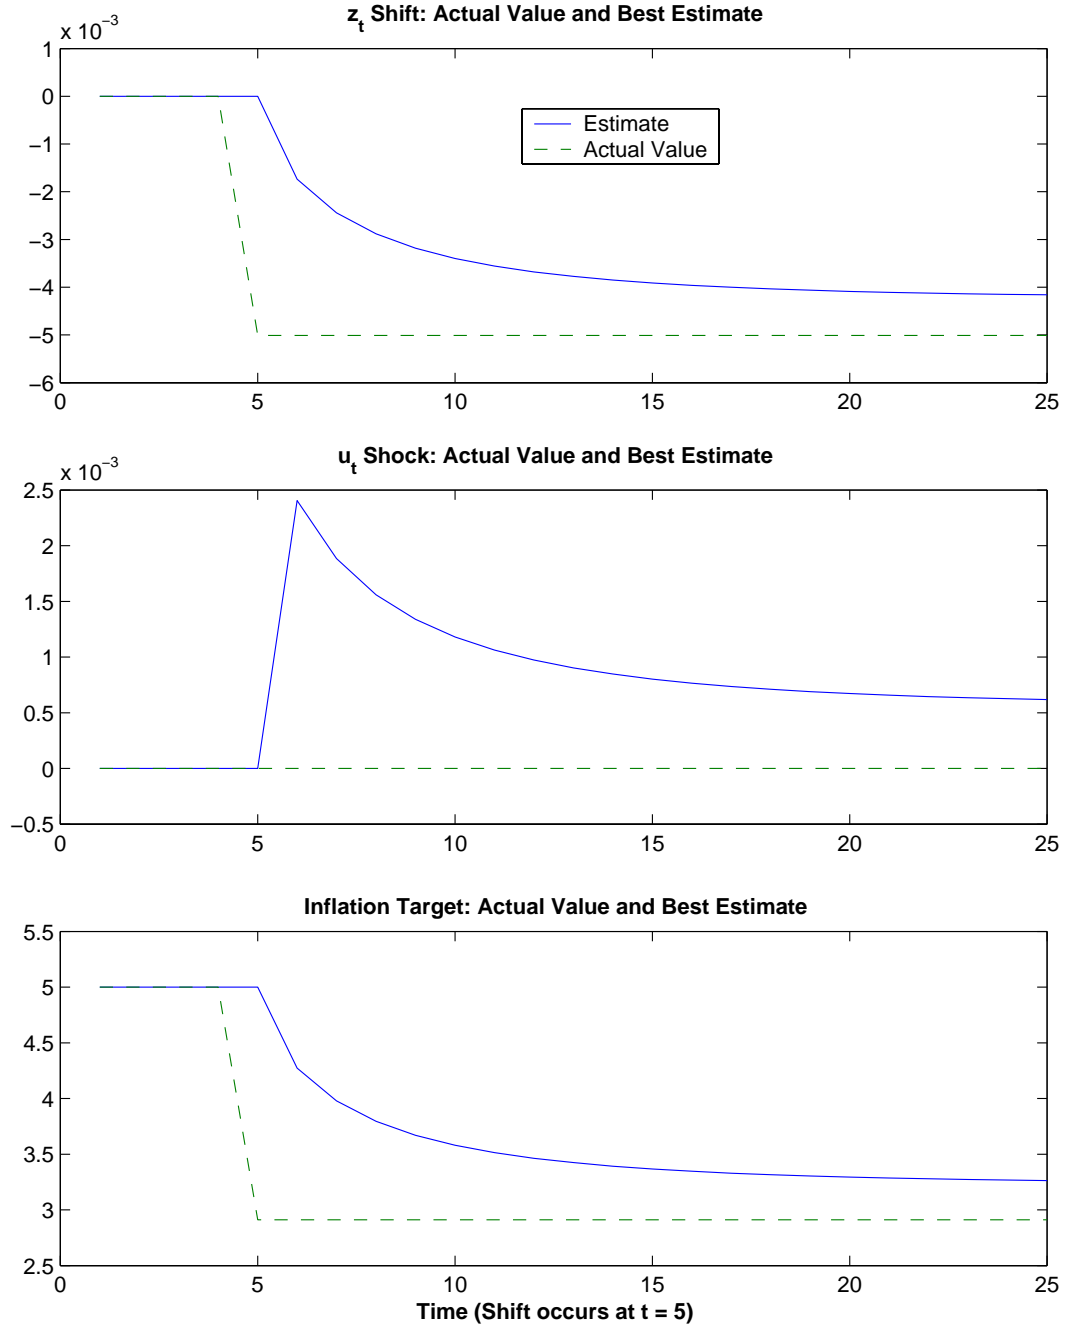
Figure 2: The Learning Mechanism: The Benchmark Calibration

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

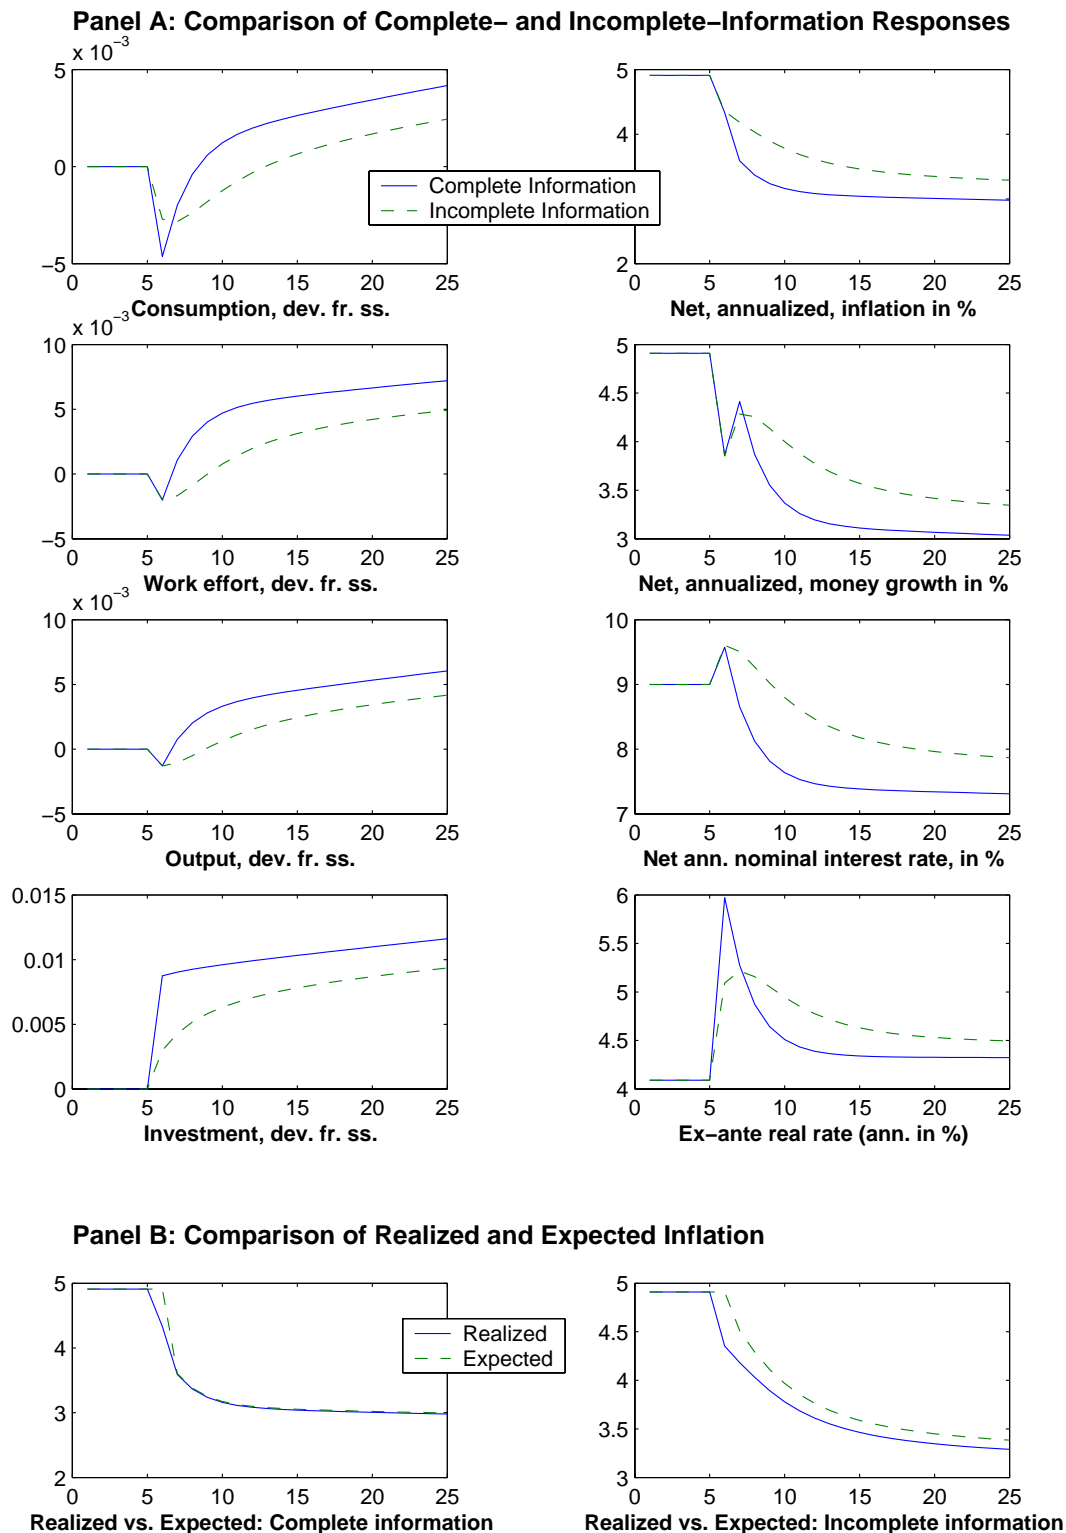


Figure 4: Results of Monte Carlo Simulations: One-Quarter-Ahead Expectations, Complete Information

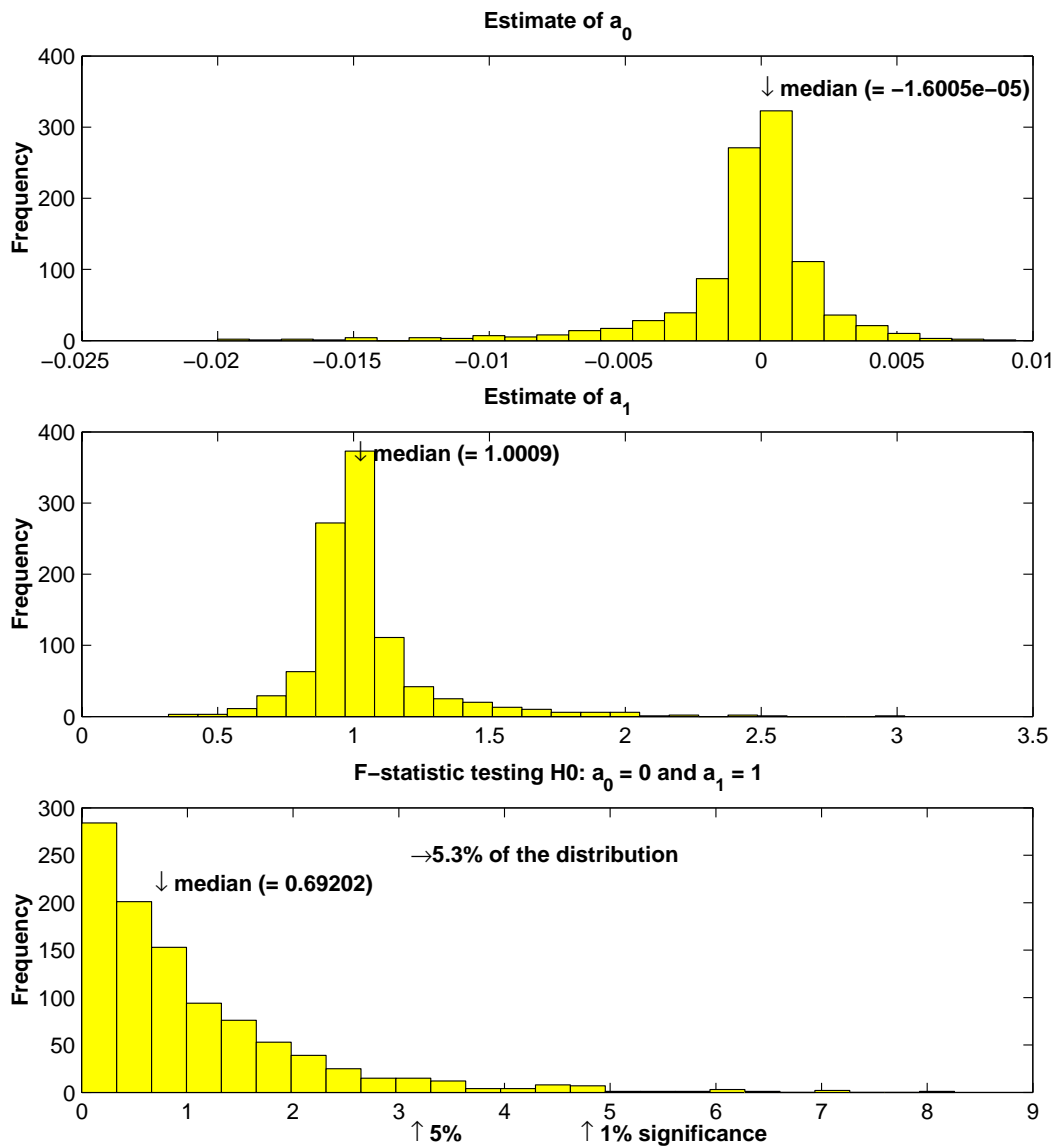


Figure 5: Results of Monte Carlo Simulations: Four-Quarters-Ahead Expectations, Complete Information




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

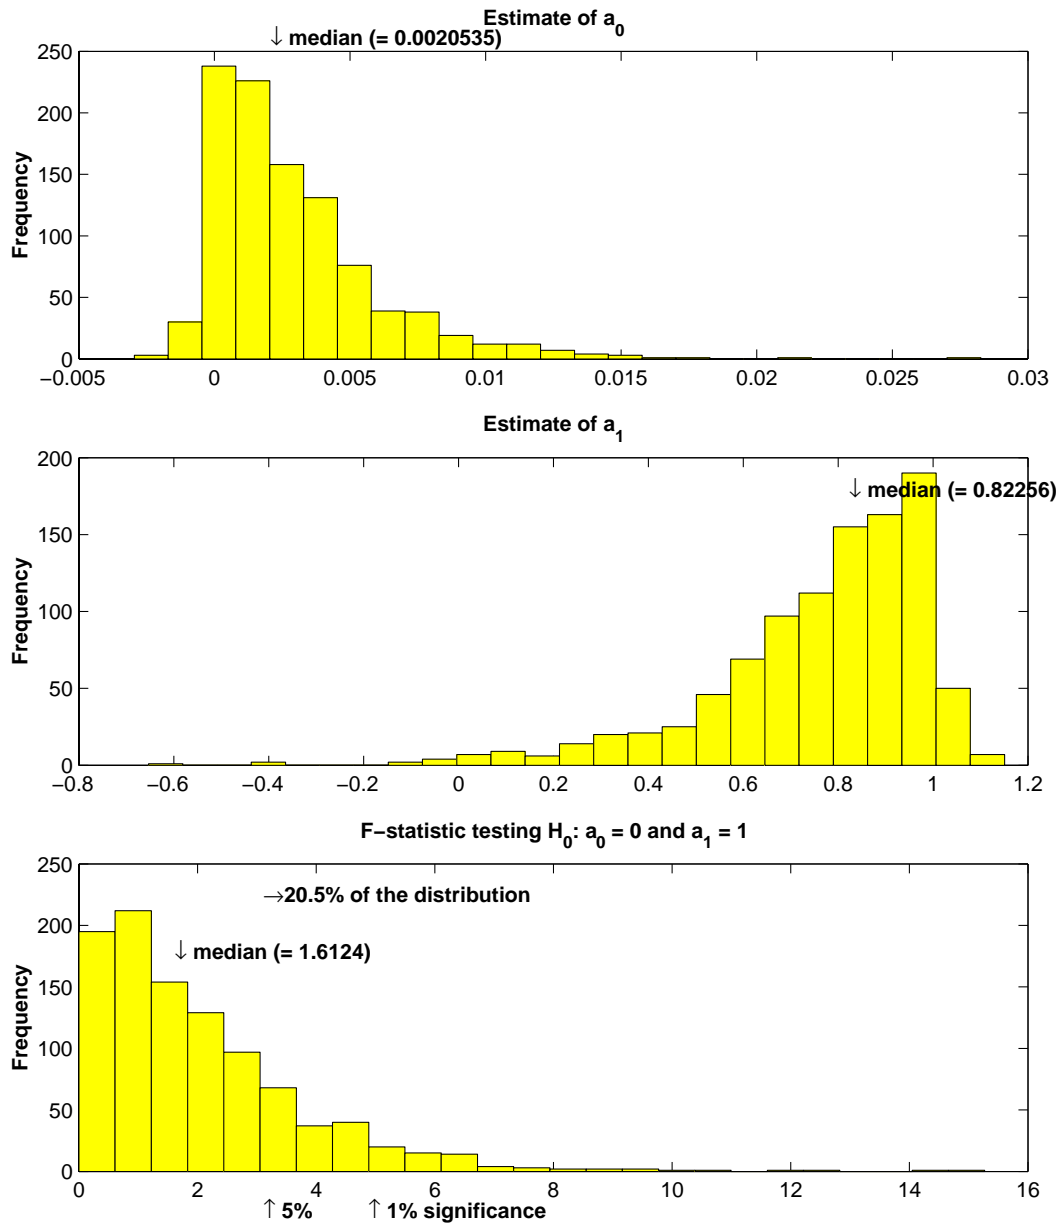
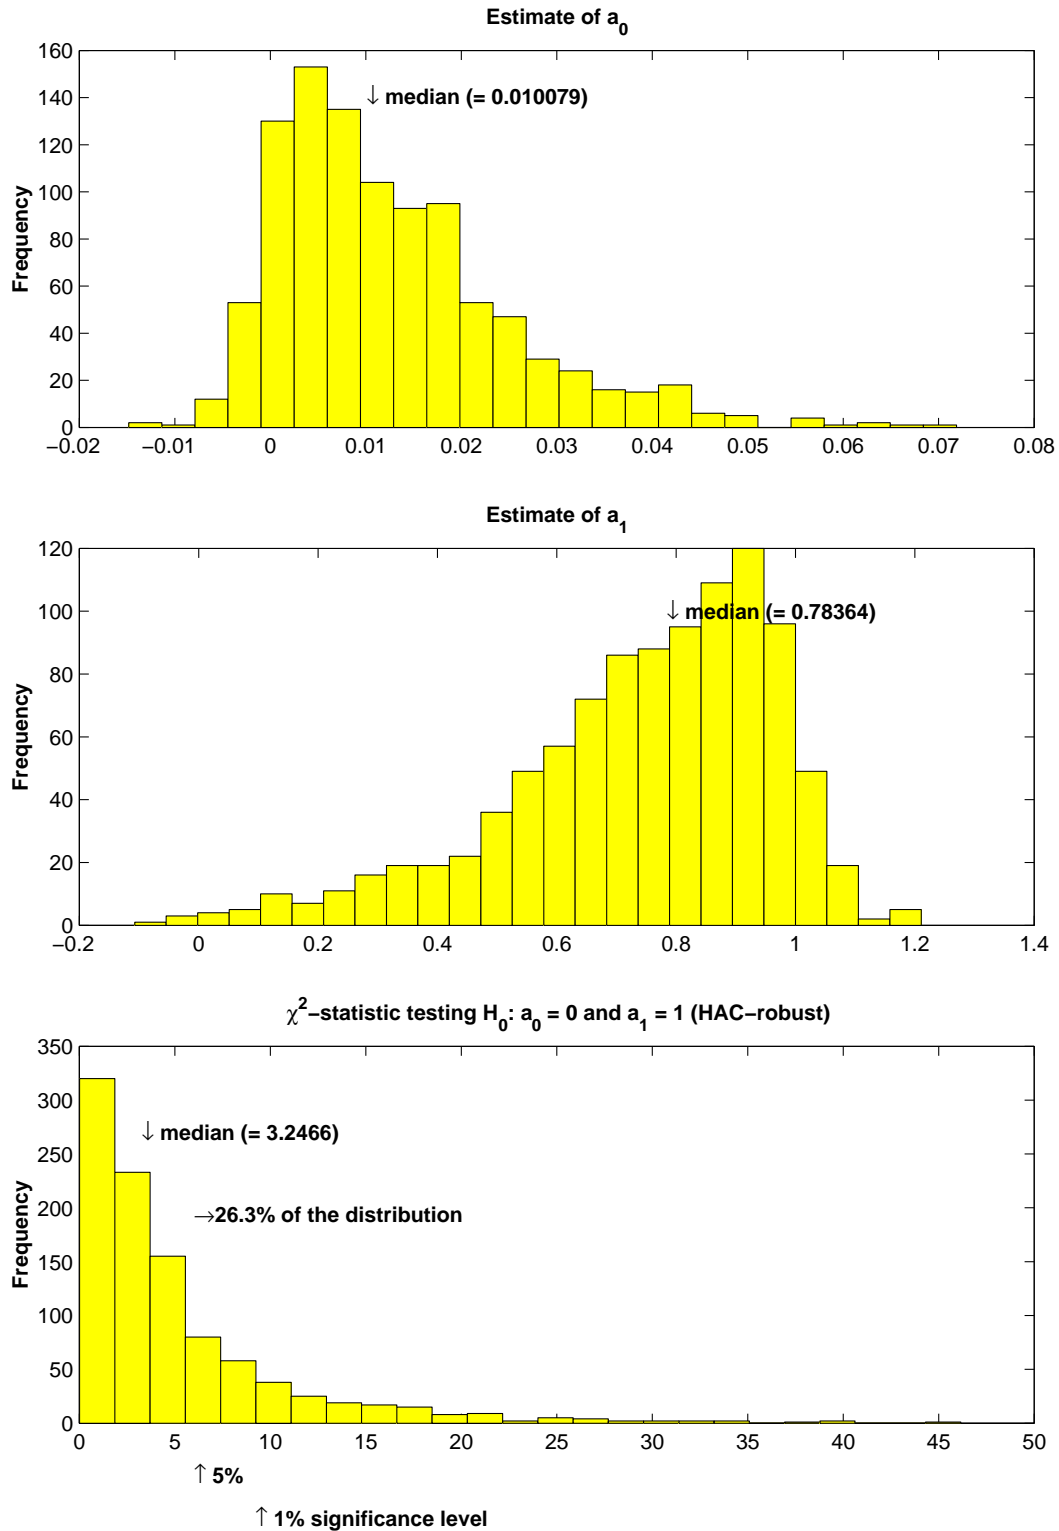


Figure 7: Results of Monte Carlo Simulations: Four-Quarters-Ahead Expectations, Incomplete Information



Appendix A: Kalman Filter

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:

$$N_{t+1} = \begin{cases} (1-\phi_2)z_t & \text{with probability } \phi_2; \\ g_{t+1} - \phi_2 z_t & \text{with probability } 1-\phi_2; \end{cases} \quad (\text{A.3})$$

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 with the one described in Hamilton's (1994, chapter 13) discussion of state-space models and the Kalman filter:

$$\begin{aligned} y_t &= \mathbf{A}' \cdot \mathbf{x}_t + \mathbf{H}' \cdot \xi_t + \mathbf{w}_t; \\ \xi_{t+1} &= \mathbf{F} \cdot \xi_t + \mathbf{v}_{t+1}; \\ E(\mathbf{v}_t \mathbf{v}_t') &= \mathbf{Q}; \\ E(\mathbf{w}_t \mathbf{w}_t') &= \mathbf{R}. \end{aligned} \quad (\text{A.4})$$

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 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 $P_1|_0$ ³⁰, the following recursive structure that describes the evolution of $\hat{\xi}_{t+1|t}$ and $P_{t+1|t}$ emerges:

$$\mathbf{K}_t = \mathbf{F} \mathbf{P}_t|_{t-1} \mathbf{H} (\mathbf{H}' \mathbf{P}_t|_{t-1} \mathbf{H})^{-1}; \quad (\text{A.5})$$

$$\hat{\xi}_{t+1|t} = \mathbf{F} \hat{\xi}_t|_{t-1} + \mathbf{K}_t (y_t - \mathbf{H}' \hat{\xi}_t|_{t-1}); \quad (\text{A.6})$$

$$\mathbf{P}_{t+1|t} = (\mathbf{F} - \mathbf{K}_t \mathbf{H}') \mathbf{P}_t|_{t-1} (\mathbf{F}' - \mathbf{H} \mathbf{K}_t') + \mathbf{Q}. \quad (\text{A.7})$$

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

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

³⁰We use the unconditional expectations.

states, ξ_t . The mechanics of this updating take the form of linear projection and are detailed in Hamilton.

Under the assumption of normality of both processes (for z_t and u_t), the sequence $(\hat{\xi}_{t+1|t})_{t=1}^T$ represents the optimal one-step-ahead forecasts of the unobserved states.³¹

³¹Even without the assumption of normality, the sequence of filtered estimates remains the best linear forecasts of the unobserved states conditional on time- t information.

Bank of Canada Working Papers

Documents de travail de la Banque du Canada

Working papers are generally published in the language of the author, with an abstract in both official languages. *Les documents de travail sont publiés généralement dans la langue utilisée par les auteurs; ils sont cependant précédés d'un résumé bilingue.*

2002

2002-29	Exponentials, Polynomials, and Fourier Series: More Yield Curve Modelling at the Bank of Canada	D.J. Bolder and S. Gusba
2002-28	Filtering for Current Analysis	S. van Norden
2002-27	Habit Formation and the Persistence of Monetary Shocks	H. Bouakez, E. Cardia, and F.J. Ruge-Murcia
2002-26	Nominal Rigidity, Desired Markup Variations, and Real Exchange Rate Persistence	H. Bouakez
2002-25	Nominal Rigidities and Monetary Policy in Canada Since 1981	A. Dib
2002-24	Financial Structure and Economic Growth: A Non- Technical Survey	V. Dolar and C. Meh
2002-23	How to Improve Inflation Targeting at the Bank of Canada	N. Rowe
2002-22	The Usefulness of Consumer Confidence Indexes in the United States	B. Desroches and M-A. Gosselin
2002-21	Entrepreneurial Risk, Credit Constraints, and the Corporate Income Tax: A Quantitative Exploration	C.A. Meh
2002-20	Evaluating the Quarterly Projection Model: A Preliminary Investigation	R. Amano, K. McPhail, H. Pioro, and A. Rennison
2002-19	Estimates of the Sticky-Information Phillips Curve for the United States, Canada, and the United Kingdom	H. Khan and Z. Zhu
2002-18	Estimated DGE Models and Forecasting Accuracy: A Preliminary Investigation with Canadian Data	K. Moran and V. Dolar
2002-17	Does Exchange Rate Policy Matter for Growth?	J. Bailliu, R. Lafrance, and J.-F. Perrault
2002-16	A Market Microstructure Analysis of Foreign Exchange Intervention in Canada	C. D'Souza
2002-15	Corporate Bond Spreads and the Business Cycle	Z. Zhang
2002-14	Entrepreneurship, Inequality, and Taxation	C.A. Meh

Copies and a complete list of working papers are available from:

Pour obtenir des exemplaires et une liste complète des documents de travail, prière de s'adresser à :

Publications Distribution, Bank of Canada
234 Wellington Street, Ottawa, Ontario K1A 0G9
E-mail: publications@bankofcanada.ca
Web site: <http://www.bankofcanada.ca>

Diffusion des publications, Banque du Canada
234, rue Wellington, Ottawa (Ontario) K1A 0G9
Adresse électronique : publications@banqueducanada.ca
Site Web : <http://www.banqueducanada.ca>