Forecasting Canadian Time Series with the New Keynesian Model *

Ali Dib
Bank of Canada

Mohamed Gammoudi
Bank of Canada

Kevin Moran
Université Laval

February 7, 2006

Abstract

The authors document the out-of-sample forecasting accuracy of the New Keynesian model for Canada. They estimate their variant of the model on a series of rolling subsamples, computing out-of-sample forecasts one to eight quarters ahead at each step. They compare these forecasts with those arising from simple vector autoregression (VAR) models, using econometric tests of forecasting accuracy. Their results show that the forecasting accuracy of the New Keynesian model compares favourably with that of the benchmarks, particularly as the forecasting horizon increases. These results suggest that the model could become a useful forecasting tool for Canadian time series. The authors invoke the principle of parsimony to explain their findings.

JEL classification: E32, E37, C12

Bank classification: Business fluctuations and cycles; Economic models; Econometric and statistical methods

*The authors wish to thank Takashi Kano, Lynda Khalaf, Pierre St-Amant, Greg Tkacz, Carolyn Wilkins, and seminar participants at Laval University, the Bank of Canada, and the 2005 SCE conference 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

New Keynesian models are becoming standard tools in applied macroeconomic analysis.[1] They are used widely to study the impact of shocks on economic activity and inform the decisions of monetary policy-makers in several central banks worldwide. These models are relevant because their optimizing environment coherently determines the time paths of aggregate variables in a framework suitable for monetary policy analysis. It has become common to estimate, rather than calibrate, the parameters of these models, using aggregate time series and standard econometric techniques.[2] The models, however, are seldom used to generate out-of-sample forecasts: evidence on the quality of these forecasts thus remains scarce.

To contribute to this evidence, this paper documents the out-of-sample forecasting properties of a New Keynesian model for Canada. Specifically, we develop a variant of the model, estimate it on a series of rolling subsamples, and compute out-of-sample forecasts one to eight quarters ahead at each step. We then compare these forecasts with those arising from vector autoregressions (VARs), using econometric tests of forecasting accuracy.

We find that the model's forecasting accuracy compares favourably with that of the VAR benchmarks, particularly as the forecasting horizon increases. Specifically, the model can forecast output, interest rates, and money as well as or better than the benchmarks, and its forecasts for inflation are no worse. Our results also suggest that a combination of the two sets of forecasts may have forecasting power that is superior to each set alone. Overall, our findings indicate that the New Keynesian class of models has the potential to become a useful forecasting tool for Canadian time series.

Using VARs as the benchmarks for comparing forecasts is natural, because the New Keynesian model itself can be written as a VAR whose parameters are restricted by non-linear

[1] New Keynesian models are dynamic stochastic general-equilibrium (DSGE) environments where monopolistically competitive firms set prices subject to various adjustment costs. They are built around a core that consists of a price-setting equation (the 'New Phillips curve'), an equation linked to intertemporal consumption smoothing, and a monetary policy rule. Although derived from the methodology for the real business cycle, their emphasis on nominal rigidities and monetary features makes them well suited to monetary policy analysis. Woodford (2003) provides a synthesis of the model's implications for monetary policy analysis.

[2] For example, Ireland (1997, 2001a, 2003, 2004) and Dib (2003a, b, 2006) estimate parameters using maximum likelihood; Christiano, Eichenbaum, and Evans (2005) do so by minimizing the distance between the model's impulse responses following monetary policy shocks and those computed with VARs; Smets and Wouters (2003) and Del Negro and Schorfheide (2004) employ a Bayesian strategy to compute the posterior distribution for the parameters.

constraints linked to the model's structure. Our forecasting experiments thus compare the out-of-sample forecasting properties of a restricted model with those of an unrestricted counterpart. Clements and Hendry (1998) discuss conditions under which better forecasting accuracy may be attained by the restricted, or parsimonious, model. This requires a trade-off between squared inconsistency (how 'wrong' the restrictions are) and sampling uncertainty (estimating a large number of parameters lowers precision) to favour the parsimonious specification.[3] This situation is more likely when the sample size for estimation is small and the forecasting horizon is high, as in monetary policy practice.

Evidence is emerging about the practical value of parsimony for forecasting. For example, Doan, Litterman, and Sims (1984), Ingram and Whiteman (1994), and, more recently, Del Negro and Schorfheide (2004) demonstrate that constraining the estimation of a VAR by employing a Bayesian strategy and priors linked to structural models[4] improves the VAR's forecasting accuracy. Working within the classical perspective, Ireland (2004) shows that a version of the real business cycle model estimated with maximum likelihood can have more accurate forecasts than simple VARs; Dolar and Moran (2002) verify that Ireland's results hold for Canada. Recent papers by Smets and Wouters (2004) and Boivin and Giannoni (2005) contribute to the emerging evidence about the good forecasting properties of New Keynesian models.

The rest of this paper is organized as follows. Section 2 develops our variant of the New Keynesian model. Section 3 discusses the model's estimation and provides estimation results for the first subsample. Section 4 describes our forecasting experiment and reports the results. Section 5 offers some conclusions.

2. Model

In this section we develop our variant of the New Keynesain class of models. The structure of the model is similar to that in Ireland (2003) and Dib (2006). Time is discrete and one model period represents a quarter. There are two sectors of production. The first sector, producing final goods, is competitive: firms take input prices as given and produce a homogeneous good

[3]In addition, problems associated with the small-samples properties of the more complex non-linear estimation must be relatively small.

[4]While Doan, Litterman, and Sims (1984) use the 'Minnesota prior' (random walks for all variables), Ingram and Whiteman (1994) derive priors from the basic real business cycle model, while those in Del Negro and Schorfheide (2004) arise from a simple New Keynesian model.

that they sell at flexible prices. Final-good production is divided between consumption and investment. Capital-adjustment costs restrict the accumulation of capital and thus influence investment choices. The firms in the second sector, which produce intermediate goods, operate under monopolistic competition. Each firm produces a distinct good for which it chooses the market price. Changes to the price of these goods are constrained by the Calvo (1983) mechanism, so that these prices are 'sticky.' Intermediate-good production requires capital and labour services, inputs for which the firms act as price-takers. The economy is closed.[5] The monetary authority's policy rule manages movements in the short-term nominal interest rate to respond to inflation deviations from its target, as well as to deviations of output and money growth from their trends.

2.1 Household

There exists a continuum of identical, infinitely lived households that derive utility from consumption, C_t, detention of real money balances, M_t/P_t, and leisure $(1-h_t)$, where h_t represents hours worked. A representative household's expected lifetime utility is described as follows:

$$U_0 = E_0 \sum_{t=0}^{\infty} \beta^t u(C_t, M_t/P_t, h_t), \tag{1}$$

where $\beta \in (0,1)$ is the discount factor and the single-period utility function is specified as:

$$u(\centerdot) = \frac{\gamma z_t}{\gamma - 1} \log \left(C_t^{\frac{\gamma-1}{\gamma}} + b_t^{\frac{1}{\gamma}} (M_t/P_t)^{\frac{\gamma-1}{\gamma}} \right) + \zeta \log(1-h_t), \tag{2}$$

where γ and ζ are positive structural parameters, and z_t and b_t are serially correlated shocks. As McCallum and Nelson (1999) show, the preference shock, z_t, resembles, in equilibrium, a shock to the IS curve of more traditional Keynesian analysis. On the other hand, b_t is interpreted as a shock to money demand. These shocks follow the first-order autoregressive processes:

$$\log(z_t) = \rho_z \log(z_{t-1}) + \varepsilon_{zt}, \tag{3}$$

[5] While this assumption deprives the model of the ability to capture information related to external factors, we believe that it does not invalidate our forecasting experiments, for two reasons. First, the VAR benchmarks that we compare the model to are also run using only Canadian time series. Second, Dib (2003b) shows that most estimates unrelated to open-economy features do not change when a model extended to comprise open-economy features is estimated using Canadian data.

and

$$\log(b_t) = (1 - \rho_b)\log(b) + \rho_b \log(b_{t-1}) + \varepsilon_{bt}, \tag{4}$$

where $\rho_z, \rho_b \in (-1, 1)$ and the serially uncorrelated innovations ε_{zt} and ε_{bt} are normally distributed with zero mean and standard deviations σ_z and σ_b, respectively.

The representative household enters period t with K_t units of physical capital, M_{t-1} units of nominal money balances, and B_{t-1} units of bonds. During period t, the household supplies labour and capital to the intermediate-good-producing firms, for which it receives total factor payment $R_{kt}K_t + W_t h_t$, where R_{kt} is the nominal rental rate for capital and W_t is the economy-wide nominal wage. Further, the household receives a lump-sum transfer from the monetary authority, X_t, as well as dividend payments, D_t, from intermediate-good-producing firms.[6] The household allocates these funds to consumption purchases, C_t, and investment in capital goods, I_t, (both priced at P_t) to money holdings, M_t, and to bond holdings, B_t, priced at $1/R_t$, where R_t denotes the gross nominal interest rate between t and $t+1$. The following budget constraint therefore applies:

$$P_t (C_t + I_t) + M_t + B_t/R_t \leq R_{kt}K_t + W_t h_t + M_{t-1} + B_{t-1} + X_t + D_t. \tag{5}$$

Investment increases the capital stock over time according to

$$K_{t+1} = (1 - \delta)K_t + I_t - \Psi(K_{t+1}, K_t), \tag{6}$$

where $\delta \in (0, 1)$ is the constant capital depreciation rate and $\Psi(.,.)$ is a capital-adjustment cost function specified as $\frac{\psi}{2}\left(\frac{K_{t+1}}{K_t} - \eta\right)^2 K_t$, where $\psi > 0$ is the capital-adjustment cost parameter and $\eta > 1$ is the growth rate of the economy. With this specification, both the total and marginal costs of adjusting capital are zero in the steady-state equilibrium.

The representative household chooses C_t, M_t, h_t, K_{t+1}, and B_t in order to maximize expected lifetime utility (1) subject to the budget constraint (5) and the investment constraint

[6] The transfer, X_t, is related to the monetary authority's management of short-term interest rates through its policy rule.

(6). The first-order conditions for this problem are as follows:

$$\frac{z_t C_t^{-\frac{1}{\gamma}}}{C_t^{\frac{\gamma-1}{\gamma}} + b_t^{\frac{1}{\gamma}} (M_t/P_t)^{\frac{\gamma-1}{\gamma}}} = \Lambda_t; \tag{7}$$

$$\frac{z_t b_t^{\frac{1}{\gamma}} (M_t/P_t)^{-\frac{1}{\gamma}}}{C_t^{\frac{\gamma-1}{\gamma}} + b_t^{\frac{1}{\gamma}} (M_t/P_t)^{\frac{\gamma-1}{\gamma}}} = \Lambda_t - \beta E_t \left(\frac{P_t \Lambda_{t+1}}{P_{t+1}} \right); \tag{8}$$

$$\frac{\zeta}{1-h_t} = \Lambda_t \frac{W_t}{P_t}; \tag{9}$$

$$\beta E_t \left[\frac{\Lambda_{t+1}}{\Lambda_t} \left(\frac{R_{kt+1}}{P_{t+1}} + 1 - \delta + \psi \left(\frac{K_{t+2}}{K_{t+1}} - \eta \right) \frac{K_{t+2}}{K_{t+1}} \right) \right] = \psi \left(\frac{K_{t+1}}{K_t} - \eta \right) + 1; \tag{10}$$

$$\frac{\Lambda_t}{R_t} = \beta E_t \left[\frac{P_t \Lambda_{t+1}}{P_{t+1}} \right]; \tag{11}$$

where Λ_t is the Lagrange multiplier associated with constraint (5).

As Ireland (1997) shows, combining conditions (7), (8), and (11) yields the following optimization-based money-demand equation:

$$\log(M_t/P_t) \simeq \log(C_t) - \gamma \log(r_t) + \log(b_t), \tag{12}$$

where $r_t = R_t - 1$ denotes the net nominal interest rate between t and $t+1$, γ is the interest rate elasticity of money demand, and b_t is the serially correlated money-demand shock described earlier.

2.2 The final-good-producing firm

The final good, Y_t, is produced by assembling a continuum of intermediate goods, $Y_{jt}, j \in (0,1)$, that are imperfect substitutes with a constant elasticity of substitution, θ. The aggregation function is defined as

$$Y_t \leq \left(\int_0^1 Y_{jt}^{\frac{\theta-1}{\theta}} \mathrm{d}j \right)^{\frac{\theta}{\theta-1}}, \qquad \theta > 1. \tag{13}$$

Final-good-producing firms behave competitively, maximizing profits and taking the market price of the final good, P_t, as well as the intermediate-good prices, $P_{jt}, j \in (0,1)$, as given. The maximization problem of a representative final-good-producing firm is therefore

$$\max_{\{Y_{jt}\}_{j=0}^1} \left[P_t Y_t - \int_0^1 P_{jt} Y_{jt} \mathrm{d}j \right],$$

subject to (13). The resulting input-demand function for the intermediate good, j, is

$$Y_{jt} = \left(\frac{P_{jt}}{P_t}\right)^{-\theta} Y_t. \tag{14}$$

Equation (14) represents the economy-wide demand for good j as a function of its relative price and of the economy's total output of final good Y_t. Competition in the sector and the constant-returns-to-scale production (13) imply that these firms make zero profits. Imposing the zero-profit condition leads to the following description of the final-good price index, P_t:

$$P_t = \left(\int_0^1 {P_{jt}}^{1-\theta} \mathrm{d}j\right)^{\frac{1}{1-\theta}}. \tag{15}$$

2.3 The intermediate-good-producing firm

The intermediate-good-producing firm j uses capital and labour services, K_{jt} and h_{jt}, respectively, to produce Y_{jt} units of good j, according to the following constant-returns-to-scale technology:

$$Y_{jt} \leq K_{jt}^{\alpha} \left(A_t \eta^t h_{jt}\right)^{1-\alpha}, \quad \alpha \in (0,1), \tag{16}$$

where $\eta > 1$ denotes the gross rate of labour-augmenting technological progress.[7] The presence of such growth implies a balanced growth path, so that output, investment, consumption, the real wage, capital, and real money balances all grow at the same rate, η. Thus, these variables must be linearly trended.[8] A_t describes an aggregate technology shock common to all firms. This shock follows a stationary first-order autoregressive process:

$$\log A_t = (1-\rho_A)\log(A) + \rho_A \log(A_{t-1}) + \varepsilon_{At}, \tag{17}$$

where $\rho_A \in (-1,1)$ is an autoregressive coefficient, $A > 0$ is a constant, and ε_{At} is normally distributed with mean zero and standard deviation σ_A.

Each intermediate-good-producing firm sells its output under monopolistic competition; the economy-wide demand for the good produced by producer j is given by (14). Following Calvo (1983), we assume that each firm is allowed to reoptimize its output price only at specific times. Specifically, with probability ϕ, the firm must charge the price that was in effect in the preceding period, indexed by the steady-state rate of inflation, π; with probability $1-\phi$, the

[7] Similarly, η^t represents the effect of trend productivity growth.

[8] It will be crucial to consider stochastic trends in these variables in future work.

firm is free to reoptimize and choose an unrestricted new price. On average, each firm therefore reoptimizes every $1/(1-\phi)$ periods.[9]

At time t, if firm j receives the signal to reoptimize, it chooses prices $\widetilde{P}_{jt}$, h_{jt}, and K_{jt} that maximize its discounted, expected (real) total profit flows for the period where it will not be able to reoptimize. The profit-maximization problem is as follows (θ^k represents the probability that $\widetilde{P}_{jt}$ remains in effect at $t+k$):

$$\max_{\{K_{jt},h_{jt},\widetilde{P}_{jt}\}} E_0 \left[\sum_{k=0}^{\infty} (\beta\phi)^k \Lambda_{t+k} D_{jt+k}/P_{t+k} \right],$$

with D_{jt+k}/P_{t+k}, the real profit flow at time $t+k$, and

$$D_{jt+k} = \widetilde{P}_{jt}\pi^k Y_{jt+k} - R_{kt+k}K_{jt+k} - W_{t+k}h_{jt+k}. \tag{18}$$

Profit maximization is subject to the demand for good j (14) and the production function (16) (to which the Lagrange multiplier $\Xi_t > 0$ is associated). The first-order conditions for K_{jt+k}, h_{jt+k}, and $\widetilde{P}_{jt}$ are:

$$\frac{R_{kt}}{P_t} = \alpha q_t \frac{Y_{jt}}{K_{jt}}; \tag{19}$$

$$\frac{W_t}{P_t} = (1-\alpha) q_t \frac{Y_{jt}}{h_{jt}}; \tag{20}$$

$$\widetilde{P}_{jt} = \frac{\theta}{\theta-1} \frac{E_t \sum_{k=0}^{\infty} (\beta\phi\pi^{-\theta})^k \Lambda_{t+k} Y_{t+k} q_{t+k} P_{t+k}^{\theta}}{E_t \sum_{k=0}^{\infty} (\beta\phi\pi^{1-\theta})^k \Lambda_{t+k} Y_{t+k} P_{t+k}^{\theta-1}}; \tag{21}$$

where $q_t \equiv \Xi_t/\Lambda_t$ is the real marginal cost of the firm.

The symmetry in the demand for their good implies that all firms that are allowed to reoptimize choose the same price, $\widetilde{P}_{jt}$, which we denote $\widetilde{P}_t$. Considering the definition of the price index in (15) and the fact that, at the economy's level, a fraction $(1-\phi)$ of intermediate-good-producing firms reoptimize, the aggregate price index, P_t, evolves according to

$$P_t^{1-\theta} = \phi(\pi P_{t-1})^{1-\theta} + (1-\phi)(\widetilde{P}_t)^{1-\theta}. \tag{22}$$

Equations (19) and (20) state that firms choose production inputs in order for their costs to equal the marginal product weighted times real marginal costs. Equation (21) relates the

[9] This specification of the Calvo mechanism follows Yun (1996). Alternatively, Christiano, Eichenbaum, and Evans (2005) assume that when the reoptimization signal is not received, the price is increased by the *preceding period*'s rate of inflation. Smets and Wouters (2003) implement a flexible specification that nests the two cases.

optimal price to the expected future price of the final good and to expected future marginal costs. Taking a first-order approximation of this condition and of (22), and then combining them, gives the model's New Keynesian Phillips curve:

$$\widehat{\pi}_t = \beta\widehat{\pi}_{t+1} + \frac{(1-\phi)(1-\beta\phi)}{\phi}\widehat{q}_t, \tag{23}$$

where a hatted variable denotes its deviation from the steady-state value. This expression relates the current period's inflation rate to its expected future value, as well as to its current marginal costs, an indicator of the strength of economic activity.

2.4 The monetary authority

As in Ireland (2003) and Dib (2006), we assume that the monetary authority manages the short-term nominal interest rate, R_t, to respond to deviations of inflation, $\pi_t \equiv P_t/P_{t-1}$, output, $y_t = Y_t/\eta^t$, and money growth, $\mu_t \equiv M_t/m_{t-1}$, from their steady-state equilibrium values.[10] This monetary policy rule is given by:

$$\log(R_t/R) = \varrho_\pi \log(\pi_t/\pi) + \varrho_y \log(y_t/y) + \varrho_\mu \log(\mu_t/\mu) + \log(v_t), \tag{24}$$

where R, π, y, and μ are the steady-state values of R_t, π_t, y_t, and μ_t, respectively. Further, v_t is a monetary policy shock that evolves according to

$$\log(v_t) = \rho_v \log(v_{t-1}) + \varepsilon_{vt}, \tag{25}$$

where $\rho_v \in [0,1)$ is an autoregressive coefficient and ε_{vt} is a zero-mean, serially uncorrelated shock with standard deviation σ_v. The monetary authority implements this rule with the appropriate lump-sum injection/withdrawal of money, X_t.

The policy coefficients ϱ_π, ϱ_y, and ϱ_μ are chosen by the monetary authorities. When $\varrho_\pi > 0$, $\varrho_y > 0$, and $\varrho_\mu = 0$, monetary policy follows a Taylor (1993) rule, in which nominal interest rates increase in response to deviations of inflation and output from their steady-state values.

In contrast, (24) states that monetary policy follows a modified Taylor (1993) rule that adjusts short-term nominal interest rates in response to changes in money growth as well as to deviations of inflation and output. In this case, a unique equilibrium exists as long as the sum of ϱ_π and ϱ_μ exceeds one.

[10] $y_t = Y_t/\eta^t$ is stationarized (lineary detrended) output.

Ireland (2003) interprets such a rule as a combination policy that influences a linear combination of the interest rate and the money-growth rate to control inflation. Alternatively, the money growth rate can be interpreted as an indicator of expected inflation or as a proxy for some omitted variables, such as the exchange rate or financial variables, to which monetary policy responds.

2.5 Symmetric equilibrium

In a symmetric equilibrium, all intermediate-goods-producing firms are identical. They make the same decisions, so $Y_{jt} = Y_t$, $\widetilde{P}_{jt} = \widetilde{P}_t$, $K_{jt} = K_t$, $h_{jt} = h_t$, $D_{jt} = D_t$. Let $\tilde{r}_{kt} \equiv R_{kt}/P_t$, $\tilde{w}_t \equiv W_t/P_t$, and $\tilde{m}_t \equiv M_t/P_t$ denote the real capital rental rate, the real wage, and real money balances, respectively. A symmetric equilibrium for this economy consists in a sequence of allocations $\{Y_t, C_t, I_t, \tilde{m}_t, h_t, K_t\}_{t=0}^{\infty}$, a sequence of prices and co-state variables $\{\tilde{w}_t, \tilde{r}_{kt}, R_t, \pi_t, \lambda_t, q_t\}_{t=0}^{\infty}$, and the stochastic processes for preference, money demand, technology, and monetary policy shocks. These allocations, prices, and shocks are such that (i) households, final-good-producing firms, and intermediate-good-producing firms optimize, (ii) the monetary policy rule (24) is satisfied, and (iii) the following market-clearing conditions are satisfied:

$$K_t = \int_0^1 K_{jt}\, dj; \tag{26}$$

$$h_t = \int_0^1 h_{jt}\, dj; \tag{27}$$

$$M_t = M_{t-1} + X_t; \tag{28}$$

$$B_t = 0; \tag{29}$$

$$Y_t = C_t + I_t. \tag{30}$$

Allowing for trend productivity growth in the production process (13) implies that Y_t, C_t, I_t, K_t, w_t, and m_t all grow at the same rate η in equilibrium. This parameter is estimated among the other model's structural parameters. In the equilibrium, most of the model's real variables inherit a deterministic trend, so we transform them by dividing by η^t to induce stationarity.[11]

Next, the steady-state of the system is computed, a first-order linear approximation of the equilibrium system around the steady-state values is formed, and Blanchard and Kahn's

[11] The transformed variables are: $y_t = Y_t/\eta^t$, $c_t = C_t/\eta^t$, $i_t = I_t/\eta^t$, $k_t = K_t/\eta^t$, $r_{kt} = \tilde{r}_{kt}/\eta^t$, $w_t = \tilde{w}_t/\eta^t$, $m_t = \tilde{m}_t/\eta^t$, $\lambda_t = \Lambda_t\eta^t$.

(1980) procedure is used to transform this forward-looking model into the following state-space solution:

$$\widehat{s}_{t+1} = \mathbf{\Phi}_1 \widehat{s}_t + \mathbf{\Phi}_2\, \varepsilon_{t+1}, \tag{31}$$

$$\hat{d}_t = \mathbf{\Phi}_3 \widehat{s}_t, \tag{32}$$

where $\widehat{s}_t$ is a vector of state variables that includes predetermined and exogenous variables; $\hat{d}_t$ is the vector of control variables; and the vector ε_{t+1} contains the random innovations.[12] The elements of matrices $\mathbf{\Phi}_1, \mathbf{\Phi}_2$, and $\mathbf{\Phi}_3$ depend on the model's structural parameters.

3. Estimation

3.1 Methodology

It is usual in the literature to calibrate the values of some of the model's parameters before estimating the values of the remaining ones, because the data used contain weak or no information about them. In light of this fact, we set the weight on leisure in the utility function, ζ, to 1.35, which implies that households spend around one-third of their non-sleeping time in market activities (work). The share of capital in production, α, and the depreciation rate, δ, are assigned values of 0.33 and 0.025, respectively; these values are commonly used in the literature. The degree of monopoly power in intermediate-goods markets, θ, is equal to 6, which implies a markup of 20 per cent in steady state: this matches values usually used in similar studies. Both Ireland (2001a) and Dib (2003b) remark that the capital-adjustment parameter, ψ, is difficult to estimate without data on capital stock. We fix this parameter to 15, as in Dib (2003b).[13]

The remaining 18 parameters are estimated using the maximum-likelihood procedure.[14] This requires that we select a subset of the control variables, $\hat{d}_t$, in (32) for which data are available, and select the appropriate rows of $\mathbf{\Phi}_3$. Next, the likelihood of the sample $\{\hat{d}_t\}_{t=1}^T$

[12]For any stationary variable x_t, $\hat{x}_t = \log(x_t/x)$ denotes the deviation of x_t from its steady-state value, x. In our specification, $\widehat{s}_t = \left(\hat{k}_t, \hat{m}_{t-1}, \hat{z}_t, \hat{b}_t, \hat{A}_t, \widehat{v}_t,\right)'$, $\hat{d}_t = \left(\hat{\lambda}_t, \hat{q}_t, \hat{m}_t, \hat{y}_t, \hat{R}_t, \hat{r}_{kt}, \hat{c}_t, \hat{\imath}_t, \hat{\pi}_t, \hat{w}_t, \hat{h}_t, \hat{\mu}_t\right)'$, and $\varepsilon_{t+1} = (\varepsilon_{zt+1}, \varepsilon_{bt+1}, \varepsilon_{At+1}, \varepsilon_{vt+1})'$. Appendix A lists the equilibrium conditions of the model, the steps involved in finding the steady state, and the linearized equations introduced into Blanchard and Kahn's (1980) algorithm.

[13]The estimated values are very robust and marginally affected by changing the calibrated value of ψ.

[14]These are $\beta, \gamma, \varrho_\pi, \varrho_\mu, \varrho_y, \rho_v, \sigma_v, \phi, A, \rho_A, \sigma_A, b, \rho_b, \sigma_b, \rho_z, \sigma_z, \pi$, and η.

is computed recursively using the Kalman filter (Hamilton 1994, chapter 13). The parameter values that maximize the likelihood are found using standard numerical procedures.[15]

Since the model is driven by four shocks, we estimate the model using data for four series, to avoid problems of stochastic singularity. We use Canadian data on output, inflation, a short-term interest rate, and real money balances. Output is measured by real final domestic demand that includes only personal consumption expenditures and gross private investment. Inflation is the gross rate of increase in the GDP deflator. The nominal interest rate is the rate of the three-month treasury bill. Real money balances are measured by dividing the M2 money stock by the GDP deflator. Output and real money balances are expressed in per-capita terms using the civilian population age 15 and over.[16]

Following Ireland (1997, 2004), we directly estimate the parameter η, which describes the growth rate of output and real money balances. This trend, however, is not shared by inflation and the nominal interest rate. We assume that these variables are trendless (stationary). Our treatment of trends therefore differs from that of Smets and Wouters (2003), who render data stationary before estimation by linearly detrending all series.[17]

We believe that the strategy pursued in this paper is particularly relevant in the context of a forecasting exercise. It enables us to produce forecasts for the log levels of the data directly, rather than forecasts for detrended series that must be transformed into forecasts for log levels.

3.2 Estimation results

3.2.1 Parameter estimates

The model's parameters are estimated using Canadian data from 1981Q1 to 2004Q1.[18] We estimate the model using several subsamples, from 1981Q3 to 1995Q4, 1981Q3 to 1996Q1, and 1981Q1 to 2004Q1. Since the results are similar for the different estimates, we report only those of the first subsample, from 1981Q1 to 1995Q4. Table 1 reports the maximum-likelihood estimates of the parameters, with their standard deviations and t-statistics. Almost all of the

[15]In addition to Dib (2003a,b, 2006) and Ireland (2003, 2004), this estimation method is used by Bergin (2003), Bouakez, Cardia, and Ruge-Murcia (2005), and several others. Ireland (2004) provides some of the details about the estimation procedure. We use the simplex algorithm, as implemented by *Matlab*.

[16]Appendix B provides additional details, notably the mnemonics, about the data.

[17]In future work, we aim to adopt a common stochastic trend that provides a possible explanation for changes in the low-frequency movements of macroeconomic variables.

[18]The sample starts at 1981Q3, since the Bank of Canada officially abandoned targeting the M1 growth rate in mid-1981.

estimated parameters are statistically significant and economically meaningful. The estimate of the discount rate, β, is 0.99, which implies an annual steady-state real interest rate of just over 4 per cent. The estimates of b, determining the steady-state ratio of real balances to consumption, is 0.5, whereas the constant elasticity of substitution between consumption and real balances, γ, is around 0.06, similar to that estimated by Dib (2003a) for the Canadian economy. The estimate of ϕ, the probability of not adjusting prices in the next period, is 0.63. Thus, on average, firms keep their prices unchanged, except for indexation, for about two quarters and a half. This estimate is very close to the closed-economy estimates of Dib (2003b).

The estimates of the monetary policy parameters are statistically significant, with the exception of ϱ_y. Specifically, the responses of monetary policy to inflation, output, and money growth (ϱ_π, ϱ_y, and ϱ_μ) are 0.75, 0.02, and 0.48, respectively.[19] The estimates of ρ_v and σ_v, the persistence coefficient and standard deviation of monetary policy shocks, are 0.20 and 0.006, respectively. Overall, the estimates of monetary policy parameters are similar to the estimates of Dib (2003b, 2006) for the Canadian economy. They indicate that, to achieve its objectives, the Canadian monetary authorities have responded significantly to inflation and money growth, and scarcely (if at all) to output deviations from its trend.

The autoregressive coefficient estimates indicate that the technology, money-demand, and preference shocks are relatively persistent, with the money-demand shock being the most persistent ($\rho_b = 0.994$). The standard deviation estimates suggest that the aggregate demand-side shocks (money demand and preferences) are the most volatile.

3.2.2 Impulse-response functions

To assess the model's performance, we briefly analyze the impulse-response functions drawn from the estimated model and its variance decomposition.[20] Figures 1 to 4 show the economy's responses to the four types of exogenous shocks, at the estimated parameter values. The response of output is measured as a deviation from its steady-state value, whereas the responses of the other variables are in net (annualized) percentage points.

Figure 1 plots the economy's response to monetary policy tightening; i.e., setting the innovation, ε_{vt}, to 0.01, a value close to its estimated standard deviation. Following the tightening, interest rates increase and return to steady state moderately fast (recall that the estimated

[19] Indeterminate equilibria do not occur as long as $\varrho_\pi + \varrho_\mu > 1$.

[20] Similar analysis is available elsewhere; see Dib (2003a), for example.

serial correlation in monetary policy shocks, ρ_v, is 0.20). Output, inflation, and money growth, in contrast, fall sharply on impact. Output and inflation return gradually to steady state, while money growth overshoots slightly in the following periods, and then converges back to steady state. This gradual return to steady state reflects the actions of the Calvo (1983) mechanism and the serial correlation of the shock. Notice that the negative, contemporaneous correlation between interest rates and money growth – the liquidity effect – is consistent with the evidence.[21]

Figure 2 shows the economy's responses to a money-demand shock (setting the innovation, ε_{bt}, to 0.01). The shock causes output and inflation to decrease only slightly on impact. Money growth increases sharply, however, to accommodate the increase in demand. Since the rule followed by the monetary authority includes a response to increases in money growth, the nominal interest rate increases slightly, which results in the slight decreases in output. These responses roughly match Poole's (1970) classic analysis, in which the monetary policy authority changes the short-term nominal interest rate in response to exogenous demand-side disturbances.

Figure 3 shows responses following a shock to technology (an increase in ε_{At} of 0.01). Output jumps on impact, while the nominal interest rate and inflation fall below their steady-state levels. Money growth responds positively to the shock before falling below its steady-state level after two quarters. The deflationary pressure brought about by the shock leads to a sustained easing of monetary policy; recall the monetary policy rule in (24). This mechanism serves to accommodate the shock and gradually increase output, which peaks three quarters after the shock. Therefore, the monetary authority's response helps the economy adjust to the supply-side disturbances.

Figure 4 shows the impulse responses to a 1 per cent increase in the preference shock, a disturbance to households' marginal utility of consumption. In response to this shock, output, the nominal interest rate, inflation, and money growth jump immediately above their steady-state levels before returning gradually to those levels. Because the estimates of the preference autoregressive coefficient, ρ_z, are relatively large, the computed impulse responses are highly persistent. To control the rises in output and inflation, the monetary authority increases short-

[21] Evidence also suggests that the responses of inflation and output to monetary policy shocks should be characterized by hump-shaped patterns, where the maximum impact on the variables is attained several periods after the shock. Christiano, Eichenbaum, and Evans (2005) show that adding several additional features to the model enables it to display these patterns. Because our emphasis is on the model's out-of-sample forecasting ability and we want to keep the model parsimonious, we do not use such a model in our experiments.

term interest rates slightly but persistently.

3.2.3 Variance decomposition

Table 2 reports the standard deviations, expressed in percentage terms, of output, real balances, inflation, and the nominal interest rate as computed from the data and the estimated model. In the data, output and real balances have standard deviations of 3.44 and 2.78 per cent, respectively. Inflation and the short-term nominal interest rate are less volatile; their standard deviations are less than 0.6 per cent. The table shows that the model (i) underpredicts the volatility of output, (ii) generates real balance volatility close to that observed in the data, and (iii) slightly overpredicts the volatility of inflation and the nominal interest rate.

To understand which of the four shocks are driving the results, Table 3 decomposes the forecast-error variances of output, real balances, inflation, and the nominal interest rate into the fractions that can be attributed to each of the shocks. The table shows that preference and technology shocks are the most important sources of fluctuations in output, both in the short and the long term. Monetary policy shocks also account for a smaller but significant fraction of output fluctuations in the short term. Monetary policy and technology shocks are the most important factors determining fluctuations in the inflation rate. Together, they account for around 80 per cent of fluctuations at the one-quarter-ahead horizon. Preference shocks do contribute to some inflation volatility, particularly at longer horizons. The great majority of fluctuations in interest rates are attributable to preferences shocks; the contribution of the other shocks, and of monetary policy shocks in particular, is not significant. Technology and money-demand shocks explain more than 90 per cent of the fluctuations in real money balances, while monetary policy shocks explain about 10 per cent of the short-term fluctuations.

Overall, the estimation results indicate that the New Keynesian model can provide a coherent explanation for how several types of shocks affect the economy. Next, we assess the model's out-of-sample forecasting properties.

4. The Model's Forecasting Properties

4.1 The experiment

We compute out-of-sample forecasts for the New Keynesian (NK), simple VAR, and Bayesian VAR (BVAR) models. The VAR(2) and BVAR(2) models are used as benchmarks.[22] The VAR model includes linear deterministic trends for output, real money balances, the nominal interest rate, and inflation. BVAR imposes the Minnesota prior of Doan, Litterman, and Sims (1984). We begin by estimating the models using data from 1981Q3 to 1995Q4. These estimates are used to produce forecasts one- to eight-quarters-ahead (i.e., for 1996Q1 to 1997Q4), for the four variables used. We next use data from 1981Q4 to 1996Q1 to update the estimates, and then produce another set of forecasts for 1996Q2 to 1998Q1. Estimates and forecasts are updated in this manner until the end of the available sample; to date, we have time series for one- to eight-quarter-ahead forecasts from 1996Q1 to 2004Q1. Table 4 reports the results of the forecasting experiment.

Figures 5 to 8 compare the forecasts with actual data for the period. Figure 5 compares the NK model's forecasts with actual data, and shows that the model provides what appears to be a relatively good characterization of output fluctuations for the one-quarter-ahead, four-quarter-ahead, and eight-quarter-ahead horizons. The model also maintains a reasonably balanced forecast for inflation, although the actual data exhibit some transitory fluctuations that are not well captured by the model. Further, the model is slow to incorporate the interest rate decreases of 2001 in its forecasts. The model's forecasts for money are reasonably accurate.

Figures 6 to 8 show the forecasting errors of the model (the solid line) with those arising from the VAR benchmark[23] (the dotted lines) for the case of one-quarter-ahead (Figure 6), four-quarter-ahead (Figure 7), and eight-quarter-ahead forecasts (Figure 8). In Figure 6, the two models appear to give forecasts that are roughly equivalent, except for output, where the VAR benchmark may produce smaller errors. At the four-quarter-ahead horizon (Figure 7), the NK model seems to outperform the benchmark for output, interest rates, and real money balances, whereas the inflation forecasts appear to be very close. At the eight-quarter-ahead horizon (Figure 8), the model's forecasting is better than the benchmark for output, interest rates, and money, while the inflation forecasts remain close.

[22]The number of lags in the VAR model are selected using Akaike's information criterion.

[23]This VAR contains a constant and a trend, as well as two lags for each variable; the Minnesota prior is not used.

The first column of Table 5 synthesizes the information contained in Figures 6 to 8. It reports the mean square error (MSE) of the New Keynesian model, relative to that of the VAR benchmark. Values smaller than one suggest that the NK model has superior forecasting accuracy, whereas values bigger than one favour the VAR benchmark. As suggested earlier, the MSEs tend to favour the NK model, particularly as the forecasting horizon increases. In particular, at the eight-quarter-ahead horizon, the model's MSE is around only 17 per cent of the VAR benchmark for output and less than 30 per cent for interest rates. Note, however, that this favourable performance does not apply to inflation, for which the VAR benchmark has slightly lower MSEs. This result may be explained by the fact that the Bank of Canada adopted inflation targeting in 1993. Table 5 also shows that, for very short-term horizons, the advantage for the NK models vanishes: the VAR benchmark appears to be more accurate in forecasting one-quarter-ahead output.

4.2 Econometric tests of forecasting accuracy

To test whether these improvements in MSE are statistically significant, we first use Diebold and Mariano's (1995) test. To compute the test, we define the forecast errors of the NK model as $\{e_t^M\}_{t=1}^T$ and those of the VAR(2) benchmark as $\{e_t^B\}_{t=1}^T$. Further, we define a sequence of 'loss differentials,' $\{l_t\}_{t=1}^T$, where $l_t = (e_t^B)^2 - (e_t^M)^2$. If the NK model is a better forecasting tool, one would expect that, on average, the loss differentials, l_t, would be positive. Conversely, one would expect negative values if the VAR benchmark is superior. Following this intuition, the Diebold and Mariano (1995) test considers the null hypothesis $H_0 : E[l_t] = 0$; positive values of the statistic suggest that the forecasts from the New Keynesian model have lower MSEs, while negative values favour the VAR benchmark. The test statistic (denoted DM) is asympotically normal and standard critical values are used.[24] Harvey, Leybourne, and Newbold (1997) propose a corrected Diebold and Mariano (1995) statistic, to reduce size distortions that might be significant in small samples. The corrected statistic is compared with a Student's t distribution with $N - 1$ degrees of freedom, where N is the number of forecast data.

The last two columns of Table 5 report the Diebold and Mariano (1995) and Harvey, Leybourne, and Newbold (1997) statistics, as well as their p-values in parentheses. Due to the small number of forecasts available (30 for the one-quarter-ahead forecasts, and 22 for the eight-quarter-ahead), it is not surprising that many test statistics are not significant. Nevertheless,

[24]The statistic is computed as $DM = l/\hat{\sigma}(l)$ where l is the sample average of l_t and $\hat{\sigma}(l)$ is a heteroscedastic and autocorrelation consistent estimate of the standard deviation of l.

Table 5 and Figures 6 to 8 show that the NK model's forecasting accuracy compares very favourably with that of the VAR(2) benchmark, performing better for output, interest rates, and money at longer-term horizons.[25] The New Keynesian model performs less well when forecasting inflation, probably because, as estimated, the model does not allow trends to affect inflation, when ample anecdotal or econometric evidence suggests that structural breaks have affected inflation over the past two decades. Section 5 discusses possible future research on the New Keynesian model to address this important issue.

Even if all the tests in Table 5 conclusively identified a superior model, forecasts from the lesser model could still contain information not present in the superior one; combining both forecasts could therefore further reduce the forecasting errors.[26] A more stringent examination of whether one model is superior to another in forecasting may be to test whether the second model contains *any* information not contained in the first model's forecasts.

In this context, Granger and Newbold (1973) define the forecasts from one model as "conditionally efficient" when they can be combined with those from another model. Chong and Hendry (1986) define the same situation as one in which the first set of forecasts "encompasses" those from the second model: there is no need to keep the second model's forecasts, because the information they contain is encompassed by those of the first model.

To implement the test for forecast encompassing, we follow Harvey, Leybourne, and Newbold (1998), who propose test statistics similar to those in Diebold and Mariano (1995) and the Harvey, Leybourne, and Newbold (1997) correction. The null hypothesis is that the New Keynesian model's forecasts contain no information that isn't already contained in those from the VAR.[27]

Table 6 reports the results. The first column shows the test statistics as proposed by

[25]This favourable performance is also obtained when the New Keynesian model is compared with a VAR with one lag in each variable. The results are available from the authors on request.

[26]For example, the lesser model might outperform the first in specific periods, such as when the economy is in recession.

[27]Assume the following regression equation:

$$e_t^B = \gamma(e_t^B - e_t^M) + \epsilon_t,$$

where e_t^B and e_t^M represent the forecasting errors from the VAR benchmark model and the NK model, respectively. The null hypothesis is $H_0 : \gamma = 0$. Under the null, the errors made by the VAR benchmark cannot be explained (and thus potentially reduced) by information arising from the NK model. Conversely, one can test whether the NK forecasts encompass those from the VAR; i.e., whether there is *any* information in the VAR forecasts that is not present in the NK forecasts.

Diebold and Mariano (1995) and the second the correction proposed by Harvey, Leybourne, and Newbold (1997). Recall that high values of the test statistics reject the hypothesis that no value can be gained from using the NK forecasts when the VAR model is available. As with Table 5, the results in Table 6 suggest that VAR forecasts for output, interest rates, and money are improved when they are combined with those from the NK model, whereas one cannot reject the hypothesis of no value from the NK model for inflation.

Researchers have often pointed out that imposing the Minnesota prior – all variables follow simple random walks – on the Bayesian estimation of simple VARs results in superior forecasting accuracy. In this context, Table 7 repeats the results of Table 5, but with a BVAR(2) serving as the benchmark with which to compare the NK model. A comparison of the two tables shows that the forecasting accuracy of the BVAR is indeed often superior to what it was without the priors (the MSE of the NK model is often higher than it was in Table 5). Nevertheless, similar observations about the model's forecasting properties can be made: in particular, the NK model's forecasting accuracy for output, interest rates, and money compares very favourably with that for the BVAR benchmark, and, as the forecasting horizon increases, several of these differences become statistically significant.

Note that the comparisons conducted have been between a model in which inflation and interest rates are restricted to having no trends (the NK model) and a model in which such trends are present (the VAR benchmark). A better comparison might be between two models for which inflation and interest rates are restricted to contain no significant linear trends. Repeating the analysis using a VAR model where the deterministic trend components have been taken out of the equations for inflation and interest rates reinforces the results reported in Tables 5 to 7.[28]

5. Conclusion

Since the introduction of the real business cycle models, researchers have often identified dimensions along which these structural models seem at odds with features of observed data.[29] Further, researchers who extend the simple real business cycle structure to New Keynesian models that feature nominal rigidities and multiple sources of volatility have had difficulties

[28]Further, using the sign test to compare the forecasts from the two models does not modify our overall conclusions.

[29]For example, Cogley and Nason (1995) show that the simple real business cycle model could not match the autocorrelation function of output or the impulse responses of Blanchard and Quah (1989).

replicating features of the data, such as the strong autocorrelation properties of inflation or output.

In this context, the evidence that structural New Keynesian models may have comparable or even better out-of-sample forecasting accuracy than unrestricted VARs is surprising.[30] The evidence suggests that restricted or parsimonious specifications, although at odds with some features of the data, may often outperform unrestricted alternatives in out-of-sample exercises. Clements and Hendry (1998, 1999) assess this conjecture. The main trade-off they identify is that of sampling variability (introduced in the unrestricted specification by estimating numerous parameters) versus inconsistency (introduced in parsimonious models by imposing possibly false restrictions).

Overall, our results suggest that this trade-off may be favourable to parsimonious specifications similar to the New Keynesian model. Such findings are encouraging for researchers who work with models of this type. The econometric tests we report make clear that, at a minimum, restricting a VAR by appealing to the structure of the New Keynesian model has no negative impact on its forecasting performance. In the case of output, interest rates, and money, the restricted model may in fact have superior forecasting accuracy, particularly as the horizon one is interested in increases. As the results show, the forecasting properties of the model for inflation are not as strong, although they are not significantly worse than those of the simple benchmark VARs.

The forecasting properties of the NK model for inflation could likely be improved if deterministic trends in the inflation rate could be introduced. This would allow the model to better track the apparent downward trend in inflation over the past 20 years. It would also be interesting to study whether using the common stochastic trend, rather than the linear trend, would improve the model's forecasting ability. Positive results in this area would imply that real variables such as output would be differentiated to make them stationary. The natural benchmark with which to compare forecasts would then be a VAR in differences.[31]

[30] As noted in the introduction, Ingram and Whiteman (1994) and DeJong, Ingram, and Whiteman (2000) provide such evidence. In an earlier paper, Ireland (1995) reports that, once translated into a bivariate VAR, the simple version of the permanent-income theory is rejected within-sample but helps the model to better forecast out-of-sample.

[31] In the experiments of Stock and Watson (1999, 2002) the variables used are, for the most part, differentiated. Ireland (2001b), however, reports, in a formal comparison between estimating trend-stationary and difference-stationary real business cycle models, that the trend-stationary specification has the better out-of-sample forecasting accuracy.

Finally, using an open-economy specification would allow the model to capture information related to external (principally American) data and the various channels by which they affect the Canadian economy, which could prove particularly useful for forecasting Canadian time series.[32]

[32]It is doubtful, however, whether such an open-economy specification would modify significantly the estimation of other parameters: Dib (2003b) shows that most estimates unrelated to open-economy features do not change when a model extended to comprise open-economy features is estimated using Canadian data.

References

Bergin, P.R. 2003. "Putting the "New Open Economy Macroeconomics" to a Test." *Journal of International Economics* 60: 3–34.

Blanchard, O. and C. Kahn. 1980. "The Solution of Linear Difference Models under Rational Expectations." *Econometrica* 48: 1305–11.

Blanchard, O. and D. Quah. 1989. "The Dynamic Effects of Aggregate Demand and Aggregate Supply Shocks." *American Economic Review* 79: 655–73.

Boivin, J. and M. Giannoni. 2005. "DSGE Models in a Data-Rich Environment." Computing in Economics and Finance Working Paper No. 431.

Bouakez, H., E. Cardia, and F. Ruge-Murcia. 2005. "Habit Formation and the Persistence of Monetary Shocks." *Journal of Monetary Economics* 52: 1073–88.

Calvo, G.A. 1983. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics* 12: 383–98.

Chong, Y. and D.F. Hendry. 1986. "Econometric Evaluation of Linear Macro-Economic Models." *Review of Economics and Statistics* 113: 1–45.

Christiano, L.J., M. Eichenbaum, and C. Evans. 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy* 103: 51–78.

Clements, M. and F. Hendry. 1998. *Forecasting Economic Time Series.* Cambridge: Cambridge University Press.

________. 1999. "On Winning Forecasting Competitions in Economics." *Spanish Economic Review* 1: 123–60.

Cogley, T. and J. Nason. 1995. "Output Dynamics in Real Business Cycle Models." *American Economic Review* 85: 492–511.

Del Negro, M. and F. Schorfheide. 2004. "On the Fit and Forecasting Performance of New-Keynesian Models." CEPR Working Paper No. 4848.

DeJong, D.N., B.F. Ingram, and C.H. Whiteman. 2000. "Keynesian Impulses Versus Solow Residuals: Identifying Sources of Business Cycle Fluctuations." *Journal of Applied Econometrics* 15: 311–29.

Dib, A. 2003a. "An Estimated Canadian DSGE Model with Nominal and Real Rigidities." *Canadian Journal of Economics* 36: 949–72.

________. 2003b. "Monetary Policy in Estimated Models of Small Open and Closed Economies." Bank of Canada Working Paper No. 2003-26.

Dib, A. 2006. "Nominal Rigidities and Monetary Policy in Canada." *Journal of Macroeconomics* 28: 285–304.

Diebold, F.X. and R.S. Mariano. 1995. "Comparing Predictive Accuracy." *Journal of Business and Economic Statistics* 13: 253–63.

Doan, T., R. Litterman, and C. Sims. 1984. "Forecasting and Conditional Projections using Realistic Prior Distributions." *Econometric Review* 3: 1–100.

Dolar, V. and K. Moran. 2002. "Estimated DGE Models and Forecasting Accuracy: A Preliminary Investigation with Canadian Data." Bank of Canada Working Paper No. 2002–18.

Granger, C.W.J. and P. Newbold. 1973. "Some Comments on the Evaluation of Economic Forecasts." *Applied Economics* 5: 35–47.

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

Harvey, D., S. Leybourne, and P. Newbold. 1997. "Testing the Equality of Prediction Mean Squared Errors." *International Journal of Forecasting* 13: 281–91.

________. 1998. "Tests for Forecast Encompassing." *Journal of Business and Economic Statistics* 15: 254–59.

Ingram, B.F. and C.H. Whiteman. 1994. "Supplanting the Minnesota Prior: Forecasting Macroeconomic Times Series using Real Business Cycle Model Priors." *Journal of Monetary Economics* 34: 497–510.

Ireland, P.N. 1995. "Optimal Disinflationary Paths." *Journal of Economic Dynamic and Control* 19: 1429–48.

________. 1997. "A Small, Structural, Quarterly Model for Monetary Policy Evaluation." *Carnegie-Rochester Conference Series on Public Policy* 47: 83–108.

________. 2001a. "Sticky-Price Models of the Business Cycle: Specification and Stability." *Journal of Monetary Economics* 47: 3–18.

________. 2001b. "Technology Shocks and the Business Cycle: An Empirical Investigation." *Journal of Dynamic Economic and Control* 25: 703–19.

________. 2003. "Endogenous Money or Sticky Prices." *Journal of Monetary Economics* 47: 3–18.

________. 2004. "A Method for Taking Models to the Data." *Journal of Dynamic Economic and Control* 34: 497–510.

McCallum, B.T. and E. Nelson. 1999. "Nominal Income Targeting in an Open-Economy Optimizing Model." *Journal of Monetary Economics* 49: 947–81.

Poole, W. 1970. “Optimal Choice of Monetary Policy Instrument in a Simple Stochastic Macro Model.” *Quarterly Journal of Economics* 84: 197–216.

Smets, F. and R. Wouters. 2003. “An Estimated Stochastic Dynamic General Equilibrium Model of the Euro Area.” *Journal of the European Economic Association* 1: 123–75.

________. 2004. “Forecasting with a Bayesian DSGE Model: An Application to the Euro Area.” European Central Bank Working Paper No. 389.

Stock, J.H. and M.W. Watson. 1999. “Forecasting Inflation.” *Journal of Monetary Economics* 44: 293–35.

________. 2002. “Macroeconomic Forecasting Using Diffusion Indexes.” *Carnegie-Rochester Conference Series on Public Policy* 39: 195–214.

Taylor, J.B. 1993. “Discretion Versus Policy Rules in Practice.” *Journal of Monetary Economics* 47: 3–18.

Woodford, M. 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy.* Princeton: Princeton University Press.

Yun, T. 1996. “Nominal Price Rigidity, Money Supply Endogeneity, and Business Cycles.” *Journal of Monetary Economics* 37: 345–70.

Table 1: Maximum-Likelihood Estimates and Standard Deviations (1981Q3 to 1995Q4)

Parameter	Estimate	Std. deviation	t-statistic
β	0.990	0.002	642.0
γ	0.055	0.022	2.45
ϱ_π	0.754	0.165	4.56
ϱ_μ	0.481	0.105	4.57
ϱ_y	0.018	0.026	0.68
ρ_v	0.194	0.094	2.06
σ_v	0.006	0.001	7.46
ϕ	0.630	0.060	10.47
A	3.532	0.200	18.13
ρ_A	0.950	0.062	15.38
σ_A	0.013	0.002	5.51
b	0.498	0.060	8.40
ρ_b	0.994	0.010	99.22
σ_b	0.012	0.001	8.35
ρ_z	0.917	0.050	18.48
σ_z	0.017	0.004	4.5
π	1.010	0.004	221.13
η	1.002	0.002	550.98
LL	898.7011		

Note: LL is the maximum log-likelihood value.

Table 2: Volatility (1981Q3 to 1995Q4)

Variable	$\hat{y}_t$	$\hat{\pi}_t$	$\hat{R}_t$	$\hat{m}_t$
Data	3.44	0.60	0.48	2.78
Model	2.33	0.74	0.61	2.87

Table 3: Decomposition of Forecast-Error Variances

		Percentage owing to:			
Quarters	Variance (%)	Tech.	Mon.dem.	Policy	Pref.
		A. Output			
1	0.0121	49.71	6.94	12.09	31.27
2	0.0248	58.18	3.84	7.29	30.68
4	0.0476	66.09	2.06	3.98	27.88
10	0.0904	75.56	1.10	2.10	21.24
		B. Inflation			
1	0.0023	31.52	17.20	48.61	2.66
2	0.0029	30.46	15.12	46.08	8.34
4	0.0034	28.56	13.03	40.47	17.94
10	0.0040	26.91	11.09	34.43	27.56
		C. Nominal interest rate			
1	0.0010	6.30	4.10	0.5	95.22
2	0.0017	2.72	2.78	2.5	94.25
4	0.0026	5.89	5.89	2.2	92.03
10	0.0038	10.72	10.72	1.5	87.84
		D. Real money balances			
1	0.0125	41.25	41.83	7.93	9.00
2	0.0324	42.56	44.80	4.26	8.38
4	0.0757	42.70	48.42	1.98	6.91
10	0.1887	39.69	55.49	0.80	4.01

Table 4: The Forecasting Experiment (1996Q1 to 2004Q1)

Estimate	Forecast k periods ahead				
	$k=1$	$k=2$	$k=3$	$\cdots$	$k=8$
1981Q3$\longrightarrow$ 1995Q4	1996Q1	1996Q2	1996Q3	$\cdots$	1997Q4
1981Q3$\longrightarrow$ 1996Q1	1996Q2	1996Q3	1996Q4	$\cdots$	1998Q1
1981Q3$\longrightarrow$ 1996Q2	1996Q3	1996Q4	1997Q1	$\cdots$	1998Q2
1981Q3$\longrightarrow$ 1996Q3	1996Q4	1997Q1	1997Q2	$\cdots$	1998Q3
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
1981Q3$\longrightarrow$ 2003Q2	2003Q3	2003Q4	2004Q1	$---$	$---$
1981Q3$\longrightarrow$ 2003Q3	2003Q4	2004Q1	$---$	$---$	$---$
1981Q3$\longrightarrow$ 2003Q4	2004Q1	$---$	$---$	$---$	$---$

Table 5: Testing for Equal Forecasting Accuracy: NK Model and VAR(2) Benchmark (1997Q1 to 2004Q1)

Variable	Relative MSE (NK model)[a]	DM stat.[b] (p-value)	HLN stat.[c] (p-value)
	Forecasting One Period Ahead		
Output	1.73	-2.18(0.03)	-2.14(0.04)
Inflation	1.21	-0.51(0.61)	-0.50(0.62)
Interest Rate	0.73	1.57(0.12)	1.54(0.13)
Money	0.67	2.29(0.02)	2.25(0.03)
	Forecasting Two Periods Ahead		
Output	1.26	-0.74(0.45)	-0.70(0.49)
Inflation	1.28	-0.63(0.53)	-0.60(0.56)
Interest Rate	0.60	2.78(0.00)	2.64(0.01)
Money	0.49	2.44(0.02)	2.31(0.03)
	Forecasting Four Periods Ahead		
Output	0.57	0.70(0.48)	0.61(0.55)
Inflation	1.28	-0.46(0.64)	-0.40(0.69)
Interest Rate	0.53	2.72(0.00)	2.36(0.03)
Money	0.30	1.87(0.06)	1.62(0.12)
	Forecasting Six Periods Ahead		
Output	0.33	1.22(0.22)	0.94(0.36)
Inflation	1.17	-0.12(0.91)	-0.09(0.93)
Interest Rate	0.45	5.07(0.00)	3.91(0.00)
Money	0.18	1.41(0.16)	1.08(0.29)
	Forecasting Eight Periods Ahead		
Output	0.26	1.20(0.23)	0.79(0.44)
Inflation	1.32	-0.81(0.42)	-0.54(0.60)
Interest Rate	0.40	NA(NA)	NA(NA)
Money	0.13	1.24(0.22)	0.82(0.42)

[a]MSE (NK model)/MSE (VAR benchmark); values smaller than 1 suggest that the NK model has superior forecasting accuracy.

[b]Test statistics from the Diebold and Mariano (1995) test. The null hypothesis is of equal forecasting accuracy between the two models. The statistics are asymptotically standard normal.

[c]Harvey, Leybourne, and Newbold's (1997) correction of the Diebold and Mariano (1995) test. The statistics follow a t_{N-1} distribution, with N denoting the number of forecasts.

Table 6: Forecast Encompassing: Does the NK model provide *any* information not contained in the VAR Benchmark?

Variable	DM stat.[a] (p-value)	HLN stat.[b] (p-value)
Forecasting One Period Ahead		
Output	2.94 (0.00)	2.88 (0.00)
Inflation	0.54 (0.30)	0.53 (0.30)
Interest Rate	3.42 (0.00)	3.36 (0.00)
Money	2.40 (0.01)	2.36 (0.01)
Forecasting Two Periods Ahead		
Output	3.03 (0.00)	2.87 (0.00)
Inflation	0.41 (0.34)	0.39 (0.35)
Interest Rate	4.31 (0.00)	4.09 (0.00)
Money	2.52 (0.01)	2.38 (0.01)
Forecasting Four Periods Ahead		
Output	1.74 (0.04)	1.51 (0.07)
Inflation	0.60 (0.28)	0.52 (0.31)
Interest Rate	6.00 (0.00)	5.18 (0.00)
Money	1.90 (0.03)	1.56 (0.04)
Forecasting Six Periods Ahead		
Output	1.42 (0.08)	1.09 (0.14)
Inflation	1.27 (0.10)	0.98 (0.17)
Interest Rate	26.0 (0.00)	20.7 (0.00)
Money	1.58 (0.06)	1.22 (0.12)
Forecasting Eight Periods Ahead		
Output	1.29 (0.10)	0.85 (0.20)
Inflation	0.61 (0.27)	0.40 (0.35)
Interest Rate	NA (NA)	NA (NA)
Money	1.39 (0.08)	0.91 (0.19)

[a]Test statistics proposed by Harvey, Leybourne, and Newbold (1998). The null hypothesis is that the forecasts from the NK model provide no information not already contained in those from the VAR benchmark.

[b]Harvey, Leybourne, and Newbold's (1997) correction of the Diebold and Mariano (1995) test. The statistics follow a t_{N-1} distribution, with N denoting the number of forecasts.

Table 7: Testing for Equal Forecasting Accuracy: NK Model and BVAR(2) Benchmark (1997Q1 to 2004Q1)

Variable	Relative MSE (NK model)[a]	DM stat.[b] (p-value)	HLN stat.[c] (p-value)
	Forecasting One Period Ahead		
Output	2.31	-2.44 (0.02)	-2.40 (0.02)
Inflation	1.01	-0.04 (0.97)	-0.04 (0.97)
Interest Rate	0.90	0.70 (0.49)	0.68 (0.50)
Money	0.78	1.64 (0.10)	1.61 (0.12)
	Forecasting Two Periods Ahead		
Output	1.56	-0.91 (0.36)	-0.86 (0.40)
Inflation	0.93	0.36 (0.72)	0.34 (0.74)
Interest Rate	0.79	0.93 (0.35)	0.88 (0.39)
Money	0.72	1.53 (0.13)	1.45 (0.16)
	Forecasting Four Periods Ahead		
Output	0.69	0.69 (0.49)	0.60 (0.58)
Inflation	0.94	0.32 (0.75)	0.28 (0.78)
Interest Rate	0.65	1.24 (0.22)	1.07 (0.30)
Money	0.56	1.07 (0.29)	0.93 (0.36)
	Forecasting Six Periods Ahead		
Output	0.35	1.87 (0.06)	1.44 (0.16)
Inflation	0.80	2.45 (0.01)	1.89 (0.07)
Interest Rate	0.48	2.35 (0.02)	1.81 (0.08)
Money	0.42	1.18 (0.24)	0.91 (0.27)
	Forecasting Eight Periods Ahead		
Output	0.24	2.17 (0.03)	1.43 (0.17)
Inflation	1.06	-0.42 (0.68)	-0.27 (0.79)
Interest Rate	0.38	7.80 (0.00)	5.14 (0.00)
Money	0.30	1.44 (0.15)	0.95 (0.35)

[a]MSE (NK model)/MSE (VAR benchmark); values smaller than 1 suggest that the NK model has superior forecasting accuracy.

[b]Test statistics from the Diebold and Mariano (1995) test. The null hypothesis is of equal forecasting accuracy between the two models. The statistics are asymptotically standard normal.

[c]Harvey, Leybourne, and Newbold's (1997) correction of the Diebold and Mariano (1995) test. The statistics follow a t_{N-1} distribution, with N denoting the number of forecasts.

Figure 1: The Economy's Response to Monetary Policy Tightening

(Shock occurs at $t = 5$)

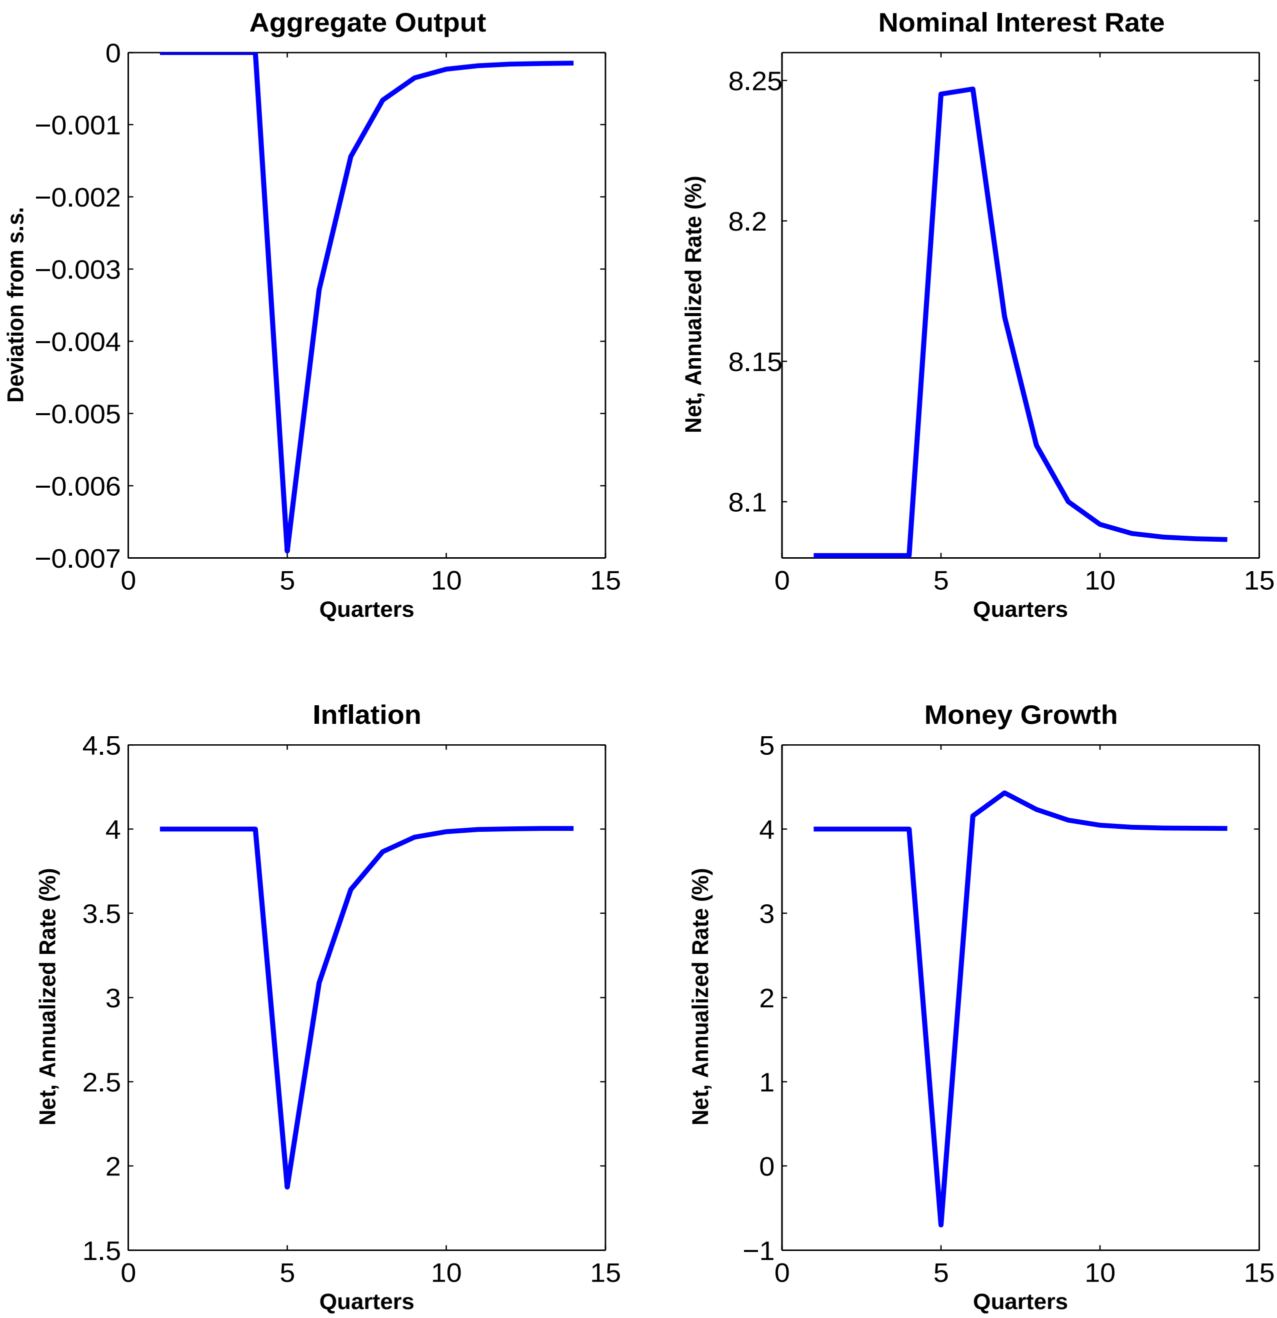

Figure 2: The Economy's Response to a Positive Money-Demand Shock
(Shock occurs at $t = 5$)

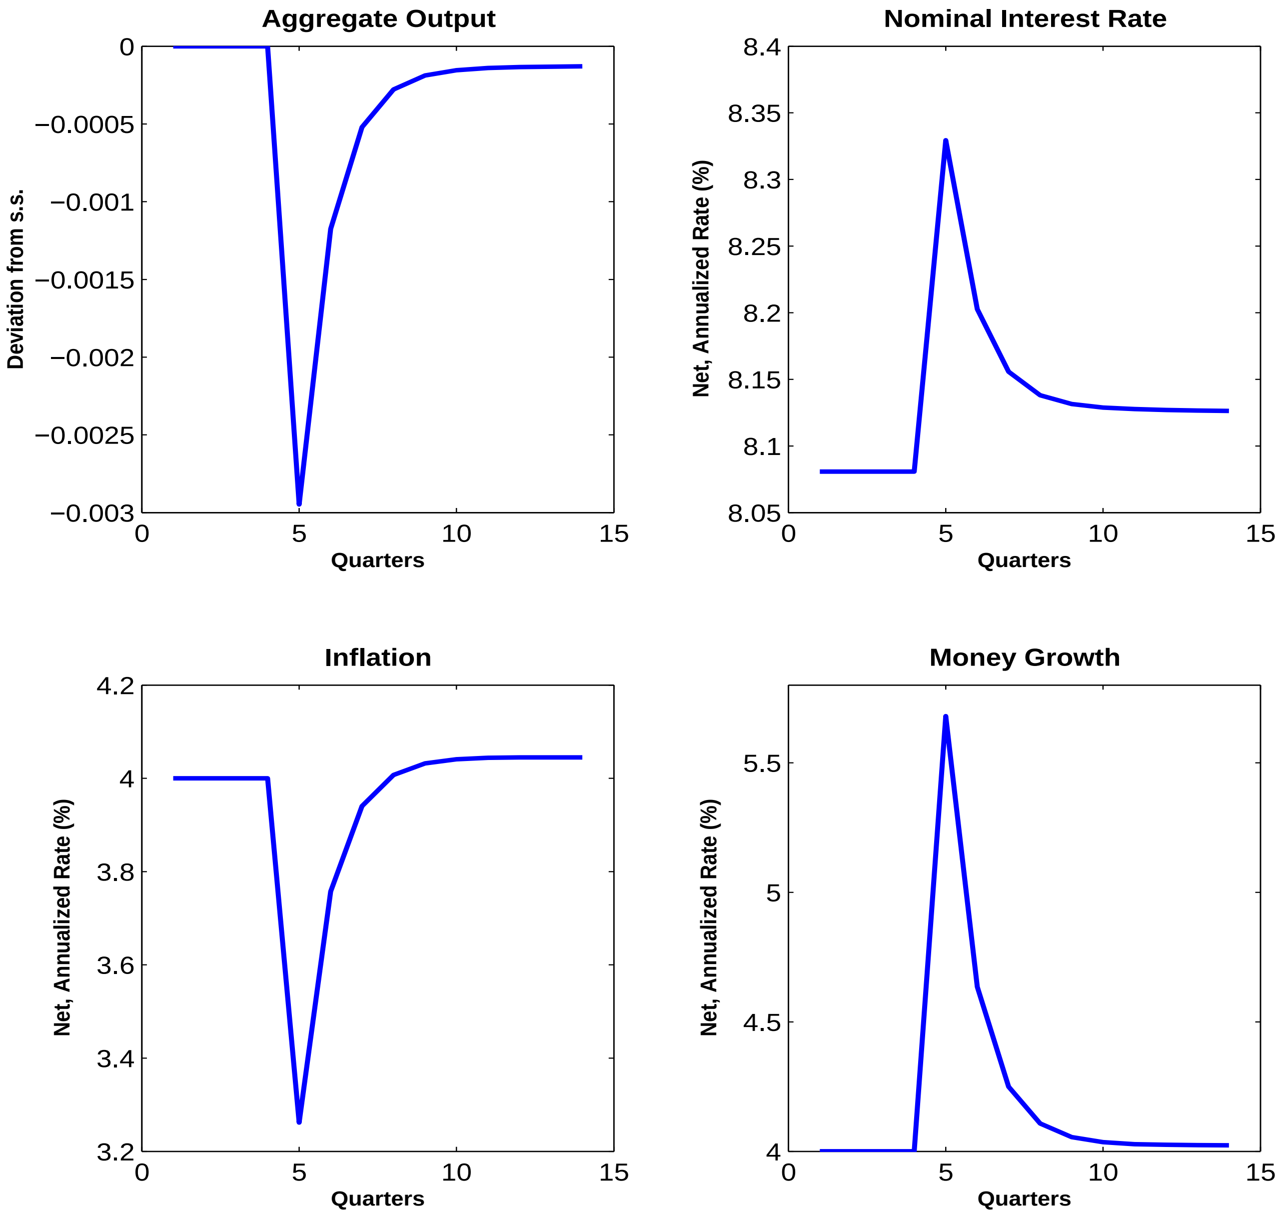

Figure 3: The Economy's Response to a Positive Technology Shock
(Shock occurs at $t = 5$)

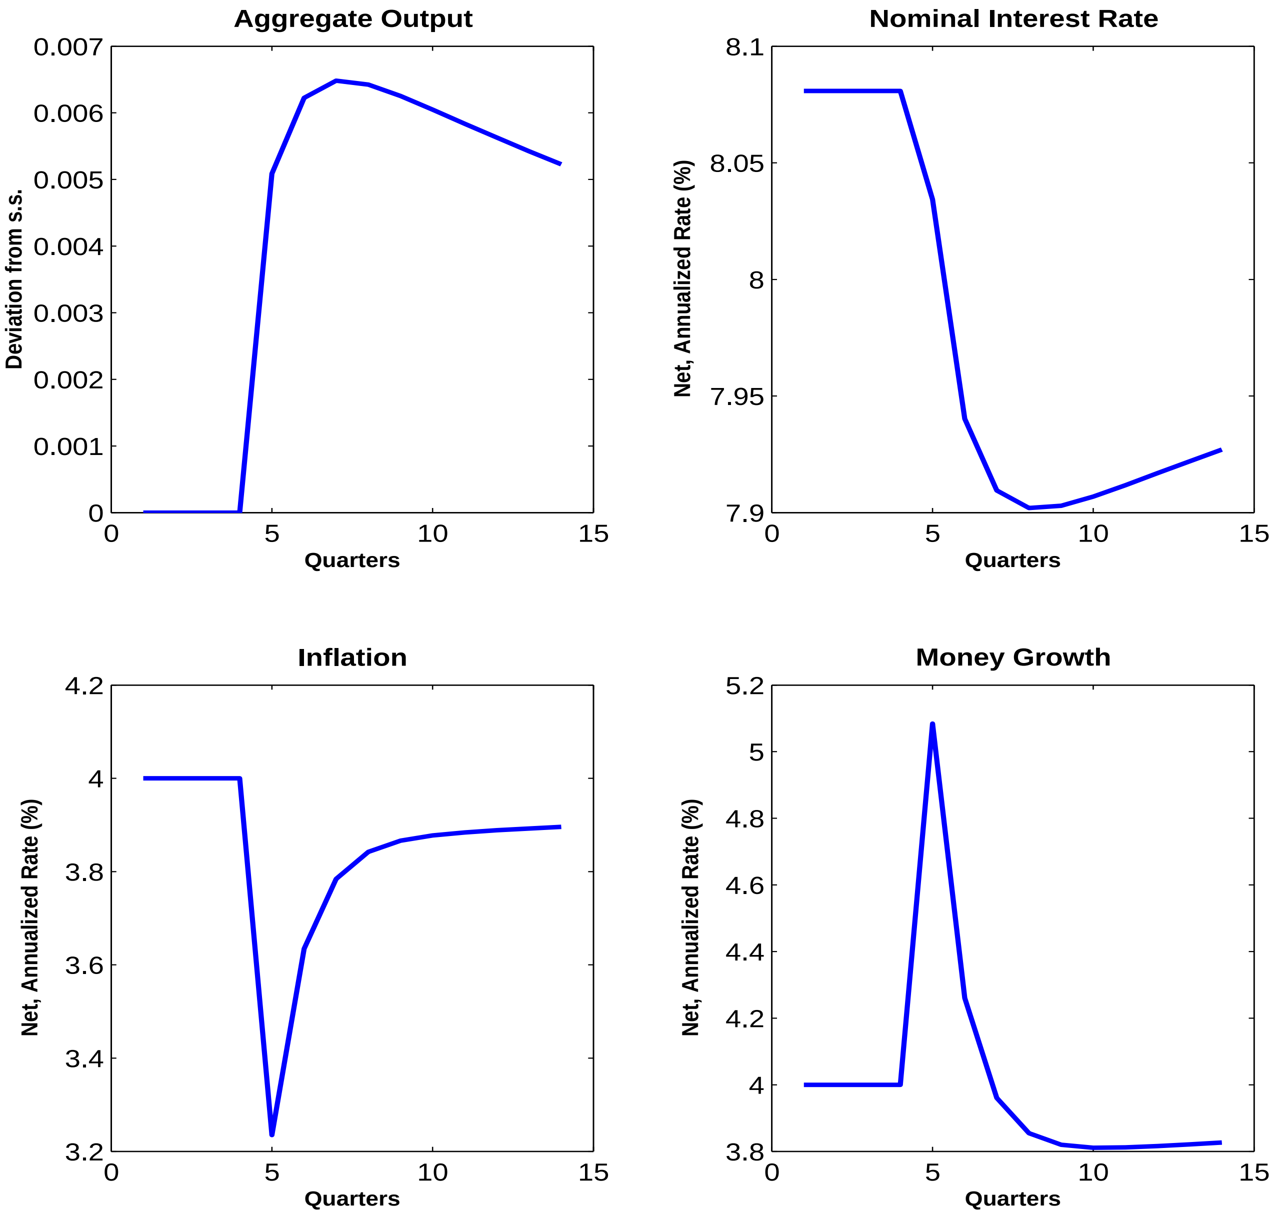

Figure 4: The Economy's Response to a Positive Preference Shock

(Shock occurs at $t = 5$)

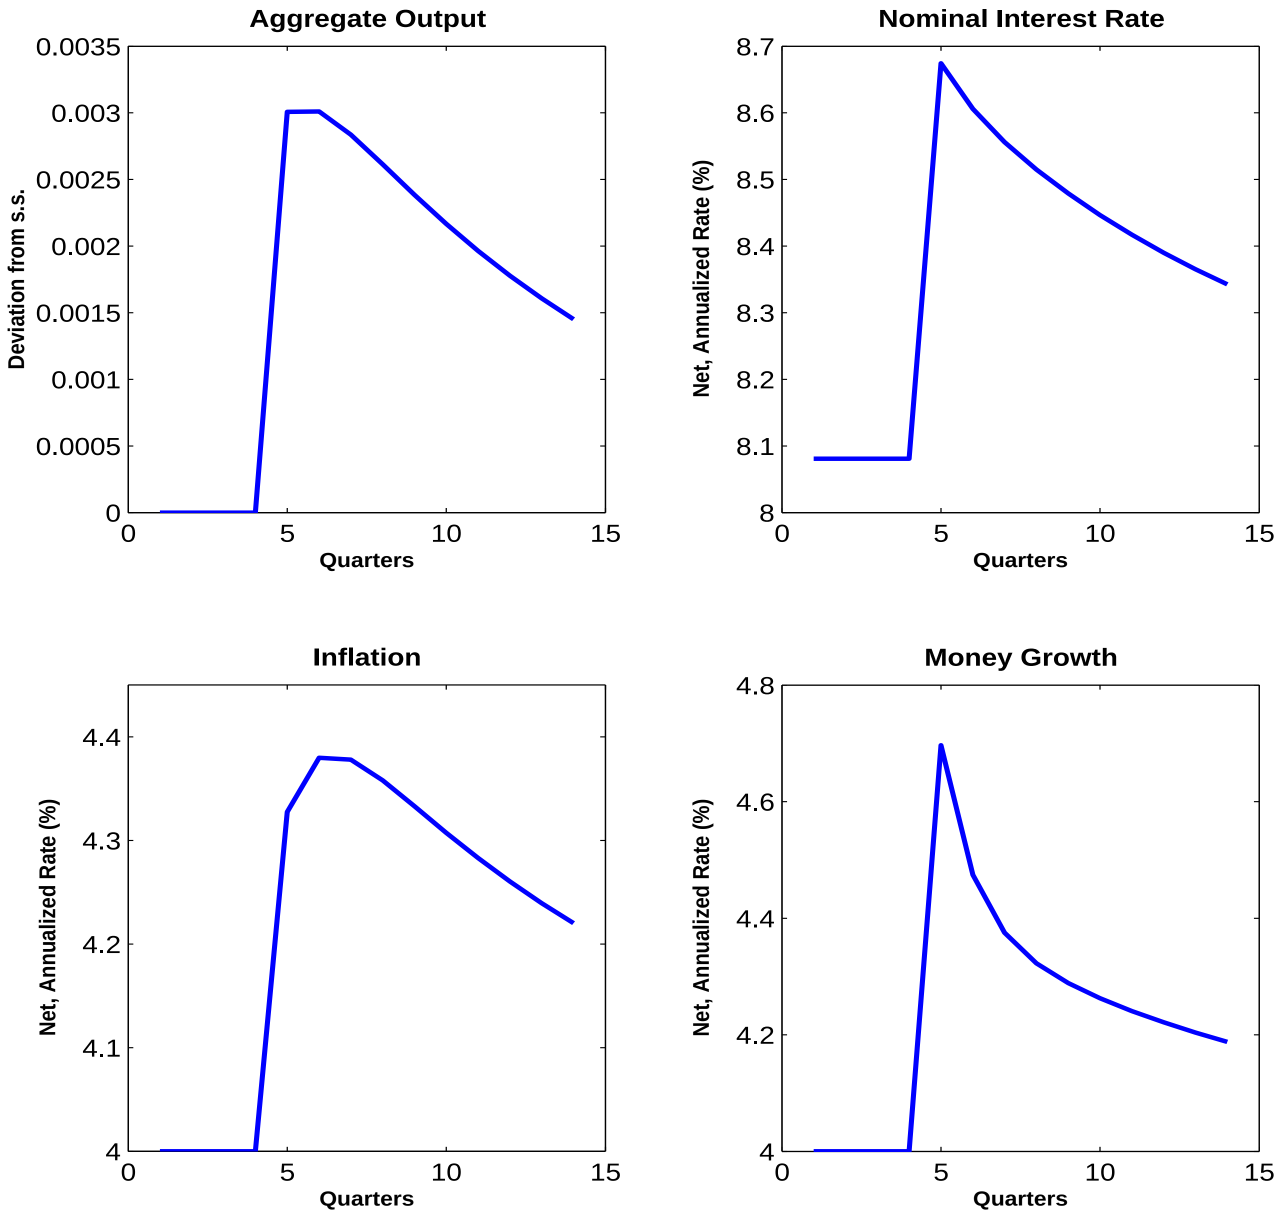

Figure 5: Actual Data and Forecasts from the New Keynesian Model

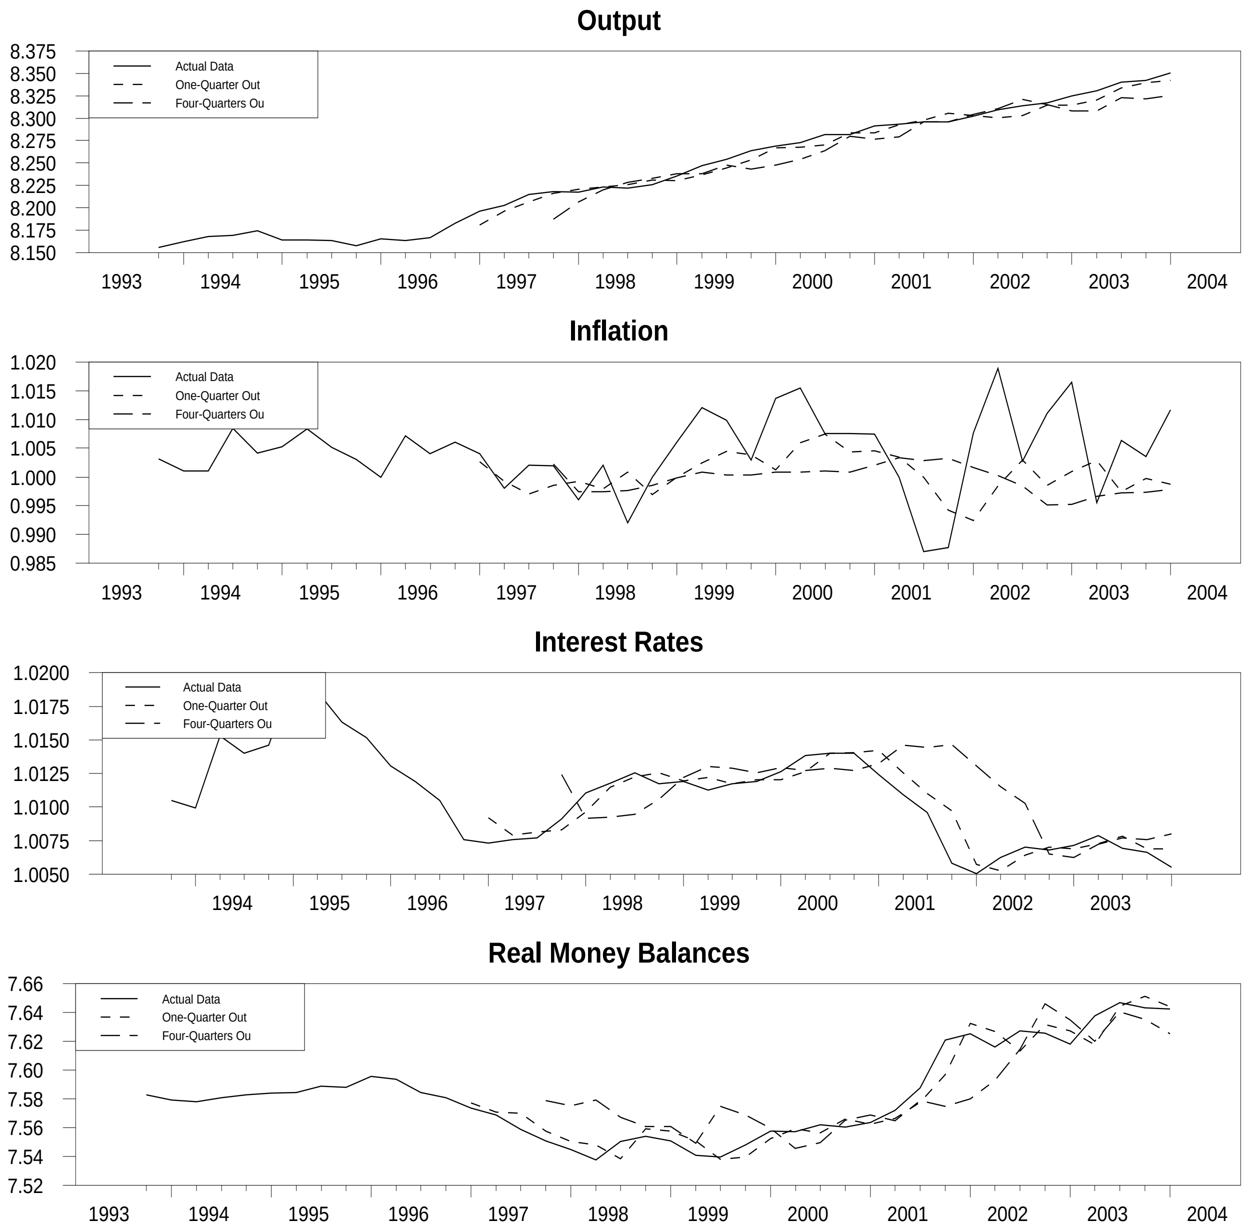

Figure 6: Forecast Errors, One Quarter Ahead

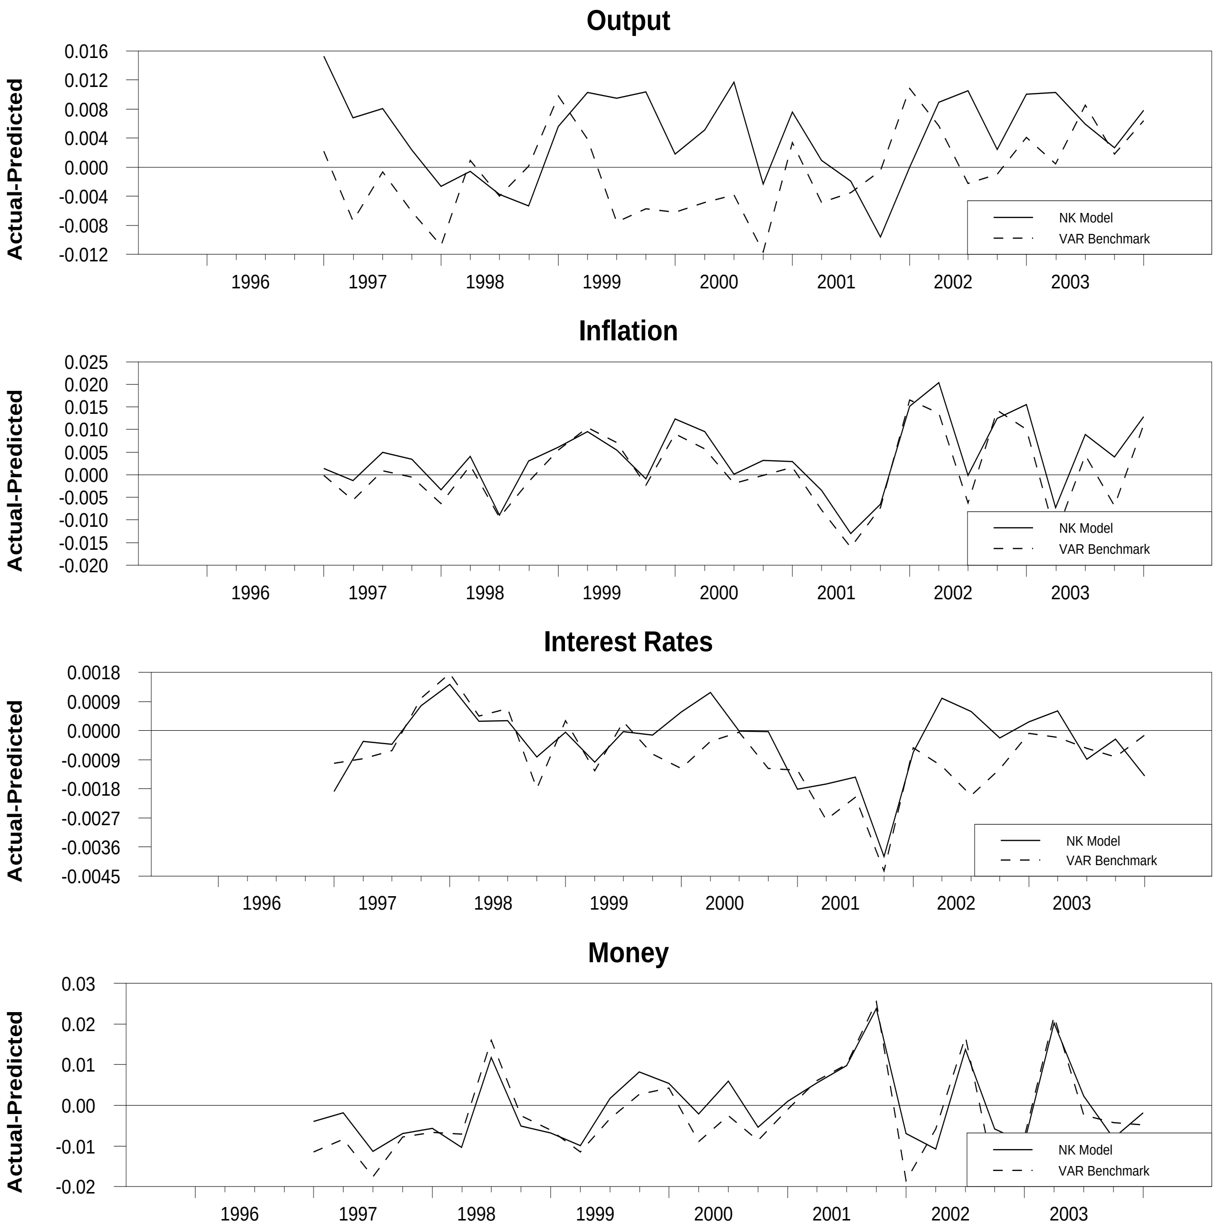

Figure 7: Forecast Errors, Four Quarters Ahead

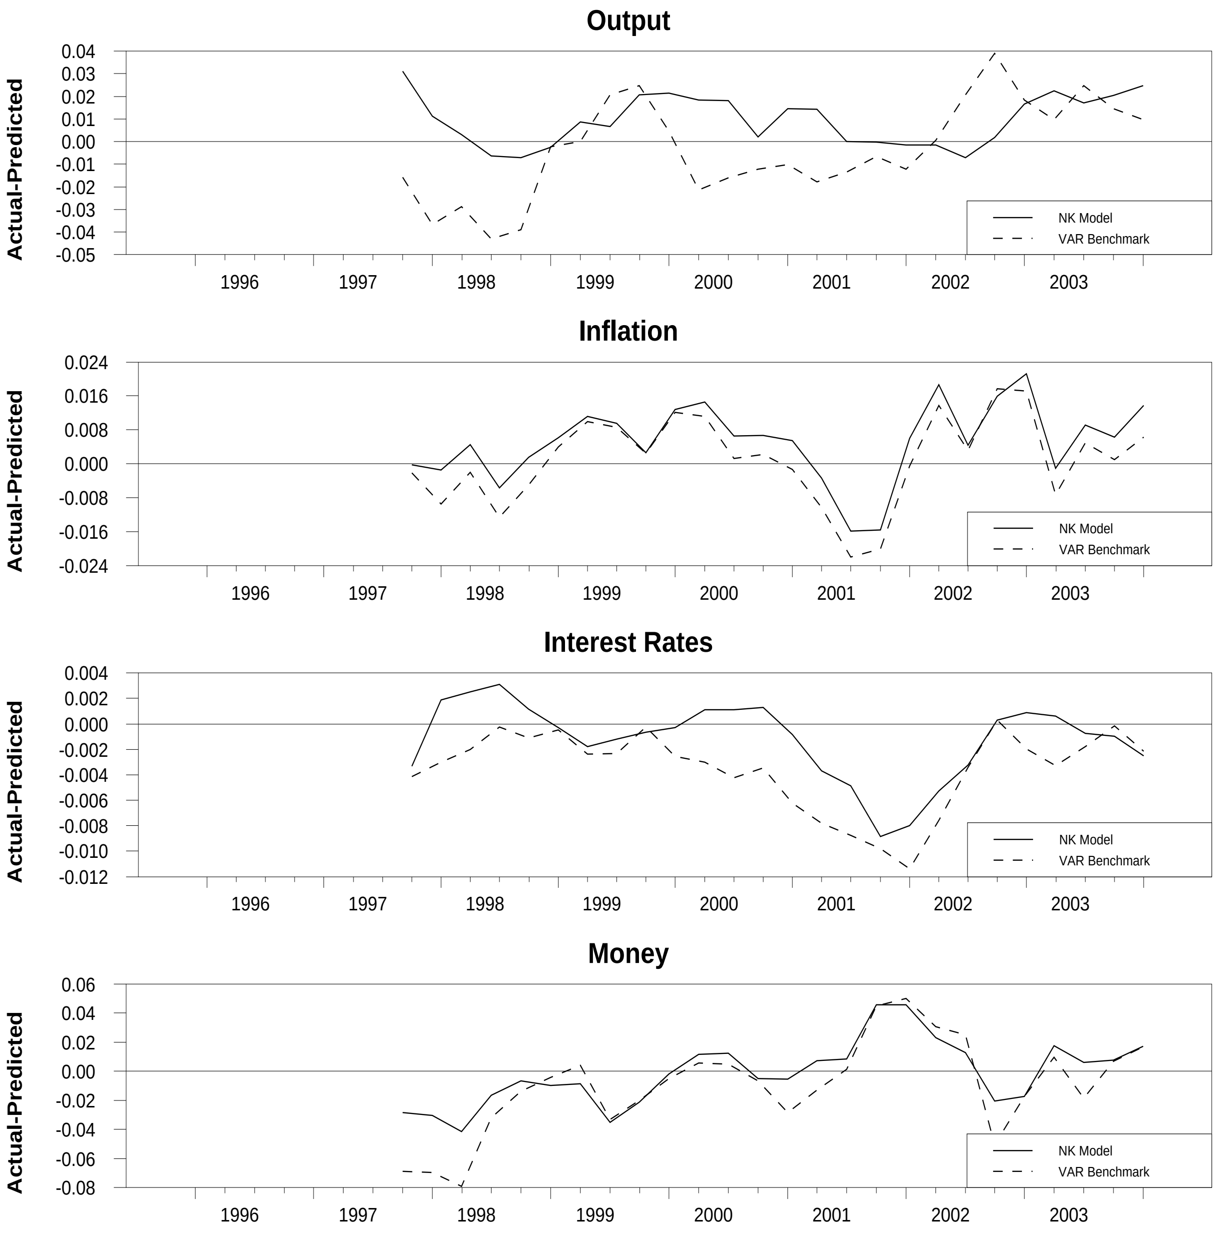

Figure 8: Forecast Errors, Eight Quarters Ahead

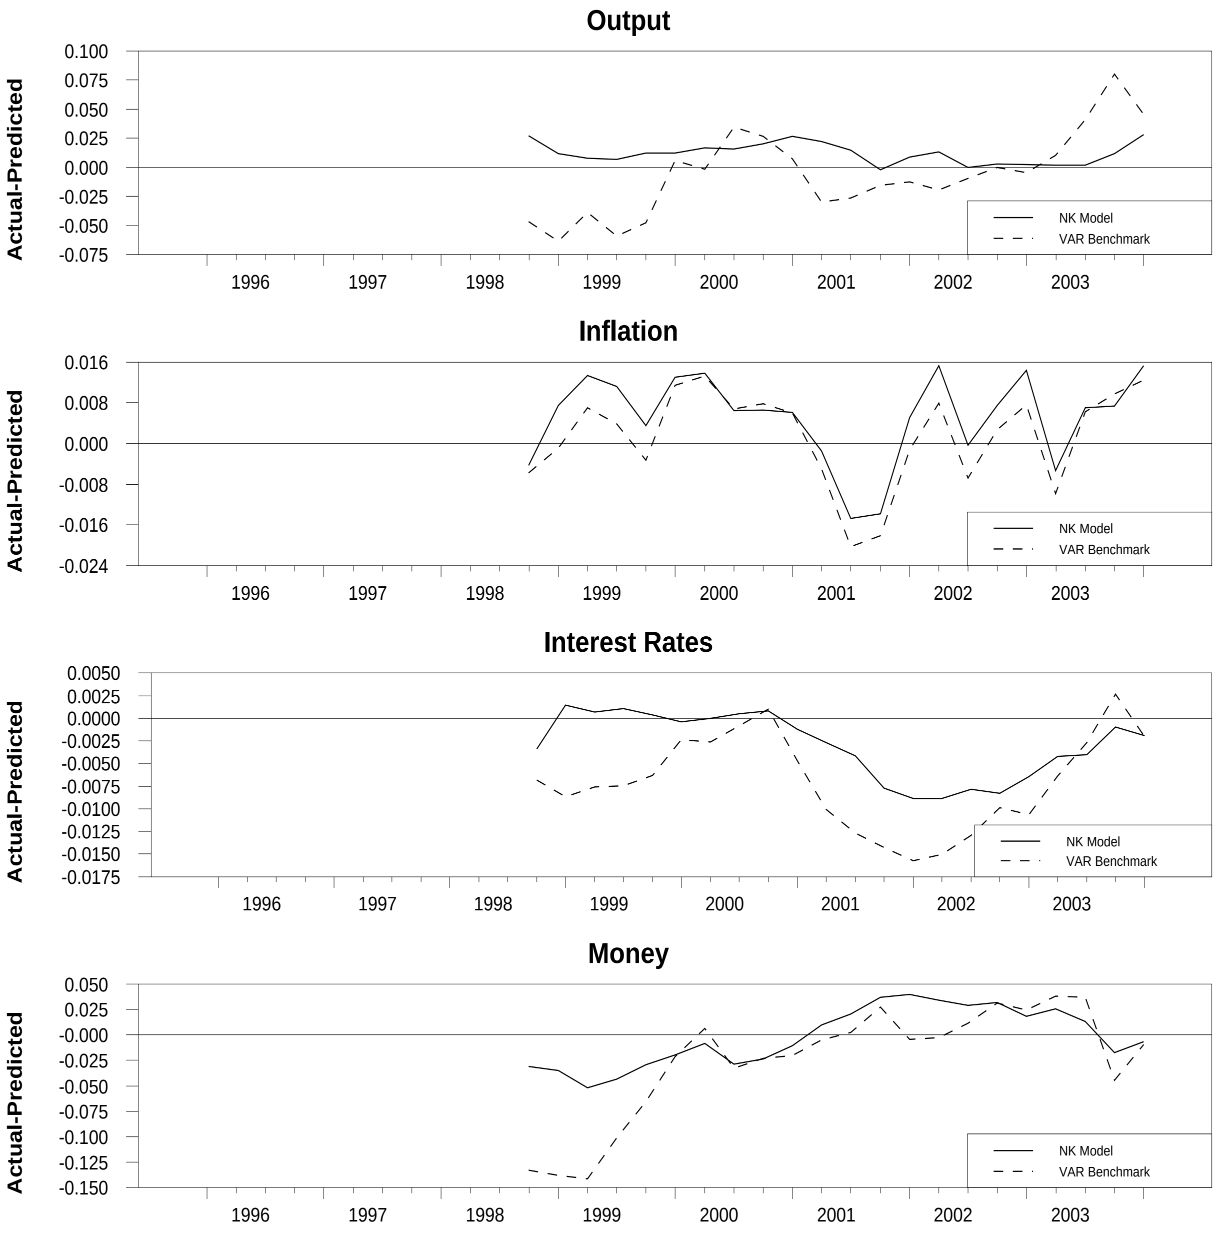

Appendix A: The Transformed Equilibrium System

$$\frac{z_t c_t^{-\frac{1}{\gamma}}}{c_t^{\frac{\gamma-1}{\gamma}} + b_t^{1/\gamma} m_t^{\frac{\gamma-1}{\gamma}}} = \lambda_t; \tag{A.1}$$

$$\frac{z_t b_t^{1/\gamma} m_t^{-\frac{1}{\gamma}}}{c_t^{\frac{\gamma-1}{\gamma}} + b_t^{1/\gamma} m_t^{\frac{\gamma-1}{\gamma}}} = \lambda_t - \beta E_t \left(\frac{\lambda_{t+1}}{\eta \pi_{t+1}} \right); \tag{A.2}$$

$$\frac{\zeta}{1 - h_t} = \lambda_t w_t; \tag{A.3}$$

$$\beta E_t \left[\frac{\lambda_{t+1}}{\eta \lambda_t} \left(r_{kt+1} + 1 - \delta + \psi \eta \left(\frac{k_{t+2}}{k_{t+1}} - 1 \right) \frac{k_{t+2}}{k_{t+1}} - \frac{\psi \eta}{2} (\frac{k_{t+2}}{k_{t+1}} - 1)^2 \right) \right]$$
$$= \psi \eta \left(\frac{k_{t+1}}{k_t} - 1 \right) + 1; \tag{A.4}$$

$$\frac{1}{R_t} = \beta E_t \left[\frac{\lambda_{t+1}}{\eta \lambda_t \pi_{t+1}} \right] \tag{A.5}$$

$$y_t = k_t^{\alpha} (A_t h_t)^{1-\alpha}; \tag{A.6}$$

$$\frac{\alpha y_t}{k_t} = q_t r_t; \tag{A.7}$$

$$\frac{(1-\alpha) y_t}{h_t} = q_t w_t; \tag{A.8}$$

$$\widetilde{p}_t = \frac{\theta \pi_t}{\theta - 1} \frac{E_t \sum_{k=0}^{\infty} (\beta \phi \pi^{-\theta})^k \lambda_{t+k} y_{t+k} q_{t+k} (\prod_{s=1}^{k} \pi_{t+s}^{\theta})}{E_t \sum_{k=0}^{\infty} (\beta \phi \pi^{1-\theta})^k \lambda_{t+k} y_{t+k} (\prod_{s=1}^{k} \pi_{t+s}^{\theta-1})}; \tag{A.9}$$

$$1 = \phi \left(\frac{\pi}{\pi_t} \right)^{1-\theta} + (1-\phi) \widetilde{p}_t^{1-\theta}; \tag{A.10}$$

$$y_t = c_t + \eta k_{t+1} - (1-\delta) k_t - \frac{\psi \eta}{2} (\frac{k_{t+1}}{k_t} - 1)^2 k_t; \tag{A.11}$$

$$\mu_t = \frac{\eta m_t \pi_t}{m_{t-1}}; \tag{A.12}$$

$$\log(R_t/R) = \varrho_\pi \log(\pi_t/\pi) + \varrho_y \log(y_t/y) + \varrho_\mu \log(\mu_t/\mu) + \log(v_t); \tag{A.13}$$

$$\log(A_t) = (1-\rho_A) \log(A) + \rho_A \log(A_{t-1}) + \varepsilon_{At}; \tag{A.14}$$

$$\log(b_t) = (1-\rho_b) \log(b) + \rho_b \log(b_{t-1}) + \varepsilon_{bt}; \tag{A.15}$$

$$\log(z_t) = \rho_z \log(z_{t-1}) + \varepsilon_{zt}; \tag{A.16}$$

$$\log(v_t) = \rho_v \log(v_{t-1}) + \varepsilon_{vt}. \tag{A.17}$$

Appendix B: Data

The model is estimated using data from 1981Q3 to 2004Q1. The data are taken from Statistics Canada's CANSIM database, for which we list the associated mnemonics. Output, Y_t, is *final domestic demand* [V1992068], of quarterly frequency and in chained 1987 dollars. We convert this series into per-capita terms using the population age 15 and over.

The interest rate, R_t, is the three-month treasury bill rate [V122531], a series of daily frequency, for which we take a quarterly average.

The money stock, M_t, is *M2* [B1630], which is of monthly frequency; we take a quarterly average and convert the resulting series into real per-capita terms by dividing it with the GDP implicit price deflator [D100465] and the population age 15 and over. Output and money data are logged before estimation.