Optimal Monetary Policy and Model Selection in a Real-Time Learning Environment*

Federico Ravenna[†]

Institute of Applied Economics, HEC Montreal

Abstract

We investigate how real-time parameter learning, optimal policies and shocks' volatility affect the policymakers' ability to distinguish across competing models of the economy. The detection speed of model misspecification depends only on the relative volatility of supply and demand shocks.

Keywords: Learning, Optimal Monetary Policy, Model Misspecification, Business Cycle.

JEL Classification: E58.

*I would like to thank Richard Dennis, Bart Hobijn, Oscar Jorda, David Lopez-Salido and Carl Walsh for useful comments and suggestions. Richard Dennis provided the algorithm to solve for the optimal policy and the rational expectations equilibrium.

[†]Corresponding author: Federico Ravenna, Institute of Applied Economics, HEC Montreal, Montreal, QC, H3T 2A7, Canada. Email: federico.ravenna@hec.ca. Tel: + 1-514-3406668. Fax: +1-514-3406469.

1 Introduction

Central banks look with caution at the prescriptions of optimal policy rules because they tend to be highly model-specific. Model uncertainty is sufficiently severe that alternative models offer widely different policy prescriptions, for example by including alternative formulations of expectations or price-setting mechanisms. In addition, policymakers also face uncertainty over the parameters of any given model.

A growing literature has suggested several approaches to account for model uncertainty in optimal policy design (Hansen and Sargent, 2003, Cogley and Sargent, 2005, Cogley et al., 2011). These methodologies aim at allowing policymakers to explicitly account for the likelihood of competing models to correctly describe the economy when choosing policy. Nevertheless, policymakers may need to engage in model selection, and take a stand on what is the 'true' model of the economy. For example, central banks need to produce forecasts of the economy, which constitute an important input in the policymaking process.

In this paper we study the problem of a central bank which must identify the true model across competing alternatives while conducting policy according to its beliefs on the model of the economy. We assume the policymaker tries to detect the existence of misspecification in its believed model against a proposed alternative, while estimating in real time the deep parameters of the model needed to compute the optimal policy. We also allow for the private sector to learn over time the law of motion of the economy, as in Evans and Honkapohja (2003b). In this framework, we can assess whether the objective of selecting the correct model across competing alternatives provides an incentive to deviate from the full-information optimal policy, as in Wieland (2000).

Our results show that a central bank using an optimal policy updated in real-time can detect the misspecification in its believed model as fast as a non-optimizing policymaker setting policy using a given Taylor-type rule. Moreover, real-time learning of the model parameters by the policymaker or of the law of motion by the private sector does not affect the speed of misspecification detection. Crucially, our results depend on the relative volatility of supply and demand shocks. When supply shocks are sufficiently more volatile than demand shocks, detection of the model misspecification does not occur even after the central bank has accumulated a 30-year long sample of data.

In our analysis we focus on competing models for the inflation process within a new Keynesian business cycle framework. Our choice of competing models reflects a long-standing debate on the nature of inflation inertia - whether inflation is driven by inertial exogenous shocks, or by an endogenous inertial mechanism for price updating.

Evans and Honkapohja (2003) and Dennis and Ravenna (2008) examine the case of a central bank estimating the model's parameters in real time, and updating its estimate when formulating the optimal policy. Within the real-time learning framework, we examine the problem of a central bank which must also select across non-nested model alternatives. Cogley and Sargent (2005) and Cogley et al. (2011) are closer in spirit to our work, assuming that the central bank updates the probability assigned to alternative reference models when formulating policy. These contributions rely on model averaging in a Bayesian learning framework, thus the optimal policy is affected by the possibility of large losses in competing models. Our analysis focuses on the econometric problem of distinguishing across model alternatives in a real-time learning environment. As in most of the learning literature, we assume a boundedly rational policymaker does not account for uncertainty in the parameters estimate when implementing policy, or for the learning process of the private sector.

2 Competing Reference Models and the Detection of Misspecifications

2.1 Economic Environment

The economy is described by a log-linear, two-equations new Keynesian model. The true model is given by the Gali and Gertler (1999) hybrid-inflation model:

$$\textit{True model} \qquad \pi_t = (1-\theta)\,\beta E_t\pi_{t+1} + \theta\pi_{t-1} + \lambda x_t + \varepsilon_{u_t} \tag{1}$$

$$x_t = \delta x_{t-1} + (1-\delta)E_t x_{t+1} - \phi\,(i_t - E_t\pi_{t+1}) + \varepsilon_{g_t} \tag{2}$$

where x_t, π_t and i_t are the output gap, inflation, and the nominal interest rate. The parameters δ and θ depend on the degree of habit persistence in households' preferences, and on the share of backward-looking price-setters among firms. The exogenous supply and demand shocks ε_{u_t}, ε_{g_t} are *iid* random processes.

The central bank believes are given by:

$$\textit{Central bank} \qquad \pi_t = \beta E_t\pi_{t+1} + \lambda x_t + u_t \tag{3}$$

$$\textit{misspecified model} \qquad x_t = \delta x_{t-1} + (1-\delta)E_t x_{t+1} - \phi\,(i_t - E_t\pi_{t+1}) + \varepsilon_{g_t} \tag{4}$$

$$u_t = \rho u_{t-1} + \varepsilon_{u_t}$$

Thus the central bank believes that the economy is impacted by serially correlated supply shocks and inflation is completely forward-looking.

The central bank chooses the time-consistent policy to minimize

$$\textit{Loss}\,(t,\infty) = E_t \sum_{j=0}^{\infty} \beta^j \left[\pi_{t+j}^2 + \alpha x_{t+j}^2 + \nu\,(\Delta i_{t+j})^2\right]. \tag{5}$$

The central bank formulates the optimal policy conditional on eqs. (3), (4) and on the real-time estimate of the unknown parameters, $\lambda,\ \phi,\ \rho$. The other parameters are known. The observable state variables at t (in both the true and misspecified model) are $x_{t-1}, \pi_{t-1}, i_{t-1}$. We specify policy as a robust optimal explicit rule (Giannoni and Woodford, 2005). This targeting rule is expressed in terms of endogenous variables and is robust to misspecification of the shocks process.

Replacing expected values with realizations, eq. (4) becomes

$$\begin{aligned} s_t &= -\phi\left(i_t - \pi_{t+1}\right) + g_t - (1-\delta)\varepsilon^x_{t+1} + \phi\varepsilon^\pi_{t+1}, \\ s_t &\equiv x_t - \delta x_{t-1} - (1-\delta)x_{t+1} \end{aligned} \tag{6}$$

where $\pi_{t+1} - E_t\pi_{t+1} = \varepsilon^\pi_{t+1}$ and $x_{t+1} - E_t x_{t+1} = \varepsilon^x_{t+1}$. Similarly, eq. (3) can be expressed as

$$\begin{aligned} p_t &= \lambda x_t + u_t - \varepsilon^\pi_{t+1} \\ p_t &\equiv \pi_t - \beta\pi_{t+1}, \end{aligned} \tag{7}$$

Using the Cochrane-Orcutt transform:

$$\begin{aligned} p_t(1-\rho L) &= \lambda(1-\rho L)x_t + (1-\rho L)(u_t - \varepsilon^\pi_{t+1}) \\ p_t &= \rho p_{t-1} + \lambda x_t - \rho\lambda x_{t-1} + (\varepsilon^u_t - \varepsilon^\pi_{t+1} + \rho\varepsilon^\pi_t). \end{aligned} \tag{8}$$

We assume the central bank estimates eqs. (6) and (8) using a GIV estimator without imposing any non-linear parameter restriction. Details on this estimation problem are in Dennis and Ravenna (2008).

Private agents estimate the economy's law of motion and use it to form expectations of future output and inflation. The REE takes the form:

$$z_t = Az_{t-1} + Bv_t \tag{9}$$

where $z_t = \left[\begin{array}{ccc} \pi_t & x_t & i_t \end{array}\right]'$ and $v_t = \left[\begin{array}{cc} \varepsilon_{u_t} & \varepsilon_{g_t} \end{array}\right]'$. The private sector forms expectations $E^*_t z_{t+1}$ according to the Perceived Law of Motion (PLM)

$$z_t = A^* z_{t-1} + B^* v_t \tag{10}$$

where the matrices $A^*,\ B^*$ are estimated through OLS. Then:

$$E^*_t z_{t+1} = A^* z_t + B^* E_t v_{t+1} \tag{11}$$

Thus the structural model can be rewritten as:

$$A_0 z_t = A_1 z_{t-1} + A_2 E^*_t z_{t+1} + A_3 v_t$$

The reduced-form dynamics of the economy, or Actual Law of Motion (ALM), is then:

$$z_t = A^{ALM} z_{t-1} + B^{ALM} v_t \tag{12}$$

where

$$\begin{aligned} A^{ALM} &= (A_0 - A_2 A^*)^{-1} A_1 \\ B^{ALM} &= (A_0 - A_2 A^*)^{-1} A_3 \end{aligned} \tag{13}$$

The true model parameter values are given by $\delta = 0.5$, $\theta = 0.30$, $\beta = 0.99$, $\phi = -0.15$, $\lambda = 0.20$, $\alpha = 1.00$ and $\nu = 0.5$. We set the standard deviations σ_{ε_u}, σ_{ε_g} of the demand and supply shocks equal to 1.5. In every period t, a new observation is generated using the ALM which includes the central bank's optimal policy conditional on its estimates. Then, the central bank and the private sector update the estimate of ϕ, λ and A^*, B^*, which enter into the ALM at $t+1$.

2.2 Learning the True Model

We ran 500 simulations to determine the average speed of detection of the misspecification under various policies and learning environment, assuming that the central bank has available an initial sample of twenty observations, generated by the true model conditional on the rule $i_t = 1.5\pi_t + 0.5x_t$.

The policy maker runs an econometric test of the believed model against the alternative, given by the true data-generating process. The testing strategy of the central bank is as follows. Eqs. (3) and (1) can be rewritten as

$$\pi_{t+1} = \frac{1}{\beta}\left[-\rho\pi_{t-1} - \lambda x_t + \rho\lambda x_{t-1} + (1+\beta\rho)\pi_t\right] + \varepsilon'_{t+1} \tag{14}$$

$$\pi_{t+1} = \frac{1}{(1-\theta)\beta}\left[-\theta\pi_{t-1} - \lambda x_t + \pi_t\right] + \varepsilon'_{t+1} \tag{15}$$

where we replaced expected values with actual realizations and the variable ε'_t is a linear combination of *iid* exogenous shocks and forecast errors. While neither equation nests the alternative, the reduced-form, linear unrestricted coefficient vector of eq. (14) nests the reduced-form alternative (15). The central bank can distinguish the two models by testing whether the coefficient on x_{t-1} in (14) is significantly different from zero.

We assume the policy maker runs two Wald tests on the estimated equation:

$$\pi_{t+1} = c + \beta_1\pi_{t-1} + \beta_2 x_t + \beta_3 x_{t-1} + \beta_4\pi_t + \varepsilon'_{t+1} \tag{16}$$

The first test compares the hypotheses:

$$
\begin{aligned}
H_0 &: \quad \beta_3 = 0 \\
H_1 &: \quad \beta_3 \neq 0
\end{aligned}
$$

If H_0 cannot be rejected, doubt is cast on the central bank's model. Since the acceptance region for H_0 can be very large, implying the possibility of a high probability of type-II error, we assume the misspecification detection occurs when the policy maker cannot reject H_0 and at the same time can reject H_2 in the test:

$$
\begin{aligned}
H_2 &: \quad \beta_3 = k \\
H_3 &: \quad \beta_3 \neq k
\end{aligned}
$$

The parameter k is set equal to 0.5 - a very large value, which will bias the test towards detection.

Through Monte-Carlo experiments comparing the performance of an IV and OLS estimator for β_3 in eq. (16) we established that the OLS estimator has negligible bias, and that in our test it provides superior power in detecting the mis-specification. Table 1 reports the average detection time for the misspecification in the central bank's model, measured using the average number of quarters after which H_0 cannot be rejected and H_2 is rejected at 10% significance level for four consecutive samples.

Does the real-time learning behaviour of the central bank or the private sector impact the speed of detection? Interestingly, very little. The first and second line of Table 1 shows that in the full information case the detection time is about the same as in the case when the central bank is estimating the values of λ and ϕ in real time.

The third and fourth lines of Table 1 show that when the private sector is learning - either with or without the central bank also estimating the model's parameters - misspecification detection happens at about the same speed as in the full information case.

Wieland (2000) argues that a central bank weighing the advantages of learning fast against the advantages of using an optimal policy conditional on partial information may find desirable to deviate from the full-information optimal policy to speed up the learning process. Within our framework, using an optimal policy updated in real time does not penalize significantly the learning speed of the policy-maker. Table 2 shows a similar detection time for the (boundedly rational) optimizing policymaker, and for the one achieved by a policymaker using a Taylor rule $i_t = 1.5\pi_t + 0.5x_t$.

The speed of model misspecification detection is heavily impacted by the shocks realization. Table 3 shows that relative to our baseline case, where $\sigma_g^2 = \sigma_u^2 = 1.5$, lowering the volatility of the cost-push shock increases the detection time by nearly

40%. The learning speed changes dramatically if the volatility of demand shocks decreases, so that $\sigma_g^2/\sigma_u^2 = 1/5$. In this case, learning has not happened even after 30 years.

3 Conclusions

This paper uses a New Keynesian model of the business cycle to investigate how real-time parameter learning, policy choices and shocks' volatility affect the policymakers' ability to distinguish across competing models of the economy

Our results show that a central bank using an optimal policy while trying at the same time to learn the deep parameters of the model and to identify a misspecification in its believed model of inflation, will detect the misspecification at about the same speed as a policymaker using a given Taylor rule. While there may exist an incentive for experimentation, using the boundedly-rational optimal policy does not put the policy-maker at a clear disadvantage in the context of the our model. This result crucially depends on the relative volatility of the supply and demand shocks. When the supply shock volatility is five times as large as the demand shock, model misspecification can go undetected for over 30 years.

References

[1] Cogley, T., De Paoli, B., Mattes, C., Nikolov, K. and Yates, T., (2011), "A Bayesian approach to optimal monetary policy with parameter and model uncertainty", mimeo.

[2] Cogley, T. and Sargent, T., (2005), "The Conquest of U.S. Inflation: Learning and Robustness to Model Uncertainty", *Review of Economic Dynamics,* 8, 210-261.

[3] Dennis, R. and Ravenna, F., (2008), "Learning and Optimal Monetary Policy", *Journal of Economic Dynamics and Control,* 32, 1964-1994.

[4] Evans, G., Honkapohja, S., (2003), Expectations and the stability problem for optimal monetary policies. Review of Economic Studies 70, 807-824.

[5] Galí, J., and M. Gertler, (1999), "Inflation Dynamics: A Structural Econometric Analysis," *Journal of Monetary Economics*, 44, 195-222.

[6] Giannoni, M., Woodford, M., (2005), Optimal inflation targeting rules, in: Bernanke, B., Woodford, M. (Eds.), The Inflation Targeting Debate. University of Chicago Press, Chicago.

[7] Hansen, L., and T. Sargent, (2003), "Robust Control of Forward-Looking Models," *Journal of Monetary Economics*, 50, 581-604.

[8] Wieland, V., (2000), "Learning-by-doing and the Value of Optimal Experimentation", *Journal of Economic Dynamics and Control,* 24 (4).

Appendix: Monte-Carlo Assessment of the Misspecification Detection Wald Test Based on the OLS and IV estimator

To check the power of the Wald test described in section 2.2, we ran a Monte-Carlo experiment using an IV estimator for eq. (16) and the four instruments π_{t-1}, x_{t-1}, i_{t-1}, π_{t-2}. The experiment assumes all agents have full information, and no learning takes place. The large variance of the IV estimator for β_3 makes the Wald test very inefficient in detecting model misspecification. Our experiment over 500 simulation of 20 to 120-period sample-length shows that the number of rejections of the false null hypothesis H_2 increases very slowly with the sample size, and after a 30 year-long span still amounts to less than 20% of the simulations.

Since a Hausman test could not reject the null hypothesis that the OLS estimator of β_3 in eq. (16) is consistent, we performed the same Monte-Carlo experiment using an OLS estimator. We ascertained that the bias of the OLS estimate for β_3 is negligible over the sample sizes considered, while the Wald test could reject in at least 90% of the samples the false null hypothesis H_2 (and reject in less than 10% of the samples the true null hypothesis H_0) with a sample size as small as 50 observations. The results in the paper adopt the OLS estimator.

Table 1: Model selection and the impact of learning

	Average Detection Time (quarters)
Full Information	35
Central Bank Learning	36
Private Sector Learning	36
All Agents Learning	36

Table 2: Model selection and the impact of policy

	Average Detection Time (quarters)
All Agents Learning - optimal policy	36
Private sector Learning - Taylor rule	38
Full Information - Taylor rule	36

Table 3: Model selection and the impact of shocks

All Agents Learning	*Average Detection Time (quarters)*
$\sigma_{\varepsilon_g} = \sigma_{\varepsilon_u} = 1.5$	36
$\sigma_{\varepsilon_g} = 1.5$; $\sigma_{\varepsilon_u} = 0.3$	50
$\sigma_{\varepsilon_g} = 0.3$; $\sigma_{\varepsilon_u} = 1.5$	> 120

Note: Average Detection Time is the average number of quarters after which H_0 cannot be rejected and H_2 is rejected at 10% significance level for four consecutive samples, computed over 500 simulations.