



Contents lists available at ScienceDirect

Journal of Economic Dynamics & Control

journal homepage: www.elsevier.com/locate/jedc

Learning under fear of floating

Saki Bigio*

Department of Economics, New York University, 19 w 4th Street, New York, NY 10012, USA

ARTICLE INFO

Article history:

Received 3 May 2009

Accepted 29 March 2010

Available online 12 May 2010

JEL classification:

C11

E44

E58

F31

F33

Keywords:

Balance-sheet effect

Fear of floating

Model uncertainty

Bayesian learning

Robustness

ABSTRACT

In recent years a large fraction of economies overcame the fear of floating. We study a model that describes the policy of a Central Bank uncertain about whether currency depreciations cause output to expand (textbook model) or contract (balance-sheet model).

We conclude that the movement away from fear of floating may not be explained by Bayesian or Robust policies. When the private sector anticipates the Central Bank's policy and endogenously determines the model, Central Banks may fall in a learning trap. An increase in financial volatility provides an escape to such trap that replicates patterns in the data.

© 2010 Elsevier B.V. All rights reserved.

1. Introduction

Cross-country evidence suggests that in recent years a large number of developing countries seem to have moved towards more flexible exchange rates. Central Bankers in these economies often think of exchange rates¹ as having two potential effects that work in opposite directions. On the one hand, high unexpected depreciations might be beneficial to agents that generate some fraction of their income in foreign currency (e.g., the exporting sector or agents holding dollar-denominated assets). Under nominal rigidities, a sudden shift in the value of the nominal exchange rate translates to a sudden shift in the real exchange rate and therefore produces beneficial wealth and competitiveness effects. This is a standard economics textbook effect.²

On the other hand, the balance-sheet effect is a harmful after-effect that might result from high unexpected exchange rate depreciations if it induces imbalances in the asset-liability positions of economic agents. In underdeveloped financial markets, it is likely that firms have low coverage and that the balance-sheet effect may have aggregate implications. This may occur when a large fraction of agents hold assets denominated in domestic currency while holding liabilities denominated in US Dollars. Without insurance to cover them from these losses, exposed firms might face a credit crunch

* Tel.: +1 212 998 8900; fax: +1 212 995 4186.

E-mail address: sbigio@nyu.edu¹ We define the exchange rate as the domestic currency price of the US Dollar.² See, for example, Clarida et al. (2001) or Gali and Monacelli (2005) for versions of open economy models with nominal rigidities in which the textbook effect is present.

that leads to an aggregate downturn. Theoretical and empirical work to date formalized these ideas and have documented evidence on the balance-sheet effect.³

In this paper, we ask if learning about either effect can explain the trend towards more flexible exchange rates. We abstract from the main mechanisms driving either model. We characterize the problem of a Central Bank that has uncertainty over a single parameter that encompasses both effects. Both sub-models are used to guide policy, but each of them delivers opposite effects from unexpected exchange rate movements. In this environment, policy makers fear missing the true model because outcomes in terms of the data-generating process might render high losses.

We first study a Bayesian Central Bank that acts in the spirit of Brainard (1967). Policy makers react according to prior beliefs about either model and update these beliefs through Bayes rule. We then explain how this policy and its learning properties are altered when the Central Bank behaves according to a policy consistent with fear of misspecification. That is, we analyze the behavior of a Central Bank that behaves optimally according to multiplier preferences as described in Hansen and Sargent (2006, 2007).

The tested hypothesis in this case is that Central Banks around the world assigned a strong prior during the early 1990s to the balance-sheet model, and that it took them a decade to learn that they were wrong. In the context of the model, slow learning could have been caused by a fear of floating policy: due to initial uncertainty, no important exchange rate swings were allowed but these swings are key to a quick detection of the true economic structure. Hence, fear of floating policies may have slowed down learning which further sustained the fear-of-floating policy. In this context, these economies would have slowly moved towards more flexible exchange rates.

In spite of being a reasonable explanation, this story is inconsistent with our model's predictions. When we calibrate the model to reasonable parameters, we find that learning is expected to occur in about 5 years for both Bayesian and Robust policies. For a large range of plausible parameterizations, the learning rate is faster than what we found in cross-country data. This conclusion remains unaltered even when the Central Bank assigns less weight to past information (constant gain).

Later in the paper, we study the learning process in which the Central Bank is uncertain about whether private agents are passive or active. Active agents take hedging decisions based on a cost parameter and the Central Bank's prior on their type. When the economy is populated by active agents, the balance-sheet effect is determined by their decisions to hedge. Passive agents, on the other hand, are unable to hedge and, therefore, always induce a balance-sheet effect. With this, we study a simple mechanism where the structure of the economy depends on the government's policy in the spirit of Chang and Velasco (2006).

In this context, the Central Bank that strongly believes that private agents are passive, falls into a learning trap. Because it has this strong belief, it suffers from fear of floating. With low exchange rate volatility, active agents that would otherwise hedge do not, and the Central Bank loses the ability to learn.

We use this model to show how a short temporal shock in this model could help explain the learning pattern found in the data. We provide an example in which an increase in the financial volatility, perhaps due to financial crisis of the late 1990s, could have provided some Central Banks with a window of opportunity to learn. We discuss several explanations for such shocks.

The title is suggestive: this paper tries to explain learning under fear of floating. For this purpose, we proceed as follows: In the following section we describe how the symptoms of fear of floating across countries have been systematically weaker during the last 4 years of the data. In Section 2 we introduce a benchmark small open-economy model and in Section 3 we present an analytical solution to the optimal Bayesian Central Bank policy under passive learning. In Section 4 we explain how multiplier preferences alter this policy and how Robust policies are equivalent to Bayesian policy with a distorted prior. In Section 5 we discuss the learning properties of both policies when the uncertainty is about the balance-sheet effect. Section 6 introduces the model where the private sector reacts endogenously and describes learning in this context. We conclude in Section 7.

1.1. Evidence on learning under fear

Hausmann et al. (2001) and Calvo and Reinhart (2002) provide a thorough study of exchange rate interventions in different countries. Both papers agree that during the 1990s decades, many developing countries had lower relative variance of the exchange rate depreciations over the interest rates than developed economies. These studies suggest that countries that shared this feature were countries in which the Central Banks often intervened in the exchange rate market. Calvo and Reinhart (2002) call this phenomenon fear of floating and is rationalized by balance-sheet effects.

What if Central Banks were unsure about the presence of the balance-sheet effect? Would they have been able to learn the structure of their economies over time? Two graphs based on the statistic built by Hausmann et al. (2001) (hereafter HPS-statistic) are suggestive. Fig. 1 compares the relative volatility of exchange rate depreciations and interest rates for 40 countries.⁴ The HPS-statistic reflects the degree of intervention in exchange rate markets using nominal interest rates. A low statistic shows *high* intervention. For a Bayesian Central Bank, relatively low HPS-statistics are consistent with placing a strong prior on the balance-sheet model.

³ See Aghion et al. (2000, 2001), Caballero and Krishnamurthy (2001), Krugman (1999), Cespedes et al. (2002), Chang and Velasco (2006), Schneider and Tornell (2004) and Frankel (2005) for theoretical work. See Hausmann et al. (2001) and Calvo and Reinhart (2002) for empirical evidence.

⁴ The data were obtained from the International Financial Statistics provided by the International Monetary Fund. The panel includes 40 countries. Exchange rates correspond to the money markets. The interest rates are the Central Bank's discount rate when available which are replaced by the inter-bank rate when not available. The data cover the months between January 1994 and December 2008 as the broadest range. The author provides the data set upon request.

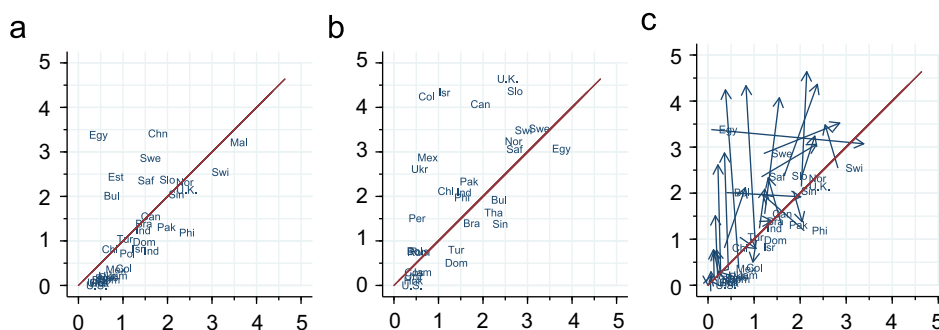


Fig. 1. Cross-country HPS-statistic: the figure plots the HPS-statistic for countries with a low statistic (high intervention) for different sample periods. Panel (a) plots the statistic computed for 1994–1998 (x -axis) against 1999–2003 (y -axis). Panel (b) plots the statistic computed for 1999–2003 (x -axis) against 2004–2008 (y -axis). Panel (c) shows the change from Panel (a) to Panel (b). The solid line depicts the 45° line. Countries close to the line had a stable statistic.




Fig. 2. Cross-country HPS-statistic for HP-filtered data. The figure plots the HPS-statistic for countries with a low statistic (high intervention) for different sample periods. The data were filtered using the Hodrick–Prescott filter with parameter 40000. Panel (a) plots the statistic computed for 1994–1998 (x -axis) against 1999–2003 (y -axis). Panel (b) plots the statistic computed for 1999–2003 (x -axis) against 2004–2008 (y -axis). Panel (c) shows the change from Panel (a) to Panel (b). The solid line depicts the 45° line. Countries close to the line had a stable statistic.

The sample of countries is restricted to a sample of economies for which the statistic is, at most, 6. There are three panels in the figure. Panel (a) is a cross-country scatter plot of the statistic for two different periods: the x -axis shows the average statistic for the period 1994–1998, and the y -axis is for 1999–2004. The solid line is a 45° line. Countries close to the 45° line had a stable HPS-statistic over both periods. The closer the points to this line, the more stable were the policies. Panel (a) is consistent with Hausmann et al. (2001) and Calvo and Reinhart (2002), in that the ranking of exchange rate interventions remained constant through time: countries like Australia, Canada or Japan intervened less than Peru, Ukraine or Uruguay.

Panel (b) is the same plot as Panel (a) for the 1999–2004 sample against the 2005–2008 sample. Panel (c) depicts, with arrows, the change in each country's position from Panel (a) to Panel (b). Panel (c) shows an international trend: a movement out of the 45° line. Note that the majority of developing countries in the beginning sample are clustered around the 45° line in Panel (a). This suggests that during the 1990s developing countries constantly intervened in exchange rate markets although across countries the implied statistic was different. Surprisingly, the pattern is not repeated in the center panel. The bulk of developing countries in the sample jumped to a higher statistic revealing a trend towards more flexible regimes. That is, most of the countries in the sample show a change from maintaining a low statistic (high intervention) for the first 2 sub-samples, and suddenly allowed more volatility in the exchange rate at the end of the sample.

We interpret this as an international trend towards more flexible exchange rates. It is plausible to think more of a change in policy rather than having the data driven by common shocks as the sample periods are long, of about 50 months, on average, per period. On the other hand, when the sample is restricted to the group of least intervention, the pattern is lost, suggesting that the pattern is only constant among countries suffering fear of floating at the beginning.⁵

To check if the evidence is Robust to a common component, we detrended the data. Detrending would eliminate any artificial increase in the variance of exchange rates induced by systematic changes in the mean. Fig. 2 mimics Fig. 1 for HP-filtered data. The overall picture that the pattern is still present though the evidence is not as strong.

For the rest of the paper, we test the hypothesis that many developing countries assigned a strong prior to the balance-sheet model at the beginning of the sample and that they learned, towards the 2000s, that their fear was unfounded. Before proceeding to the model though, we discuss related papers.

⁵ In fact, the 1990s were a period in which exchange rate interventions by Central Banks were constantly discussed by policy makers and researchers. See for example: Chang and Velasco (2000, 2006), Ize and Levy-Yeyati (2003), Levy-Yeyati and Sturzenegger (2003a, 2003b), or Goodhart (2005).

1.2. Related papers

The nature of the problem we present here is closely related to other recent works that study the behavior of Central Banks under model uncertainty. This paper follows the work by Wieland (2000a, 2000b), Ellison (2006), Cogley et al. (2007, 2008) and Ellison et al. (2007). Most of these papers focused on uncertainty about the sacrifice ratio implied by the Phillips curve and the benefits of policy experimentation in closed economies. Among these papers, Ellison et al. (2007) is the closest to ours. That paper studies uncertainty about the exchange rate in a two-country setup in which devaluations have effects through the relative price of imports to exports. Our paper differs from theirs not only by the nature of uncertainty, but also because we pay special attention to the pace at which the Central Bank under Bayesian and Robust policies learn using a model as close as possible to the benchmark used by many Central Banks.

Cho and Kasa (2008) is another related paper.⁶ That paper explains how recurrent currency crises occur when a Central Bank misinterprets the forward-looking balance-sheet model of Aghion et al. (2001) from a backward-looking open economy Keynesian model. In the balance-sheet model that they study, unexpected nominal devaluations dampen output. The Central Bank only believes the Keynesian model but ignores the parameter that governs the trade-offs between exchange rate and output. When an unexpected appreciation occurs, the Central Bank will observe a surprising increase in output. Because it believes in the wrong model, it will misinterpret the data and estimate that a higher exchange rate leads to higher output. Because private agents incorporate a higher exchange rate into their expectations, this phenomenon reinforces the Central Bank's beliefs.⁷ Therefore, the combination of large shocks and discounted learning lead their economy to experience recurrent swings from one self-confirming equilibrium path to another.

Cho and Kasa (2008) and our paper share the common feature of a Central Bank ignoring whether the balance-sheet effect is present. Nevertheless, whereas Cho and Kasa (2008) is suited to explain currency crises, it is not suited to explain the trend towards less interventions. By construction, in that model, the relative conditional volatility of exchange rates to the interest rate is constant. As in our model, the Central Bank controls the exchange rate via the uncovered interest rate parity but only in our paper does the interest rate affect output through a direct effect and indirectly via the effect of exchange rates. This allows the HPS-statistic to vary depending on the Central Bank's beliefs. In Cho and Kasa (2008) the Central Banks sets the interest rate to target a particular exchange rate so the HPS-statistic should be constant over time.

In Section 6, we study a model where the private sector determines the balance-sheet effect. This exercise is inspired by Chang and Velasco (2006). In that paper, firms can issue debt in foreign or domestic currency. The Central Bank observes that decision and decides between allowing a completely flexible exchange rate or fixing it. They show that both flexible and fixed exchange rate equilibria may coexist.

Though the exercise we perform is based on a similar insight, it differs from that paper in several respects. Most importantly, the timing is different. Here we assume that the Central Bank decides on a policy first and then agents decide on hedging. As opposed to Chang and Velasco (2006), we assume that agents can be of multiple types and the Central Bank may learn the composition in the population. This uncertainty eventually determines the Central Bank's policy.

Like Caplin and Leahy (1996), in our paper, the Central Bank faces uncertainty about a feature of private agents. In that paper, private agents belong to a continuum of types and the Central Bank updates beliefs which ultimately determine policy. Given these beliefs, private agents make decisions to carry out or postpone investment. Unlike that paper, learning here will not always be possible.

2. A standard small-open economy model

We use a standard new-Keynesian small-open economy model as our workhorse since it is a benchmark for Central Banks around the world. The structure here is close to Ball (1999) and Gali and Monacelli (2005). The model of Gali and Monacelli (2005) begins from proper micro-foundations that include monopolistic competition and nominal rigidities which is the reason why nominal variables have real effects.

As in all new-Keynesian models, we have a Phillips curve:

$$\pi_t = \beta E_t[\pi_{t+1}] + \gamma y_t + \varepsilon_{\pi,t} \quad (1)$$

where π_t is the inflation rate, y_t is the output gap and β is the period discount factor. Inflation depends on the expectation of inflation $E_t[\pi_{t+1}]$ and the output gap y_t .⁸ The term $\varepsilon_{\pi,t}$ represents a cost-push shock.

In addition, there is an aggregate demand equation:

$$y_t = E_t[y_{t+1}] - \chi(i_t - E_t[\pi_{t+1}] - r_t^e) + \theta E_t[\Delta s_{t+1}] + \varepsilon_{y,t} \quad (2)$$

⁶ That paper belongs to the literature on adaptive learning. This literature typically uses two ingredients to explain recurrent changes in policy: a Central Bank that ignores a forward-looking component in the true model, and discounting of past information. Work in this area includes Sargent (1999), Cho et al. (2002) and Sargent and Williams (2005).

⁷ Indeed the Central Bank is observing a higher exchange rate, with less effects on output. This leads to dynamics in which an initial large unexpected appreciation propels a large expected devaluation path.

⁸ Notice that the small-open economy Phillips curve presented here ignores the pass-through mechanism. This mechanism would not add much to a discussion on whether the textbook effect or the balance-sheet effect dominates.

This equation describes how the output gap depends on its own one-period ahead expectation term, the gap between the real interest rate and its natural level r_t^n , the expected nominal depreciation $E_t[\Delta s_{t+1}]$ and a demand shock $\varepsilon_{y,t}$. The textbook model will have θ to be positive. Therefore, nominal devaluations will expand output. The balance-sheet model works the opposite way with θ being negative. The nominal interest rate i_t is controlled by the Central Bank, so we will refer to it as its policy instrument.

The expected nominal exchange rate depreciation $E_t[s_{t+1} - s_t]$ is obtained via an uncovered interest rate parity equation:

$$E_t[\Delta s_{t+1}] = i_t - i_t^* - \varepsilon_{et} \quad (3)$$

This is a non-arbitrage condition between domestic i_t and foreign interest rates i_t^* . We call the shock to this equation a financial shock $\varepsilon_{e,t}$.

The natural interest rate r_t^n evolves according to the following law of motion:

$$r_t^n = \rho_n r_t^n + \varepsilon_t^r \quad (4)$$

and ε_t^r is a shock to this autoregressive process. The system is characterized by the following set of parameters which are constant over time: $[\beta, \gamma, \chi, \theta, \{\rho_s, \mu_s, \sigma_s^2\}_{s=\pi, y, e, i, r}]$.

Notice that the model does not include terms of trade as some new-Keynesian models do. In fact, Lubik and Schorfheide (2007) have shown that at least for Australia, Canada, New Zealand and the UK, the terms of trade in the model of Gali and Monacelli (2005) are not important for the business cycle. Hence, one can be comfortable that this simplification will not alter the results.

The model is very similar to its closed economy counterpart as suggested by Clarida et al. (2001). Following Clarida et al. (2002), the Central Bank will seek to minimize a standard quadratic loss function⁹:

$$L_\tau = E_\tau \left[\sum_{t=\tau}^{\infty} \beta^{t-\tau} (\pi_t^2 + w y_t^2) \right] \quad (5)$$

The minimization is constrained by Eqs. (1)–(4). In the new-Keynesian framework, Central Banks are concerned about stabilizing output because it puts pressure on inflation and because deviations from its steady state are inefficient when the market is not competitive and there are nominal frictions.

The timing protocol of the models is as follows: all shock processes $\varepsilon_{s,t}$ for $s = \pi, y, e, i, r$, follow an AR(1) process with corresponding set of parameters $\{\rho_s, \mu_s, \sigma_s^2\}$. We denote the innovations to these processes by $v_{s,t}$ correspondingly. We further assume that all the $\varepsilon_{s,t}$ innovations are observed only with a one-period lag except for the ε_{et} risk premium shock that affects the depreciation expectations and the world interest rate shock $\varepsilon_{i,t}$. This assumption is supported by the fact that spot exchange rates, as well as nominal domestic and world interest rates, are observed immediately.

Under this protocol, optimal policies with certainty will be the limit cases of the Bayesian policies we describe below. We can obtain the optimal discretionary policy (without commitment) as a function of parameters and observed shocks. In the rest of the paper, we focus on the optimal discretionary policies as opposed to commitment policies. Credible policies under commitment for a Central Bank remains an unresolved question in economics. The optimal policy, steady state equilibrium and other properties are described in Appendix A.

3. Optimal Bayesian policy

A Bayesian Central Bank that ignores whether data are generated by Model A or B, will act in accordance to a model prior probability assigned to each, and maximize expected utility. In dynamic setups, the Central Banker may incorporate the fact that his policies will affect his ability to learn. Here, Central Banks are assumed to be myopic, in the sense that they will not take into account the effect that their actions will have on their ability to learn.¹⁰ Active learning would not affect the results since, as we will show, learning would have been too fast even without experimentation.

3.1. Knowledge assumptions

The only particular feature about the knowledge assumption is that θ has an unknown value to both policy makers and agents within the model. Recall that θ is the parameter that translates exchange rate depreciations to an effect in the output gap. In addition, we adopt the common prior assumption. By doing so, uncertainty becomes common knowledge

⁹ This loss function is analogous to the closed-economy version of this model. Recent papers have provided micro-foundations of a welfare function in the small-open economy that differ from this one. Faia and Monacelli (2008) show that home bias in consumption is a sufficient condition for introducing the exchange rate into the objective. De Paoli (2009) also argues that the exchange rates should be included when the economy faces monopolistic competition in domestic and foreign goods.

¹⁰ In a related paper, Bigio and Vega (2006) study whether intentional policy experimentation could be beneficial. In their model, the balance-sheet effect has a non-linear impact on output. In spite of this non-linearity, no substantial gains from intentional policy experimentation are found. This is a result in line with the one obtained by Cogley et al. (2007, 2008). We keep the focus of the discussion here to passive learning policies. In contrast Wieland (2000a, 2000b) show that intentional policy experimentation is beneficial and may avoid policy traps.

and does not require solving the problem of ‘forecasting the forecast of others’.¹¹ A theorem in [Aumann \(1979\)](#) asserts that under this assumption, posteriors coincide too.

3.2. The Bayesian Central Bank’s problem

When Central Bankers are uncertain about which of the two possible models is the actual model driving the economy, they assign a probability p_t to Model A and a probability $(1-p_t)$ to Model B. This probability evolves according to the odds ratio, i.e., a combination of model fit and their shocks and an initial prior belief p_0 .

Note that the uncertainty is constrained to only two values of parameters. In fact, we could have chosen to model this source of uncertainty as a density over the values of θ . In fact, the optimal policy in the case of two parameter values is the same optimal policy for a continuum of parameters if we choose an appropriate prior density (although using Bayes rules as an updating rule could change these relation after one period). The advantage of restricting the parameter space is that each parameter has the interpretation of a different model: the textbook and the balance-sheet model.¹²

By only taking into account this simple version of uncertainty, the expected loss function for period t becomes

$$L(p_t)_t = p_t L_{t|A} + (1-p_t)L_{t|B} \quad (6)$$

where $L_{t|A}$ and $L_{t|B}$ represent the value functions conditional on Model A or B being the true models. Note that because in the passive learning environment the priors are fixed, the prior can be taken as a parameter. To avoid notation, Model A will refer to the textbook model and Model B to the balance-sheet model. The loss function (6) is consistent with expected utility theory. The optimal policy will be a function of the prior belief and current available information. The following proposition describes this optimal policy:

Proposition 1 (Optimal Bayesian policy). *The optimal Bayesian monetary policy without commitment is given by the following rule:*

$$i_t^{BAY}(p_t) = \frac{1}{\Psi(p_t)} \left[\left(\chi_\pi + \frac{\gamma_y}{w} \left(\frac{1}{\rho_\pi} - 1 \right) \right) A \rho_\pi \varepsilon_{\pi,t-1} + E_t(\mu_{y,t}|p_t) \right]$$

where

$$\Psi(p_t) = \chi - (p_t \theta^H + (1-p_t) \theta^L)$$

and

$$\mu_{y,t} = \chi r_t^n - \theta i_t^* - \theta \varepsilon_{et} + \varepsilon_{y,t}$$

The detail of the solution to the Central Bank’s problem is presented in the Appendix. The optimal policy reacts to all of the shocks to the model by weighting their effects into the loss function. The reaction to shocks depends on the sign of $\Psi(p_t)$, which in turn depends on the prior. In fact, the sign of the policy reaction may even change depending on the prior (the next section provides detail for a particular calibration). Also note that because the optimal policy depends on expectations of shocks that are not observable, the set of state variables to compute the optimal policy includes all the shocks to the system while including two demand shocks: one for each model. Hence, computing optimal policies requires computing the demand shocks for both models.

The update of Model A probability in the Bayesian approach depends on the Bayesian odds ratio. Let Z_t be the vector of observables at time t . The Bayesian odds ratio is written recursively as

$$p_t = \frac{P(Z_t|M^A, Z_{t-1})p_{t-1}}{P(Z_t|M^B, Z_{t-1})(1-p_{t-1}) + P(Z_t|M^A, Z_{t-1})p_{t-1}} \quad (7)$$

The intuition behind this ratio is that the probability that the policy-maker will assign to a given model for being true will depend on a weighted average (with weights given by the previous period’s model probabilities) of the likelihood of the data at period t , conditional to each model and the data. In Section 5 we describe how the priors converge to the true model, i.e., how Central Banks learn.

By holding priors fixed, we can understand some things about the learning dynamics of this policy. We can observe that the smaller the variance of expected devaluation, $E_t[\Delta s_{t+1}]$, the smaller will be the difference between the implied residuals of Eq. (2) according to Models A and B. Thus the likelihood ratio will be close and the prior will evolve more slowly. When the balance-sheet effect is severe, the variance of expected devaluation will be smaller when the prior p_t is closer to 0, because the balance-sheet model will suggest stronger stabilization of the financial shock. This feature is consistent with fear-of-floating economies as explained in [Calvo and Reinhart \(2002\)](#).

¹¹ [Ellison \(2006\)](#) makes the same assumption for the same reason.

¹² For example, one can generalize the two-parameter structure here, by using a Gaussian prior. The drawback of this approach is that it introduces an additional free parameter: the initial prior precision. The initial prior precision could be accommodated to make the learning process as slow as one wants. In fact, we can also slow down the learning pattern with a two-parameter structure, provided that we begin with a prior arbitrary close to the balance-sheet model. Since a Gaussian prior precision is harder to interpret, there is no further gain in imposing such priors.

No matter what the prior is, it can be shown that under this policy, P_t converges almost surely to the true model. Kasa (1999) provides a formal treatment on conditions for these results. Intuitively, the fact that agents will always learn the true model in this context occurs because fixed exchange rates are never optimal in this setup. Because this is the case, on average, the true model will yield a lower error term in the aggregate demand equation, making its likelihood function higher. This property comes from the fact that uncertainty is described over a single parameter. In general, with uncertainty over more parameters, this assertion is not necessarily true. Therefore, the focus of our description is on the speed of learning rather than conditions that ensure learning at all. The next section describes the benchmark calibration of the model and its main features.

3.3. Calibration and dynamics

Our calibration is designed so that the textbook model and the balance-sheet model prescribe the policy-maker to react in opposite directions after a financial shock. This property is stressed by earlier studies on the balance-sheet model, e.g., Aghion et al. (2000). The Bayesian optimal policy is a convex combination of both policies. We illustrate the model's dynamics implied by different priors. In particular, we want to stress how asymmetric the responses are according to the prior. This has important consequences on learning. In this exercise we fix the prior in order to have a clear idea on what is the expected behavior of Central Banks after a shock to the uncovered interest rate parity (or exchange rate shock).

We calibrate the model according to the following approach: we take the parameters in Eqs. (1)–(3), according to the estimation done by Lubik and Schorfheide (2007) for the Canadian economy. The parameters that determine the structure of the shock are chosen to mimic the Bank of Canada's Quarterly Projection Model described in Murchison and Rennison (2006). We chose Canada as our benchmark small-open economy for Model A because it represents a small-open economy that has developed financially in such a way that the balance-sheet is not present. On the other hand, the HPS-statistic for Canada was not as high as for Australia or New Zealand, but close enough to countries that seem to suffer from fear of floating at the beginning of the sample but not at the end. We calibrated θ_A and the variance of the process to match the HPS-statistic of Canada, and the volatility of devaluations and interest rates. To parameterize Model B, we chose θ_A to match the HPS-statistic of Peru, an example of a fear-of-floating economy. The summary is presented in Table 1.

Under this parametrization we can have a good idea of the exchange rate policy behind the optimal policy by observing the impulse responses functions shown in Figs. 3 and 4. The impulse we analyze is a negative financial shock (exchange rate shock). A negative shock of this type will put pressure towards appreciation. The plots show a continuum of impulse responses for each prior. The darkest lines indicate the behavior when the Central Bank was closer to certainty, that is, when the prior was either close to 1 or close to 0. As the lines get lighter, this means that the prior was of more uncertainty about either model.

Arrows in the figure point out the outcome under the optimal policy for each model (a prior of 1 or 0). One can see in both figures that if the policy-maker knows the true model he is able to stabilize both output and inflation perfectly. Under Model A the optimal policy makes the Central Bank reduce interest rates. It does so because interest rates will have two

Table 1
Calibration.

Parameter	Value	Notes
<i>Reduced form parameters</i>		
β	0.99	Specifying a quarterly model with 4% steady-state real interest rate
γ	0.86	Following Lubik and Schorfheide (2007)
χ	0.37	Following Lubik and Schorfheide (2007)
w	0.5	Putting 2/3 of weight on inflation and 1/3 on output
<i>Shock process parameters</i>		
ρ_π	0.7	Guided by Bank of Canada's Quarterly Projection Model price mark-up shock
ρ_y	0.2	Guided by Bank of Canada's Quarterly Projection Model demand shock
ρ_r	0.3	Following Lubik and Schorfheide (2007) estimate of terms-of-trade estimate
ρ_i	0.8	Consistent with Libor data
ρ_e	0.35	Guided by Bank of Canada's Quarterly Projection Model risk premium (exchange rate) shock
σ_π	0.15	Guided by Bank of Canada's Quarterly Projection Model price mark-up shock
σ_y	0.18	To match volatility of interest rates and devaluations in Canada
σ_r	0.5	Following Lubik and Schorfheide (2007) estimate of terms-of-trade estimate
σ_i	0.25	Consistent with Libor data
σ_e	0.08	To match volatility of interest rates and devaluations in Canada
<i>Models A and B</i>		
θ_A	0.1	To match Canada's HPS-statistic
θ_B	−0.6	To match Peru's HPS-statistic

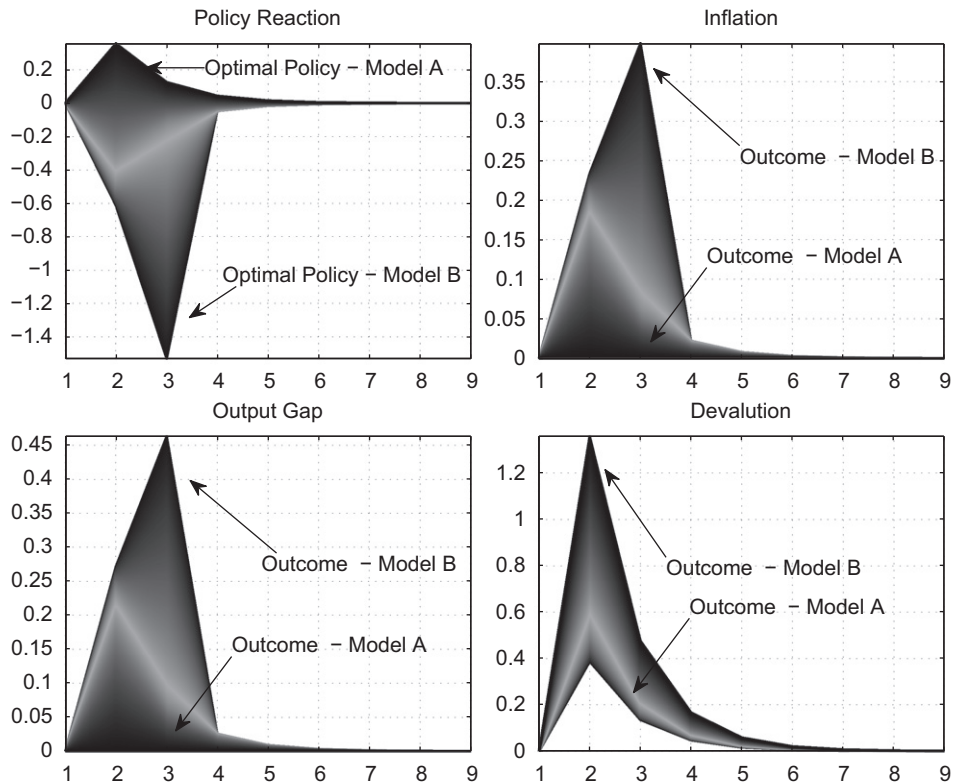


Fig. 3. Bayesian impulse response—true Model A: the figure shows the reaction of all the endogenous variables after a positive financial shock. This shock puts pressure towards devaluation. Darker colors reflect priors closer to certainty. The lighter color is put on the prior with 0.5 probability.

effects that offset each other exactly stabilizing the output gap. On the one hand, a reduction in the interest rates provokes a further appreciation according to the uncovered parity (Eq. (3)). Because under Model A, the aggregate demand (Eq. (2)) reacts negatively to the appreciation but positively to the reduction in interest rates, the policy is designed in such a way that both effects exactly offset each other.

Under Model B, the effects work the opposite way. Appreciations cause output to expand, but the belief that output will contract prescribes to lower rates. This pushes the depreciation even further on, and deviations are even stronger.

Under certainty, there is no tension between stabilizing output and inflation because the model lacks a pass-through rate. A 0 pass-through rate is unlikely in an open economy but for many countries pass-through rates are usually not high.¹³ Our assumption implies that the costs of nominal devaluations are due to model uncertainty rather than to a sacrifice ratio trade-off caused by pass-through.

When the Central Bank assigns a positive prior to both models, the policies are the same during the first period after the shock is realized. In subsequent periods, the policies differ because the Central Bank will believe that there was a non-zero demand shock under the false model. Because it assigns a non-zero prior to that model, and shocks have serial correlation, in subsequent periods, the policy reacts partially to an expected shock in aggregated demand too.

The cost of missing the true model manifests itself through deviations from the steady state inflation and output (upper right and lower left panel). Under Model A, the textbook model, a Central Bank assigning a low prior will increase interest rates inducing a fall in output and a deflation. Fear of floating induces the wrong policy prescription.

When the true model is Model B, the balance-sheet model, and the prior assigned to this model is low, the Central Bank acts by reducing rates, thus causing an output expansion and an undesired inflation.

One additional feature of missing the correct prior is that interest rates show an acceleration effect: the peak increase (decrease) in rates is not immediately after the shock occurs. Rather, when beliefs are far from the truth, after the shock is realized, the Central Bank puts a high prior on the event that the economy was also hit by a demand shock. The Central Bank estimates the shock according to the wrong model and finds it was positive when in fact it did not exist. Because shocks have memory, in subsequent periods they react to the exchange rate shock and a demand shock that never happened. This phenomenon translates into the hump-shape of impulse responses when the prior is wrong. Along the lines

¹³ Lubik and Schorfheide (2007) shows that it is low for industrialized economies. Monacelli (2005) studies optimal policies in low pass-through environments.

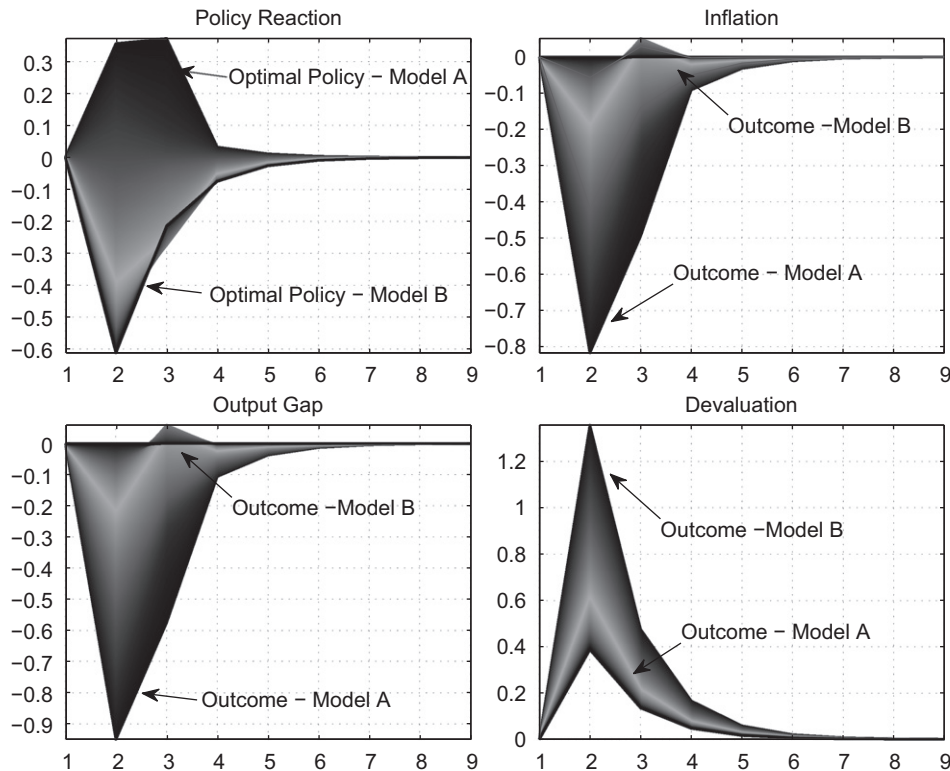


Fig. 4. Bayesian impulse response—true Model B: the figure shows the reaction of all the endogenous variables after a positive financial shock. This shock puts pressure towards devaluation. Darker colors reflect priors closer to certainty. The lighter color is put on the prior with 0.5 probability.

of Chari et al. (2002), this property of the model adds additional persistence of real exchange rates through monetary policy reactions.

Because losses are weighted differently depending on the model, we observe that the Bayesian policy is asymmetric. For a prior of 0.5, the reaction after the shock is closer to the reaction under the balance-sheet model because outcomes are much worse under this model. This condition will have effects over Robust policies.

The differences of the optimal policy under each model will affect the HPS-statistic. Fig. 5 shows the expected statistic for different priors for both models. When a high prior is assigned to the balance-sheet model, the figure shows that the statistic is much smaller. This image is consistent with the fear-of-floating literature. The figure also shows an additional perverse consequence of the balance-sheet model. Volatility is extremely high when the balance-sheet effect is disregarded. When we solve for Robust policies, this feature will have important effects on how the Central Bank will distort its priors. The figures also include the statistic for Canada and Peru. The arrows show the direction of the increase of the statistic from the first sub-sample to the last sub-sample. Though the statistic is itself a random variable, the figure is consistent with the argument that an increase in the statistic would show signs of learning about Model A. Peru could have faced more shocks in the last sample, but a potential alternative story is that it has slowly started to discover that the textbook model applies to its economy.

4. Robust policy

In Section 3 we described how learning about the effects of exchange rate depreciations may occur in less than a decade when the Central Bank is Bayesian under our calibration. In this section, we study the Central Bank that has fear of misspecification of its prior probabilities. Hence, it optimizes in accordance to multiplier preferences instead of expected utility theory. Hansen and Sargent (2006, 2007) provide substantial support to use these preferences as a modeling device that summarizes fear of misspecification. The idea here is that Central Banks can fear that their updating rules are misspecified. They do not fully trust Bayes rule. There are two reasons that may suggest that Central Bankers in countries that face a potential balance-sheet effect may be particularly concerned about misspecification. The first reason is that in spite of observing aggregate data, Central Bankers constantly may receive information from the private sector. Microeconomic evidence on the balance position of firms with foreign lenders and the banking system may be inconsistent with aggregate effects but the Central Bank does not know how to incorporate this evidence. On the other hand, interest groups may have incentives to complain about strong exchange rate movements. Finally, perceived financial fragility from

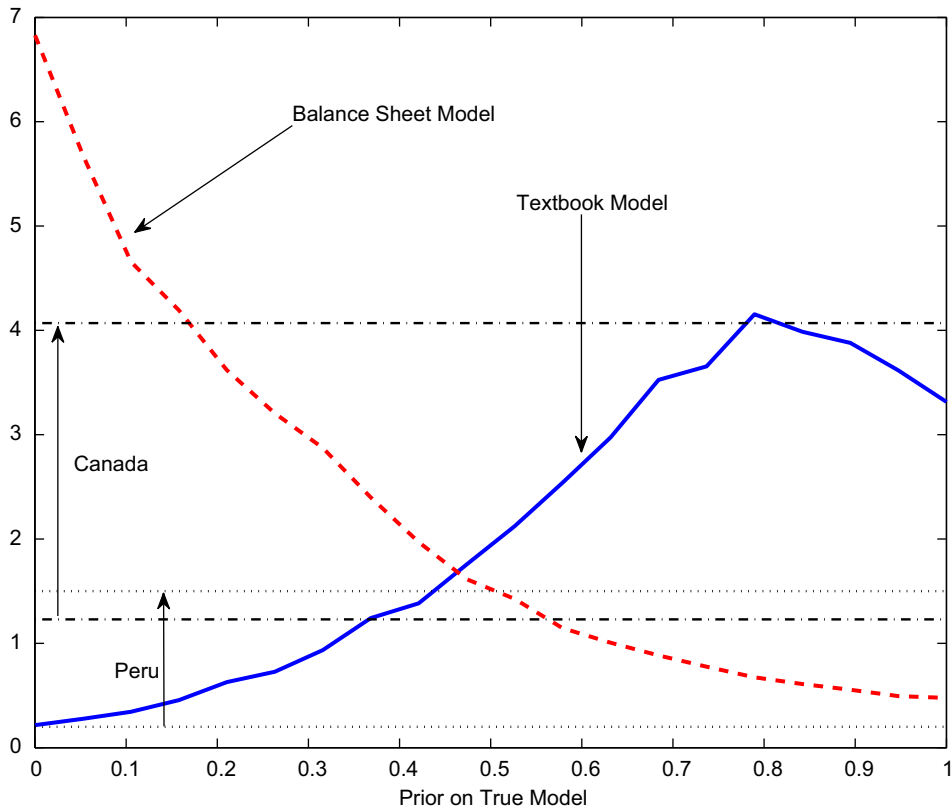


Fig. 5. HPS-statistic implied by each model.

balance-sheet’s solely may neglect the fact that firms can renegotiate their loans to imply the same real burden as before. In fact, firms may be covered against unforeseen exchange rate movements in several ways that include forward-back operations. Coverage mechanisms may be hard to observe by Central Banks.

The second reason is purely statistical: Bayes rule requires additional knowledge. It requires Central Banks to put *a priori* knowledge on the distribution of economic shocks. Even if shocks are believed to be Gaussian, Central Bankers are required to know the true variance in advance.

By using Robust policies, the Central Bank will use Bayes rule as a pivotal mechanism to assess risks, but will minimize loss for a set of priors close to the one computed by the data and Bayes rule. Hansen and Sargent’s framework deals with two sources of misspecification.¹⁴

Robust preferences are characterized by a single parameter, Θ .¹⁵ The higher this parameter is, the more the Central Bank trusts its prior. The Robust policy-maker solves the following problem:

$$L(p_t)_t = \min_{\{\bar{p}_t\}} \max_{\{t_t\}} \left\{ \bar{p}_t \left[-L_{\bar{p}_t|A} + \Theta \log \left(\frac{\bar{p}_t}{p_t} \right) \right] + (1 - \bar{p}_t) \left[-L_{\bar{p}_t|B} + \Theta \log \left(\frac{1 - \bar{p}_t}{1 - p_t} \right) \right] \right\} \tag{8}$$

Here p_t is the same prior probability as in the problem defined in Section 3. The Central Banker acts as if there exists a malevolent agent that attempts to distort the prior probabilities assigned to each of the models. This artificial agent is not entirely free to choose any value, but is constrained by a multiplier, Θ in this case, that summarizes the Central Banker’s fear of misspecification. Without this restriction, the artificial evil agent would choose the balance-sheet model always for the reasons discussed in the previous section. Hence, higher values of Θ allow the evil agent lower distortions to the prior, $\Theta = \inf$ allows no distortions at all. A Central Bank with strong concerns about Robustness will set Θ very low, and the evil agent will distort probabilities accordingly.

¹⁴ For example, in Cogley et al. (2008) two operators are defined to deal with both forms of fear of misspecification. In this paper, we concentrate only on the latter. The following problem is close to what they define as their T^2 -operator only case.

¹⁵ To avoid confusion, this parameter is θ in Hansen and Sargent’s notation. Here θ refers to the exchange rate parameter in the aggregate demand equation.

The game represented by (8) is played simultaneously. The minimizing agent will therefore take a sequence $\{i_t\}$ of policy decisions as given and minimize the welfare function. The first order and sufficient condition for the artificial minimizing agent is

$$\left[-L_{\bar{p}_t|A} + \Theta \log\left(\frac{\bar{p}_t}{p_t}\right)\right] - \left[-L_{\bar{p}_t|B} + \Theta \log\left(\frac{1-\bar{p}_t}{1-p_t}\right)\right] + \Theta = 0$$

so regrouping and clearing out this equation yields

$$\left(\frac{\bar{p}_t}{1-\bar{p}_t}\right) = \left(\frac{p_t}{1-p_t}\right) \exp\left[\frac{L_{\bar{p}_t|A} - L_{\bar{p}_t|B}}{\Theta}\right] \quad (9)$$

This equality says that if the loss originated from Model A is bigger than the loss caused by Model B, the Central Banker will act according to $\bar{p}_t > p_t$. Note that for at least some value of the preference parameter, there is well defined mapping from p_t to \bar{p}_t . Once we obtain \bar{p}_t the problem for the Robust Central Bank is the same as in Section 4. This condition is also used in Cogley et al. (2008).¹⁶ The main distinction here is that the Phillips curve and output gap equations depend on the agent's expectations. Dealing with Robustness in contexts in which agent beliefs affect the law of motion of variables has not yet been dealt with in the literature. In the next subsections we describe an assumption that simplifies this complication and allows one to compute the value function for each model for a given prior.

4.1. Knowledge assumptions and optimal policy

We maintain the assumptions of Section 3. In addition, we assume that expectations of agents in Eqs. (2) and (3) are constructed in such a way that they are based upon the same prior model belief that assigns the same distorted probability \bar{p}_t that the Central Bank does. This assumption indirectly suggests that when agents derive pricing strategies that give rise to the Phillips curve, they do so in such a way that the distorted probabilities are identical to the Central Bank.

4.2. Distorted probabilities and policies

Robust policies will depend on a distorted probability calculated from Eq. (9). Once these probabilities are computed¹⁷ the optimal Robust policy replaces the standard prior by the distorted prior in the optimal policy given in Proposition 1. The state space has multiple dimensions. Therefore, in this section we describe the distortions of probabilities for a given financial shock as an example. Nevertheless, it is clear that other realizations of shocks will generate different distortions.

The value functions conditional on the true model depend on the current shock realization. Because for some priors, distinct realizations of shocks may induce a worse outcome from one model or the other, it will be the case that Robust policies may tilt the prior towards one model and tilt it towards another model for a different realization. Fig. 6 plots the distortion over the prior over Model A on the vertical axis for each value of the prior over Model A in the horizontal axis. The shock is 1 standard deviation devaluation of the exchange rate. All alternative shocks are set to 0 in this exercise, to stress the effects of the exchange rate. Each curve around the 45° lines shows distortions under different values of Θ .

The figure shows that when Model A is the true model, upon a devaluation, the Robust Central Banker will react as if its Bayesian prior were closer to Model B. This implies an interest policy that reacts with more strength to exchange rate devaluations. From the analysis carried out in the previous sections, we could interpret the points in which the distorted probabilities tilt the prior towards Model B as points in which learning will occur at a slower pace. Nevertheless, Robustness will work in the opposite direction for a shock of the opposite sign. Thus, it is clear that the effect of fear of misspecification will affect learning depending on the nature of the shock and its sign. The effect of Robust policies will be described in the following section when we present some simulations based on Bayesian and Robust policies.

The reason for this ambiguity is hinted at Hansen et al. (1999) and Adam (2004). In the latter paper, it is shown that for concerns about the nature of the shocks affecting the economy, for well defined Robust preferences, there is an equivalent Bayesian problem. In our framework, the nature of Robustness differs. It only operates over the prior and not over the perceived law of motion for the shocks. Nevertheless, the distorted prior, by shifting the weight from one model to the other, seems to operate in an analogous way to a change in the distribution of the shocks. However, Robustness concerns over the priors are a more restrictive concern than over the nature of the shocks. For the set of models they analyze, these authors find that the optimal Robust policy will have an equivalent optimal Bayesian problem if the risk-aversion parameter is perturbed. Because our model has the property that the model that yields the worst outcomes depends on the shock, this seems to suggest that Robustness will not tilt the prior consistently towards one model or the other. This then suggests that learning under Robust policies should not be dramatically different from the Bayesian counterpart. We explore this in the next section.

¹⁶ In particular, these preferences are consistent with Cogley et al. (2008) setting $\theta_1 = \infty$ and $\theta_2 = \Theta$.

¹⁷ See Appendix C for the numerical algorithm.




Fig. 6. Distorted prior from Robust policy and Bayesian prior. After a devaluation shock of 1 standard deviation.

5. Learning dynamics

5.1. Learning under Bayesian policies

For the calibration described in Table 1, Fig. 7 describes the speed of convergence of the priors for different initial values and for each model. Solid lines report the expected¹⁸ path for priors depending on the correct model: Model A (solid) and Model B (dashed). On the horizontal axis of the graph we have the time elapsed, where the scale is quarters. The vertical axis measures the value of the prior so that the curves represent the mean of the prior model evolution.

When the true model is A, we find that beginning from a prior of 10%, the expected waiting time for convergence to a prior probability of 90% is slightly above 5 years, roughly, the time interval of each of our cross-country sub-samples. The learning process is very similar, independent of the model. The dashed and solid lines are very close to each other.

The conclusion from this graph is that following a strictly optimal Bayesian policy, the Central Bank needs about 4–5 years to detect the true model when the economy is initiated with the wrong prior. Four to five years seem faster than what we observe in the data. If the textbook effect had been the true structure of the economy, we would have expected to see the systematic increase in volatility at the Panel (a) of Fig. 1, not in Panel (b). In the next section we ask if concerns for model misspecification can explain the pattern.

5.2. Learning under Robust policies

Fig. 8 describes the speed of convergence to the true model by replicating (7), when Robust policies are implemented. The figures are based on a value of $\Theta = 2$.¹⁹ The main lesson obtained from the graph is that learning is not modified substantially as expected. That is, the result is consistent with the conjecture that this is due to the fact that Robustness works as a change in risk aversion. If this is the case, because outcomes are worse in one model depending on the realization of shocks, Robustness would not tilt the prior systematically towards either model modifying the propensity to learn. However, a detailed inspection suggests that learning seems to occur slightly faster under Robust policies.

¹⁸ Expectations are approximated by 20 000 replications. The simulations are initiated with a shock to the exchange rate and the other shocks set to 0.

¹⁹ We could have chosen a tighter parameter but results did not change much for small perturbations. Moreover, concerns for Robustness cannot be taken much further. The problem becomes ill-posed for values of $\Theta = 0.8$.




Fig. 7. Prior evolution for Bayesian policy: solid lines show the evolution of priors towards Model A when Model A is the true model. Dashed lines show the evolution of priors towards Model B when Model B is the true model. Simulations are carried out considering Bayesian policies based on updated priors.

5.3. Learning under constant gain

We now study the sensitivity of the learning process when a constant gain algorithm is used instead of the ordinary Bayes rule. The implied assumption is that Central Banks discount data from further in the past. Sargent and Williams (2005) motivate this assumption when Central Banks are concerned about time-varying parameters. We take the following approach: the Central Bank discounts previous observations as if they had been generated by data generated by a similar process with a larger variance. In particular, the Central Bank computes the likelihood function by assuming that the standard deviation of the cost-push shock, at time $t - T$ was

$$\sigma_{x,t-T} = \sigma_x \kappa^{-T/2}$$

where $\kappa \leq 1$ is a parameter governing the constant gain. In addition, the initial prior at time t is also punished over time because the Central Bank associates it with information from further in the past. The Central Bank therefore assumes that its initial prior at time t takes the form of $P(M^A)^{1/\kappa^t}$.²⁰ The constant gain recursive Bayesian odds ratio is

$$\hat{p}_t = \frac{P(Z_t|M^A, Z_{t-1})\hat{p}_{t-1}^\kappa}{P(Z_t|M^A, Z_{t-1})(1-\hat{p}_{t-1})^\kappa + P(Z_t|M^B, Z_{t-1})\hat{p}_{t-1}^\kappa} \tag{10}$$

As opposed to Eq. (7), the Bayesian odds ratio (10) with constant gains is slightly modified.²¹ When updating posteriors, the prior is penalized by raising the priors to the κ -th power. Since priors are probabilities, raising $(1-\hat{p}_{t-1})$ and \hat{p}_{t-1} to the κ -th power, makes the weights closer to each other.

The consequences for learning are interesting in this context. The simulations show that the expected path for the prior will not converge to the true model. Rather, as it approaches 1, the constant gain algorithm places more weight on the wrong model than Bayes rule. At some point, this effect more than compensates the effect of the likelihood ratio. Therefore, posteriors no longer settle down.

²⁰ Note that this is not a probability since the priors assigned to each model do not sum up to one. One can nevertheless, divide both $P(M^A)^{1/\kappa^t}$ priors by their sum, to obtain an appropriate prior probability. This transformation does not alter the posteriors in any way.

²¹ The derivation of this odds ratio is carried out in the Appendix.




Fig. 8. Prior evolution for Robust policy: solid lines show the evolution of priors towards Model A when Model A is the true model. Dashed lines show the evolution of priors towards Model B when Model B is the true model. Simulations are carried out considering Robust policies based on updated priors.

Fig. 9 plots the expected path of the average simulation. In fact, posteriors converge to some probability which is independent of initial prior. However, as shown in the simulations, the speed of convergence of this process is very quick. Moreover, since the prior may wander around, only common shocks can explain the pattern in the data. The conclusion is valid for virtually all values of κ . Hence, introducing a constant gain does not explain the slow pattern shown in the data.²²

5.4. Determinants of the speed of convergence

A natural question is how large can concerns over the balance-sheet effect be, to prevent a Central Bank from learning? One is tempted to think that for a sufficiently low value of θ_B , the learning dynamics may be slowed down. Interestingly, whereas the pair (θ_A, θ_B) determines the HPS-statistic, it has a minor impact on the speed of convergence to the true model. The reason is that two effects countervail each other. On the one hand, a larger θ_B increases the Central Bank's concern over exchange rate depreciations. This leads to a lower volatility in the exchange rate. On the other hand, when θ_A and θ_B are further apart, both models are more easily distinguished. That is, the signal to noise ratio in the output gap equation is not very sensitive to changes in θ_A and θ_B .

To affect the signal to noise ratio significantly, the key parameter is the value of σ_y , the standard deviation of the innovations to the output equation (which determines the noise). This parameter is not free in our calibration. Moreover, to postpone learning, the value of σ_y should increase by 100%. Reaching the 95th percentile takes a double of the time. Nevertheless, as shown in Fig. 7, learning occurs faster at the beginning of the sample. Even though reaching the 95th percentile takes much longer, the simulations suggest an increase in the HPS-statistic, again faster than seen in the data. We could prevent this by increasing σ_y further, but we would lose the discipline of the calibration exercise.

So far, we have been unsuccessful in explaining the pattern in the data with learning. The next section explores an alternative explanation.

²² There is a difference between the model studied here and models of adaptive learning that follow from Sargent (1999). In particular, these papers share in common a Central Bank that ignores that agents are forward-looking. These papers have shown that by introducing constant gain algorithms, the economy may cycle from one equilibria to another. These equilibria are characterized by a Central Bank believing in a wrong model but in which the true structure reinforces such beliefs. In the context of this paper, constant gain algorithms do not have this property because the government knows the forward looking structure of the true model.




Fig. 9. Prior evolution for constant gain: solid lines show the evolution of priors towards Model A when Model A is the true model. Dashed lines show the evolution of priors towards Model B when Model B is the true model. Simulations are carried out considering Bayesian policies based on prior updates according to the constant gain Bayes rule. For this simulations $\kappa = 1.1$.

6. Endogenous balance sheets

The previous sections show that neither Bayesian optimal nor Robust policies can account for a learning pattern that fits the international trend towards more flexible exchange rates. Moreover, this conclusion is maintained under a constant gain algorithm and, more so, does not depend much on the parameters θ_A and θ_B . For this reason, a successful alternative hypothesis must involve some sort of learning trap that was active at least at the beginning of the sample. On the other hand, it must provide a mechanism that randomly allowed learning. This section presents an alternative hypothesis with these features.

The model presented in the previous sections is now modified to allow some endogeneity in the determination of the true model. Precisely, instead of assuming θ as given deterministically, we assume that it is the outcome of an endogenous process. This section follows a similar route as [Chang and Velasco \(2006\)](#) who study endogenous hedging decisions for the case in which the Central Bank chooses among pure fixed and floating regimes.

Here, we assume that private agents can belong to either of two types: active or passive. Active agents take hedging decisions based on a cost parameter and a common prior on the population of each type. Passive agents, on the other hand, are unable to hedge. The population measures are such that either active or passive agents are measure zero. Nevertheless, we also assume that both populations always coexist. Therefore, private agents know their own types but their own expectations over the average population are given by the common prior, and not by conditioning on their own type.²³ Hence, both the Central Bank and agents are unaware of the typical population type.²⁴

Active agents will hedge against exchange rate fluctuations depending on a cost \bar{c} . Therefore, if the representative agent is active, the economy behaves according to the balance-sheet model depending on its hedging decisions. The hedging decisions are determined by the active agents policy function:

$$\text{Hedge if and only if } \text{VAR}[\Delta s_t | Z_t, p_t, \sigma_{\varepsilon,t}] \geq \bar{c}$$

²³ Potentially these agents could condition their posterior distributions on their own types. Nevertheless we assume that this is not done in order to maintain the common prior—common posterior structure of the model. By imposing this structure, we are able to avoid, once more, 'forecasting the forecast of others.'

²⁴ Similar assumptions are made in models of Reputation by assuming the existence of *commitment types*. See [Mailath and Samuelson \(2006, Chapter 16\)](#).

$\sigma_{e,t}$ is the volatility of the financial shock, which is now assumed to be time varying. Whenever active agents hedge, θ takes the value θ^A .

Therefore, the parameter θ is determined endogenously whenever agents are active. We refer to a model in which the representative agent is active as model C (where C is for *contingent* balance-sheet model). M^C is the event in which the true model is Model C in which agents hedge. The relevant posterior to determine the Central Bank's policy is $p_t^* = \Pr[\theta_t = \theta^A | Z^t, \bar{c}]$. p_t^* is the probability that the Central Bank has that θ at period t is θ^A .

The law of motion for the parameter θ_t is

$$\theta_t = \theta^A \quad \text{if } \text{VAR}[\Delta s_t | Z_{t-1}, p_t^*, \sigma_{e,t}] \geq \bar{c} \text{ and } M^C$$

$$\theta_t = \theta^B \quad \text{otherwise}$$

Here the conditional variance of the exchange rate $\text{VAR}[\Delta s_t | Z_{t-1}, p_t^*, \sigma_{e,t}]$ is a function of the prior p_t^* through the government's policy instrument. The decision to hedge is carried out before the financial shock is realized, so the relevant variance term is the one conditional on Z_{t-1} . To compute p_t^* the Central Bank must assign a prior on the probability that the true model is model C. We label call probability $q_t \equiv \Pr(M^C | Z^t)$. Both p_t and q_t are jointly determined and given by the following formula:

$$p_t^* = q_t \quad \text{if } \text{VAR}[\Delta s_t | Z_{t-1}, p_t^*, \sigma_{e,t}] \geq \bar{c}$$

$$0 \quad \text{otherwise} \tag{11}$$

and

$$q_{t+1} = q_{t-1} \quad \text{if } \text{VAR}[\Delta s_t | Z_{t-1}, p_t^*, \sigma_{e,t}] < \bar{c}$$

$$= \frac{\Pr(Z_t | \theta^A, Z_{t-1}) q_t}{\Pr(Z_t | \theta^B, Z_{t-1})(1 - q_t) + \Pr(Z_t | \theta^A, Z_{t-1}) q_t} \quad \text{otherwise} \tag{12}$$

A Central Bank knows that if agents are active, they will hedge for a sufficiently high volatility. Therefore, following Bayes rule, it assigns the same probability to the event $\theta_A = \theta_t$ and to the event that agents are active. This explains Eq. (11). Eq. (12) also follows from Bayes Rule. If volatility is insufficient, the Central Bank knows that active agents would have acted as passive agents and therefore has no information to update its prior about the population type. On the other hand, if volatility is large enough, active agents will hedge, and the Central Bank has information to update its prior. This posterior is interesting because learning is activated only when there is a sufficiently high expected exchange rate volatility. The Central Bank induces this volatility by choosing its own policy i_t^{BAY} (p_t^*) which depends on its prior q_t via (11). Depending on the calibration of parameters, this posterior may place the Central Bank on a different trap as the one discussed in the previous sections. If the Central Bank has a strong prior about the balance-sheet model, active agents will behave just as passive agents. Knowing this, the Central Bank is always stuck with the same posterior.

We need to compute $\text{VAR}[\Delta s_t | Z_t, p_t^*, \sigma_{e,t}]$ to determine the thresholds $\bar{p}(\sigma_{e,t})$ that trigger the change in θ . Plugging in the closed form for i_t^{BAY} (p_t^*) into Eq. (3), and computing the conditional variance, allows us to pin down the formula for the conditional volatility:

$$\text{VAR}(\Delta s_{t+1} | Z_{t-1}, p_t^*) = \left[\frac{\chi}{p_t^* \theta^H + (1 - p_t^*) \theta^L - \chi} \right]^2 \sigma_i^2 + \left[\frac{p_t^* \theta^H + (1 - p_t^*) \theta^L}{p_t^* \theta^H + (1 - p_t^*) \theta^L - \chi} \right]^2 \sigma_{e,t}^2 \tag{13}$$

Thus, Eq. (13) shows that conditioned on the policy function, the volatility of the exchange rate is determined by the volatility of the international interest rate σ_i^2 and the volatility of financial shocks $\sigma_{e,t}^2$.

Fig. 10 plots the expected volatility for $\sigma_e = 0.08$ (dashed line) and $\sigma_e = 0.16$ (solid line). The figure shows how, the larger the prior on θ^A , the larger is the implied volatility of the exchange rate. The horizontal broken line is an example of threshold \bar{c} that determines the active agent's decision to hedge. When the financial shock is small, there is no learning for priors below 0.4. An increase in the volatility of the financial shock would enable learning for all priors, since the threshold volatility drops below the volatility implied for all priors. This figure provides an example where learning is possible for low priors only when the financial volatility is high.

The next section explores the possibility that changes in the volatility of the financial shock, only for a small period of time, provided a window of opportunity to learn. Without a shock like this, countries would either fall in a learning trap forever or learn quickly.

6.1. Learning under endogenous balance sheets

Here we study the learning process when private agents behave as described in the previous section. We perform an experiment and assume that the economy is populated by active agents. Again, we maintain the assumption that $\sigma_{e,t}$ is observable in spite of being time varying. We calibrate the model as in Section 3 but we allow a transitory twofold increase in the volatility at the middle of the sample for four periods. We have in mind episodes in which financial crises induce greater exchange rate volatility but there are multiple reasons that could trigger a learning process as we explain below.

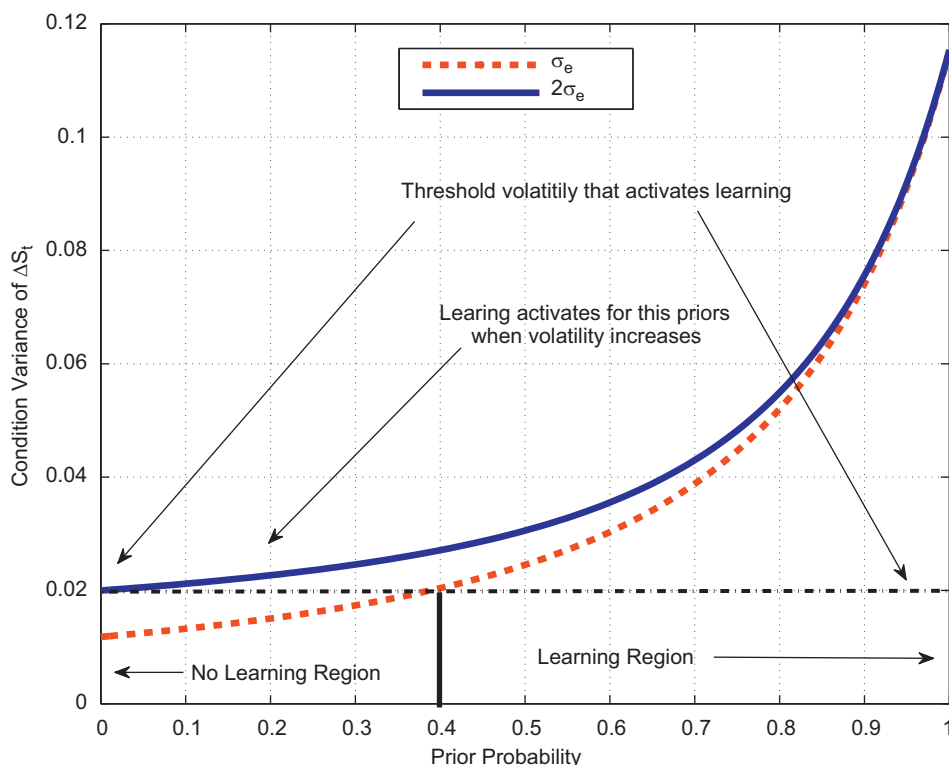


Fig. 10. Exchange rate volatility as a function of model prior: the figure plots the exchange rate volatility induced by Central Bank policies. Each curve corresponds to two levels of financial volatility. The horizontal (dash-dot) line represents the threshold volatility that activates learning. For low priors, an increase in financial volatility causes the implied exchange rate volatility to jump from the dashed curve to the solid curve activating the learning process.

With this intuition in mind, Fig. 11 shows the expected behavior (thick line) of the prior when it is started at 0.15 along with 200 different sample paths (thin lines). The threshold prior upon which learning is activated is set to 0.4, as in Fig. 10. Since the prior starts at a level below the threshold, no learning is possible until the increase in volatility occurs at the middle of the sample.²⁵ At this stage, depending on the shocks, Central Banks begin to learn. The volatility is then reset to its baseline level. Most of the simulated economies are lucky enough to learn their way out of the trap. These countries have the right combination of shocks during the volatile period. Another group continues to suffer from fear of floating. Their posteriors remain in the trap region. Some of these countries may be stuck forever, as no further increase would allow them to learn their way out.

Fig. 12 shows a replica of Panel C in Fig. 2. The replica simulates the HPS-statistic for 29 countries. The countries are subject to an increase in volatility for four periods at the middle of the sample. Fig. 12 replicates the pattern away from the 45° line. This shows that the trend towards more flexible exchange-rates could be explained by an exogenous process that affected the decision to hedge by active agents.

6.2. Discussion on endogenous balance sheets

In the previous section we presented an experiment in which an endogenous decision by private agents helps replicate the pattern observed in the data. The experiment shows that the endogeneity of the balance-sheet effect is key in explaining the data. It introduces a mechanism under which learning is impossible if Central Banks start with a prior that renders low exchange rate volatility. We showed that a temporal increase in the volatility of the exchange rate may trigger learning. An increase in the volatility of the exchange rate, without an endogenous balance-sheet mechanism, would have made learning faster in the model of Section 3.

Whether a country follows a path towards more flexible regimes depends on its initial priors, the cost of hedging and luck. The main insight of this section is that explaining the data requires a mechanism that prevented learning at the begging of the sample. A temporal disturbance in the decision to hedge would allow learning thereafter.

²⁵ The timing of the volatility episodes does not matter. That is, shocks could occur in separate periods and the results would not change in expectation. On the other hand, the magnitude of the increase in the volatility is important. If the increase is not sufficiently high, then the Central Bank remains stuck in its learning trap.

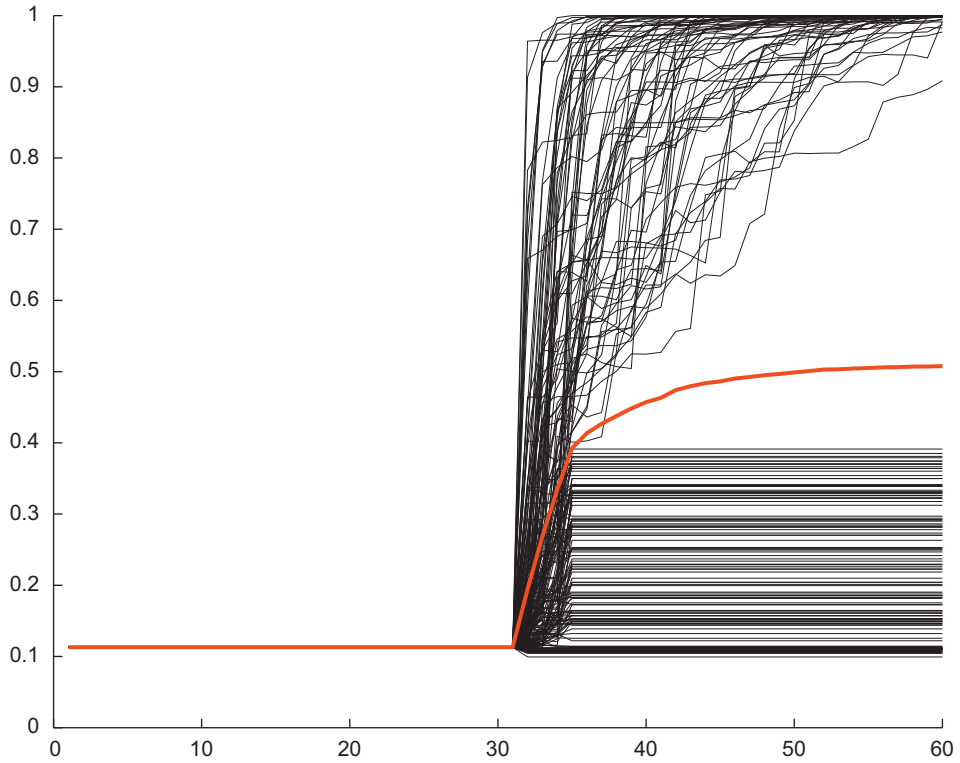


Fig. 11. Learning simulations under endogenous balance-sheet: the thick line shows the mean path for 200 simulations of the posteriors (thin lines). Financial volatility is increased twofold for four periods at the middle of the sample.




Fig. 12. Simulation of cross-country evidence: simulations of Fig. 1 corresponding to a twofold increase in financial volatility for four periods at the middle of the sample in the model with endogenous balance sheets.

Whereas the main lesson from this experiment is clear, there are multiple sources for a mechanism that triggered learning. We provided only one: financial volatility, that we associate with the financial crises of the late 1990s. Yet, any disturbance in the decision to hedge would do the job that more volatility does our example. Because learning seems to have been activated for many countries simultaneously, we think of these financial crises as the underlying triggering events. An increase in the volatility of the financial shock could follow from instability in capital markets or a systematic loss in reserves which disrupted the Central Bank's ability to smooth financial shocks. On the other hand, the literature has associated the severeness of the financial crises to the balance-sheet effect. Hence, we should take our experiment with caution and interpret it as a temporal change the private sector's trade-offs.

Many other explanations associated with the crises episodes may have played a similar role. If foreign lending became more sensitive to the balance-sheet structure, the opportunity cost of remaining exposed to the exchange rate risk may have increased. This effect would lower the value of \bar{c} , that triggered hedging decisions in our model. This would have the same effects as in our exercise.

Alternatively, crises may have led to an exogenous structural change due to more regulation both by international lenders or domestic borrowers. On the other hand, many hedging instruments not available in the past may have developed post-crises. In the context of our model, these explanations would reflect an exogenous parameter change which could also explain the data.

We have abstracted from the possibility that countries learned from one another. This could explain the co-movement in the data without common shocks. Nevertheless, we would need to describe a mechanism that prevented countries from learning at the beginning of the sample as done here. Without it, learning from each other would made countries learn faster. Uncovering which effects explains the data best is left for further research.

7. Concluding remarks

This paper was motivated by a conversation between the Chairman of a Central Bank in a developing country and a member of his modeling unit:

- *Econometrician*: 'If you don't let it varyate you'll never learn what its impact is!'
- *Central Banker*: 'I don't want to learn!'

Indeed, this dialogue may have often occurred in various countries. Cross-country data, on the other hand, suggest that floating was less fearsome for many countries. We tested the hypothesis that by believing in the balance-sheet model, countries intervened substantially in their exchange rate markets and by doing this, they lost the ability to learn about the balance-sheet effect.

We studied this hypothesis under Bayesian and Robust optimal policies. Under this structure, our calibration suggests that the econometrician was wrong. Regardless of the fact that the exchange rate variations relative to output variations were low, a Bayesian or Robust Central Bank would have discovered the true model in a shorter time than in the data.

We then tested a model in which the relevance of the balance-sheet effect depends on the private agents' decision to hedge. We modeled private agents who choose to hedge whenever they anticipate sufficient exchange rate volatility. This volatility depends on the Central Bank's prior on the private sector's ability to hedge. Together with an increase in the financial volatility, this version was capable of replicating the data.

Behind the scenes, this finding suggests that the international currency crises of the late 1990s provided a window of opportunity to learn. In this case, the econometrician was right, more volatility would have enabled the Central Bank's to learn.

We conclude this paper asking what characteristic of the currency crises could have triggered learning? Whether the phenomenon responds to a change in the private sector's decision, or whether it is due to a change in policy, remains an open empirical question. Yet, we are more confident than not that initial wrong beliefs on their own cannot account for an international trend towards more flexible exchange rates.

Acknowledgments

I wish to thank David Backus, Francisco Barillas, Ryan Booth, Roberto Chang, Tim Cogley, Jonathan Eaton, Mark Gertler, Boyan Jovanovic, Virgiliu Midrigan, Carlos Montoro, Taisuke Nakata, Kris Nimark, Albert Queralto, Thomas Sargent, Chris Tonetti, Vicente Tuesta, Marco Vega, Vivian Yue and Eduardo Zilberman for insightful comments. I wish to especially thank James Bullard and an anonymous referee for comments that substantially shaped this paper. The usual disclaimer applies.

Appendix A. Optimal discretionary policy under certainty

The behavior of private agents described by Eqs. (1)–(3) can be reduced to a simple two-equation system by direct substitution:

$$\pi_t = \beta E_t[\pi_{t+1}] + \gamma y_t + \varepsilon_{\pi,t} \quad (14)$$

$$y_t = E_t[y_{t+1}] - (\chi - \theta)i_t + \chi E_t[\pi_{t+1}] + \mu_{y,t} \quad (15)$$

where we have replaced (3) into (2), and set $\mu_{y,t} = \chi r_t^n - \theta i_t^* - \theta \varepsilon_{et} + \varepsilon_{y,t}$. We use the Phillips curve to solve the minimization problem and the IS equation to infer the policy rate that implements the solution:

$$\varrho_\tau = E_\tau \left[\sum_{t=\tau}^{\infty} \beta^{t-\tau} \{ \pi_t^2 + w y_t^2 + \lambda_t (\beta \pi_{t+1} + \gamma y_t + \varepsilon_{\pi,t} - \pi_t) \} \right]$$

The solution under discretion implies that the following first-order conditions with respect to $E_t[\pi_t]$ and $E_t[y_t]$ hold every period²⁶:

$$2E_t[\pi_t] = E_t[\lambda_t]$$

and

$$2wE_t[y_t] + \gamma E_t[\lambda_t] = 0$$

These two variables are controlled using the relation:

$$-\frac{\gamma}{w} E_t[\pi_t] = E_t[y_t] \quad (16)$$

In order to implement this rule through the policy instrument i_t , we replace the above condition in the Phillips curve to obtain

$$\left(1 + \frac{\gamma^2}{w} \right) E_t[\pi_t] = \beta E_\tau[\pi_{t+1}] + \rho_\pi \varepsilon_{\pi,t-1} \quad (17)$$

This equation can be solved by guessing a solution for the expectations operator $E_t[\pi_t] = A \varepsilon_{\pi,t-1}$, where A is a parameter to be determined. This guess, in turn, implies $E_t[\pi_{t+1}] = A \rho_\pi \varepsilon_{\pi,t-1}$. Substituting the guess back into the Phillips curve allows us to obtain the solution to the unknown parameter A :

$$A = \frac{\rho_\pi}{1 + \frac{\gamma^2}{w} - \beta \rho_\pi} \quad (18)$$

Given expectations, the interest rate can be backed out through Eq. (15):

$$i_t = \frac{1}{(\chi - \theta)} \{ E_t[y_{t+1}] - E_t[y_t] + \chi E_t[\pi_{t+1}] + E_t[\mu_{y,t}] \}$$

and replacing the initial guess into this functional leads to

$$i_t = \frac{1}{(\chi - \theta t)} \left\{ \left(\frac{\gamma}{w} \left(\frac{1}{\rho_\pi} - 1 \right) + \chi \right) A \rho_\pi \varepsilon_{\pi,t-1} + E_t[\mu_{y,t}] \right\} \quad (19)$$

which completes the proof of the optimal policy in the perfect certainty case.

For the special case of $\rho_\pi \rightarrow 0$ we have that

$$i_t \rightarrow \frac{1}{(\chi - \theta)} E_t[\mu_{y,t}]$$

which is a simple rule that attains zero inflation by reacting to the observed shock to the output gap. Given the timing protocol, the solution to the expectation of exogenous shocks is

$$E_t[\mu_{y,t}] = \chi(\rho_r \varepsilon_{r,t-1}) - \theta \varepsilon_{i,t} - \theta \varepsilon_{e,t} + \rho_{y,t} \varepsilon_{y,t-1}$$

where the expectation takes that form given the assumption that ε_{et} and $\varepsilon_{i,t}$ are observable. We then substitute this result Eq. (15) to obtain

$$y_t = E_t[y_{t+1}] - \left(\frac{\gamma}{w} \left(\frac{1}{\rho_\pi} - 1 \right) + \chi \right) A \rho_\pi \varepsilon_{\pi,t-1} + \chi E_t[\pi_{t+1}] + \mu_{y,t} - E_t[\mu_{y,t}] \quad (20)$$

with the initial guess that this equation further simplifies to

$$y_t = -\frac{\gamma}{w} A \varepsilon_{\pi,t-1} + \mu_{y,t} - E_t[\mu_{y,t}]$$

This confirms our guess for the law of motion y_t 's expectation.

In terms of the primitives we obtain

$$y_t = -\frac{\gamma}{w} A \varepsilon_{\pi,t-1} + \chi v_{r,t} + v_{y,t} \quad (21)$$

²⁶ Implicitly we are assuming that the Central Bank controls expectations directly.

To determine the law of motion of inflation from Eq. (14):

$$\pi_t = \beta A \rho_\pi \varepsilon_{\pi,t-1} + \gamma \left(-\frac{\gamma}{W} A \varepsilon_{\pi,t-1} + \chi v_{r,t} + v_{y,t} \right) + \varepsilon_{\pi,t} = A \left(\beta \rho_\pi - \frac{\gamma^2}{W} \right) \varepsilon_{\pi,t-1} + \gamma \chi v_{r,t} + \gamma v_{y,t} + \rho_p \varepsilon_{\pi,t-1} + v_{\pi,t}$$

so synthesizing

$$\pi_t = \left((A\beta + 1) \rho_\pi - \frac{\gamma^2}{W} A \right) \varepsilon_{\pi,t-1} + \gamma \chi v_{r,t} + \gamma v_{y,t} + v_{\pi,t} \tag{22}$$

The value of the exchange rate is obtained directly from the UIP equation (3)

$$E_t[\Delta s_{t+1}] = i_t - i_t^* - \varepsilon_{et} \tag{23}$$

A.1. Steady state of the economy

In the absence of shocks we obtain from Eq. (19) the steady state value of the interest rate, \bar{i} , in terms of the steady state of the stochastic process, $\bar{\mu}_{y,t}$, and inflation, $\bar{\pi}_t$. The steady state for the stochastic process is

$$\bar{\mu}_y = \chi \bar{r} + \theta \bar{i}^*$$

so re-writing this

$$\bar{i} = \frac{1}{(\chi - \theta)} [\bar{\mu}_y]$$

From the steady state version of the Phillips curve and the gap condition we obtain

$$0 = -(\chi - \theta) \bar{i} + \chi \bar{\pi} + \bar{\mu}_y$$

which in turn implies $\bar{\pi} = 0$.

The non-commitment policy yields the desired result of $\bar{y} = 0$. Thus, in the absence of shocks, the policy induces the lowest cost possible.

A.2. State-space representation

A.2.1. Exogenous state block

The exogenous state vector-equation for this model is

$$S_t = AS_{t-1} + Bw_t$$

where it is explicitly expressed in the matrix notation matching our parameters according to

$$\begin{bmatrix} \varepsilon_{\pi,t} \\ \varepsilon_{\pi,t-1} \\ \varepsilon_{y,t} \\ \varepsilon_{y,t-1} \\ r_t^n \\ r_{t-1}^n \\ i_t^* \\ \varepsilon_{e,t} \\ 1 \end{bmatrix} = \begin{bmatrix} \rho_\pi & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \rho_y & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \rho_r & 0 & 0 & 0 & r^n \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \rho_i & 0 & r^n \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & \rho_e & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \varepsilon_{\pi,t-1} \\ \varepsilon_{\pi,t-2} \\ \varepsilon_{y,t-1} \\ \varepsilon_{y,t-2} \\ r_{t-1}^n \\ r_{t-2}^n \\ i_{t-1}^* \\ \varepsilon_{e,t-1} \\ 1 \end{bmatrix} + \begin{bmatrix} \sigma_\pi & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & \sigma_y & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \sigma_r & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \sigma_i & 0 \\ 0 & 0 & 0 & 0 & \sigma_e \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \begin{matrix} v_{\pi,t} \\ v_{y,t} \\ v_{r,t} \\ v_{i,t+1} \\ v_{s,t+1} \end{matrix}$$

Recall that by expositional motives we adopt the convention that $i_t^* = \varepsilon_{i,t}$.

A.2.2. Observable state equation

The observable vector-equation for this model is

$$Z_t = CS_t + Du_t$$

We use the following parameters as auxiliary variables to find a parsimonious representation.

Policy rule: For the policy reaction we can define the following set of auxiliary parameters:

$$F_1 = \frac{1}{(\chi - \theta)} \left(\frac{\gamma}{W} \left(\frac{1}{\rho_\pi} - 1 \right) + \chi \right) A \rho_\pi$$

$$F_2 = \frac{1}{(\chi - \theta)} \rho_{y,t}$$

$$F_3 = \frac{1}{(\chi - \theta)} \chi \rho_r$$

$$F_4 = -\frac{1}{(\chi - \theta)} \theta$$

$$F_5 = -\frac{1}{(\chi - \theta)} \theta$$

and we have that

$$i_t = F_1 \varepsilon_{\pi,t-1} + F_2 \varepsilon_{y,t-1} + F_3 \varepsilon_{r,t-1} + F_4 \varepsilon_{i,t} + F_5 \varepsilon_{e,t}$$

Phillips curve: Redefining (22) we obtain that

$$\pi_t = \left((\lambda\beta + 1) \rho_\pi - \frac{\gamma^2}{w} \lambda \right) \varepsilon_{\pi,t-1} + \gamma \chi (r_t^n - \rho_r r_{t-1}^n) + \gamma (\varepsilon_{y,t} - \rho_y \varepsilon_{y,t-1}) + (\varepsilon_{\pi,t} - \rho_\pi \varepsilon_{\pi,t-1})$$

and using the auxiliary variables:

$$G_1 = 1$$

$$G_2 = \lambda \left(\beta - \frac{\gamma^2}{w} \right)$$

$$G_3 = \gamma$$

$$G_4 = -\gamma \rho_y$$

$$G_5 = \gamma \chi$$

$$G_6 = -\gamma \chi \rho_r$$

Output gap: Recall that we can summarize the output gap by Eq. (21) so regrouping it we obtain

$$y_t = -\frac{\gamma}{w} \lambda \varepsilon_{\pi,t-1} + \chi (r_t^n - \rho_r r_{t-1}^n) + \varepsilon_{y,t} - \rho_y \varepsilon_{y,t-1}$$

so we define the following auxiliary variables:

$$H_1 = -\frac{\gamma}{w} \lambda$$

$$H_2 = 1$$

$$H_3 = -\rho_y$$

$$H_4 = \chi$$

$$H_5 = -\chi \rho_r$$

Nominal depreciation and final matrix form: We are looking for a system of the form:

$$Z_t = CS_t$$

To obtain this matrix equation we can write this by using Eq. (23) and noting that it can be written as a function of the exogenous processes that affect the policy instrument and the observable shocks that affect these equation directly. Summing up, the matrix form for the set of observables should look like

$$\begin{bmatrix} i_t \\ \pi_t \\ y_t \\ E[\Delta s_{t+1}] \end{bmatrix} = \begin{bmatrix} 0 & F_1 & 0 & F_2 & 0 & F_3 & F_4 & F_5 & 0 \\ G_1 & G_2 & G_3 & G_4 & G_5 & G_6 & 0 & 0 & 0 \\ 0 & H_1 & H_2 & H_3 & H_4 & H_5 & 0 & 0 & 0 \\ 0 & F_1 & 0 & F_2 & 0 & F_3 & F_4 - 1 & F_5 - 1 & 0 \end{bmatrix} \begin{bmatrix} \varepsilon_{\pi,t} \\ \varepsilon_{\pi,t-1} \\ \varepsilon_{y,t} \\ \varepsilon_{y,t-1} \\ r_t^n \\ r_{t-1}^n \\ i_t^* \\ \varepsilon_{e,t} \\ 1 \end{bmatrix}$$

Appendix B. Solution to the passive learning policy without commitment under uncertainty

This appendix describes the solution to the Central Bank’s problem when it behaves as a Bayesian optimizer.

B.1. Setting up the Lagrangian

We solve the Central Bank problem without commitment under model uncertainty following similar procedures as we use for the case with model certainty. The main difference corresponds to the way in which agents and the Central Bank form their expectations.

Hence, the objective is to solve for the following, prior-conditional Lagrangian:

$$\mathcal{L}_\tau(p_\tau) = L_\tau(p_\tau) + \sum_{t=\tau}^{\infty} \lambda_t \beta^t E_t[(\pi_{t+1} + \gamma_y y_t + \varepsilon_\pi - \pi_t) | p_t]$$

so that the first order conditions, with respect to $\{\pi_t\}_{t=\tau}^{\infty}$ and $\{y_t\}_{t=\tau}^{\infty}$, are the same as in the case with model certainty, except for the fact that expectations take the prior into consideration. Recall from Section 3 that agent expectations and the Central Bank’s expectations on inflation regarding the prior are the same by assumption. This assumption allows us to regroup the expectation operator of the Central Bank as well as the agents’ expectations present in the Phillips curve. The Bayesian counterpart of Eq. (16) is

$$-\frac{\gamma}{w} E_t[\pi_t | p_t] = E_t[y_t | p_t] \tag{24}$$

We can apply the expectations operator conditional on the prior to the Phillips curve and obtain a counterpart to (17):

$$\left(1 + \frac{\gamma}{w}\right) E_t(\pi_t | p_t) = \beta E_t(\pi_{t+1} | p_t) + \rho_\pi \varepsilon_{\pi,t-1}$$

The main complication here is how to deal with the conditional expectations $E_t(\pi_{t+1} | p_t)$. Again we take the guess of a linear functional in terms of the previous shocks. We use a linear combination of both these shocks to make a guess:

$$E_\tau(\pi_t | p_t) = A^{BAY} \rho_\pi \varepsilon_{\pi,t-1}$$

where A^{BAY} is the parameter yet to be determined.

The solution to A^{BAY} is the same as for Eq. (18):

$$A^{BAY} = \frac{\rho_\pi}{1 + \frac{\gamma}{w} - \beta \rho_\pi} \tag{25}$$

this is the same result as the one obtained in the case under certainty. Thus, the intuition behind this result is that the Central Bank’s goal is the same, both under certainty and uncertainty and, given what it knows, it will try to maintain inflation as if it knew the true model. Consequently, regardless of the definition of expectations, the Central Bank will try to set these expectations to the same number. The difference will be in the action it takes.

The law of motion for $\{E_\tau(\pi_t | p_t)\}_{t=\tau}^{\infty}$ will be an analog to the perfect certainty case and thus

$$E_\tau[\pi_t | p_t] = A^{BAY} \rho_\pi^{t+1-\tau} \varepsilon_{\pi,\tau-1}$$

Having sorted out what $\{E_\tau(\pi_t)\}_{t=\tau}^{\infty}$ is, we may obtain the corresponding expected path for $\{E_\tau(y_t)\}_{t=\tau}^{\infty}$ with the use of Eq. (15):

$$E_\tau[y_t] = -\frac{\gamma}{w} E_\tau[\pi_t]$$

We obtain the result for the optimal no-commitment policy by taking the Bayesian expectation in the aggregate demand equation (15) and clearing out the interest rate:

$$i_\tau^{BAY} = \frac{1}{\Psi(p_t)} [E_\tau(y_{t+1} - y_t | p_t) + \chi E_\tau(\pi_{t+1} | p_t) + E_\tau(\mu_{y,\tau} | p_t)]$$

where $\Psi(p_t) = \chi - (p_t \theta^H + (1 - p_t) \theta^L)$.

Replacing the previous results we obtain a similar result as before:

$$i_\tau^{BAY} = \frac{1}{\Psi(p_t)} \left[\left(\chi + \frac{\gamma}{w} \left(\frac{1}{\rho_\pi} - 1 \right) \right) A \rho_\pi \varepsilon_{\pi,\tau-1} + E_\tau(\mu_{y,\tau} | p_t) \right] \tag{26}$$

For the special case of $\rho_\pi = 0$ we have that

$$i_\tau^{BAY} = \frac{1}{\Psi(p_t)} E_t[\mu_{y,t} | p_t]$$

which is a simple rule that reacts only to the shocks that affect the output gap. Given the data set available at τ , $E_\tau(\mu_{y,\tau} | p_t)$ may be easily computed by taking the weighted average of the shocks that affect the model. Notice that $\varepsilon_{y,t-1}$ is not identified and depends on the model as it is obtained as a residual to the output gap equation, that, in the end, depends on the parameter over which the uncertainty is about.

B.2. State space representation

B.2.1. Exogenous state block

The exogenous state vector-equation for this model is

$$S_t = A(p_t)S_{t-1} + Bw_t \tag{27}$$

where it is explicitly expressed in matrix notation matching our parameters according to

$$\begin{bmatrix} \varepsilon_{\pi,t} \\ \varepsilon_{\pi,t-1} \\ \varepsilon_{y,t} \\ \varepsilon_{y,t-1} \\ r_t^n \\ r_{t-1}^n \\ \tilde{i}_t^* \\ \varepsilon_{e,t} \\ 1 \\ \tilde{\varepsilon}_{y,t} \\ \tilde{\varepsilon}_{y,t-1} \end{bmatrix} = \begin{bmatrix} \rho_\pi & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \rho_y & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \rho_r & 0 & 0 & 0 & r^n & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \rho_i & 0 & r^n & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & \rho_e & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ J_1 & 0 & J_2 & 0 & J_3 & 0 & J_4 & J_5 & 0 & J_6 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \varepsilon_{\pi,t-1} \\ \varepsilon_{\pi,t-2} \\ \varepsilon_{y,t-1} \\ \varepsilon_{y,t-2} \\ r_{t-1}^n \\ r_{t-2}^n \\ \tilde{i}_{t-1}^* \\ \varepsilon_{e,t-1} \\ 1 \\ \tilde{\varepsilon}_{y,t-1} \\ \tilde{\varepsilon}_{y,t-2} \end{bmatrix} + \begin{bmatrix} \sigma_\pi & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & \sigma_y & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \sigma_r & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \sigma_i & 0 \\ 0 & 0 & 0 & 0 & \sigma_e \\ 0 & 0 & 0 & 0 & 0 \\ J_7 & 0 & 0 & J_8 & J_9 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \begin{matrix} v_{\pi,t} \\ v_{y,t} \\ v_{r,t} \\ v_{i,t} \\ v_{s,t} \end{matrix}$$

As opposed to the system that we defined in the previous appendix, we have introduced two extra columns. The columns refer to the residual of the false model that will be an endogenous outcome of the model. Note that we have not yet solved for the unknowns $\{J_s\}_{s=1}^9$. To do so, we will use a result from the next section.

B.2.2. Endogenous state block

Take the policy reaction function described in Eq. (26) and our guess for the formation of expectations, and replace these results correspondingly in Eq. (15) for each model. Assume that θ^T is the true value and θ^F the false value. We will obtain that the output gap will be

$$y_t = \left(\chi - \frac{\gamma}{w}\right) A^{BAY} \rho_\pi \varepsilon_{\pi,t-1} - (\chi - \theta^T) \tilde{i}_t + \mu_{y,t}$$

and the according outcome of the misspecified model would be

$$\tilde{y}_t = \left(\chi - \frac{\gamma}{w}\right) A^{BAY} \rho_\pi \varepsilon_{\pi,t-1} - (\chi - \theta^F) \tilde{i}_t + \tilde{\mu}_{y,t}$$

The $\tilde{\cdot}$ terms in the variables of this equation are the outcomes under the false model. For the same observation y_t , each model implies a different innovation. To obtain the false model's output gap residual as a function of the innovations to the true model we subtract both equations and equate the left-hand side to 0.

$$\mu_{y,t} + (\theta^T - \theta^F) \tilde{i}_t - \tilde{\mu}_{y,t} = 0$$

We now use the definition of both terms and find that

$$\varepsilon_{y,t} + (\theta^T - \theta^F) \tilde{i}_t = \tilde{\varepsilon}_{y,t}$$

which in turn implies that

$$\rho \varepsilon_{y,t-1} + v_{y,t} + (\theta^T - \theta^F) \tilde{i}_t = \rho_y \tilde{\varepsilon}_{y,t-1} + \tilde{v}_{y,t} \tag{28}$$

Recall that $\varepsilon_{y,t-1}$ and $\tilde{\varepsilon}_{y,t-1}$ are state variables for the Central Bank. Then, $\tilde{v}_{y,t}$ is an endogenous outcome of the model. To obtain an explicit solution to the false innovation we use a linear form of the Central Bank's Policy as a function of the innovations to the system.

Interest rate: The solution to the optimal policy equation (26) is

$$\tilde{i}_t^{BAY} = \frac{1}{\Psi(p_t)} \left[\left(\chi + \frac{\gamma_y}{w} \left(\frac{1}{\rho_\pi} - 1 \right) \right) A \rho_\pi \varepsilon_{\pi,t-1} + E_t(\mu_{y,t} | p_t) \right]$$

For the policy reaction we can define the following set of auxiliary variables:

$$F_1(p_t) = \frac{1}{\Psi(p_t)} \left(\frac{\gamma_y}{w} \left(\frac{1}{\rho_\pi} - 1 \right) - \chi \right) A \rho_\pi$$

$$F_2(p_t) = \frac{1}{\Psi(p_t)} (p_t) \rho_{y,t}$$

$$F_3(p_t) = \frac{1}{\Psi(p_t)} \chi \rho_r$$

$$F_4(p_t) = -\frac{1}{\Psi(p_t)}(p_t\theta^T + (1-p_t)\theta^F)$$

$$F_5(p_t) = -\frac{1}{\Psi(p_t)}(p_t\theta^T + (1-p_t)\theta^F)$$

$$F_6(p_t) = \frac{1}{\Psi(p_t)}(1-p_t)\rho_{y,t}$$

and we have that

$$\dot{i}_t = F_1\varepsilon_{\pi,t-1} + F_2\varepsilon_{y,t-1} + F_3\varepsilon_{r,t-1} + F_4\varepsilon_{i,t} + F_5\varepsilon_{e,t} + F_6\tilde{\varepsilon}_{y,t-1}$$

Replacing this result in Eq. (28) allows us to compute the value of the $\{J_s\}_{s=1}^5$ since we can regroup these terms to obtain the value of $\tilde{v}_{y,t}$ and simply add to the corresponding row of the evolution of the exogenous matrices. We obtain the following result:

$$\begin{aligned} \tilde{v}_{y,t} = & \rho\varepsilon_{y,t-1} + v_{y,t} + (\theta^T - \theta^F)\dot{i}_t - \rho_y\tilde{\varepsilon}_{y,t-1} = (\theta^T - \theta^F)F_1\varepsilon_{\pi,t-1} + ((\theta^T - \theta^F)F_2 + \rho)\varepsilon_{y,t-1} + (\theta^T - \theta^F)F_3\varepsilon_{r,t-1} + (\theta^T - \theta^F)F_4\rho_i\varepsilon_{i,t-1} \\ & + (\theta^T - \theta^F)F_5\rho_e\varepsilon_{e,t-1} + ((\theta^T - \theta^F)F_6 - \rho_y)\tilde{\varepsilon}_{y,t-1} + v_{y,t} + (\theta^T - \theta^F)F_4v_{i,t} + (\theta^T - \theta^F)F_5v_{e,t} \end{aligned}$$

Therefore we summarize this equation by

$$\tilde{\varepsilon}_{y,t} = J_1\varepsilon_{\pi,t-1} + J_2\varepsilon_{y,t-1} + J_3\varepsilon_{r,t-1} + J_4\varepsilon_{i,t-1} + J_5\varepsilon_{e,t-1} + J_6\tilde{\varepsilon}_{y,t-1} + J_7v_{y,t} + J_8v_{i,t} + J_8v_{e,t}$$

and the solution to these parameters is

$$J_1 = (\theta^T - \theta^F)F_1$$

$$J_2 = ((\theta^T - \theta^F)F_2 + \rho_y)$$

$$J_3 = (\theta^T - \theta^F)F_3$$

$$J_4 = (\theta^T - \theta^F)F_4\rho_i$$

$$J_5 = (\theta^T - \theta^F)F_5\rho_e$$

$$J_6 = (\theta^T - \theta^F)F_6$$

$$J_7 = 1$$

$$J_8 = (\theta^T - \theta^F)F_4$$

$$J_9 = (\theta^T - \theta^F)F_5$$

Note that because the terms $\varepsilon_{r,t-1}$ and $\varepsilon_{i,t-1}$ do not have a zero mean, *a priori* we can expect the wrong model to have a bias (unless $\theta^T = \theta^F$ are trivially equal) whichever that is.

Inflation equation: We can use Eq. (22) to proceed in the same manner as in Appendix A so the set $\{G_s\}_{s=\pi,y,r,i,e}$ remains the same.

Output gap equation: The output gap's matrix form is obtained from Eq. (21) so we have

$$y_t = \left(\chi - \frac{\gamma}{w}\right)A^{BAY}\rho_\pi\varepsilon_{\pi,t-1} - (\chi - \theta)\dot{i}_t + \chi r_t^n - \theta i_t^* - \theta\varepsilon_{et} + \varepsilon_{y,t}$$

Therefore, use the following auxiliary variables:

$$H_1(p_t) = \left(\chi - \frac{\gamma}{w}\right)A^{BAY}\rho_\pi - (\chi - \theta^T)F_1(p_t)$$

$$H_2(p_t) = 1$$

$$H_3(p_t) = -(\chi - \theta^T)F_2(p_t)$$

$$H_4(p_t) = \chi$$

$$H_5(p_t) = -(\chi - \theta^T)F_3$$

$$H_6(p_t) = -\theta^T - (\chi - \theta^T)F_4$$

$$H_7(p_t) = -\theta^T - (\chi - \theta^T)F_5$$

$$H_8(p_t) = -(\chi - \theta^T)F_6$$

Nominal depreciation and matrix form: Nominal depreciations equation remains the same. Therefore, the system that summarizes the economy takes the form

$$Z_t = C(p_t)S_t \tag{29}$$

where the difference in relation to the result in Appendix A is the appearance of parameter p_t .

Summing up, we can rewrite it as

$$\begin{bmatrix} \dot{i}_t \\ \pi_t \\ y_t \\ E[\Delta s_{t+1}] \end{bmatrix} = \begin{bmatrix} 0 & F_1 & 0 & F_2 & 0 & F_3 & F_4 & F_5 & 0 & 0 & F_6 \\ G_1 & G_2 & G_3 & G_4 & G_5 & G_6 & 0 & 0 & 0 & 0 & 0 \\ 0 & H_1 & H_2 & H_3 & H_4 & H_5 & H_6 & H_7 & 0 & 0 & H_8 \\ 0 & F_1 & 0 & F_2 & 0 & F_3 & F_4 - 1 & F_5 - 1 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \varepsilon_{\pi,t} \\ \varepsilon_{\pi,t-1} \\ \varepsilon_{y,t} \\ \varepsilon_{y,t-1} \\ r_t^n \\ r_{t-1}^n \\ \dot{i}_t^* \\ \varepsilon_{e,t} \\ 1 \\ \tilde{\varepsilon}_{y,t} \\ \tilde{\varepsilon}_{y,t-1} \end{bmatrix}$$

B.3. Conditional moments

Unconditional moments may be obtained via numerical simulations of the system described above. On the contrary, we can use the matrix forms in the above section here to obtain an analytical expression though not in close-form. The system is summarized by two blocks of linear equations and a single non-linear equation computed by

$$S_t = A(p_t)S_{t-1} + Bw_t$$

$$Z_t = C(p_t)S_t$$

and Bayes's updating rule

$$p_{t+1} = \frac{p_t P(Data_t | Data_{t-1}, M^A)}{p_t P(Data_t | Data_{t-1}, M^A) + (1 - p_t) P(Data_t | Data_{t-1}, M^B)}$$

Cogley et al. (2007) show that p_{t+1} is a martingale. Conditional on the true model, it can be shown that p_{t+1} is either a super-martingale or sub-martingale depending on which is the true model. With this, one can also show that the model will converge to the true model.

B.3.1. Moments conditional on the true model and state

First moments: Conditional expectations over the exogenous states are

$$E_t[S_{t+1} | S_t, p_{t+1}] = A(p_{t+1})S_t$$

Substituting this into the observable states we have

$$E_t[Z_{t+1} | S_t, p_{t+1}] = C(p_{t+1})[A(p_{t+1})S_t]$$

where for simplicity we defined

$$M(p_{t+1}) = C(p_{t+1})[A(p_{t+1})]$$

Second moments: Recall the following equivalence:

$$S_{t+1} = A(p_{t+1})S_t + Bw_{t+1}$$

and therefore

$$Z_{t+1} = C(p_{t+1})[A(p_{t+1})S_t + Bw_{t+1}]$$

The conditional variance-covariance matrix for the observable states is then

$$VAR[Z_{t+1}] = E[[C(p_{t+1})Bw_{t+1}][C(p_{t+1})Bw_{t+1}]']$$

and is a function of the innovations only. To compute a particular element one can simply use a selector matrix.

Instantaneous loss: The loss at period $t+1$ is

$$Z'_{t+1} W Z_{t+1}$$

where

$$W = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & w & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

We can substitute in the components of Z_{t+1} to obtain

$$[A(p_{t+1})S_t + Bw_{t+1}]'C(p_{t+1})'WC(p_{t+1})[A(p_{t+1})S_t + Bw_{t+1}]$$

Taking expectations over the above equation and defining it as a value V we obtain

$$V(S_t, p_{t+1}) = S_t' M(p_{t+1})' W M(p_{t+1}) S_t + E[w'_{t+1} B' C(p_{t+1})' W C(p_{t+1}) B w'_{t+1}]$$

where we used the fact that w_{t+1} are i.i.d. shocks to eliminate the cross terms. Note that the p_{t+2} here will depend on the outcome of w'_{t+1} and the true model.

B.3.2. Value conditional on true model

Here we use p_{t+1} to refer to the probability of the event that A is the true model. The value of the Central Bank's problem conditional on the event that the true model is A is given by

$$L_{t|A}(S_t, p_{t+1}) = V(S_t, p_{t+1}|A) + \beta E[(p_{t+1})L_{t+1|A}(S_{t+1}, p_{t+2}) + (1-p_{t+1})L_{t+1|B}(S_{t+1}, p_{t+2})|A]$$

and

$$L_{t|B}(S_t, p_{t+1}) = V(S_t, p_{t+1}|B) + \beta E[(p_{t+1})L_{t+1|A}(S_{t+1}, p_{t+2}) + (1-p_{t+1})L_{t+1|B}(S_{t+1}, p_{t+2})|B]$$

Both value functions can be obtained via standard methods of computation of value functions described in Appendix C.

Appendix C. Approximation of Robust policies—value function method

In Appendix B we described $L_{\bar{p}_{t|A}}$ and $L_{\bar{p}_{t|B}}$. In this section we explain the numerical procedure to compute this value functions.

Algorithm.

1. First, discretize the exogenous state space using Tauchen's method. The relevant space is $S^* = \{\varepsilon_{\pi,t-1}, \varepsilon_{y,t-1}, r_{t-1}^n, i_{t-1}^*, \varepsilon_{e,t-1}, \varepsilon_{y,t-1}\}$.
2. We model prior space according to Chebyshev nodes.
3. Use an initial guess for the two value functions indexed by their corresponding model priors: $L_{t|A}(\cdot, \cdot)$ and $L_{t|B}(\cdot, \cdot)$.
4. Define a closed loop to satisfy a convergence condition for the guess in $L_{t|A}(\cdot, \cdot)$ and $L_{t|B}(\cdot, \cdot)$.
5. Define an inner loop over all model prior probabilities and exogenous state sock grids.
6. Define an open loop for draws in w_{t+1} , using 1000 draws.
7. Draw a random sample of w_{t+1} where this variable is a five dimensional normal standard vector.
8. Compute $w'_{t+1} B' C(p_{t+1})' W C(p_{t+1}) B w'_{t+1}$, use w'_{t+1} to update $p(t+2)$ given $p(t+1)$, for Model A and Model B and the point in state S^* .
9. Update the point S^* , and use updated points in $[0, 1] \times S^*$ and evaluate them at $L_{t|A}(\cdot, \cdot)$ and $L_{t|B}(\cdot, \cdot)$ using an interpolation method. Save the outcome.
10. Compute the mean of the outcomes in 8–9, and use these to update the guess values in the grid space for $L_{t|A}(S_t, p_{t+1}) \times L_{t|B}(S_t, p_{t+1})$.
11. Repeat 5–10 until convergence.

Once the value functions are computed, we use them to compute a solution to the distorted prior in Eq. (9).

Appendix D. Derivation of recursive odd's ratio

Let Z_t be the vector of observables at time t . Accordingly, we define as $Z^t \equiv \{Z_0, Z_1, \dots, Z_t\}$ the history of all observations up to time t . The standard Bayes rule is derived in the following way:

$$p_t = P(M^A | Z^t) = \frac{P(Z_t | M^A, Z_{t-1}) P(M^A | Z^{t-1})}{P(Z_t | M^B) P(M^B | Z^{t-1}) + P(Z_t | M^A, Z_{t-1}) P(M^A | Z^{t-1})} = \frac{P(Z_t | M^A, Z_{t-1}) p_{t-1}}{P(Z_t | M^B, Z_{t-1}) (1 - p_{t-1}) + P(Z_t | M^A, Z_{t-1}) p_{t-1}} \quad (30)$$

Under constant gain, we have that the standard deviation of the output gap equation takes the form

$$\sigma_{x,t-T} = \sigma_x \kappa^{-T/2}$$

where $\kappa \leq 1$ is a parameter governing the constant gain. In addition, the initial prior at time t is also altered and becomes $P(M^A)^{1/\kappa^t} < P(M^A)$.

Define $M_y \equiv E(y_{t-T} | Z^{t-T-1}, M_A)$. The likelihood at time t , of y_{t-T} , conditional on history Z^{t-T-1} and Model A, is therefore given by

$$\hat{P}(Z_{t-T} | Z^{t-T-1}, M_A, t) = N(Z_{t-T}, M_y, \sigma_{x,t-T}) = \frac{1}{\sqrt{2\pi}} \kappa^{-T/2} \exp\left(-\frac{1}{2} \frac{(Z_{t-T} - M_y)^2}{\sigma_x^2 \kappa^{-T}}\right) = (2\pi)^{(\kappa-1)/2} \kappa^{(\kappa-(\kappa-1)T)/2} [N(Z_{t-T}, M_y, t-1)]^\kappa$$

The recursive formulation as in Eq. (7) is obtained by using this formula to obtain the posterior:

$$\begin{aligned} \hat{p}_t &= \hat{P}(M^A | Z^t) = \frac{\hat{P}(Z_t | M^A, Z_{t-1}, 1) \hat{P}(Z_{t-1} | M^A, Z_{t-2}, 2) \cdots \hat{P}(Z_{t-T} | M^A, t) P(M^A)^{t/\kappa}}{\sum_{M \in \{M^A, M^B\}} \hat{P}(Z_t | M, Z_{t-1}, 1) \hat{P}(Z_{t-1} | M, Z_{t-2}, 2) \cdots \hat{P}(Z_{t-T} | M, t) P(M)^{t/\kappa}} \\ &= \frac{P(Z_t | M^A, Z_{t-1}) \hat{p}_{t-1}^\kappa}{P(Z_t | M^B, Z_{t-1}) (1 - \hat{p}_{t-1})^\kappa + P(Z_t | M^A, Z_{t-1}) \hat{p}_{t-1}^\kappa} \end{aligned}$$

where the first line follows from the Markov property of the model and the second is obtained by dividing the numerator and the denominator by $(2\pi)^{(\kappa-1)/2} \kappa^{(\kappa-(\kappa-1)T)/2}$. Finally, the term \hat{p}_{t-1}^κ is obtained by factoring the likelihood functions.

References

- Adam, K., 2004. On the relation between robust and Bayesian decision making. *Journal of Economic Dynamics and Control* 28, 2050–2117.
- Aghion, P., Bacchetta, P., Banerjee, A., 2000. A simple model of monetary policy and currency crises. *European Economic Review* 44, 728–738.
- Aghion, P., Bacchetta, P., Banerjee, A., 2001. Currency crises and monetary policy in an economy with credit constraints. *European Economic Review* 45, 1121–1150.
- Aumann, R., 1979. Agreeing to disagree. *Annals of Statistics* 4, 1236–1239.
- Ball, L., 1999. *Policy Rules for Open Economies, Monetary Policy Rules*. University of Chicago Press, Chicago.
- Bigio, S., Vega, M., 2006. Optimal monetary policy under balance sheet uncertainty. Unpublished manuscript.
- Brainard, W., 1967. Uncertainty and the effectiveness of policy. *American Economic Review Papers and Proceedings* 57, 411–425.
- Caballero, R.J., Krishnamurthy, A., 2001. International and domestic collateral constraints in a model of emerging market crises. *Journal of Monetary Economics* 48, 513–548.
- Calvo, G., Reinhart, C., 2002. Fear of floating. *Quarterly Journal of Economics* 117, 379–408.
- Caplin, A., Leahy, J., 1996. Monetary policy as a process of search. *American Economic Review* 86, 689–702.
- Céspedes, L.F., Chang, R., Velasco, A., 2002. Balance sheets and exchange rate policy. *American Economic Review* 94, 1183–1193.
- Chang, R., Velasco, A., 2000. Exchange rate policy for developing countries. *American Economic Review* 90, 71–75.
- Chang, R., Velasco, A., 2006. Monetary policy and the currency denomination of debt: a tale of two equilibria. *Journal of International Economics* 69, 150–175.
- Chari, V.V., Kehoe, P.J., McGrattan, E.R., 2002. Can sticky price models generate volatile and persistent real exchange rates? *Review of Economic Studies* 69, 533–563.
- Cho, I.-K., Kasa, K., 2008. Learning dynamics and endogenous currency crises. *Macroeconomic Dynamics* 12, 257–285.
- Cho, I.-K., Williams, N., Sargent, T.J., 2002. Escaping Nash inflation. *Review of Economic Studies* 69, 1–40.
- Clarida, R., Gali, J., Gertler, M., 2001. Optimal monetary policy in open versus closed economies: an integrated approach. *American Economic Review Papers and Proceedings* 91, 248–252.
- Clarida, R., Gali, J., Gertler, M., 2002. A simple framework for international monetary policy analysis. *Journal of Monetary Economics* 49, 879–904.
- Cogley, T., Colacito, R., Sargent, T., 2007. Benefits from U.S. monetary policy experimentation in the days of Samuelson and Solow. *Journal of Money, Credit and Banking* 39, 67–99.
- Cogley, T., Colacito, R., Hansen, L.P., Sargent, T., 2008. Robustness and U.S. monetary policy experimentation. Unpublished manuscript.
- De Paoli, B., 2009. Monetary policy and welfare in a small open economy. *Journal of International Economics* 77, 11–22.
- Ellison, M., 2006. The learning cost of interest rate reversals. *Journal of Monetary Economics* 53, 1895–1907.
- Ellison, M., Sarno, L., Vilmunen, J., 2007. Caution or activism? Monetary policy strategies in an open economy. *Macroeconomic Dynamics* 11, 519–541.
- Faia, E., Monacelli, T., 2008. Optimal monetary policy in a small open economy with home bias. *Journal of Money, Credit and Banking* 40, 721–750.
- Frankel, J., 2005. Contractionary currency crashes in developing countries. *IMF Staff Papers*, 52, No. 2.
- Gali, J., Monacelli, T., 2005. Monetary policy and exchange rate volatility in a small open economy. *Review of Economic Studies* 72, 707–734.
- Goodhart, C., 2005. The future of central banking. *FMG Special Paper No. 162*.
- Hansen, L.P., Sargent, T.J., 2006. Wanting robustness in macroeconomics. Unpublished manuscript.
- Hansen, L.P., Sargent, T.J., 2007. *Robustness*. Princeton University Press, Princeton.
- Hansen, L.P., Sargent, T.J., Tallarini, T.D., 1999. Robust permanent income and pricing. *Review of Economic Studies* 66, 873–907.
- Hausmann, R., Panizza, U., Stein, E., 2001. Why do countries float the way they float? *Journal of Development Economics* 66, 387–414.
- Ize, A., Levy-Yeyati, E., 2003. Financial dollarization: causes and consequences. *Journal of International Economics* 59, 323–347.
- Kasa, K., 1999. Will the fed ever learn? *Journal of Macroeconomics* 21, 279–292.
- Krugman, P., 1999. Balance sheets effects, the transfer problem and financial crises. In: Isard, P., Razin, A., Rose, A.K. (Eds.), *International Finance and Financial Crises: Essays in Honour of Robert P. Flood, Jr.*. Kluwer Academic Publishers, New York, pp. 31–44.
- Levy-Yeyati, E., Sturzenegger, F., 2003a. To float or to fix: evidence of the impact of exchange rate regimes on economic growth. *American Economic Review* 93, 1173–1193.
- Levy-Yeyati, E., Sturzenegger, F., 2003b. Classifying exchange rate regimes: deeds vs. words. *European Economic Review* 49, 1603–1635.
- Lubik, T.A., Schorfheide, F., 2007. *Journal of Monetary Economics* 54, 1069–1087.
- Mailath, G.J., Samuelson, L., 2006. *Repeated Games and Reputation*. Oxford University Press, New York, New York.
- Monacelli, T., 2005. Monetary policy in a low pass-through environment. *Journal of Money, Credit and Banking* 37, 1047–1066.
- Murchison, S., Rennison, A., 2006. ToTEM: the Bank of Canada's new quarterly projection model. *Bank of Canada Technical Report No. 97*.
- Sargent, T.J., 1999. *The Conquest of American Inflation*. Princeton University Press, Princeton, NJ.
- Sargent, T.J., Williams, N., 2005. Impacts of priors on convergence and escapes from Nash inflation. *Review of Economic Dynamics* 8, 360–391.
- Schneider, M., Tornell, A., 2004. Balance sheet effects, bailout guarantees and financial crises. *Review of Economic Studies* 71, 883–913.
- Wieland, V., 2000a. Learning by doing and the value of optimal experimentation. *Journal of Economic Dynamics and Control* 24, 501–534.