DISCUSSION PAPER SERIES

No. 4766

**CAUTION OR ACTIVISM?
MONETARY POLICY STRATEGIES
IN AN OPEN ECONOMY**

Martin Ellison, Lucio Sarno
and Jouko Vilmunen

INTERNATIONAL MACROECONOMICS

Centre for Economic Policy Research

www.cepr.org

Available online at: **www.cepr.org/pubs/dps/DP4766.asp**

ISSN 0265-8003

CAUTION OR ACTIVISM? MONETARY POLICY STRATEGIES IN AN OPEN ECONOMY

Martin Ellison, University of Warwick and CEPR
Lucio Sarno, University of Warwick and CEPR
Jouko Vilmunen, Bank of Finland

Discussion Paper No. 4766
November 2004

Centre for Economic Policy Research
90–98 Goswell Rd, London EC1V 7RR, UK
Tel: (44 20) 7878 2900, Fax: (44 20) 7878 2999
Email: cepr@cepr.org, Website: www.cepr.org

CEPR Discussion Paper No. 4766

November 2004

ABSTRACT

Caution or Activism? Monetary Policy Strategies in an Open Economy*

We examine optimal policy in a two-country model with uncertainty and learning, where monetary policy actions affect the real economy through the real exchange rate channel. Our results show that whether policy should be cautious or activist depends on the size of one country relative to another. If one country is small relative to the other then activism is optimal. In contrast, if the two countries are equal sized then caution prevails. Caution is induced in the latter case because of the interaction between the home and foreign central banks. In a two-country symmetric equilibrium, learning is shown to be detrimental to welfare, implying that optimal policy is cautious.

JEL Classification: E52, E58 and F41
Keywords: learning, monetary policy and open economy

Martin Ellison
Department of Economics
University of Warwick
Coventry
CV4 7AL
Tel: 7652 3472
Email: martin.ellison@warwick.ac.uk

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=146947

Lucio Sarno
Finance Group
Warwick Business School
University of Warwick
Coventry
CV4 7AL
Tel: (44 24) 7652 8219
Fax: (44 24) 7657 2871
Email: lucio.sarno@warwick.ac.uk

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=138519

Jouko Vilmunen
Bank of Finland
Research Department
PO Box 160
SF-00101, Helsinki
FINLAND
Tel: (358 9) 183 2526
Fax: (358 9) 174 872
Email: jouko.vilmunen@bof.fi

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=117786

*This Paper was partly written while Martin Ellison was visiting the Bank of Finland and Lucio Sarno was visiting the International Monetary Fund and the Central Bank of Norway. The authors are grateful for helpful remarks and suggestions to James Bullard, Jagjit Chadha, Seppo Honkapohja, Patrick Minford, Kaushik Mitra, Mark Salmon, Mark Taylor, Dan Thornton, David Vestin and Volker Wieland, in addition to participants at presentations held at the 2003 Anglo-French Macroeconomics Workshop, University of Oxford; the MMF Workshop on ‘Dynamics and Policy in the Open Economy’, London; and the CEPR/Bank of Finland Workshop on ‘Heterogeneous Information and Modeling of Monetary Policy’. The authors remain solely responsible for the views expressed in the Paper and any errors that may remain.

Submitted 11 November 2004

1 Introduction

Since the rational expectations revolution of the 1970s, the literature on optimal monetary policy has increasingly emphasized the importance of learning by private agents about government behavior and policy. However, it seems equally important to understand the process by which the government itself learns in the conduct of policy. The central bank learning process, whose importance is witnessed by the large staffs employed to produce information of practical value for monetary policy, attempts to understand the behavior of both private agents and other governments or central banks. In a highly interdependent world, the actions of other governments and central banks play a crucial role in determining the efficacy of policy.

This paper contributes to a growing literature that takes the learning processes of central banks seriously. Specifically, we examine the design of optimal monetary policy under uncertainty and learning in a two-country open-economy environment. We emphasize the open-economy aspect by focussing on the exchange rate channel of monetary policy, and assume that monetary actions affect the real economy through their impact on the real exchange rate. In our model, central banks implement monetary policy to achieve inflation and output targets, whilst simultaneously learning about the link between the real exchange rate and the real economy.

Our research builds on previous studies that have analyzed optimal monetary policy and learning in closed economy models, which have generally studied a world where the central bank has less knowledge than possible about the policy environment.[1] In this context, it has been argued that an 'activist' policy is optimal because strong policy actions generate information useful for learning.[2] In our open-economy framework, the result does not necessarily hold and activist policy may not be optimal. Hence, our generalization to the open-economy case is not a mere technical extension of existing models to cover both learning and open

[1]See Basar and Salmon (1990), Bertocchi and Spagat (1993), Balvers and Cosimano (1994) and Wieland (2000b).

[2]Whilst most academic papers argue in favour of activism, central bankers have often stated that policy should predominantly be cautious due to uncertainty, following the result of Brainard (1967).

economy considerations: it potentially leads to very different policy recommendations.

We analyze two different scenarios with our two-country open economy model. In the first scenario, one country is dominant (large) relative to the other. With a large and a small country, we replicate the result obtained in the closed economy literature and find an incentive for the small country to follow an activist monetary policy. Intuitively, this is because the central bank in the large country does not respond to the monetary policy actions of the small country.

For the second scenario, we assume that the two countries are of equal size. In this case, each central bank has to set policy based in part on projections of the behavior of the other central bank. The way the economy responds becomes closely connected to the policy itself. In particular, an activist policy leads to stronger actions from the other country in the future, which translate back into unwanted future output or inflation fluctuations in the original country. For this reason, activist polices are likely to be suboptimal: the existence of a foreign central bank in a symmetric open-economy setting serves to neutralize the gains from activism. The result, which calls for more caution in optimal monetary policy under uncertainty and learning than previously suggested by the literature, derives from the interaction between countries that is central to our open economy framework.

The structure of the paper is as follows. In Section 2 we present our two-country open-economy model and discuss the roles of uncertainty and learning. The optimal monetary policy problem under learning is defined in Section 3. Section 4 describes our calibration exercise and reports numerical results for scenarios with i) a large and a small country and ii) two countries of equal size. We also perform sensitivity analysis to check the robustness of our results. Conclusions are presented in Section 5. A Technical Appendix provides details of the numerical methods employed.

2 The model

The salient features of our two-country model are the structure of the open economy, the preferences and beliefs of each central bank, the learning mechanism through which beliefs are updated, and the definition of equilibrium.

2.1 Structure of the open economy

The structure of the open economy is represented by twin open-economy Phillips curve relationships (1) and (2) that emphasize the exchange rate channel of monetary policy, following the stylized version of the Aghion, Bacchetta and Banerjee (2001) model suggested by Cho and Kasa (2004).

$$y_t = y_0 + \theta_t(s_t - s_t^e) + z_t + \varepsilon_t \tag{1}$$

$$y_t^* = y_0^* - \theta_t^*(s_t - s_t^e) + z_t^* + \varepsilon_t^*. \tag{2}$$

Outputs y_t and y_t^* in the home and foreign countries are primarily determined by unexpected changes in the bilateral real exchange rate s_t, defined as the price of home goods in terms of foreign goods. The parameters θ_t and θ_t^* are assumed to be positive because of strong balance sheet effects, so an unexpected appreciation in the real exchange rate stimulates home output but depresses foreign output. Output in each country is also subject to observable and unobservable shocks, (z_t, z_t^*) and $(\varepsilon_t, \varepsilon_t^*)$, which are normally distributed with variance-covariance matrices Σ_z and Σ_ε. We define the activism of monetary policy as the degree to which central banks react to the observable output shocks z_t and z_t^*.

The bilateral real exchange rate is defined in logarithms by $s_t = e_t + p_t - p_t^*$, where e_t is the nominal exchange rate, and (p_t, p_t^*) denote the domestic and foreign price levels. The unexpected change in the real exchange rate may then be written as equation (3), in which unexpected real exchange movements are caused by either nominal exchange rate shocks w_t or unexpected relative inflation rate differentials. In our analysis, the main focus is on the last two terms in equation (3) and the incentives for central banks to use surprise inflation to affect output by creating unexpected real exchange rate movements through surprise inflation differentials:

$$
\begin{aligned}
s_t - s_t^e &= (e_t - e_t^e) + (p_t - p_t^e) - (p_t^* - p_t^{*e}) \\
&= (e_t - e_t^e) + (p_t - p_{t-1} - p_t^e + p_{t-1}) - \left(p_t^* - p_{t-1}^* - p_t^{*e} + p_{t-1}^*\right) \\
&= w_t + (\pi_t - \pi_t^e) - (\pi_t^* - \pi_t^{*e}) \,. \qquad (3)
\end{aligned}
$$

To introduce a role for learning in the model, we assume that the parameters θ_t and θ_t^* cannot be observed directly by either the home or foreign central bank. For simplicity of the learning process, (θ_t, θ_t^*) are restricted to be either both high (θ_H, θ_H^*) or both low (θ_L, θ_L^*), with switches between high and low values occurring according to a two-state Markov process. In other words, the economy switches between periods in which the exchange rate channel of monetary policy is strong or weak. The condition probabilities of not switching, i.e. $\rho_H = P[(\theta_{t+1}, \theta_{t+1}^*) = (\theta_H, \theta_H^*) \,|(\theta_t, \theta_t^*) = (\theta_H, \theta_H^*)]$ and $\rho_L = P[(\theta_{t+1}, \theta_{t+1}^*) = (\theta_L, \theta_L^*) \,|(\theta_t, \theta_t^*) = (\theta_L, \theta_L^*)]$ are treated as exogenous.

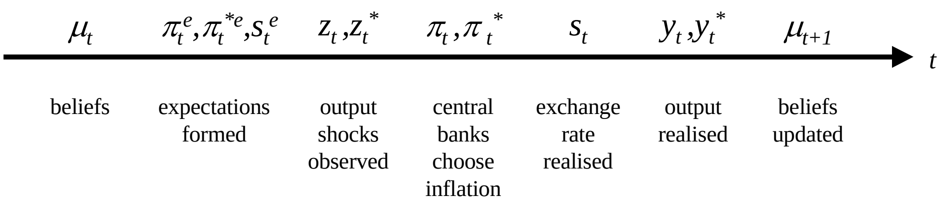

Figure 1: Timing of the model

The timing of the model is shown in Figure 1. At the beginning of the period, private agents form expectations of home inflation, foreign inflation and the real exchange rate. The observable output shocks are then revealed to both the home and foreign central banks. The two central banks next set inflation, the instrument of monetary policy. Activism or caution is reflected in the degree to which inflation reacts to observable output shocks. The real exchange rate and outputs are revealed at the end of the period. There is no asymmetric information in the model since each central bank always knows both observable output shocks.

2.2 Central bank loss function

The loss function of each central bank is assumed to be quadratic in the deviations of its own output and inflation from target. We assume that the targets are set consistent with the natural rates of output (y_0, y_0^*) so there is no systematic inflation bias. In the terminology of Svensson (1999), each central bank has a flexible inflation target. Equation (4) shows the algebraic form of the home central bank loss function, with the parameter χ measuring the relative weight of home inflation and output and β being the subjective discount factor:

$$\mathcal{L}_t = \sum_{n=0}^{\infty} \beta^n \left[(y_{t+n} - y_0)^2 + \chi \pi_{t+n}^2 \right]. \tag{4}$$

The loss function of the foreign central bank is given analogously in terms of foreign output and inflation by equation (5). We assume that both central banks care about inflation and output with the same relative weight:

$$\mathcal{L}_t^* = \sum_{n=0}^{\infty} \beta^n \left[(y_{t+n}^* - y_0^*)^2 + \chi \pi_{t+n}^{*2} \right]. \tag{5}$$

2.3 Beliefs

The central banks cannot observe the parameters θ_t and θ_t^* directly and so form beliefs about their current values. Since the parameters are restricted to switch simultaneously between two sets of values, the beliefs of the home central bank can be conveniently represented by a single variable, $\mu_t = P[(\theta_t, \theta_t^*) = (\theta_H, \theta_H^*)]$, the belief at time t that the exchange rate channel of monetary policy is currently strong. If $\mu_t = 1$ the home central bank is certain that the channel is strong. Conversely, if $\mu_t = 0$ there is certainty that the channel is weak. The beliefs of the foreign central bank can similarly be summarized as $\mu_t^* = P[(\theta_t, \theta_t^*) = (\theta_H, \theta_H^*)]$. Since all information in the model is symmetric, the beliefs of the home and foreign central banks always coincide and $\mu_t = \mu_t^*$ for all t.

2.4 Learning

Beliefs are dynamic in the model, being updated as the central banks receive information about the current values of the parameters in the Phillips curves. In Figure 1, initial beliefs μ_t are updated to μ_{t+1} at the end of the period on the basis of private sector expectations, observable output shocks, central bank inflation choices, and the realizations of the real exchange rate and output. This forms the link between policy activism and learning in the model. In general, an activist policy involves making inflation choices that are more informative, leading to faster learning and quicker updating of beliefs.

The central banks update their beliefs at the end of the period according to whether the outputs realized are more consistent with high or low parameters in the Phillips curves. Equations (6) and (7) show the predicted distributions of output, conditional on prior information $\mathcal{I}_t \equiv (\pi_t^e, \pi_t^{*e}, s_t^e, z_t, z_t^*, \pi_t, \pi_t^*, s_t)$ and the two sets of possible values of the Phillips curve parameters, (θ_H, θ_H^*) and (θ_L, θ_L^*). The central banks have to infer which of these two distributions has most likely generated y_t and y_t^*. The role for policy activism and interaction in learning is apparent since it is inflation choices that determine the unexpected change in the real exchange rate, hence how the distributions (6) and (7) differ and how easy it is for the central bank to learn:

$$\left.\begin{pmatrix} y_t \\ y_t^* \end{pmatrix}\right|_{\mathcal{I}_t,(\theta_t,\theta_t^*)=(\theta_H,\theta_H^*)} \sim N\left[\begin{pmatrix} y_0 + \theta_H(s_t - s_t^e) + z_t \\ y_0^* - \theta_H^*(s_t - s_t^e) + z_t^* \end{pmatrix}; \Sigma_\varepsilon\right] \tag{6}$$

$$\left.\begin{pmatrix} y_t \\ y_t^* \end{pmatrix}\right|_{\mathcal{I}_t,(\theta_t,\theta_t^*)=(\theta_L,\theta_L^*)} \sim N\left[\begin{pmatrix} y_0 + \theta_L(s_t - s_t^e) + z_t \\ y_0^* - \theta_L^*(s_t - s_t^e) + z_t^* \end{pmatrix}; \Sigma_\varepsilon\right]. \tag{7}$$

A simple application of Bayes rule solves the inference problem of the home central bank. Equation (8) shows how initial beliefs of the home central bank, μ_t, are updated to μ_t^+ at the end of the period on the basis of the information contained in $\mathcal{I}_t$ and realized outputs. Under such Bayesian learning, μ_t^+ depends on the relative probability of observing the outcome $\mathbf{y}_t \equiv (y_t \; y_t^*)'$ under high and low values of the Phillips curve parameters:

$$\mu_t^+ = \frac{\mu_t P(\mathbf{y}_t|_{\mathcal{I}_t,(\theta_t,\theta_t^*)=(\theta_H,\theta_H^*)})}{\mu_t(\mathbf{y}_t|_{\mathcal{I}_t,(\theta_t,\theta_t^*)=(\theta_H,\theta_H^*)}) + (1-\mu_t)P(\mathbf{y}_t|_{\mathcal{I}_t,(\theta_t,\theta_t^*)=(\theta_L,\theta_L^*)})}. \tag{8}$$

μ_t^+ is the optimal inference for the home central bank of the current value of the Phillips curve parameters given expectations, observed output shocks, the unexpected exchange rate movement and realized outputs. The home central bank is consequently able to make a prediction μ_{t+1} of whether the parameters will be high in the next period, taking into account the possibility that the parameters may shift before then. In equation (9), the prediction is calculated as a weighted average of the probability of keeping high values and the probability of switching back from low to high values.

$$\mu_{t+1} = \mu_t^+ \rho_H + (1-\mu_t^+)(1-\rho_L). \tag{9}$$

Equations (8) and (9), when combined with the conditional distributions (6) and (7) for $\mathbf{y}_t$, define the nonlinear equation (10) for updating the beliefs of the home central bank:

$$\mu_{t+1} = \mathcal{B}(\mu_t, \pi_t^e, \pi_t^{*e}, s_t^e, z_t, z_t^*, \pi_t, \pi_t^*, s_t, y_t, y_t^*) \tag{10}$$

According to equation (10), updated beliefs are a function of current beliefs, expectations, observed output shocks, inflation rates, the realized real exchange rate, and realized outputs. $\mathcal{B}(\cdot)$ represents the Bayesian operator modified to take into account Markov-switching effects.

The symmetric nature of information in the model means that the beliefs of the foreign central bank, μ_t^*, are updated using exactly the same information and Bayesian formula (10) as the home central bank. In the model, there is always joint learning and the beliefs of the home and foreign central banks are updated simultaneously and identically. With such joint learning, $d\mu_t = d\mu_t^*$ for all t.

2.5 Equilibrium

We assume that the home and foreign central banks play a non-cooperative Nash game, in which each central bank takes private sector expectations and the actions of the other central bank as given. Both central banks therefore follow "beggar thy neighbor" policies. In equilibrium, each central bank chooses inflation in its own country, taking expectations and inflation in the other country as given.

3 Optimal policy under learning

To analyze optimal monetary policy in the model, we derive the policy of the home central bank under two alternative assumptions about how learning issues are taken into account. With the *passive learning* policy, both central banks learn but neither consciously attempts to influence the speed of learning by adjusting the degree of activism or caution in policy. This policy forms our baseline case since, although each central bank is learning, neither takes into account that current actions affect learning. In contrast, under the *active learning* policy, each central bank does internalize the consequences of its actions for learning. From the viewpoint of the central bank, we refer to this as the optimal policy.[3]

3.1 Passive learning policy

In the passive learning policy, each central bank optimally accounts for current uncertainty but fails to realize that current policy actions also affect expected future losses. Learning is not internalized. Since learning is the only source of dynamics in the model, the problem of the home central bank reduces to that of minimizing the expected one-period loss function each period, subject to expectations, foreign inflation, the observable home output shock and the Phillips curve. The loss minimization problem of the home central bank is given in equation (11):

$$
\begin{aligned}
&\min_{\pi_t} E_t(\mathcal{L}_t(y_t,\pi_t)|_{\mu_t,\pi_t^e,\pi_t^{*e},s_t^e,z_t,\pi_t^*}) \\
&\text{s.t.} \\
&y_t = y_0 + \theta_t(s_t - s_t^e) + z_t + \varepsilon_t \\
&s_t - s_t^e = w_t + (\pi_t - \pi_t^e) - (\pi_t^* - \pi_t^{*e})
\end{aligned}
\tag{11}
$$

[3]The policy is optimal in the absence of a commitment technology by which the central banks can coordinate their actions.

By substituting the Phillips curve and the definition of the real exchange rate into the loss function, the problem can be written as equation (12). The expected one-period loss is a weighted average of the expected losses conditional on the true values of the Phillips curve parameters, with weights equal to the probabilities the home central bank places on each set of parameters values:

$$\min_{\pi_t} \left\{ \begin{array}{c} \mu_t E_t \left[(\theta_H(w_t + (\pi_t - \pi_t^e) - (\pi_t^* - \pi_t^{*e})) + z_t + \varepsilon_t)^2 + \chi \pi_t^2 \big|_{\pi_t^e, \pi_t^{*e}, z_t, \pi_t^*} \right] \\ +(1-\mu_t) E_t \left[(\theta_L(w_t + (\pi_t - \pi_t^e) - (\pi_t^* - \pi_t^{*e})) + z_t + \varepsilon_t)^2 + \chi \pi_t^2 \big|_{\pi_t^e, \pi_t^{*e}, z_t, \pi_t^*} \right] \end{array} \right\} \tag{12}$$

Solving the first order condition for loss minimization, the passive learning policy of the home central bank policy is given by equation (13):

$$\pi_t = \frac{\mu_t \theta_H^2 + (1-\mu_t)\theta_L^2}{\mu_t \theta_H^2 + (1-\mu_t)\theta_L^2 + \chi} \left[\pi_t^e + (\pi_t^* - \pi_t^{*e}) \right] - \frac{\mu_t \theta_H + (1-\mu_t)\theta_L}{\mu_t \theta_H^2 + (1-\mu_t)\theta_L^2 + \chi} z_t \tag{13}$$

The private sector takes expectations of equation (13) at time *t-1*, which implies $\pi_t^e = 0$ and expected inflation is identically equal to zero at all times. Taking expectations of the corresponding passive learning policy of the foreign central bank leads to $\pi_t^{*e} = 0$, so inflation expectations in both countries are zero. The passive learning policy of the home central bank can then be simplified to equation (14):

$$\pi_t = \frac{\mu_t \theta_H^2 + (1-\mu_t)\theta_L^2}{\mu_t \theta_H^2 + (1-\mu_t)\theta_L^2 + \chi} \pi_t^* - \frac{\mu_t \theta_H + (1-\mu_t)\theta_L}{\mu_t \theta_H^2 + (1-\mu_t)\theta_L^2 + \chi} z_t. \tag{14}$$

Under the passive learning policy, home inflation reacts linearly to foreign inflation and the observable home output shock. The extent to which the home central bank reacts depends on its beliefs, μ_t, and the distaste-for-inflation parameter, χ. In general, the home central bank is more activist in responding if the distaste for inflation is low. An identical derivation defines the policy followed by the foreign central bank under passive learning. In equation (15), foreign inflation is a linear reaction to home inflation and the observable foreign output shock. In this case, the size of the reaction depends on the beliefs of the foreign central bank, μ_t^*:

$$\pi_t^* = \frac{\mu_t^* \theta_H^2 + (1-\mu_t^*)\theta_L^2}{\mu_t^* \theta_H^2 + (1-\mu_t^*)\theta_L^2 + \chi} \pi_t - \frac{\mu_t^* \theta_H + (1-\mu_t^*)\theta_L}{\mu_t^* \theta_H^2 + (1-\mu_t^*)\theta_L^2 + \chi} z_t^*. \tag{15}$$

Equations (14) and (15) simultaneously define the non-cooperative Nash equilibrium in the model with passive learning.

3.2 Active learning policy

The passive learning policy is not fully optimal because it does not internalize the benefits of learning. This forms the basis for the argument that central bank policy should be more activist. Activism is beneficial because a stronger reaction to the observed output shocks creates valuable information about the state of the economy. By increased activism, the central bank learns quicker about the economy and so can stabilize the economy more effectively in the future when faced with shocks.

To assess this argument in our open economy model, we calculate the active learning policy followed by the home central bank when it takes all the costs and benefits of learning into account, and then ask whether it is more activist or cautious than the passive learning policy. The active learning problem of the home central bank solves the dynamic loss-minimization problem (16), in which the home central bank minimizes the net present value of expected losses, subject to four constraints. The first two constraints are the Phillips curve and the updating equation (10) for the beliefs of the home central bank. The third constraint recognizes that all learning is joint with the foreign central bank so future beliefs of the home and foreign central banks will coincide. The fourth constraint is that the foreign central bank follows the active learning policy in the future. The minimization problem is intertemporal because future beliefs depend on current actions:

$$
\begin{aligned}
& \min_{\{\pi_t\}} E_t \sum_{n=0}^{\infty} \beta^n \left[\mathcal{L}_{t+n}(y_{t+n}, \pi_{t+n}|_{\mu_t, \pi_t^e, \pi_t^{*e}, s_t^e, z_t, \pi_t^*}) \right] \\
& \text{s.t.} \\
y_t &= y_0 + \theta_t(w_t + (\pi_t - \pi_t^e) - (\pi_t^* - \pi_t^{*e})) + z_t + \varepsilon_t \\
\mu_{t+1} &= \mathcal{B}(\mu_t, \pi_t^e, \pi_t^{*e}, s_t^e, z_t, z_t^*, \pi_t, \pi_t^*, s_t, y_t, y_t^*) \\
\mu_{t+1}^* &= \mu_{t+1} \\
\pi_{t+1}^* &= \pi_{t+1}^*(\pi_{t+1}, z_{t+1}^*, z_{t+1}). \qquad (16)
\end{aligned}
$$

This problem has a recursive nature so the active learning policy must satisfy the Bellman equation (17). As there is no systematic inflation bias in the model, all expected values are zero as in the passive learning case:

$$
V(\mu_t, z_t, \pi_t^*) = \min_{\pi_t} E_t \left[\mathcal{L}(\mu_t, z_t, \pi_t^*, \pi_t, w_t) + \beta V(\mu_{t+1}, z_{t+1}, \pi_{t+1}^*) \right]. \qquad (17)
$$

It is not possible to derive a closed-form solution to this problem because of the nonlinearity in the equation for updating beliefs. The presence of the nonlinear learning constraint also complicates any proof of existence and uniqueness of equilibrium. However, Wieland (2000a) shows that a unique equilibrium does exist so standard dynamic programming algorithms can be applied to obtain a numerical approximation to the Bellman equation and the active learning policy.[4]

4 Numerical results

4.1 Calibration

The model is calibrated to match monthly data, roughly reflecting the frequency with which monetary policy decisions are made. Table 1 shows our baseline calibrations for the large

[4]Blackwell's sufficiency conditions are satisfied for this class of problems (see Kiefer and Nyarko, 1989), so it is possible to define a contraction mapping that converges to a unique fixed point. Repeated iterations over the Bellman equation will therefore converge to the stationary optimal policy and value function. More details about this solution technique are given in the Technical Appendix.

and small country scenario and the two equally sized countries scenario. The first eight parameters refer to the structure of the twin Phillips curve relationships (1) and (2). The natural levels of output are normalized to zero. The high and low Phillips curve parameters are taken from the upper and lower bounds of the range suggested by Cho and Kasa (2004). In the large and small country scenario, the parameters (θ_H^*, θ_L^*) are set at $(0, 0)$ so unexpected exchange rate movements have no effect on output in the large (foreign) country. The remaining parameters of the Phillips curve cannot be estimated from the data. We therefore set the variance-covariance matrix of the observable shocks, Σ_z, to be the identity matrix. This implies that there are no common observable output shocks and the variance of idiosyncratic shocks is the same in both countries. The variance-covariance matrix of the unobservable output shocks, Σ_ε, is calibrated so that the ratio of unobservable to observable shocks in each country is 0.15. In addition, there is no correlation between the unobservable home and foreign shocks.

Parameter	Small and large country scenario	Two equal sized countries scenario
y_0	0	0
y_0^*	0	0
θ_H	0.3	0.3
θ_L	0.1	0.1
θ_H^*	0	0.3
θ_L^*	0	0.1
Σ_z	$\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$
Σ_ε	$\begin{pmatrix} 0.15 & 0 \\ 0 & 0.15 \end{pmatrix}$	$\begin{pmatrix} 0.15 & 0 \\ 0 & 0.15 \end{pmatrix}$
σ_w^2	0	0
ρ_H	0.975	0.975
ρ_L	0.975	0.975
χ	0.1	0.1
β	0.997	0.997

Table 1: Baseline calibrated values

The last five parameters in Table 1 concern nominal exchange rate shocks, the stochastic process governing switches in the Phillips curve parameters, and policy preferences. We set σ_w^2 to zero to sharpen our focus on the unexpected inflation differential as the primary determinant of unexpected movements in the real exchange rate. The process determining switching of the Phillips curve parameters is calibrated as symmetric, with a persistence parameter of ρ=0.975 corresponding to an average duration between switches of 1/(1-0.975)=40 months. χ reflects the relative weight that each each central bank places on inflation and output deviations from target. It is calibrated for monthly data to give inflation and output volatility the same order of magnitude. The choice of the discount factor, β, implies a quarterly discount rate of 1%.

4.2 Results with a small and large country

In the scenario with a small and large country, the foreign (large) central bank has no exploitable Phillips curve trade-off. Optimal policy in the foreign country therefore sets $\pi_t^* = 0$ for all t, and only the behavior of the home (small) central bank is of interest. Figure 2 shows numerically the equilibrium inflation choices made by the passive learning home central bank, as a function of prior beliefs, μ_t, and the observable home output shock, z_t. The upper line in Figure 2 is the response to a one standard deviation negative shock, $z_t = -1$, while the lower line refers to the case of a one standard deviation positive shock, $z_t = +1$. The vertical distance between the upper and lower lines measures the extent to which the home central bank reacts to the observable home output shock for a given prior belief. The figure shows that changes in beliefs have a significant effect on the extent to which the home central bank reacts to the observable home output shock. An increased belief that the Phillips curve parameter is high is associated with a stronger reaction.

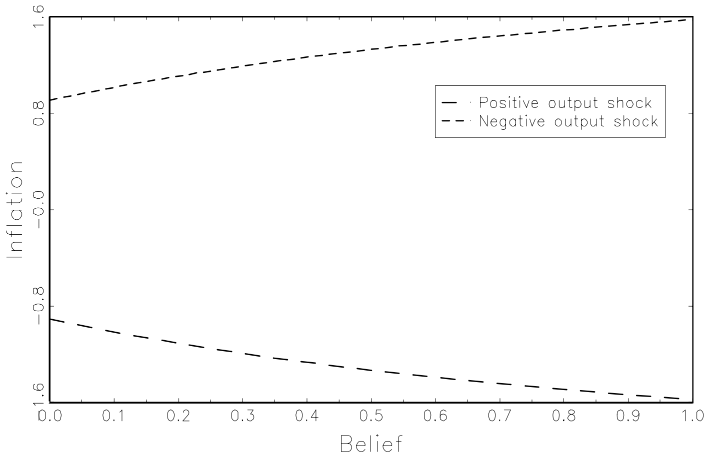

Figure 2: Inflation choices under the passive learning policy: small and large country

Although learning issues are not internalized under the passive learning policy, there will still be learning in the economy. As time passes, both central banks receive information which helps them to infer the current value of the Phillips curve parameter. The central banks update their beliefs using the learning procedure discussed in Section 2.4. Since equation (10) for the Bayesian updating of beliefs is nonlinear, simulations are needed to gain insights into the dynamic behavior of the economy. Table 2 shows some stylized facts for the calibrated economy, based on 100 simulations of 250 periods each. In the table, σ^2 denotes variance. Since there is always joint learning in the model, we only report the variance of the beliefs of the home central bank.

		Passive learning policy
Home inflation	σ^2_π	1.73
Home output	σ^2_y	0.70
Inflation differential	$\sigma^2_{\pi-\pi^*}$	1.73
Beliefs	σ^2_μ	0.0996

Table 2: Stylized facts of the passive learning policy: small country

Figure 3 shows the policies followed by the home central bank under passive and active learning. As the figure demonstrates, there is a slight difference in the degree of activism between the two policies. In contrast to the passive learning policy, the active learning policy involves a stronger reaction to the observable home output shock. In other words, the home central bank's reaction to the observable home output shock is increased. The active learning policy is marginally more activist than the passive learning policy.

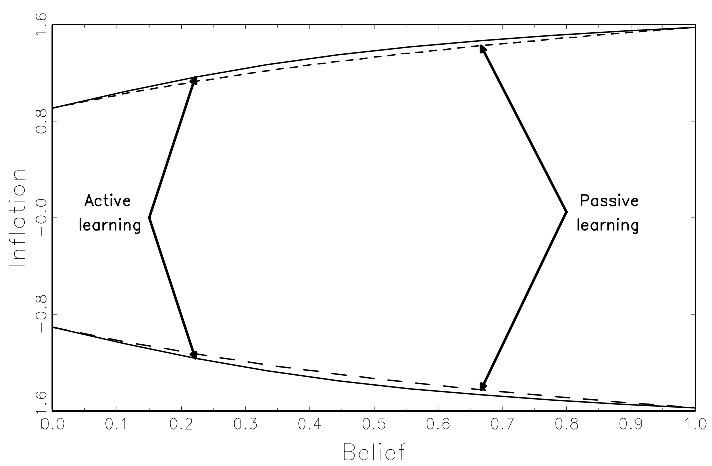

Figure 3: Inflation choices under the active learning policy: small country

The dynamic properties of the active and passive learning policies are compared in Table 3. Under the active learning policy, the volatility of home inflation naturally rises as the home central bank becomes more activist. As policy is more activist, the inflation choices of the central banks are more informative and beliefs are updated faster. Indeed, the variance of beliefs rises by approximately 2%.

		Passive learning policy	Active learning policy
Home inflation	σ_{π}^2	1.73	1.78
Home output	σ_{y}^2	0.70	0.70
Inflation differential	$\sigma_{\pi-\pi^*}^2$	1.73	1.78
Beliefs	σ_{μ}^2	0.0996	0.1015

Table 3: Stylized facts of the active learning policy: small country

The increased volatility of beliefs has a positive but small effect on the welfare of the home central bank. It therefore appears that there are incentives for activist policy in our open economy model when the country in question is small relative to the other. In this case, an active learning policy that internalizes the costs and benefits of learning is more activist than a passive learning policy.

4.3 Results with two equally sized countries

The policies followed by the home central bank when both countries are of equal size are shown in Figure 4. The dashed lines refer to the passive learning policy, with the upper line corresponding to the (z_t, z_t^*) shock combination $(1, 1)$ and the lower line corresponding to $(-1, -1)$. We highlight these particular combinations of shocks because they create the greatest conflict between the two countries, with the home central bank wanting a depreciation of the real exchange rate to offset its positive output shock, and the foreign central bank wanting an appreciation to offset its positive shock. The solid lines are for the response of the active learning policy to the same shocks. According to the figure, there is a marked difference in the degree of activism between the two policies. Compared to the passive learning policy, the active learning policy reacts less to the observable output shocks. In this scenario with two equal sized countries, the central banks' reactions to output shocks is dampened. The active learning policy becomes more cautious than the passive learning policy.

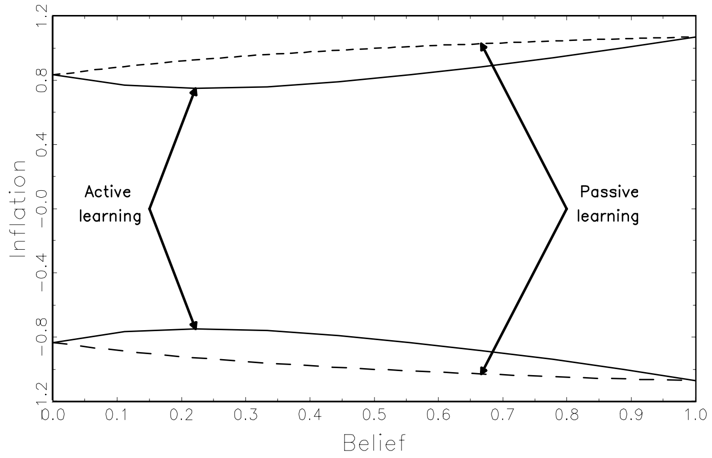

Figure 4: Inflation choices under the active learning policy: equal sized countries

The dynamics of the passive and active learning policies in the two country scenario are compared in Table 4. We report statistics only for the home country as the model is completely symmetric and volatilities in each country are the same. The cautious nature of the active learning policy means that volatility of inflation falls by 12% and the volatility of the inflation differential falls by 40%. Reduced volatility in inflation implies that the policy is less informative, which is reflected in the reduced variance of beliefs as learning is slower.

		Passive learning policy	Active learning policy
Home inflation	σ^2_π	2.94	2.60
Home output	σ^2_y	0.82	0.87
Inflation differential	$\sigma^2_{\pi-\pi^*}$	2.65	1.59
Beliefs	σ^2_μ	0.1250	0.0939

Table 4: Stylized facts of the passive learning policy: equal sized countries

Reducing the speed of learning improves the welfare of both countries, albeit only slightly. When both countries are of equal size, the incentives appear to be in favour of cautious policy. An active learning policy that internalizes learning is more cautious than its passive learning counterpart.

4.4 Sensitivity analysis

To establish the robustness of our results, we show that our conclusions hold for a wide range of calibrated parameter values. We do this by examining whether taking learning into account creates an incentive to decrease or increase activism relative to the passive learning policy in each scenario. If the incentive is to decrease activism then the active learning policy will be more cautious than the passive learning policy.

The incentive to decrease or increase activism in turn depends on the convexity or concavity of the expected central bank loss function with respect to beliefs when the home and foreign central banks follow the passive learning policies.[5] If the expected loss function is convex with respect to beliefs then the central bank strictly prefers a cautious policy that does not reduce uncertainty. For example, if the expected loss function is convex, then the central bank *ex ante* prefers uncertainty (e.g. $\mu_t = 0.5$) to certainty (i.e., $\mu_t = 0$ or $\mu_t = 1$) about the parameters (θ_t, θ_t^*). If the expected loss function is concave with respect to beliefs then the opposite is true and the central bank strictly prefers an activist policy which decreases uncertainty.

The expected central bank loss function can be written as a function, $\mathcal{L}^e(\mu_t, \mu_t^*)$, of the beliefs of the home and foreign central banks by substituting the passive learning policies (14) and (15) into the objective function (4) and taking expectations. Its convexity or concavity depends on the sign of its second derivative with respect to beliefs. Total differentiation of the expected loss function gives equation (18), in which the time subscripts on beliefs have been dropped for ease of notation:

$$d^2\mathcal{L}^e(\mu_t, \mu_t^*) = \mathcal{L}_{\mu\mu} d\mu^2 + \mathcal{L}_{\mu^*\mu^*} d\mu^{*2} + 2\mathcal{L}_{\mu\mu^*} d\mu d\mu^*. \tag{18}$$

Due to the symmetric nature of the model there is always joint learning of the home and foreign central banks, so $d\mu_t = d\mu_t^*$ for all t. We therefore evaluate the convexity or concavity of the expected loss function by looking at the sign of the second derivative

[5] See DeGroot (1962) for more details of this approach to calculating the incentives for policy experimentation under learning.

along the path of joint learning. Equation (19) shows that this second derivative has three components. The first, $\mathcal{L}_{\mu\mu}$, is the second partial derivative with respect to the beliefs of the home central bank. The second, $\mathcal{L}_{\mu^*\mu^*}$, is the second partial derivative with respect to the beliefs of the foreign central bank. The third, $2\mathcal{L}_{\mu\mu^*}$, depends on the cross partial derivative with respect to the joint beliefs of both the home and foreign central banks:

$$\left.\frac{d^2\mathcal{L}^e(\mu_t,\mu_t^*)}{d\mu^2}\right|_{d\mu=d\mu^*} = \mathcal{L}_{\mu\mu} + \mathcal{L}_{\mu^*\mu^*} + 2\mathcal{L}_{\mu\mu^*}. \tag{19}$$

The components of the second derivative of the expected loss function for the two scenarios, evaluated at $\mu_t = \mu_t^* = 0.5$, are shown in Table 4. For the small and large country scenario, the expected loss of the home (small) central bank is concave with respect to the beliefs of the home central bank but independent of the beliefs of the foreign (large) central bank. Overall, the loss function is concave, so that informative, activist policy is optimal.

θ_H	θ_L	θ_H^*	θ_L^*	$\mathcal{L}_{\mu\mu}$	$\mathcal{L}_{\mu^*\mu^*}$	$\mathcal{L}_{\mu\mu^*}$	$\left.\frac{d^2\mathcal{L}^e}{d\mu^2}\right\|_{d\mu=d\mu^*}$	Active learning policy
0.3	0.1	0	0	-0.116	0	0	-0.116	activist
0.3	0.1	0.3	0.1	-0.218	-0.134	0.384	0.412	cautious

Table 4: Activism of monetary policy in the baseline calibrations

When the two countries are of equal size, the expected central bank loss function is concave with respect to the beliefs of both the home and foreign central banks since $\mathcal{L}_{\mu\mu}$ and $\mathcal{L}_{\mu^*\mu^*}$ are both negative. If the home central bank were able to learn in isolation then it would *ex ante* strictly prefer the certainty of knowing the parameter (θ_t, θ_t^*) to the uncertainty represented by $\mu_t = 0.5$. Similarly, if the foreign central bank were able to learn in isolation then the home central bank would be better off. Hence, if we consider the learning of the home and foreign central banks in isolation there will be an incentive for increased activism to reduce uncertainty. However, the home and foreign central banks do not learn in isolation in our model. There is always joint learning so we also have to consider whether the expected loss function is convex or concave with respect to the joint beliefs of the home and foreign central banks. Since $\mathcal{L}_{\mu\mu^*}$ is positive, we see that the expected loss function turns out to be convex in the joint beliefs and therefore there is an incentive for more caution in monetary policy. This effect dominates the effects of learning in isolation so overall the active learning policy is cautious. The interaction in our model is so strong that even though learning in isolation is beneficial, the costs of joint learning are sufficiently high that the active learning policy is more cautious than the passive learning policy.[6]

Table 5 examines whether our results are still valid under a wide range of alternative calibrations for the parameters (θ_H, θ_H^*) and (θ_L, θ_L^*) for the small and large country scenario, $\mathcal{L}_{\mu\mu}$ is always negative and optimal policy is cautious. In all cases of the two equally sized countries scenario, the value of $\mathcal{L}_{\mu\mu}$ is positive and $\mathcal{L}_{\mu\mu^*}$ is negative so learning of the home central bank in isolation is beneficial but joint learning by both central banks is costly. The latter effect dominates for all but one of the calibrations and so the active learning policy is typically more cautious than the passive learning policy. Only if (θ_H, θ_H^*) and (θ_L, θ_L^*) are small or of very similar magnitude is it possible for the active learning policy to be more activist than its passive learning counterpart. Since these cases are of limited practical interest, we conclude that our result is fairly robust to alternative calibrations.

[6] Similar results can be obtained in closed economy models if there is a strong interaction between the central bank and private agents, see Ellison and Valla (2001), Rosal and Spagat (2002).

θ_H	θ_L	θ_H^*	θ_L^*	$\mathcal{L}_{\mu\mu}$	$\mathcal{L}_{\mu^*\mu^*}$	$\mathcal{L}_{\mu\mu^*}$	$\left.\frac{d^2\mathcal{L}^e}{d\mu^2}\right\|_{d\mu=d\mu^*}$	Active learning policy
0.6	0	0	0	-0.338	0	0	-0.338	activist
0.6	0.2	0	0	-0.005	0	0	-0.005	activist
0.6	0.4	0	0	-0.034	0	0	-0.034	activist
0.3	0	0	0	-0.590	0	0	-0.590	activist
0.3	0.1	0	0	-0.116	0	0	-0.116	activist
0.3	0.2	0	0	-0.007	0	0	-0.007	activist
0.15	0	0	0	-0.327	0	0	-0.327	activist
0.15	0.05	0	0	-0.120	0	0	-0.120	activist
0.15	0.1	0	0	-0.023	0	0	-0.023	activist
0.6	0	0.6	0	-1.076	-0.736	2.694	3.576	cautious
0.6	0.2	0.6	0.2	-0.704	-0.884	1.543	1.499	cautious
0.6	0.4	0.6	0.4	-0.198	-0.277	0.402	0.327	cautious
0.3	0	0.3	0	-0.526	-0.027	0.623	0.693	cautious
0.3	0.1	0.3	0.1	-0.218	-0.138	0.384	0.412	cautious
0.3	0.2	0.3	0.2	-0.053	-0.061	0.105	0.096	cautious
0.15	0	0.15	0	-0.281	0.019	0.094	-0.075	activist
0.15	0.05	0.15	0.05	-0.109	0.00001	0.066	0.024	cautious
0.15	0.1	0.15	0.1	-0.023	-0.003	0.022	0.018	cautious

Table 5: Activism of monetary policy for alternative calibrations

5 Conclusions

In this paper, we have revisited the controversy that exists in both academic and policy circles about whether monetary policy should be activist or cautious. Our results suggest that the answer depends crucially on the wider environment in which policy is conducted, which in our open economy context includes the central bank in another country. If the world is populated by a small and a large country, then the standard result that optimal policy should be activist is obtained for the small country. However, if both countries are of equal size the conclusion is reversed and optimal policy should be cautious. Our results are found to be robust to a wide range of calibrations.

The result for two equally sized countries differs from previous results in the literature due to the strategic interaction between the home and foreign central banks, which is missing from closed economy models or the case of small and large countries. With two countries of the same size, the home central bank cannot consider its actions and learning in isolation and has to recognize that the foreign central bank is also learning in tandem. *Joint* learning of both central banks is costly in our model because they are effectively learning to play a "beggar thy neighbor" non-cooperative Nash game, in which they both lose out. The cautious nature of the active learning policy is aimed at delaying the two central banks learning how to play this game.

While this paper provides a first attempt to contribute to the ongoing debate on whether or not optimal monetary policy should be activist in the context of an open economy model, clearly our model is very simple. However, our framework appears to be a fairly natural open-economy representation of the type of models used in the relevant literature. The model could be further generalized, for example by relaxing the assumption that the two countries have identical preferences or by introducing nominal exchange rate shocks. Establishing the robustness of our conclusions to these further generalizations remains an immediate avenue for future research.

A Technical Appendix: Numerical approximation of the active learning policy

The approximation to the active learning policy is obtained by solving the Bellman equation (17) numerically. This requires expressions for the expected one-period loss, $E_t\mathcal{L}_t$, and the expected continuation value, E_tV_{t+1}, for an inflation choice π_t. The one-period loss is given by equation (4) and the expected continuation value is given by equation (A.1), where future beliefs μ_{t+1} have been substituted out using the nonlinear updating equation (10):

$$E_tV_{t+1} = E_tV(\mathcal{B}(\mu_t, z_t, z_t^*, \pi_t, \pi_t^*, s_t, y_t, y_t^*), z_{t+1}, \pi_{t+1}^*). \qquad \text{(A.1)}$$

The expectation in equation (A.1) is evaluated by the home central bank before the realization of current output, y_t and y_t^*, and the observable output shocks for next period, z_{t+1} and z_{t+1}^*. The home central bank therefore evaluates the quadruple integral in equation (A.2):

$$\begin{aligned} E_tV_{t+1} \quad = \quad & \iiiint V(\mathcal{B}(\mu_t, z_t, z_t^*, \pi_t, \pi_t^*, s_t, y_t, y_t^*), z_{t+1}, \pi_{t+1}^*) f(y_t \left|_{\mu_t, z_t, z_t^*, \pi_t, \pi_t^*, s_t}\right.) \\ & f(y_t^* \left|_{\mu_t, z_t, z_t^*, \pi_t, \pi_t^*, s_t}\right.) f(z_{t+1}) f(z_{t+1}^*) dy_t dy_t^* dz_{t+1} dz_{t+1}^* \qquad \text{(A.2)} \end{aligned}$$

where $f(z_{t+1})$ and $f(z_{t+1}^*)$ are the distributions of z_{t+1} and z_{t+1}^*; and $f(y_t \left|_{\mu_t, z_t, z_t^*, \pi_t, \pi_t^*, s_t}\right.)$ and $f(y_t^* \left|_{\mu_t, z_t, z_t^*, \pi_t, \pi_t^*, s_t}\right.)$ are the predictive distributions of y_t and y_t^*. They have independent distributions, a normal and mixture of normals respectively, as described respectively by equations (A.3) and (A.4):

$$f\begin{pmatrix} z_{t+1} \\ z_{t+1}^* \end{pmatrix} = N\left[\Sigma_z\right] \qquad \text{(A.3)}$$

$$\begin{aligned} f\left(\begin{array}{c} y_t \\ y_t^* \end{array} \Bigg|_{\mu_t, z_t, z_t^*, \pi_t, \pi_t^*, s_t} \right) &= \mu_t N\left[\begin{array}{c} y_0 + \theta_H s_t + z_t \\ y_0^* - \theta_H^* s_t + z_t^* \end{array} ; \Sigma_\varepsilon \right] \\ & + (1-\mu_t) N\left[\begin{array}{c} y_0 + \theta_H s_t + z_t \\ y_0^* - \theta_H^* s_t + z_t^* \end{array} ; \Sigma\varepsilon \right]. \qquad \text{(A.4)} \end{aligned}$$

The computational algorithm starts by defining a grid of points in the state space (μ_t, z_t, z_t^*). The grid points for beliefs, μ_t, are distributed uniformly across the interval [0,1], but grid points for the observable output shocks z_t and z_t^* are bunched around zero according to a cosine weighting function to increase accuracy. For each gridpoint, we assign an inflation choice for the home central bank, π_t, and an initial value for the value function, V_t, using the inflation choices of the passive learning central bank (14) and the expected one-period loss (4) as starting values.

An iteration of the Bellman equation involves passing through the grid point by point. For each gridpoint, the inflation choice is reoptimized by minimizing the right hand side of the Bellman equation (17), using equation (4) for the expected one-period loss and equations (A.2) to (A.4) for the expected continuation value. Numerical evaluation of the expected continuation value requires Gaussian Quadrature methods and linear interpolation of adjacent gridpoints to evaluate the quadruple integral in (A.2). After each new inflation choice has been calculated, new values of π_t and V_t are assigned to the gridpoint. Each iteration of the Bellman equation is complete when the inflation choice and value function have been updated for each gridpoint.

Repeated application of the iterative procedure outlined above converges to the active learning policy. Using a grid of $10 \times 10 \times 10$ points, we accept convergence when the values associated with each gridpoint change by less than 1×10^{-6} between successive iterations. When optimizing the inflation choice at each gridpoint we also use a convergence tolerance of 1×10^{-6}. 32 ordinates were used in the Gaussian Quadrature approximation to equation (A.2).

To calculate the non-cooperative Nash equilibrium we use a simple iterative algorithm. In the first stage, the inflation choices of the home central bank are calculated for the given inflation choices of the foreign central bank, using the procedure described above. In the next stage, the inflation choices of the foreign central bank are updated. These foreign inflation choices are then used as the basis for calculating the new inflation choices of the home central bank. This procedure is iterated until convergence to non-cooperative equilibrium is achieved. Convergence is accepted when the change in the policy of the home central bank is less than 1×10^{-6} for each gridpoint.

References

[1] Aghion, P., Bacchetta, P. and A. Banerjee, 2001, A Simple Model of Monetary Policy and Currency Crises, *European Economic Review* 44, 728-738.

[2] Balvers, R.J. and Cosimano, T.F., 1994, Inflation Variability and Gradualist Monetary Policy, *Review of Economic Studies* 61, 721-738.

[3] Basar, T. and Salmon, M., 1990, Credibility and the Value of Information Transmission in a Model of Monetary Policy and Inflation, *Journal of Economic Dynamics and Control* 14, 97-116.

[4] Bertocchi, G. and Spagat, M., 1993, Learning, Experimentation and Monetary Policy, *Journal of Monetary Economics* 32,169-183.

[5] Brainard, W., 1967, Uncertainty and the Effectiveness of Policy, *American Economic Review Papers and Proceedings* 57, 411-425.

[6] Cho, I.-K. and Kasa, K. 2004, Learning Dynamics and Endogenous Currency Crises, University of Illinois, Urbana Champaign, mimeo.

[7] DeGroot, M.H., 1962, Uncertainty, Information, and Sequential Experiments, *The Annals of Mathematical Statistics* 33, 404-419

[8] Ellison, M. and N. Valla, 2001, Learning, Uncertainty and Central Bank Activism in an Economy with Strategic Interactions, *Journal of Monetary Economics* 48, 153-171.

[9] Kiefer, N.M. and Nyarko, Y., 1989, Optimal Control of an Unknown Linear Process with Learning, *International Economic Review* 30, 571-586.

[10] Rosal, J.C. and M. Spagat, 2002, Structural Uncertainty and Central Bank Conservatism: The Ignorant Should Keep Their Eyes Shut, *CEPR Discussion Paper* No. 3568

[11] Svensson, L.E.O., 1999, Price Level Targeting versus Inflation Targeting: A Free Lunch?, *Journal of Money, Credit and Banking* 31, 277-295.

[12] Wieland, V., 2000a, Learning by Doing and the Value of Optimal Experimentation - Optimal Learning with Endogenous Information, *Journal of Economic Dynamics and Control* 24, 501-534.

[13] Wieland, V., 2000b, Monetary Policy, Parameter Uncertainty and Optimal Learning, *Journal of Monetary Economics* 46, 199-228.