

Caution or Activism?

Monetary Policy Strategies in an Open Economy*

Martin Ellison

Lucio Sarno

University of Warwick and CEPR

Warwick Business School and CEPR

Jouko Vilmunen

Bank of Finland

February 23, 2006

Abstract

We examine optimal policy in an open-economy model with uncertainty and learning, where monetary policy actions affect the economy through the real exchange rate channel. Our results show that the degree of caution or activism in optimal policy depends on whether central banks are in coordinated or uncoordinated equilibrium. If central banks coordinate their policy actions then activism is optimal. In contrast, if there is no coordination then caution prevails. In the latter case caution is optimal because it helps central banks to avoid exposing themselves to manipulative actions by other central banks.

*We thank two anonymous referees for very helpful comments and suggestions. Martin Ellison acknowledges support from an ESRC Research Fellowship, "Improving Monetary Policy for Macroeconomic Stability in the 21st Century" (RES-000-27-0126).

1 Introduction

This paper contributes to a growing literature that takes the learning processes of central banks seriously. Specifically, we examine optimal monetary policy strategies for central banks that are learning in a two-country open-economy environment. We emphasise the open-economy aspect by assuming that real exchange rate movements are an important part of the monetary transmission mechanism. Our analysis builds on previous studies of monetary policy and learning in a closed economy, which typically conclude that an activist policy is optimal since strong policy actions generate information that is useful for learning.¹ In our open-economy framework, the result does not necessarily hold and activist policy may well be suboptimal. The generalisation to the open economy case is therefore not a mere technical extension of existing models. Optimal monetary policy strategy in an open economy can be very different to that in a closed economy.

To derive our results we introduce uncertainty and learning into the textbook two-country open-economy model of Walsh (2003). Uncertainty enters through unobserved time-variation in the elasticities of aggregate supply and demand with respect to the real exchange rate. Central banks are therefore unsure whether a real exchange rate depreciation will primarily stimulate aggregate demand (through improved competitiveness) or contract aggregate supply (through more expensive imported materials). Learning arises naturally in such a framework, as central banks continually process new data and update their estimates of the relevant elasticities. We follow Walsh (2003) and draw a distinction between equilibria with and without coordination between the central banks.

The key to understanding our results is to recognise that there are informational spillovers in a two-country environment. A central bank contemplating an activist policy must therefore take into account that information generated by its own policy will also be used in the learning process of the other central bank. When there is *policy coordination* this is not a problem as informational spillovers are beneficial: it is better to coordinate policy between informed

¹See Basar and Salmon (1990), Bertocchi and Spagat (1993), Balvers and Cosimano (1994) and Wieland (2000b). Central bankers such as Blinder (1998) generally disagree and argue that policy should be predominantly cautious without any formal justification for their view.

than uninformed central banks. Optimal monetary policy strategy is activist and we replicate the closed-economy result that policy actions should be strong to promote learning. In the equilibrium *without policy coordination* the informational spillovers prove to be more problematic. Absent coordination, each central bank independently tries to manipulate the real exchange rate to its own advantage. In effect, the central banks are in a conflict equilibrium of mutual attempted exploitation. The problem with the informational spillover is that it ties together the way central banks learn about how to exploit each other. Any benefit from one central bank learning how to exploit the other then needs to be weighed against the cost of the other central bank also learning how to exploit. The threat of increased exposure to exploitative actions from the other central bank greatly attenuates the incentive to follow an activist monetary policy strategy. In many cases, the cost of the other central bank learning is sufficiently high that it completely dominates. The closed economy result is overturned and optimal monetary policy strategies are cautious rather than activist, with weak policy responses designed to retard the learning of the other central bank.

The paper is organised as follows: In Section 2 we describe the textbook two-country open-economy model and introduce a role for uncertainty and learning. Section 3 derives optimal monetary policy strategies with policy coordination. A similar analysis in Section 4 derives optimal monetary policy strategies without policy coordination. Conclusions and further discussion are presented in Section 5. Appendices provides full details of the numerical methods employed.

2 Model

Our model derives from the open-economy sticky-wage textbook model of Walsh (2003). To it, we add uncertainty in the form of measurement errors and time variation in the elasticity parameters of aggregate supply and demand. The presence of uncertainty forces central banks to form and update beliefs about the elasticities in the model. We assume Bayesian updating of beliefs, so central banks are completely rational in their learning. They are, though, hampered by the measurement errors that cloud the signals in observed variables. The final component

of the model is a definition of central bank objectives.

2.1 Structure of the open economy

The structure of the open economy is represented by AS-AD-UIP equations (1)-(5), taken directly from Section 6.3.1 in Walsh (2003). Equations (1) and (2) are open-economy aggregate supply curves, where outputs y_t and y_t^* in the home and foreign countries are determined by the real exchange rate ρ_t (the relative price of home and foreign output expressed in terms of the home currency), unexpected inflation, and aggregate supply shocks e_t and e_t^* . A real exchange rate depreciation ($\rho_t \uparrow$) contracts supply because it increases the price of imported materials and raises consumer prices relative to producer prices, which due to wage rigidity increases the real wage in terms of producer prices. The aggregate supply shocks are uncorrelated i.i.d. Gaussian distributed with mean zero and variance σ_e^2 .

$$y_t = -b_{1t}\rho_t + b_2(\pi_t - E_{t-1}\pi_t) + e_t \quad (1)$$

$$y_t^* = b_{1t}\rho_t + b_2(\pi_t^* - E_{t-1}\pi_t^*) + e_t^* \quad (2)$$

$$y_t = a_{1t}\rho_t - a_2r_t + a_3y_t^* + u_t \quad (3)$$

$$y_t^* = -a_{1t}\rho_t - a_2r_t^* + a_3y_t + u_t^* \quad (4)$$

$$\rho_t = r_t^* - r_t + E_t\rho_{t+1} \quad (5)$$

Aggregate demand curves (3) and (4) link demand to the real exchange rate, real interest rates r_t and r_t^* , foreign output, and aggregate demand shocks u_t and u_t^* . A real exchange rate depreciation ($\rho_t \uparrow$) stimulates demand through improved competitiveness as domestic goods become cheaper relative to foreign goods. The presence of foreign output in the aggregate demand curves reflects the direct spillover that arises when an increase in output in one country raises demand for goods produced in the other. The aggregate demand shocks are uncorrelated i.i.d. Gaussian with mean zero and variance σ_u^2 . Equation (5) is the uncovered interest rate parity condition. Written as $r_t^* - r_t = \rho_t - E_t\rho_{t+1}$ it shows that any real interest rate differential must be matched by an offsetting expected depreciation or appreciation in the real exchange rate. Real interest rates in the model are linked to nominal interest rates through the Fisher equations $r_t = i_t - E_t\pi_{t+1}$ and $r_t^* = i_t^* - E_t\pi_{t+1}^*$. We assume that nominal

interest rates are the instruments of monetary policy.

2.2 Uncertainty and timing protocol

To introduce uncertainty to the Walsh (2003) model we allow some of the elasticity parameters to be unobserved and time-varying, whilst making output and inflation subject to measurement errors. Since our focus is on uncertainty and learning in the open economy, we restrict unobservability and time variation to the parameters b_{1t} and a_{1t} governing the elasticities of aggregate supply and demand with respect to the real exchange rate. The other elasticity parameters (b_2, a_2, a_3) are assumed to be time-invariant and observable. For simplicity of the learning process, the time-varying parameters are further restricted to follow a hidden two-state Markov process. In other words, (b_{1t}, a_{1t}) switches between (\bar{b}_1, \bar{a}_1) and $(\underline{b}_1, \underline{a}_1)$ and there is time variation in the absolute and relative magnitude of real exchange rate effects on aggregate supply and demand.² The conditional probabilities ω of switching in the two-state Markov process are assumed to be exogenous. Uncertainty about the elasticities is magnified by assuming that observed realisations of output and inflation are subject to measurement error. This ensures that central banks only have imperfect signals of output and inflation from which to infer the current elasticities of aggregate supply and demand. The measurement errors are assumed to have an i.i.d. Gaussian distribution with mean zero.

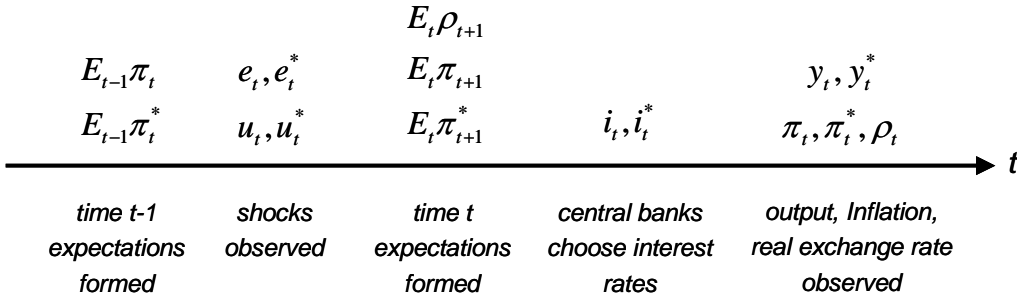


Figure 1: Timing protocol

The timing of the model is shown in Figure 1. At the beginning of the period, private

²Restricting b_{1t} and a_{1t} to switch simultaneously is done for computational reasons and is not critical for our results. Similarly, retaining the symmetry of the Walsh (2003) model by having simultaneous/identical shifts in both countries is not a prerequisite for our results. All we require is some unobserved time variation in the real exchange rate channel of the monetary policy transmission mechanism.

agents form time $t-1$ dated expectations of inflation in period t . The aggregate supply and demand shocks are then revealed to both central banks, after which private agents form time t dated expectations of inflation and the real exchange rate in period $t+1$.³ The two central banks next set nominal interest rates, the instrument of monetary policy. Activism or caution is reflected in the degree to which nominal interest rates react to the observed shocks. Inflation, output and the real exchange rate are observed (subject to measurement error) at the end of the period. There is no asymmetric information in the model since each central bank observes all shocks at the same time.

Our timing protocol retains the feature of the original textbook model that all expectations terms are zero in equilibrium. As long as monetary policy strategies do not inject systematic biases then the i.i.d. Gaussian nature of shocks is sufficient to guarantee $E_{t-1}\pi_t = E_{t-1}\pi_t^* = 0$ and $E_t\pi_{t+1} = E_t\pi_{t+1}^* = 0$. With regard to the real exchange rate, when shocks are i.i.d. Gaussian and expectations are rational it is possible to write $\rho_t = \theta E_t\rho_{t+1} + v_t$, where $0 < \theta < 1$ and v_t is white noise. It follows that $E_t\rho_{t+1} = 0$ in any no-bubbles solution.⁴

2.3 Equilibrium

A convenient representation of equilibrium can be obtained using the uncovered interest rate parity condition to substitute out for the real exchange rate in the aggregate supply and demand curves. Since all expectations terms are zero in equilibrium, the result is a structural relationship linking observed output and inflation to nominal interest rates, shocks and measurement errors. Appendix A derives the state-space form (6) of the structural relationship, where Y_t , R_t , ξ_t and η_t are vectors of observed outputs and inflations, nominal interest rates, shocks and measurement errors respectively. The matrix A_t switches between \bar{A} and \underline{A} as the elasticities in the aggregate supply and demand curves switch between (\bar{b}_1, \bar{a}_1) and $(\underline{b}_1, \underline{a}_1)$. Matrix B is time-invariant.

$$Y_t = A_t R_t + B \xi_t + \eta_t \tag{6}$$

³The precise point at which private agents form expectations of future variables is unimportant. The combination of no systematic policy biases and i.i.d. shocks means expectations of future variables will be zero throughout the period anyway.

⁴For more details on these arguments see Section 6.3.1 of Walsh (2003).

The definition of equilibrium is completed by adding monetary policy strategies to the structural relationship (6) and the stochastic processes for A_t , ξ_t and η_t . Monetary policy strategies specify nominal interest rate choices R_t by central banks, typically as functions of beliefs, shocks and nominal interest rates set by other central banks. We specify separate monetary policy strategies with and without policy coordination in Sections 3 and 4.

2.4 Beliefs

Central banks cannot observe the time-varying elasticity parameters in A_t directly - they can only form a belief about their current values. Since the elasticities are restricted to switch according to a two-state Markov process, beliefs can be represented parsimoniously by a single variable, $\mu_t = P(A_t = \bar{A})$, the belief at time t that A_t is currently equal to \bar{A} . If $\mu_t = 1$ the home central bank is certain that $A_t = \bar{A}$. Conversely, if $\mu_t = 0$ there is certainty that $A_t = \underline{A}$. The beliefs of the foreign central bank can similarly be summarised as $\mu_t^* = P(A_t = \bar{A})$. The symmetry of information in the model means that beliefs of both central banks always coincide and $\mu_t = \mu_t^*$ for all t .

2.5 Learning

Beliefs are not static in the model but are updated whenever new information becomes available. The central bank learning process updates beliefs according to which set of elasticity parameters is more consistent with new observations of output and inflation. Mathematically, the central bank has to infer whether $A_t = \bar{A}$ or $A_t = \underline{A}$ in the structural equilibrium relationship (6). The central bank observes (R_t, ξ_t) and is assumed to know all variances, so the question is which of predicted distributions (7) and (8) is most likely to have generated observed output and inflation Y_t . The role of policy activism in learning is reflected in the nominal interest rate choices R_t determining how easy it is for the central bank to distinguish between observations from the two predicted distributions. If policy is extremely cautious with $R_t = (0 \ 0)'$ then the predicted distributions coincide exactly and it is impossible for a

central bank to learn.

$$Y_t |_{A_t=\bar{A}} \sim N [\bar{A}R_t + B\xi_t; \Sigma_\eta] \quad (7)$$

$$Y_t |_{A_t=\underline{A}} \sim N [\underline{A}R_t + B\xi_t; \Sigma_\eta] \quad (8)$$

A simple application of Bayes rule describes the mechanics of the home central bank updating its beliefs and learning. Equation (9) shows how initial beliefs μ_t are updated to μ_t^+ on the basis of new observations of output and inflation. Under such Bayesian learning, μ_t^+ depends on the relative probability of observing Y_t under the two sets of elasticity parameters.

$$\mu_t^+ = \frac{\mu_t P(Y_t |_{A_t=\bar{A}})}{\mu_t P(Y_t |_{A_t=\bar{A}}) + (1 - \mu_t) P(Y_t |_{A_t=\underline{A}})} \quad (9)$$

The updated belief μ_t^+ represents optimal inference on the values of the elasticity parameters at the end of period t . From this, the home central bank forms a prediction μ_{t+1} of which elasticities will apply in period $t+1$, taking into account that the elasticities may switch before then. In equation (10), the prediction is calculated as a function of the probability of having \bar{A} at time t (and not switching from it) and the probability of having \underline{A} at time t (but switching from it). The switching probabilities ω_{11} and ω_{21} are assumed exogenous and known by the central bank.

$$\mu_{t+1} = \mu_t^+ \omega_{11} + (1 - \mu_t^+) (1 - \omega_{21}) \quad (10)$$

Equations (9) and (10), when combined with the predicted distributions for Y_t , define a nonlinear learning process (11) for updating the beliefs of the central bank.

$$\mu_{t+1} = \mathcal{B}(\mu_t, \xi_t, R_t, Y_t) \quad (11)$$

According to equation (11), updated beliefs are a function of current beliefs, shocks, nominal interest rates and observations of output and inflation. $\mathcal{B}(\cdot)$ represents the Bayesian operator modified to take account of Markov-switching effects. The symmetric nature of information means that beliefs of the foreign central bank, μ_t^* , are updated using exactly the same Bayesian formula. In the model, central banks always learn together and beliefs are updated simultaneously and identically. With such joint learning $d\mu_t = d\mu_t^*$ for all t .

2.6 Central bank objectives

The final component of the model is a definition of central bank objectives. We assume that central banks attempt to minimise a weighted average of output variability and inflation variability, following Svensson (1999) and Walsh (2003). Such an objective ensures there are no systematic biases to policy and all expectations terms are zero in i.i.d. equilibrium. Equations (12) and (13) show the algebraic form of the central bank loss functions, with λ measuring the relative weight of output variability and β being a discount factor between 0 and 1. The central banks have identical preferences.

$$V_t = \sum_{i=0}^{\infty} \beta^i (\lambda y_{t+i}^2 + \pi_{t+i}^2) \quad (12)$$

$$V_t^* = \sum_{i=0}^{\infty} \beta^i [\lambda (y_{t+i}^*)^2 + (\pi_{t+i}^*)^2] \quad (13)$$

Central bank loss functions (12) and (13) are used to derive optimal monetary policy strategies with and without policy coordination in Sections 3 and 4.

3 Policy with coordination

Our analysis of coordinated monetary policy strategies is based on deriving optimal policy under two alternative assumptions about how central banks take learning into account. With the *passive learning* policy, the central banks learn but make no conscious effort to influence the speed of learning. This forms our reference case since, although both central banks learn, neither takes into account that the degree of activism or caution in policy affects learning. In contrast, with an *active learning* policy the central banks do internalise the consequences of activism or caution for learning. The difference between monetary policy strategies under passive and active learning measures the degree of activism or caution induced by active learning.

3.1 Passive learning

With the passive learning policy, central banks optimally account for current uncertainty but fail to realise that policy actions affect future uncertainty through the learning process.

Learning is not internalised. Since learning is the only source of dynamics in the model, the coordinated policy problem reduces to that of minimising the sum of the one-period central bank loss functions each period, subject to equilibrium structural relationship (6), shocks $\xi_t = (e_t \ u_t \ e_t^* \ u_t^*)'$ and beliefs μ_t . The loss minimisation problem with policy coordination is defined in (14).

$$\begin{aligned}
& \min_{R_t} E_t(Y_t' \Lambda Y_t) \\
& \quad s.t. \\
& \quad Y_t = \bar{A}R_t + B\xi_t + \eta_t \text{ with probability } \mu_t \\
& \quad Y_t = \underline{A}R_t + B\xi_t + \eta_t \text{ with probability } 1-\mu_t \\
& \quad \mu_t, \xi_t \text{ given}
\end{aligned} \tag{14}$$

In the objective, Y_t is a vector of observations of $(y_t \ \pi_t \ y_t^* \ \pi_t^*)'$ subject to measurement errors η_t . Λ is a 4×4 matrix of preference weights with $(\lambda \ 1 \ \lambda \ 1)$ on the leading diagonal and zeros elsewhere, giving equal weight to each central bank in coordinated equilibrium. The choice variable $R_t = (i_t \ i_t^*)'$ is a vector of nominal interest rates, the instrument of monetary policy. The loss minimisation problem (14) is a simple modification of the linear-quadratic control problems studied by, for example, Sargent (1987). Standard techniques give the solution in terms of an optimal coordinated monetary policy strategy (15).

$$R_t = - \left[\mu_t \bar{A}' \Lambda \bar{A} + (1 - \mu_t) \underline{A}' \Lambda \underline{A} \right]^{-1} \left[\mu_t \bar{A}' \Lambda B + (1 - \mu_t) \underline{A}' \Lambda B \right] \xi_t \tag{15}$$

As expected in a linear-quadratic framework, optimised monetary policy strategy has nominal interest rates reacting linearly to the shocks. The extent of the reaction depends on beliefs μ_t , structural parameters \bar{A} , \underline{A} , B and preferences Λ . In general, passive learning policies tend to be cautious in order to avoid the increase in uncertainty that strong policy actions create. This call for caution under parameter uncertainty, most closely associated with Brainard (1967), resurfaces in our model as the reaction coefficient in (15) is less than its certainty equivalent value.⁵

⁵To see this note that $\mu_t \bar{A}' \Lambda \bar{A} + (1 - \mu_t) \underline{A}' \Lambda \underline{A} > (\mu_t \bar{A} + (1 - \mu_t) \underline{A})' \Lambda (\mu_t \bar{A} + (1 - \mu_t) \underline{A})$ as long as $\mu_t (1 - \mu_t) \underline{A}' \Lambda \bar{A}$ (the covariance term in the Brainard conservatism result) is small.

3.2 Active learning

The passive learning policy is not fully optimal because it does not internalise the costs and benefits of learning. This forms the basis for the closed-economy argument that monetary policy strategy should be more activist. Activism is beneficial because strong policy reactions create valuable information about the state of the economy and help the central banks to learn. To assess this argument in an open-economy coordinated policy context, we calculate the active learning policy followed by central banks when they take all learning costs and benefits into account. By definition, the coordinated active learning policy solves dynamic loss-minimisation problem (16), where the expected net present value of losses is minimised subject to the equilibrium structural relationship, shocks, beliefs and the learning process (11) by which central banks update their beliefs. The minimisation problem is intertemporal because future beliefs depend on current actions.

$$\begin{aligned}
& \min_{\{R_t\}} E_t \sum_{i=0}^{\infty} \beta^i (Y'_{t+i} \Lambda Y_{t+i}) \\
& \quad \quad \quad s.t. \\
& \quad Y_{t+i} = \bar{A}R_{t+i} + B\xi_{t+i} + \eta_{t+i} \text{ with probability } \mu_{t+i} \\
& \quad Y_{t+i} = \underline{A}R_{t+i} + B\xi_{t+i} + \eta_{t+i} \text{ with probability } 1-\mu_{t+i} \\
& \quad \mu_{t+i+1} = \mathcal{B}(\mu_{t+i}, \xi_{t+i}, R_{t+i}, Y_{t+i}) \\
& \quad \mu_t, \xi_t \text{ given}
\end{aligned} \tag{16}$$

The problem has a recursive nature so coordinated active learning policy must satisfy the Bellman equation (17).

$$V(\mu_t, \xi_t) = \min_{R_t} E_t [Y'_t \Lambda Y_t + \beta V(\mu_{t+1}, \xi_{t+1})] \tag{17}$$

It is not possible to derive a closed-form solution to this problem because of the nonlinearity in the learning process for updating beliefs. The presence of the nonlinear learning constraint also complicates any proof of existence and uniqueness of equilibrium. However, Wieland (2000a) applies results from Kiefer and Nyarko (1989) to show that a unique equilibrium policy does exist, so standard dynamic programming algorithms can be used to obtain a

numerical solution to the Bellman equation and active learning policy.⁶

3.3 Numerical example

Our baseline numerical example is designed to best highlight the incentives for activism and caution in monetary policy strategies for the open economy. To this end, we choose parameter values that emphasise the real exchange rate channel and its time variation. Table 1 shows our baseline parameterisation. In words, the parameterisation implies that the real exchange rate channel of monetary policy transmission switches between acting through aggregate supply and acting through aggregate demand. When elasticities (b_{1t}, a_{1t}) take the values $(1, 0)$ a real exchange rate depreciation contracts aggregate supply but leaves aggregate demand unchanged. With the values $(0, 1)$ the opposite is true and real exchange rate depreciations stimulate aggregate demand whilst having no supply-side effects. The Markov-switching process determining time variation in the elasticities is assumed to be symmetric with a 5% probability of switching each period. Elasticities that do not switch are less central to our analysis. We set a_2 to a small number so the direct effect of nominal interest rates on aggregate demand is weak and central banks are forced to exploit the real exchange rate channel to stabilise their economies. a_3 is set to zero to shut down the direct spillover of output in one country to aggregate demand in the other, since this spillover is not part of the real exchange rate channel. The elasticity of aggregate supply with respect to surprise inflation b_2 is normalise to unity to give a 45° Phillips curve slope.

(\bar{b}_1, \bar{a}_1)	$(\underline{b}_1, \underline{a}_1)$	ω_{11}	ω_{21}	a_2	a_3	b_2	σ_e^2	σ_u^2	σ_η^2	λ	β
(1, 0)	(0, 1)	0.95	0.05	0.1	0	1	1	0	1.5	5	0.99

Table 1: Baseline parameter values

⁶Blackwell’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 Appendix B.

The parameters of the stochastic processes for shocks and measurement errors are chosen to maintain the essence of the model whilst reducing computational complexity. To keep the dimension of the state space manageable we normalise σ_e^2 to unity and follow Walsh (2003) by setting σ_u^2 to zero, implying that only aggregate supply is subject to shocks.⁷ For similar reasons, we introduce sufficiently large measurement errors in observed inflation that it becomes completely uninformative for learning. This simplifies the learning process but does not change the incentives for activism or caution in monetary policy strategy. The value for σ_η^2 in Table 1 is the variance of measurement errors in observed output. The final parameters concern the preferences of the central bank. We set λ large since otherwise there would be too much incentive for policy to react to what are effectively inflation shocks in our model. The discount factor β is close to one as usual.

Having parameterised our numerical example, we are able to solve for optimal monetary policy strategies that map shocks and beliefs to nominal interest rates under passive and active learning. In our case the monetary policy strategies are a complex multidimensional mapping from \mathbb{R}^3 (two shocks and one belief) to \mathbb{R}^2 (two nominal interest rates). Our discussion therefore focusses on a representative part of the mapping, asking how nominal interest rates react to the combination of a positive aggregate supply shock in the home country and a negative aggregate supply shock in the foreign country. In the absence of any policy reaction, the shocks would cause inflation to fall below target in the home country and rise above target in the foreign country. The optimal policy response to such a situation depends on whether the real exchange rate is currently believed to act through the aggregate supply or the aggregate demand side of the economy. If the supply-side effect dominates there is a strong incentive for the home country to reduce its nominal interest rate and the foreign country to raise its nominal interest rate. The resulting real exchange rate depreciation (as determined by the uncovered interest rate parity condition) affects aggregate supply and so suitably offsets the aggregate supply shocks in both countries. If the demand-side effect dominates then the incentives for

⁷The precise mix of aggregate supply and demand shocks is unimportant for our recommendations vis-a-vis activism or caution in monetary policy strategy. All we require are *some* shocks that perturb the system and motivate central banks to stabilise their economies by manipulating the real exchange rate. Aggregate supply shocks provide this motivation so we can abstract from aggregate demand shocks.

central banks to react are much weaker. Engineering a real exchange rate depreciation is now less desirable since it acts on aggregate demand, which can only imperfectly offset the aggregate supply shocks.

Figure 2 presents numerically-solved optimal monetary policy strategies in the baseline parameterisation of the model. The first panel shows the reaction of nominal interest rates to a unit positive aggregate supply shock in the home country and unit negative aggregate supply shock in the foreign country. As expected, under both passive and active learning the home central bank reduces its nominal interest rate i_t and the foreign central bank raises its nominal interest rate i_t^* . Again as expected, the reactions of nominal interest rates are stronger if the central banks believe that the real exchange rate acts more through aggregate supply than aggregate demand.⁸ The second panel confirms that changes in nominal interest rates lead to a depreciation ($\rho_t \uparrow$) in the real exchange rate through the uncovered interest rate parity condition.

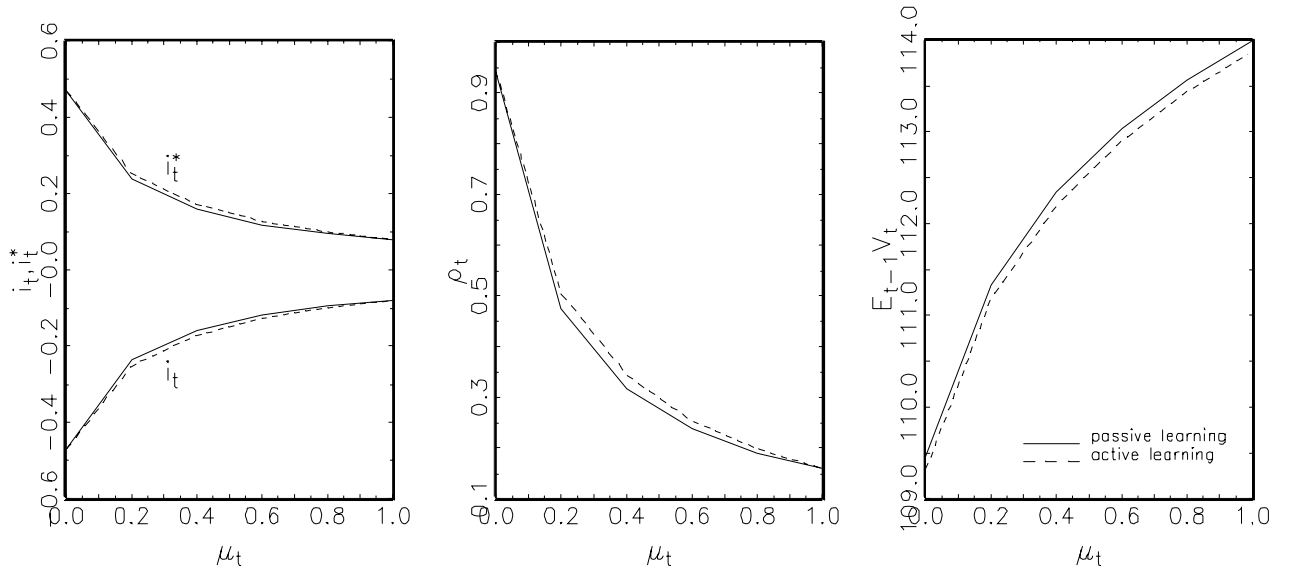


Figure 2: Policy function and unconditional value functions with policy coordination

Comparing monetary policy strategies under the different learning assumptions, it is apparent that the active learning policy involves slightly stronger policy reactions than the passive

⁸Reactions are stronger when μ_t is close to zero as the supply-side effect is believed to dominate. Note also that the reaction of nominal interest rates is muted when there is uncertainty. This reflects the Brainard (1967) result discussed in Section 3.1.

learning policy. In Figure 2, nominal interest rates react more under the active learning policy and there is a larger depreciation in the real exchange rate. We therefore conclude that optimal monetary policy strategies are *activist* when policy is coordinated by central banks in open-economy equilibrium. Our conclusion mirrors the results from the closed-economy literature. The rationale is also similar, in that activism is beneficial because it produces useful information that promotes the learning of central banks. Learning is beneficial since it puts central banks in a better position to respond to shocks in the future.

Although the policies look only slightly different in Figure 2, the benefits to pursuing an active learning policy are significant. Small differences in policy quickly cumulate into substantial differences in the dynamic behaviour of the economy. The small differences here are already sufficient to raise the correlation between beliefs and the true values of the time-varying elasticities from 0.237 to 0.260, representing a 10% improvement in the ability of central banks to track the true structure of the economy. To quantify the benefit of faster learning, the third panel of Figure 2 plots the unconditional value functions as defined by their expected values at the beginning of the period before shocks are observed.⁹ The difference in the unconditional value functions implies that the active learning policy reduces losses by approximately 0.12%

3.4 Sensitivity analysis

To establish the robustness of our results, we show they are valid for a wide range of alternative parameterisations of the model. It is computationally demanding to calculate the active learning policy for each parameterisation, so we take a short-cut and identify only the *incentives* for activism in optimal monetary policy strategy, accepting our results if we can show there are incentives to increase policy activism over and above the level found in the passive learning policy. The advantage of this approach is that incentives can be identified without

⁹The unconditional value function under passive learning is calculated by iterating on the equation $V(\mu_t, \xi_t) = E_t [Y_t' \Lambda Y_t + \beta V(\mu_{t+1}, \xi_{t+1})]$, whilst imposing learning process (11) and passive learning policy (15). The corresponding unconditional value function under active learning is taken directly from the Bellman equation (17).

iterating the Bellman equation of the active learning policy. The incentives themselves depend on whether the unconditional value function is convex or concave under passive learning. If the unconditional value function is concave then central banks strictly prefer an activist policy that reduces uncertainty.¹⁰ In Figure 2 the unconditional value function is clearly concave under passive learning, which confirms that optimal policy is activist in the baseline parameterisation. Furthermore, the first line of Table 2 shows that the expected return function $\mathcal{L}^e(\mu_t) = E_{t-1}(Y_t' \Lambda Y_t)$ is also concave under passive learning, a sufficient condition for concavity of the unconditional value function itself.

(\bar{b}_1, \bar{a}_1)	$(\underline{b}_1, \underline{a}_1)$	a_2	$\frac{d^2 \mathcal{L}^e(\mu_t)}{d\mu^2}$
{0, 1}	{1, 0}	0.1	-1.155
{0, 1}	{1, 0}	0.05	-1.163
{0, 1}	{1, 0}	0.2	-1.127
{0, 0.25}	{1, 0}	0.1	-0.191
{0, 1}	{0.25, 0}	0.1	-0.186
{0, 0.5}	{1, 0.5}	0.1	-0.037
{0.5, 1}	{0.5, 0}	0.1	-0.429

Table 2: Concavity of the expected return function at $\mu_t=0.5$ with policy coordination

In the remaining rows of Table 2 we allow parameters $\{(\bar{b}_1, \bar{a}_1), (\underline{b}_1, \underline{a}_1), a_2\}$ to change whilst keeping other model parameters at their baseline values. The fact that the expected return function is always concave implies that the unconditional value function remains concave and optimal monetary policy strategy is activist in each case. We take this as strong support for the robustness of our result. The incentives for policy activism survive a relative strengthening or weakening of the real exchange rate channel as represented by changes in a_2 in the second and third lines. Similarly, changing the relative magnitude of aggregate supply and demand elasticities in the fourth and fifth lines does not destroy the incentives for activism. Finally, it

¹⁰See DeGroot (1962) and Ellison and Vilmunen (2005) for more details of this approach to calculating the incentives for policy experimentation under learning.

is possible to remove time-variation in either elasticity in the sixth and seventh lines without overturning our result.

4 Policy without coordination

It is well known that policy recommendations are often sensitive to the assumption of whether policies are coordinated or uncoordinated.¹¹ This is particularly true in our open-economy framework. When there is no coordination, each central bank has to set its monetary policy strategy in part as a best response to the monetary policy strategy of the other central bank. The resulting non-cooperative Nash equilibrium is typically characterised by “beggar thy neighbour” policies as both central banks try to manipulate the real exchange rate to their own - rather than mutual - advantage. To disentangle the incentives for activist policy we again derive optimal monetary policy strategies under passive and active learning. For comparability with coordinated policy results we use the same parameter values in our baseline numerical example and sensitivity analysis.

4.1 Passive learning

The loss minimisation problem faced by the home passive-learning central bank is shown in (18). As in the coordinated policy case, learning is the only source of dynamics and the problem reduces to minimising the one-period central bank loss function each period. However, compared to coordination there are two important differences. Firstly, in the objective Λ_1 is a 4×4 matrix of zeros with $(\lambda \ 1 \ 0 \ 0)$ on the leading diagonal because the home central bank only cares about output and inflation in its own country. Secondly, optimisation is conditional on the monetary policy strategy of the foreign central bank, so the foreign nominal interest rate i_t^* is taken as given.

¹¹For example, Walsh (2003) shows that policy reacts less to aggregate supply shocks in his textbook model if there is no coordination.

$$\begin{aligned}
& \min_{i_t} E_t(Y_t' \Lambda_1 Y_t) \\
& \quad s.t. \\
& \quad Y_t = \bar{A}_1 i_t + \bar{A}_2 i_t^* + B \xi_t + \eta_t \text{ with probability } \mu_t \\
& \quad Y_t = \underline{A}_1 i_t + \underline{A}_2 i_t^* + B \xi_t + \eta_t \text{ with probability } 1-\mu_t \\
& \quad \mu_t, \xi_t, i_t^* \text{ given}
\end{aligned} \tag{18}$$

The constraint in the optimisation problem is the same equilibrium structural relationship (6) as before, since it holds irrespective of whether monetary policy strategies are coordinated or uncoordinated. To distinguish between i_t as a choice variable and i_t^* as given, we decompose A_t in the equilibrium structural relationship into matrices A_{1t} and A_{2t} for i_t and i_t^* respectively. Appendix A gives full details. The optimisation problem again falls in the class of linear-quadratic control problems studied by Sargent (1987), so standard solution techniques produce an optimal uncoordinated monetary policy strategy (19) for the home central bank.

$$i_t = - \left(\mu_t \bar{A}'_1 \Lambda_1 \bar{A}_1 + (1 - \mu_t) \underline{A}'_1 \Lambda_1 \underline{A}_1 \right)^{-1} \left[\begin{array}{l} \left(\mu_t \bar{A}'_1 \Lambda_1 \bar{A}_2 + (1 - \mu_t) \underline{A}'_1 \Lambda_1 \underline{A}_2 \right) i_t^* \\ + \left(\mu_t \bar{A}'_1 \Lambda_1 B + (1 - \mu_t) \underline{A}'_1 \Lambda_1 B \right) \xi_t \end{array} \right] \tag{19}$$

Optimal monetary policy strategy has the home nominal interest rate reacting linearly to shocks and the foreign nominal interest rate. The element of caution in coordinated policy is repeated here as uncertainty again dampens the reaction of the home nominal interest rate to shocks.¹² An analogous monetary policy strategy (20) is followed by the foreign central bank, making the foreign nominal interest rate a function of shocks and the home nominal interest rate.

$$i_t^* = - \left(\mu_t^* \bar{A}'_2 \Lambda_2 \bar{A}_2 + (1 - \mu_t^*) \underline{A}'_2 \Lambda_2 \underline{A}_2 \right)^{-1} \left[\begin{array}{l} \left(\mu_t^* \bar{A}'_2 \Lambda_2 \bar{A}_1 + (1 - \mu_t^*) \underline{A}'_2 \Lambda_2 \underline{A}_1 \right) i_t \\ + \left(\mu_t^* \bar{A}'_2 \Lambda_2 B + (1 - \mu_t^*) \underline{A}'_2 \Lambda_2 B \right) \xi_t \end{array} \right] \tag{20}$$

¹²We have $\mu_t \bar{A}'_1 \Lambda_1 \bar{A}_1 + (1 - \mu_t) \underline{A}'_1 \Lambda_1 \underline{A}_1 > (\mu_t \bar{A}_1 + (1 - \mu_t) \underline{A})' \Lambda_1 (\mu_t \bar{A}_1 + (1 - \mu_t) \underline{A}_1)$ as the equivalent condition to that for coordinated policy in footnote 5. The Brainard (1967) result holds as long as $\mu_t (1 - \mu_t) \underline{A}'_1 \Lambda_1 \bar{A}_1$ is not too large.

The equilibrium under passive learning with no policy coordination can now be fully specified as the structural relationship (6), monetary policy strategies (19)-(20), and the stochastic processes for A_t , ξ_t and η_t .

4.2 Active learning

When central banks follow active learning policies they internalise learning in their monetary policy strategies. To analyse such policies when there is no policy coordination we need to refine our concept of equilibrium. A natural candidate is to restrict our attention to Markov-perfect equilibria in which each central bank takes the current and future *strategies* of the other central bank as given. In practice, this translates into the home central bank taking the current foreign nominal interest rate as given, but recognising that future foreign nominal interest rates depend on future beliefs. The loss minimisation problem of the home active-learning central bank is presented in (21)

$$\begin{aligned}
& \min_{\{i_t\}} E_t \sum_{i=0}^{\infty} \beta^i (Y'_{t+i} \Lambda_1 Y_{t+i}) \\
& \quad s.t. \\
& \quad Y_{t+i} = \bar{A}_1 i_{t+i} + \bar{A}_2 i_{t+i}^* + B \xi_{t+i} + \eta_{t+i} \text{ with probability } \mu_{t+i} \\
& \quad Y_{t+i} = \underline{A}_1 i_{t+i} + \underline{A}_2 i_{t+i}^* + B \xi_{t+i} + \eta_{t+i} \text{ with probability } 1-\mu_{t+i} \\
& \quad \mu_{t+i+1} = \mathcal{B}(\mu_{t+i}, \xi_{t+i}, i_{t+i}, i_{t+i}^*, Y_{t+i}) \\
& \quad \mu_{t+i}^* = \mu_{t+i} \\
& \quad i_{t+i}^* = i_{t+i}^*(\mu_{t+i}^*, i_{t+i}, \xi_{t+i}) \\
& \quad \mu_t, \xi_t, i_t^* \text{ given}
\end{aligned} \tag{21}$$

The first two constraints are a simple restatement of the equilibrium structural relationship (6), as explained in the case of passive learning with coordination. The third constraint is the learning process by which the home central bank updates its beliefs, whereas the fourth constraint acknowledges that all learning is joint with the foreign central bank. The fifth constraint is the monetary policy strategy used by the foreign central bank to set nominal interest rates in the future. The non-linearities inherent in central bank learning preclude any

closed-form solution to the optimisation problem and we again resort to numerical techniques to characterise equilibrium. In Markov-perfect equilibrium the monetary policy strategy of one central bank depends on the monetary policy strategy of the other, so we need to iterate twin Bellman equations (21)-(22) to obtain a numerical solution. We do this sequentially until convergence of active learning policies in both central banks.¹³

$$V(\mu_t, i_t^*, \xi_t) = \min_{i_t} E_t [Y_t' \Lambda_1 Y_t + \beta V(\mu_{t+1}, i_{t+1}^*, \xi_{t+1})] \quad (22)$$

$$V^*(\mu_t^*, i_t, \xi_t) = \min_{i_t^*} E_t [Y_t' \Lambda_2 Y_t + \beta V^*(\mu_{t+1}^*, i_{t+1}, \xi_{t+1})] \quad (23)$$

The equilibrium under active learning with no policy coordination is completely specified by the structural relationship (6), monetary policy strategies that solve Bellman equations (21)-(22), and the stochastic processes for A_t , ξ_t and η_t .

4.3 Numerical example

The baseline numerical example of Section 3.3 is a natural reference point for analysing uncoordinated policies. We therefore retain the baseline parameter values in Table 1. To recap, the key feature of the baseline parameterisation is the real exchange rate switching between having supply-side and demand-side effects. We also adopt the same focus as before by looking at the reaction of nominal interest rates to a unit positive aggregate supply shock in the home country and a unit negative aggregate supply shock in the foreign country. Figure 3 presents our numerical results.

¹³An issue arises in that existence and uniqueness of the Markov-perfect equilibrium is no longer guaranteed once we move away from the linear monetary policy strategies of passive learning. Our response is to appeal to a continuity argument that a unique equilibrium still exists as long as the incentives to change policy as learning is internalised are sufficiently small. Given that the difference between passive and active learning appears marginal in our numerical solution it seems likely that this will be the case in our parameterisations of the model.

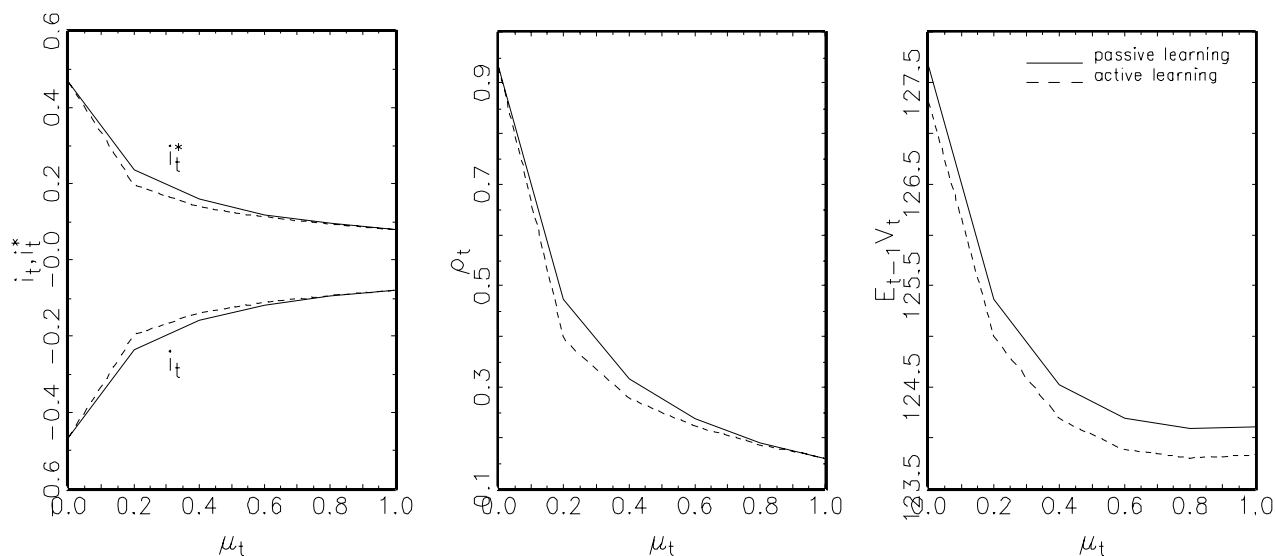


Figure 3: Policy function and unconditional value functions without policy coordination

The passive learning policy in the first panel of Figure 3 is very similar to that in Figure 2 for coordinated policies.¹⁴ The home central bank reacts to a positive shock by cutting its nominal interest rate, whilst the foreign central bank raises its nominal interest rate after a negative shock. The second panel of Figure 3 shows the resulting depreciation of the real exchange rate. As expected from the results with coordination, reactions are stronger when central banks believe the real exchange rate mostly affects the supply-side of the economy. The difference between monetary policy strategies under passive and active learning is, though, particularly striking in Figure 3. Whereas optimal coordinated monetary policy strategy was activist, we now find that optimal strategy is *cautious*. This is witnessed by a weak reactions of nominal interest rates and a smaller depreciation of the real exchange rate under active learning.

The finding that cautious monetary policy strategies can be optimal runs counter to established wisdom that policy should be activist to generate information useful for learning. If

¹⁴The similarity is partly an illusion. It is caused by focussing on the response of nominal interest rate to shocks of opposite sign in each country. Such shocks create very little conflict, as both central banks want to depreciate the real exchange rate. In contrast, if shocks have the same sign there are conflicting aims for the real exchange rate and significant differences arise in policies with and without coordination. The illusion of similarity disappears once we examine the full strategy space.

anything, it appears that optimal uncoordinated monetary policy strategies are designed to avoid generating information and learning is better retarded than promoted. This is indeed the case as the correlation between beliefs and the true values of the time-varying elasticities falls from 0.236 under passive learning to 0.216 under active learning, a deterioration of 8% in the ability of central banks to track the structure of the economy. The reason why such a seemingly counter-intuitive policy is optimal lies in the fact that central banks do not learn in isolation in the open economy, so any attempt at activist policy promotes the learning of both central banks. This is problematic when policies are not coordinated: an activist policy is desirable in that it helps a central bank to learn but undesirable in that it also helps the other central bank to learn. In our numerical example, the cost of allowing the other central bank to learn completely dominates so optimal monetary policy strategy is cautious. On balance, a central bank prefers not to learn how to manipulate the real exchange rate since in doing so the other central bank also learns and starts to manipulate the real exchange rate.¹⁵ Optimal monetary policy strategies internalise these incentives and recommend caution as a way of simultaneously slowing the learning of both central banks. The third panel of Figure 3 confirms that caution is a welfare-improving policy, with the difference in unconditional value functions representing a reduction in losses under active learning of the order of 0.25%.

4.4 Sensitivity analysis

To check the robustness of our uncoordinated policy result we subject it to the same sensitivity analysis as we did the coordinated policy result. We continue to take the short-cut of identifying the *incentives* for increasing policy activism, as this proves particularly informative when policies are uncoordinated. To proceed, we substitute passive learning policies (19)-(20) and the equilibrium structural relationship (6) into the one period return in (12) to obtain the expected return of the home central bank $\mathcal{L}^e(\mu_t, \mu_t^*)$ as a function of the beliefs of the home and foreign central banks. Writing the expected return function in terms of beliefs is useful because it admits a decomposition of incentives for activism according to which central bank

¹⁵In comparison, the problem does not arise when policies are coordinated because there is no incentive for central banks to manipulate the real exchange rate to mutually exploit each other.

is learning. The decomposition is based on the total second differential (24) of the expected return function of the home central bank. Time subscripts on beliefs have been dropped for ease of notation.

$$d^2 \mathcal{L}^e(\mu_t, \mu_t^*) = \mathcal{L}_{\mu\mu} d\mu^2 + \mathcal{L}_{\mu^*\mu^*} d\mu^{*2} + 2\mathcal{L}_{\mu\mu^*} d\mu d\mu^* \quad (24)$$

The symmetric nature of the model means that the home and foreign central banks always learn together. We therefore evaluate the convexity or concavity of the expected return function along the locus of points where $\mu_t = \mu_t^*$. Equation (25) shows that the second derivative can then be decomposed into three components. The first two terms capture convexity or concavity with respect to the separate beliefs of the home and foreign central banks, whilst the third term explicitly accounts for central banks learning together.

$$\left. \frac{d^2 \mathcal{L}^e(\mu_t, \mu_t^*)}{d\mu^2} \right|_{d\mu=d\mu^*} = \mathcal{L}_{\mu\mu} + \mathcal{L}_{\mu^*\mu^*} + 2\mathcal{L}_{\mu\mu^*} \quad (25)$$

The first line of Table 3 shows the decomposed second derivative in our baseline numerical example. The fact that $\mathcal{L}_{\mu\mu}$ is negative implies that the expected return function is concave with respect to the beliefs of the home central bank. In this respect there is an incentive to increase the level of policy activism. However, $\mathcal{L}_{\mu^*\mu^*}$ and $\mathcal{L}_{\mu\mu^*}$ are both positive so overall the expected return function is convex and the incentives for policy activism are overturned once we factor in the learning of the foreign central bank. The decomposition confirms our explanation in Section 4.3 for why optimal uncoordinated monetary policy strategies are cautious. An activist policy is desirable in that it helps the home central bank to learn ($\mathcal{L}_{\mu\mu} < 0$) but undesirable in that it also helps the foreign central bank to learn ($\mathcal{L}_{\mu^*\mu^*} > 0$). Learning together is particularly problematic ($\mathcal{L}_{\mu\mu^*} > 0$).

a_1	b_1	a_2	$\mathcal{L}_{\mu\mu}$	$\mathcal{L}_{\mu^*\mu^*}$	$\mathcal{L}_{\mu\mu^*}$	$\left. \frac{d^2 \mathcal{L}^e(\mu_t, \mu_t^*)}{d\mu^2} \right _{d\mu=d\mu^*}$
{0, 1}	{1, 0}	0.1	-0.159	0.403	0.154	0.551
{0, 1}	{1, 0}	0.05	-0.079	0.498	0.211	0.841
{0, 1}	{1, 0}	0.2	-0.248	0.273	0.081	0.187
{0, 0.25}	{1, 0}	0.1	-0.123	0.105	0.932	1.846
{0, 1}	{0.25, 0}	0.1	-0.239	0.149	0.054	0.018
{0, 0.5}	{1, 0.5}	0.1	-0.044	-0.036	0.100	0.120
{0.5, 1}	{0.5, 0}	0.1	-0.203	0.221	0.083	0.184

Table 3: Concavity of the expected return function at $\mu_t=0.5$ without policy coordination

The remaining rows of Table 3 decompose the incentives for policy activism in alternative parameterisations of the model. In all cases there is an incentive to increase policy activism if the home central bank considers its own learning in isolation, but once the learning of the foreign central bank is taken into account the optimal monetary policy strategy becomes cautious. We interpret this as evidence for the robustness of our result.

5 Conclusions

The main contribution of this paper is to show there is ambiguity in the degree of caution or activism in optimal monetary policy strategies for the open economy, a result in sharp contrast to the closed-economy literature which unambiguously recommends activist policy. The ambiguity arises because there is more than one central bank in the open economy and monetary policy strategies may or may not be coordinated. If policies are coordinated then we recover the closed-economy result that strategies should be activist to speed up learning: it is better to coordinate policy between informed than uninformed central banks. If policies are uncoordinated then the converse is true and strategies should be cautious to slow down learning: a central bank prefers to be in noncooperative equilibrium with uninformed than informed central banks. Our result is derived in a context that stresses the real exchange rate

channel of the monetary transmission mechanism, so in situations where the open economy aspect is less prominent we would expect activist policy strategies to dominate in line with the closed-economy literature. However, our result is still valid in that incentives for activist policy will be enhanced or attenuated in the open economy according to whether policies are coordinated or uncoordinated.

Our results also shed new light on the debate about the gains from international policy coordination. One of the seminal contributions in this literature is Ghosh and Masson (1991), who argue that learning can restore the benefits to international policy coordination that earlier work by Frankel and Rockett (1988) argued were lost if there is model uncertainty. The learning they consider is purely passive so our contribution takes the argument a step further and asks how *active* learning affects the gains to international policy coordination. Our results suggest two effects working in opposite directions. Firstly, active learning increases the gains to coordination by improving the “good” outcome when policies are coordinated. Active learning prompts coordinated central banks to follow activist strategies, making them learn faster and putting them in a better position to stabilise their economies. Secondly, active learning reduces the gains to coordination by improving the “bad” outcome when policies are not coordinated. Active learning in the absence of coordination gives central banks an incentive to follow cautious strategies, thereby slowing learning and avoiding conflict. The overall effect on the gains to international policy coordination is likely to be ambiguous as the “good” coordinated outcome gets better and the “bad” uncoordinated outcome gets less bad.¹⁶

¹⁶In our baseline numerical example the improvement in the unconditional value function with coordination (the third panel of Figure 2) is slightly less than the improvement without coordination (the third panel of Figure 3), so active learning marginally reduces the gains to policy coordination. We report this result for illustration purposes and make no claim about its robustness.

A Equilibrium structural relationship

The structure of the open economy (1)-(5) describes equilibrium relationships linking real interest rates to the real exchange rate and output and inflation in each country. By adding Fisher equations and substituting out for the real exchange rate, we obtain an alternative equilibrium structural relationship in the form of equation (A.1), where observations Y_t of output and inflation $(y_t \ \pi_t \ y_t^* \ \pi_t^*)'$ are a function of nominal interest rates $R_t = (i_t \ i_t^*)'$, shocks $\xi_t = (e_t \ u_t \ e_t^* \ u_t^*)'$ and measurement errors η_t . Expectation terms drop out as they are zero in i.i.d. equilibrium with no systematic biases in policy.

$$Y_t = A_t R_t + B \xi_t + \eta_t \quad (\text{A.1})$$

The elements of matrix A_t are nonlinear functions of the elasticities $(b_{1t}, b_2, a_{1t}, a_2, a_3)$ in the aggregate supply and demand curves. A_t switches between \bar{A} and \underline{A} as the time-varying elasticities switch between (\bar{b}_1, \bar{a}_1) and $(\underline{b}_1, \underline{a}_1)$. The nonlinear functions in A_t are given by equation (A.2).

$$A_t = \begin{pmatrix} \frac{a_{1t}(a_3-1)-a_2}{1-a_3^2} & -\frac{a_{1t}(a_3-1)+a_2a_3}{1-a_3^2} \\ \frac{a_{1t}(a_3-1)-a_2-b_{1t}(1-a_3^2)}{b_2(1-a_3^2)} & -\frac{a_{1t}(a_3-1)-a_2a_3-b_{1t}(1-a_3^2)}{b_2(1-a_3^2)} \\ -\frac{a_{1t}(a_3-1)+a_2a_3}{1-a_3^2} & \frac{a_{1t}(a_3-1)-a_2}{1-a_3^2} \\ -\frac{a_{1t}(a_3-1)-a_2a_3-b_{1t}(1-a_3^2)}{b_2(1-a_3^2)} & \frac{a_{1t}(a_3-1)-a_2-b_{1t}(1-a_3^2)}{b_2(1-a_3^2)} \end{pmatrix} \quad (\text{A.2})$$

The matrix B depends only on time-invariant elasticities (b_2, a_3) and is described by equation (A.3).

$$B = \begin{pmatrix} \frac{1}{1-a_3^2} & \frac{a_3}{1-a_3^2} & 0 & 0 \\ \frac{1}{b_2(1-a_3^2)} & \frac{a_3}{b_2(1-a_3^2)} & b_2^{-1} & 0 \\ \frac{a_3}{1-a_3^2} & \frac{1}{1-a_3^2} & 0 & 0 \\ \frac{a_3}{b_2(1-a_3^2)} & \frac{1}{b_2(1-a_3^2)} & 0 & b_2^{-1} \end{pmatrix} \quad (\text{A.3})$$

For the analysis with uncoordinated policies, it is useful to decompose the equilibrium structural relationship further to separate the roles of home and foreign nominal interest rate.

This is done in equation (A.4).

$$Y_t = A_{1t}i_t + A_{2t}i_t^* + B\xi_t + \eta_t \quad (\text{A.4})$$

Matrices A_{1t} and A_{2t} are defined in equation (A.5) and will switch between (\bar{A}_1, \bar{A}_2) and $(\underline{A}_1, \underline{A}_2)$ as the underlying time-varying elasticities switch between (\bar{b}_1, \bar{a}_1) and $(\underline{b}_1, \underline{a}_1)$.

$$A_{1t} = \begin{pmatrix} \frac{a_{1t}(a_3-1)-a_2}{1-a_3^2} \\ \frac{a_{1t}(a_3-1)-a_2-b_{1t}(1-a_3^2)}{b_2(1-a_3^2)} \\ -\frac{a_{1t}(a_3-1)+a_2a_3}{1-a_3^2} \\ -\frac{a_{1t}(a_3-1)-a_2a_3-b_{1t}(1-a_3^2)}{b_2(1-a_3^2)} \end{pmatrix} \quad A_{2t} = \begin{pmatrix} -\frac{a_{1t}(a_3-1)+a_2a_3}{1-a_3^2} \\ -\frac{a_{1t}(a_3-1)-a_2a_3-b_{1t}(1-a_3^2)}{b_2(1-a_3^2)} \\ \frac{a_{1t}(a_3-1)-a_2}{1-a_3^2} \\ \frac{a_{1t}(a_3-1)-a_2-b_{1t}(1-a_3^2)}{b_2(1-a_3^2)} \end{pmatrix} \quad (\text{A.5})$$

B Numerical solution of active learning policy

The numerical solution to the coordinated active learning policy is obtained by iterating on the Bellman equation (17). This requires expressions for the expected one-period return $E_t(Y_t' \Lambda Y_t)$ and continuation value $E_t V_{t+1}$ as functions of nominal interest rates R_t . The expected one-period return and continuation value can be calculated from the equilibrium structural relationship (6) and equation (B.1) respectively, where future beliefs μ_{t+1} have been substituted out using the central bank learning process (11).

$$E_t V_{t+1} = E_t V(\mathcal{B}(\mu_t, \xi_t, R_t, Y_t), \xi_{t+1}) \quad (\text{B.1})$$

The expectation in equation (B.1) is formed prior to observing output, inflation and next period shocks by evaluating the double integral in equation (B.2).

$$E_t V_{t+1} = \int \int V(\mathcal{B}(\mu_t, \xi_t, R_t, Y_t), \xi_{t+1}) f(Y_t | \mu_t, \xi_t, R_t) f(\xi_{t+1}) dY_t d\xi_{t+1} \quad (\text{B.2})$$

$f(\xi_{t+1})$ is the distribution of shocks ξ_{t+1} and $f(Y_t | \mu_t, \xi_t, R_t)$ is the predicted distribution of Y_t . They have independent multivariate distributions, normal and a mixture of normals respectively, as described by equations (B.3) and (B.4).

$$f(\xi_{t+1}) = N \left[\begin{array}{c} \left(\begin{array}{c} 0 \\ 0 \\ 0 \\ 0 \end{array} \right) \\ ; \\ \left(\begin{array}{cccc} \sigma_e^2 & 0 & 0 & 0 \\ 0 & \sigma_u^2 & 0 & 0 \\ 0 & 0 & \sigma_e^2 & 0 \\ 0 & 0 & 0 & \sigma_u^2 \end{array} \right) \end{array} \right] \quad (\text{B.3})$$

$$f(Y_t |_{\mu_t, \xi_t, R_t}) = \mu_t N [\overline{A}R_t + B\xi_t; \Sigma_\eta] + (1 - \mu_t) N [\underline{A}R_t + B\xi_t; \Sigma_\eta] \quad (\text{B.4})$$

The computational algorithm starts by defining a grid of points in the state space (μ_t, ξ_t) .¹⁷ The gridpoints for beliefs μ_t are distributed uniformly across the interval $[0,1]$, but gridpoints for the shocks ξ_t are bunched around zero according to a cosine weighting function to increase accuracy. For each gridpoint, we take starting values for nominal interest rates R_t and the value function V_t from those under the passive learning policy. An iteration of the Bellman equation involves passing through the grid point by point. At each gridpoint, the nominal interest rate choices are reoptimised by minimising the right hand side of the Bellman equation (17), using the equilibrium structural relationship (6) to calculate the expected one-period return and equations (B.2)-(B.4) to calculate the expected continuation value. Numerical evaluation of the expected continuation value requires linear interpolation of adjacent gridpoints to solve the double integral in (B.2). The reoptimised values of R_t and V_t are then assigned to the gridpoint. An iteration of the Bellman equation is complete when the nominal interest rate choices and value function have been updated for each gridpoint. Repeated application of the iterative procedure converges to the active learning policy. We accept convergence when the values associated with each gridpoint change by less than 0.0001 between successive iterations. When optimising the nominal interest rate choices at each gridpoint we use a convergence tolerance of 0.00001.

¹⁷The state space expands to (μ_t, i_t^*, ξ_t) when policies are uncoordinated.

References

- [1] Balvers, R.J. and Cosimano, T.F., 1994, Inflation Variability and Gradualist Monetary Policy, *Review of Economic Studies* 61, 721-738.
- [2] 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.
- [3] Bertocchi, G. and Spagat, M., 1993, Learning, Experimentation and Monetary Policy, *Journal of Monetary Economics* 23, 169-183.
- [4] Blinder, A., 1998, *Central Banking in Theory and Practice*, MIT Press.
- [5] Brainard, W., 1967, Uncertainty and the Effectiveness of Policy, *American Economic Review Papers and Proceedings* 57, 411-425.
- [6] DeGroot, M.H., 1962, Uncertainty, Information, and Sequential Experiments, *The Annals of Mathematical Statistics* 44, 404-419.
- [7] Ellison, M. and Vilmunen, J., 2005, A simple approach to identifying the incentives for policy experimentation, *Economics Letters* 86, 167-172.
- [8] Frankel, J., and Rockett, K., 1988, International Macroeconomic Policy Coordination when Policymakers Do Not Agree on the True Model, *American Economic Review* 78, 318-340.
- [9] Ghosh, A.R. and Masson, P.R., 1991, Model Uncertainty, Learning, and the Gains from Coordination, *American Economic Review* 81, 465-479.
- [10] Kiefer, N.M. and Nyarko, Y., 1989, Optimal Control of an Unknown Linear Process with Learning, *International Economic Review* 30, 571-586.
- [11] Sargent, T.J., 1987, *Dynamic Macroeconomic Theory*, Harvard University Press.
- [12] Svensson, L.E.O., 1999, Price Level Targeting versus Inflation Targeting: A Free Lunch?, *Journal of Money, Credit and Banking* 31, 277-295.

- [13] Walsh, C., 2003, *Monetary Theory and Policy*, MIT Press.
- [14] 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.
- [15] Wieland, V., 2000b, Monetary Policy, Parameter Uncertainty and Optimal Learning, *Journal of Monetary Economics* 46, 199-228.