DISCUSSION PAPER SERIES

No. 4135

THE LEARNING COST OF INTEREST RATE REVERSALS

Martin Ellison

INTERNATIONAL MACROECONOMICS

Centre for Economic Policy Research

www.cepr.org

ISSN 0265-8003

THE LEARNING COST OF INTEREST RATE REVERSALS

Martin Ellison, University of Warwick and CEPR

Discussion Paper No. 4135
December 2003

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. 4135

December 2003

ABSTRACT

The Learning Cost of Interest Rate Reversals*

In this Paper, we suggest a new motivation for why central banks appear averse to reversing recent changes in their interest rate. We show, in a standard monetary model with forward-looking expectations, data uncertainty and parameter uncertainty, that there is a learning cost associated with interest rate reversals. A policy that frequently reverses the interest rate makes it more difficult for the central bank and private agents to learn the key parameters of the model. Optimal monetary policy internalizes this learning cost and therefore has a lower number of interest rate reversals. The incentive to reduce the number of interest rate reversals is in addition to the optimal policy inertia created by the presence of forward-looking expectations and uncertainty in the model.

JEL Classification: D83, E52 and E58
Keywords: interest rate smoothing, learning and monetary policy

Martin Ellison
Department of Economics
University of Warwick
Coventry
CV4 7AL
Tel: (44 24) 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

*This Paper is produced as part of a CEPR Research Network on ‘The Analysis of International Capital Markets: Understanding Europe's Role in the Global Economy’, funded by the European Commission under the Research Training Network Programme (Contract No: HPRN-CT-1999-00067). and as part of a CEPR Research Network on ‘Specialization Versus Diversification: The Microeconomics of Regional Development and the Spatial Propagation of Macroeconomic Shocks in Europe’, funded by the European Commission under the Research Training Network Programme (Contract No: HPRN-CT-2000-00072). I would like to thank John Carlson, Neil Rankin, Jouko Vilmunen and Volker Wieland for helpful remarks and suggestions. The comments of seminar participants at DIW (Berlin), Bank of England, Royal Holloway (University of London), University of Southampton, London School of Economics, University of Essex, University of York and the ‘Heterogeneous Information and Modeling of Monetary Policy’ CEPR/Bank of Finland workshop, October 2003, are also gratefully acknowledged. Giacomo Carboni provided excellent research assistance.

Submitted 22 October 2003

1 Introduction

One of the most striking features of central bank behaviour is their apparent reluctance to reverse recent interest rate changes. In many countries, monetary policy tends to be characterised by runs of successive interest rate movements in the same direction, with only rare reversals during which the interest rate moves in the opposite direction to recent changes. The Federal Reserve is particularly averse to interest rate reversals. In the US, it is approximately ten times more likely that a rise in the interest rate will be followed by another rise, rather than a fall, in the interest rate. A significant literature on interest rate smoothing has developed in an attempt to rationalise the lack of reversals in central bank policy. The explanations that have been subject to formal analysis are the presence of forward-looking expectations, Woodford (1999), data uncertainty, Orphanides (1998), and parameter uncertainty, Sack (2000).[1]

The purpose of this paper is to suggest that learning provides an additional motivation for central banks to refrain from interest rate reversals. We show that, if expectations are forward looking and there is data and parameter uncertainty, then it is more difficult for the central bank and private agents to learn the key features of the monetary transmission mechanism if there are frequent reversals in monetary policy. Optimal monetary policy under learning is therefore characterised by less interest rate reversals, so that learning is promoted and welfare is improved. The incentive for interest rate smoothing due to learning is in addition to the optimal policy inertia

[1] The empirical evidence for these and other explanations is discussed in Sack and Wieland (2000).

created by the presence of forward-looking expectations and uncertainty in the model.

To establish our result, we allow the central bank and private agents to learn the parameters of a standard monetary model (the New-Keynesian framework of, *inter alia*, Clarida, Gali and Gertler (1999), McCallum and Nelson (1997) and Rotemberg and Woodford (1997)). In contrast to the adaptive learning literature associated with Evans and Honkapohja (2001), we assume that both the central bank and private agents learn rationally using all available information. This has the advantage that we can derive optimal monetary policy, assuming either that the central bank ignores learning (passive learning) or that the central bank internalises learning (active learning), without being subject to the Lucas critique. We solve the model by simulation using the parameterised expectations method of den Haan and Marcet (1990) to capture the non-linearities inherent in learning. Our results show that the active learning policy which internalises learning has less interest rate reversals than the passive learning policy that ignores learning.

The plan of the paper is as follows. In Section 2, we outline our model and derive the first order conditions for the optimal policy under passive and active learning. Section 3 discusses our calibration and solution method. The results are presented in Section 4, which examines the nature of optimal monetary policy under passive and active learning. A final section concludes.

2 Analytical framework

To obtain our results, we introduce uncertainty and learning into a standard optimising model of inflation and output determination. We adopt the forward-looking analytical framework of Woodford (1999), which can be interpreted as the log-linearised equilibrium conditions of a simple intertemporal general equilibrium model with sticky prices. A role for learning is introduced by assuming that the central bank and private agents only have imprecise estimates of the structural parameters of the model. These estimates are then updated through learning as new information is received. In what follows, we focus on the role of learning and only offer a mathematical outline of the rest of the model. Readers wanting a more detailed economic explanation of the framework should consult Woodford (*ibid.*).

2.1 Structure of the model

The structure of the model consists of an intertemporal IS equation (1) and an aggregate supply equation (2). These equations show the structural relationship between the output gap, x_t, inflation, π_t, and the short-term nominal interest rate, r_t. The short-term nominal interest rate is assumed to be the instrument of monetary policy.

$$x_t = E_t x_{t+1} - \sigma^{-1}(r_t - r_t^n - E_t \pi_{t+1}) \tag{1}$$

$$\pi_t = \kappa_t x_t + \beta E_t \pi_{t+1} + \varepsilon_t \tag{2}$$

The term r_t^n is Wicksell's natural rate of interest and corresponds to the equilibrium real rate of interest that would prevail if prices in the economy were flexible. It is given exogenously by equation (3). The task of monetary policy in this model is to respond to (possibly persistent) shocks in the Wicksellian natural rate.

$$r_{t+1}^{n} = \rho r_{t}^{n} + \epsilon_t \tag{3}$$

To introduce parameter uncertainty and a motivation for learning in the model, we assume that the structural parameter κ_t in equation (2) is time-varying but cannot be observed directly by either the central bank or private agents. Fundamentally, the parameter κ_t is assumed to follow a two-state hidden Markov process, switching between high and low values with probability $1-\gamma$. Assuming $\gamma > 0.5$, the probability of switching values is lower than the probability of continuing with the same value so the process generating κ_t has persistence.[2] However, since κ_t is unobservable, the central bank and private agents have to infer its current value of the basis of past observations of the output gap and inflation. The inference problem is complicated by the presence of an unobserved measurement error in inflation, ε_t, which corresponds to data uncertainty and is normally distributed with mean zero and standard deviation σ_ε.

2.2 Beliefs

The central bank and private agents are unable to observe the structural parameter κ_t directly and so have to form a belief about whether it is currently high or low. Since the central bank and private agents have the same

[2] The assumption of a persistent two-state Markov process can be considered as a stylised representation of other persistent processes such as an AR(1).

information, their beliefs always coincide and there is no scope for asymmetry. Under symmetry, the beliefs of the central bank and private agents can be summarised by a single variable, $p_t = P(\kappa_t = \kappa_h)$, the belief at time t that the structural parameter κ_t currently has a high value. If $p_t = 1$ then there is complete certainty that $\kappa_t = \kappa_h$. Similarly, $p_t = 0$ implies certainty that $\kappa_t = \kappa_l$.

2.3 Learning

The beliefs of the central bank and private agents are updated in the model as new information is received. To learn, the central bank and private agents must assess whether the observed output gap and inflation are most consistent with a high or low value for the structural parameter κ_t. Equations (4) and (5) show the predicted distribution of inflation, conditional on the value of the structural parameter κ_t. The central bank and private agents have to infer whether observed inflation is most likely to have come from distribution (4) or (5).

$$\pi_t\,|_{\kappa_h} \sim N[\kappa_h x_t + \beta E_t \pi_{t+1}; \sigma_\varepsilon] \tag{4}$$

$$\pi_t\,|_{\kappa_l} \sim N[\kappa_l x_t + \beta E_t \pi_{t+1}; \sigma_\varepsilon] \tag{5}$$

The assumption that there are only two possible values for κ_t, coupled with exogenous Markov switching, means that the updating of beliefs takes a particularly simple form. Under rational learning, the central bank and private agents use Bayes rules to update their beliefs in the light of new information. Equation (6) shows how initial beliefs p_t are updated to p_t^+ at the end of the period, after the realisation of π_t. With Bayesian learning, p_t^+ depends on the relative probability of observing inflation π_t under the high and low κ_t cases.

$$p_t^+ = \frac{p_t P(\pi_t |_{\kappa_h})}{p_t P(\pi_t |_{\kappa_h}) + (1 - p_t) P(\pi_t |_{\kappa_l})} \tag{6}$$

p_t^+ represents the optimal inference of the current value of the structural parameter κ_t, given the realisation of inflation π_t. The central bank and private agents then make a prediction, p_{t+1}, of which value of κ_t will apply in the next period, taking into account the probability that κ_t may switch in the meantime. Equation (7) shows how this prediction is calculated as a weighted average of the probability of κ_t being high and remaining high and the probability that κ_t was low but switches to being high for the next period.

$$p_{t+1} = p_t^+ \gamma + (1 - p_t^+)(1 - \gamma) \tag{7}$$

Equations (6) and (7), together with the predicted distributions (4) and (5), define a non-linear equation (8) for updating beliefs. Substituting for inflation using the aggregate supply equation (2), updated beliefs can be written as a function of the initial belief, the output gap and the unobserved measurement error.

$$\begin{aligned} p_{t+1} &= \frac{\gamma p_t e^{-\frac{1}{2}\left(\frac{\kappa_t x_t + \varepsilon_t - \kappa_h x_t}{\sigma_\varepsilon}\right)^2} + (1-\gamma)(1-p_t) e^{-\frac{1}{2}\left(\frac{\kappa_t x_t + \varepsilon_t - \kappa_l x_t}{\sigma_\varepsilon}\right)^2}}{p_t e^{-\frac{1}{2}\left(\frac{\kappa_t x_t + \varepsilon_t - \kappa_h x_t}{\sigma_\varepsilon}\right)^2} + (1-p_t) e^{-\frac{1}{2}\left(\frac{\kappa_t x_t + \varepsilon_t - \kappa_l x_t}{\sigma_\varepsilon}\right)^2}} \\ &= \mathcal{B}(p_t, x_t, \varepsilon_t) \end{aligned} \tag{8}$$

2.4 Central bank objective function

Following Woodford (1999), we assume that the central bank minimises a loss function of the form (9). The discount factor satisfies $0 < \beta < 1$ and the period-by-period loss function is defined by equation (10).

$$W = E_0 \left(\sum_{t=0}^{\infty} \beta^t \mathcal{L}_t \right) \tag{9}$$

$$\mathcal{L}_t = \pi_t^2 + \lambda_x x_t^2 + \lambda_r r_t^2 \tag{10}$$

The objective function quadratically penalises the central bank for any deviations in inflation, the output gap and the short-run nominal interest rate from their long-run levels. All long-run levels are normalised to zero and λ_x, λ_r represent the weights placed on output gap and short-run nominal interest deviations relative to inflation deviations. The inclusion of a penalty for deviations of the interest rate from target reflects potential problems with volatile interest rates. As Rotemberg and Woodford (1997) show, minimising a loss function such as (10) corresponds to maximising the welfare of the representative agent if interest rate volatility increases the probability that the interest rate will be distortionary or constrained by its zero bound. Note that the presence of a term in the volatility of interest rates does *not* imply an aversion to interest rate reversals on the part of the central bank. The penalty is on volatility in the level of interest rates and not on the volatility of interest rate changes.[3]

[3] See Woodford (1999) for more discussion on this point.

2.5 Optimal policy

Optimal monetary policy requires the central bank to commit to a dynamic path for short-term nominal interest rate to minimise the present discounted value of expected current and future losses. The policy will not generally be time consistent. The minimisation problem of the central bank is shown in equations (11)-(16). The central bank minimises the expected loss defined by equations (9) and (10), subject to the IS curve (1) and the aggregate supply curve (2). To ease the mathematical solution of the model, we introduce the notation π_t^h and π_t^l to denote the levels of inflation generated by the predicted distributions (4) and (5). In other words, π_t^h is the level of inflation conditional on the structural parameter κ_t taking the value κ_h. Similarly, π_t^l is the level of inflation conditional on $\kappa_t = \kappa_l$. With this notation, the aggregate supply curve constraint can be divided into two components, (13) and (14), and there is an additional constraint (15) on the consistency of expectations.

$$\min E_0 \left(\sum_{t=0}^{\infty} \beta^t [\pi_t^2 + \lambda_x {x_t}^2 + \lambda_r r_t^2] \right) \tag{11}$$

$$s.t.$$

$$x_t = E_t x_{t+1} - \sigma^{-1}(r_t - r_t^n - E_t \pi_{t+1}) \tag{12}$$

$$\pi_t^h = \kappa_h x_t + \beta E_t \pi_{t+1} + \varepsilon_t \tag{13}$$

$$\pi_t^l = \kappa_l x_t + \beta E_t \pi_{t+1} + \varepsilon_t \tag{14}$$

$$E_t \pi_t = p_t E_t \pi_t^h + (1 - p_t) E_t \pi_t^l \tag{15}$$

$$p_{t+1} = \mathcal{B}(p_t, x_t, \varepsilon_t) \tag{16}$$

The final constraint (16) is the learning mechanism by which the central bank and private agents update their beliefs when new information is received. It is fully specified by equation (8) of Section 2.3. In calculating optimal policy, we distinguish between two alternative assumptions concerning how the central bank views the learning constraint. Under a *passive learning* policy, the central bank ignores the learning constraint (16) and the problem reduces to minimising (11) subject to (12)-(15). Although the central bank and private agents do learn with a passive learning policy, the central bank does not consider that its policy actions are instrumental in determining learning. In contrast, an *active learning* policy does internalise the learning constraint (16) and the problem of the central bank is to minimise (11) subject to (12)-(16). The active learning policy takes into account that policy actions have an influence on learning.

The first order conditions for optimal passive and active learning policies can be derived using standard Lagrangian techniques. Full details of their derivation are contained in Appendix A. In what follows, we restrict ourselves to the economic interpretation of the first order conditions and a discussion of the differences between the optimal passive and active learning policies.

2.5.1 Passive learning

The optimal passive learning policy ignores the learning constraint and so only satisfies the set of first order conditions (17)-(21), obtained in Appendix A by optimising with respect to $\pi_t, \pi_t^h, \pi_t^l, x_t$ and r_t. $\phi_{1t} \cdots \phi_{4t}$ are the Lagrange multipliers associated with the IS curve, the conditional aggregate supply curves and the expectations consistency constraint respectively. We only consider bounded solutions to these equations and so avoid the issue of transversality conditions. In general, the optimal passive learning policy is not time-consistent. It involves elements of commitment due to the presence of lagged multiplier terms. We follow Giannoni and Woodford (2002) and restrict ourselves to policies that are optimal from a 'timeless perspective'.

$$0 = E_t \pi_t - (\beta\sigma)^{-1}\phi_{1t-1} - \phi_{2t-1} - \phi_{3t-1} + \phi_{4t} \tag{17}$$

$$0 = \phi_{2t} - p_t \phi_{4t} \tag{18}$$

$$0 = \phi_{3t} - (1 - p_t)\phi_{4t} \tag{19}$$

$$0 = \lambda_x x_t + \phi_{1t} - \beta^{-1}\phi_{1t-1} - \kappa_h \phi_{2t} - \kappa_l \phi_{3t} \tag{20}$$

$$0 = \lambda_r r_t + \sigma^{-1}\phi_{1t} \tag{21}$$

The first order conditions (17)-(21) are equivalent to those derived by Woodford (1999), except for additional terms due to uncertainty surrounding the structural parameter κ_t. If $p_t = 1$ or $p_t = 0$ then κ_t is know with certainty and our first order conditions exactly coincide. We therefore refer readers to Woodford (*ibid.*) for more discussion on the economic intuition of the system of equations (17)-(21). The passive learning policy that is defined by our first order conditions arises in the Woodford (*ibid.*) model if there is parameter uncertainty.

2.5.2 Active learning

The optimal active learning policy internalises the learning constraint so there is an additional equation and Lagrange multiplier in the first order conditions (22)-(27), the full derivation of which appears in Appendix A.

ϕ_{5t}^{μ} is the Lagrange multiplier associated with the learning constraint. Compared to the conditions for the optimal passive learning policy, the optimal active learning policy has extra terms in equation (25), the first order condition for the output gap. The extra terms are superscripted to indicate that they are the value of a variable conditional on beliefs. For example, x_{t+1}^{s} is the value of the output gap at time t+1, conditional on beliefs p_{t+1} being equal to s.

$$0 = E_t \pi_t - (\sigma\beta)^{-1}\phi_{1t-1} - \phi_{2t-1} - \phi_{3t-1} + \phi_{4t} \tag{22}$$

$$0 = \phi_{2t} - p_t \phi_{4t} \tag{23}$$

$$0 = \phi_{3t} - (1 - p_t)\phi_{4t} \tag{24}$$

$$\begin{aligned} 0 &= \lambda_x x_t + \phi_{1t} - \beta^{-1}\phi_{1t-1} - \kappa_h \phi_{2t} - \kappa_l \phi_{3t} \\ &+ E_t \sum_s [\phi_{1t} x_{t+1}^{s} + (\sigma^{-1}\phi_{1t} + \beta\phi_{2t} + \beta\phi_{3t})\pi_{t+1}^{s} + \tfrac{\beta}{2}\phi_{5t+1}^{s}] \frac{\partial P(\mathcal{B}(p_t, x_t, \varepsilon_t) = s)}{\partial x_t} \end{aligned} \tag{25}$$

$$0 = \lambda_r r_t + \sigma^{-1}\phi_{1t} \tag{26}$$

$$0 = E_t(\pi_t^{\mu})^2 + \lambda_x (x_t^{\mu})^2 + \lambda_r (r_t^{\mu})^2 + \beta E_t \sum_s \phi_{5t+1}^{s} P(\mathcal{B}(\mu, x_t^{\mu}, \varepsilon_t) = s) - \phi_{5t}^{\mu} \tag{27}$$

To gain an insight into the role of learning in the first order conditions, it is useful to rewrite the additional first order condition in the form of equation (28). Written in this way, the Lagrange multiplier ϕ_{5t}^{μ} has an intuitive economic interpretation.

$$\phi_{5t}^{\mu} = E_t[(\pi_t^{\mu})^2 + \lambda_x (x_t^{\mu})^2 + \lambda_r (r_t^{\mu})^2] + \beta E_t \sum_s \phi_{5t+1}^{s} P(\mathcal{B}(\mu, x_t^{\mu}, \varepsilon_t) = s) \tag{28}$$

ϕ_{5t}^{μ} is equal to the sum of the expected period-by-period loss and the discounted value of ϕ_{5t+1} expected for the next period, both expectations being conditional on current beliefs satisfying $p_t = \mu$. The additional first order condition is analogous to a value function, with a one-period return and a continuation value. It measures the value to the central bank of having a belief μ at time t.

The analogy to a value function helps explain the presence of the extra terms in the first order condition (25) for the output gap. At the margin, a change in the output gap x_t will affect the expected distribution of future beliefs through the process of learning. The extra terms therefore capture the fact that changes in x_t will have an effect on forward-looking expectations of the output gap, inflation and continuation value because the changes affect the expected distribution of future beliefs.[4]

3 Calibration and solution

We use the parameter values of Woodford (1999) as a baseline calibration of the model. The only additional parameters are the two possible values for the structural parameter κ_t, the probability $1-\gamma$ of switching between them, and the standard deviation of the measurement error, σ_ε. We set $\kappa_l = 0.024$ following Woodford (*ibid.*) and allow $\kappa_h = 0.036$. This implies that the sacrifice ratio is 50% lower when κ_t takes its high value. The probability of switching is calibrated at $1 - \gamma = 0.025$, so the average duration over which

[4] See DeGroot (1962) for a discussion of how the incentives to change the output gap x_t are determined by whether forward-looking expectations of the output gap, inflation and continuation value are convex or concave with respect to expected future beliefs.

κ_t remains unchanged is $1/0.025 = 40$ months. The choice for the standard deviation of the measurement error, σ_ε, is a compromise to preserve the incentives for active learning in the model. If σ_ε is too low, the passive learning policy is already very informative. If σ_ε is too high, even the active learning policy is not very informative. In either case, there is little incentive to manipulate learning and the passive and active learning policies are very similar. Table 1 reports the full calibration of our model. With respect to the persistence, ρ, of shocks to the Wicksellian natural rate of interest, we report our main results with the value given but also consider other values in sensitivity analysis.

Structural parameters	β	0.997
	σ	0.157
	κ_h	0.036
	κ_l	0.024
Shock processes	γ	0.975
	$sd(r^n)$	3.72
	σ_ε	0.25
	ρ	0.35
Loss function	λ_x	0.047
	λ_r	0.233

Table 1: Calibrated parameter values

To solve the model, we utilise the parameterised expectations algorithm of den Haan and Marcet (1990), which has the advantage of being able to capture the non-linearities intrinsic to rational Bayesian learning.[5] We specify a second-order polynomial in the state and co-state variables and include cross products to obtain an 'accurate solution'. To test for accuracy, we use the test statistic proposed by den Haan and Marcet (1994). For both the passive and active learning policies, the number of simulations in the upper and lower tails is in the range $2.5\% - 7.5\%$. Given the stringency of the test, we interpret this as strong support for the accuracy of our solution.

4 Results

In Sections 4.1 to 4.4 we report the results of simulating our model under the optimal passive and active learning policies. In each case, we adopt the baseline calibration presented in Section 3. In Section 4.5, we check the sensitivity of our results to changes in the calibrated parameters.

4.1 Persistence

The persistence properties of the passive and active learning policies are shown in Figure 1, which plots the estimated first-order autocorrelation coefficient for the short-term nominal interest rate, as a function of beliefs.[6]

[5] A standard iterative approach such as the one suggested by Wieland (2000) cannot be used because our commitment policies are not time-consistent and so do not satisfy Bellman's dynamic programming principle. A GAUSS program to solve our model and test its accuracy is available from the author on request.

[6] The estimates in Figure 1 are obtained by applying a standard Nadaraya-Watson kernel estimator of conditional autocorrelation to simulations of the model under the passive and active learning policies. We use a Gaussian kernel with bandwidth equal to 0.5.

For the passive learning policy, the short-term nominal interest rate shows considerable persistence, with an autocorrelation coefficient around 0.69 for all beliefs. This is partly due to persistence in the natural rate of interest itself (ρ=0.35 in the simulations), but also because the presence of forward-looking expectations in the model creates an incentive for inertia in policy, as emphasised by Woodford (1999).

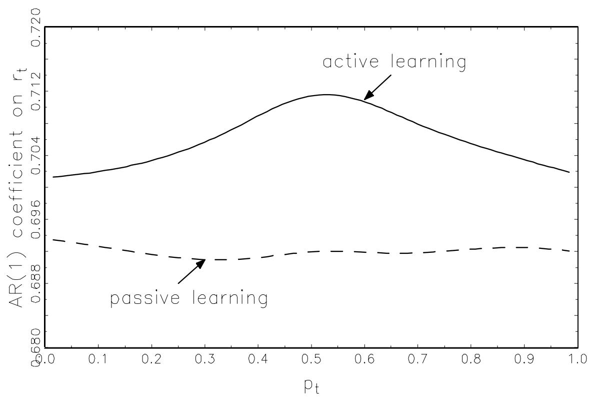

Figure 1: Autocorrelation of the short-term nominal interest rate

Turning to the active learning policy, we observe that the short-term nominal interest rate is more persistent than it was under the passive learning policy at every level of beliefs. The effect is amplified when beliefs are around 0.5, the value associated with maximum uncertainty. The increased persistence induced by the active learning policy is the central result of our paper. In the simulations, we find that internalising learning creates an additional incentive for persistence in the short-term nominal interest rate, especially when uncertainty is high. The active learning policy therefore implies more persistence in the short-term nominal interest rate and a greater degree of interest rate smoothing. The remainder of our paper is devoted to developing the economic intuition and assessing the robustness of this result.

4.2 Impulse response functions

Increased persistence in the short-term nominal rate of interest under active learning is also apparent in the reaction of the central bank to shocks in the Wicksellian natural rate of interest. Figure 2 shows the response of the short-term nominal interest rate to a one standard deviation positive shock in the natural rate, when uncertainty is initially at its highest.[7]

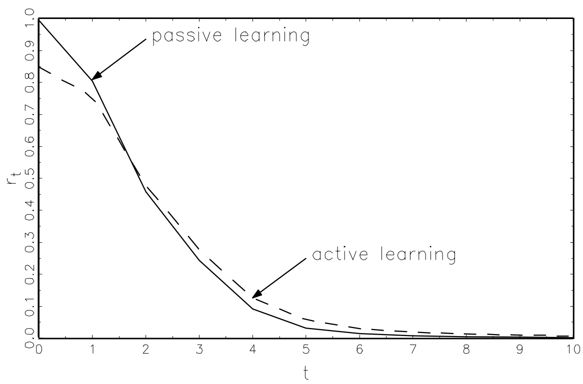

Figure 2: Response of short-term nominal interest rate to a shock in the natural rate of interest.

The response of the short-term nominal interest rate under the active learning policy is initially muted when compared to the passive learning policy. However, the response is more persistent, remaining above baseline for longer under active than passive learning.

[7] We use the same kernel estimator as in Figure 1 to estimate the impulse response function conditional on beliefs being equal to 0.5.

4.3 Beliefs and learning

To understand the economic intuition behind our result, it is useful to solve the IS curve (1) forward to obtain an expression showing how the current output gap is determined. According to equation (29), the output gap x_t in our model is a function of the sum of current and future expected deviations of the real short-term interest rate, $r_{t+j} - E_t\pi_{t+j+1}$, from the natural rate, r^n_{t+j}. In other words, the output gap is determined by the expected difference between the long-term real interest rate and the long-term natural rate.

$$x_t = -\sigma \sum_{j=0}^{\infty} (r_{t+j} - E_t\pi_{t+j+1} - r^n_{t+j}) \tag{29}$$

The persistence in short-term nominal interest rates observed under passive learning is designed to exploit the fact that the output gap is determined by long-term rates. By introducing persistence in the short rate, the central bank gains more control over the long rate and a given change in the short-term nominal interest rate will have a larger effect on the output gap. This is the subject of Woodford (1999).

Under the active learning policy, short-term nominal interest rates are even more persistent than they were with the passive learning policy. The motivation for this is that a given change in the short-term nominal interest rate will then have a larger effect on the long-term rate and hence on the output gap. Table 2 reports the standard deviations of the short-term nominal interest rate and the output gap in the model. The increased persistence of short-term nominal interest rates means that, even though short-term rates are less volatile with active learning, there is greater volatility in the output gap.

Learning	σ_r	σ_x	σ_π	σ_p
passive	1.40	12.67	0.48	0.30
active	1.34	13.21	0.48	0.31

Table 2: Standard deviation of the short-term nominal interest rate, output gap, inflation and beliefs.

The active learning policy induces greater volatility in the output gap because it is beneficial for learning. Increased volatility makes it easier for the central bank and private agents to learn the value of the structural parameter κ_t in the aggregate supply curve (2).[8] In our simulations, beliefs consequently converge faster to their true (extreme) values under active learning than passive learning. This explains the higher standard deviation of beliefs for the active learning policy in Table 2.

[8] This argument has been made by Bertocchi and Spagat (1993), Balvers and Cosimano (1994) and Wieland (2000) in the context of central bank activism and policy experimentation. If the central bank were to control the output gap directly in our model then there would be an incentive for increased volatility in policy to promote learning.

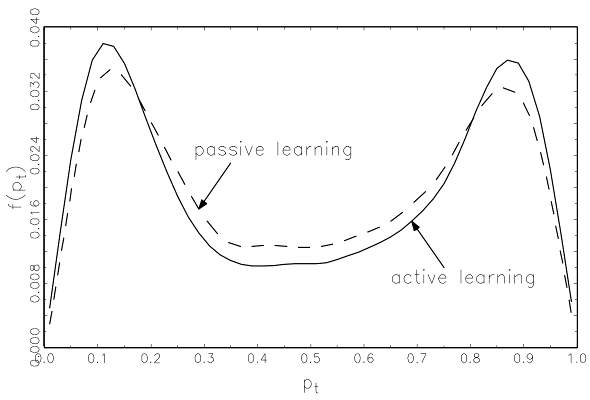

Figure 3: Distribution of beliefs.

The different behaviour of beliefs under the passive and active learning policies is illustrated in Figure 3, which shows the simulated distribution of beliefs for the two policies. The greater output gap volatility under the active learning policy means that beliefs are updated faster and consequently tend to be closer to 0 or 1, the values corresponding to certainty. In effect, the active learning policy causes a mean-preserving spread in the distribution of beliefs, with more probability mass being assigned to beliefs associated with greater certainty.

The benefit to be gained by following the active learning policy and increasing the persistence of the short-term nominal interest rate depends on the calibration of the parameter values. For the baseline calibration, the active learning policy is marginally welfare-enhancing compared to the passive learning policy. Measured in terms of the median value of the period-by-period central bank loss function, the improvement under active learning is of the order of +0.5%.

4.4 Interest rate reversals

The theme of the introduction to this paper was that learning creates an incentive to avoid interest rate reversals. In our simulations, the increased persistence in the short-run nominal interest rate does translate into less interest rate reversals under active than passive learning. Hence our claim that learning provides an additional motivation for why central banks should refrain from interest rate reversals. However, the number of interest rate reversals predicted by the model is high under both passive and active learning policies. In this respect, the predictions of the model do not match empirical data. A formal runs test rejects the hypothesis that actual US data could have been generated by the model under either passive or active learning.

4.5 Sensitivity analysis

Our result is robust to a range of alterative (sensible) calibrations for the parameters of the model. Where we do find a difference is if we increase the persistence of shocks to the natural rate of interest. With $\rho = 0.9$ our result is reversed: the active learning policy has less persistence in the short-term nominal interest rate, less volatility in the output gap, slower learning and more interest rate reversals. The improvement in welfare is of the same order of magnitude as before. The intuition behind this apparently contradictory result is that, when shocks are persistent, the strategic interplay between the central bank and the private sector becomes more important. Persistent shocks create an inflation bias as the central bank and private agents temporarily play a strategic game along the lines of Barro and Gordon (1983). The results of Ellison and Valla (2001) and Rosal and Spagat (2002) explain that, in such cases, learning can be detrimental to welfare. The optimal active learning policy is therefore designed to retard rather than promote learning.

5 Conclusions

We began this paper with the observation that central banks appear very reluctant to reverse their recent interest rate decisions. Our analysis suggests learning as an additional motivation for this behaviour. We show, in a standard monetary model with forward-looking expectations and uncertainty, that a policy which internalises learning (an active learning policy) leads to less interest rate reversals than a policy which ignores learning (a passive learning policy). The active learning policy has less interest rate reversals and more interest rate smoothing to give the central bank greater leverage over the long-term real interest rate. This greater control is used by the central bank to generate more volatility in the output gap, thereby promoting learning and improving welfare. The incentive to avoid interest rate reversals is in addition to the optimal policy inertia identified by Woodford (1999), Orphanides (1998) and Sack (2000), for models with forward-looking expectations, data uncertainty and parameter uncertainty.

References

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

[2] Barro, R.J. and G.B. Gordon, 1983, Rules, Discretion and Reputation in a Model of Monetary Policy, Journal of Monetary Economics 12, 101-121

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

[4] Clarida, R., Gali, J. and M. Gertler, 1999, The Science of Monetary Policy: A New Keynesian Perspective, Journal of Economic Literature 37, 1661-1707

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

[6] den Haan, W.J. and A. Marcet, 1990, Solving the Stochastic Growth Model by Parameterizing Expectations, Journal of Business and Economic Statistics 8, 31-34

[7] den Haan, W.J. and A. Marcet, 1994, Accuracy in Simulations, Review of Economic Studies 61, 3-17

[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] Evans, G. and S. Honkapohja, 2001, Learning and Expectations in Macroeconomics, Princeton University Press, New Jersey

[10] Giannoni, M. and M. Woodford, 2002, Optimal Interest Rate Rules I: General Theory, Unpublished, Princeton University and Columbia University

[11] McCallum, B.T. and E. Nelson, 1999, An Optimizing IS-LM Specification for Monetary Policy and Business Cycle Analysis, Journal of Money, Credit, and Banking 31:3, 296-316

[12] Orphanides, A., 1998, Monetary Policy Evaluation with Noisy Information, Board of Governors of the Federal Reserve System, Finance and Economics Discussion Paper Series: 98/50

[13] 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

[14] Rotemberg, J.J. and M. Woodford, 1997, An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy, NBER macroeconomics annual 1997, 297-346

[15] Sack, B., 2000, Does the Fed Act Gradually? A VAR Analysis, Journal of Monetary Economics 46, 229-56

[16] Sack, B. and V. Wieland, 2000, Interest-Rate Smoothing and Optimal Monetary Policy: A Review of Recent Empirical Evidence, Journal of Economics and Business 52, 205-28

[17] Wieland, V., 2000, Learning by doing and the value of optimal experimentation - optimal learning with endogenous information, Journal of Economic Dynamics and Control 24, 501-534.

[18] Woodford, M., 1999, Optimal Monetary Policy Inertia, Manchester School 67, 1-35

A First order conditions for optimal policy

We begin by deriving the first order conditions for optimality when the central bank follows an active learning policy. The Lagrangian (A.1) corresponds to the minimisation problem presented in equations (11)-(16). In equation (A.1), the inflation, output gap and short-term nominal interest rate are all indexed by the superscript $\mu \in [0,1]$ to indicate they are conditional on the belief of the central bank and private agents. For example, r^{μ}_{t+j} is the short-term nominal interest rate set by the central bank at time *t+j*, conditional on belief p_{t+j} being equal to μ. Similarly, π^{μ}_{t+j}, $\pi^{\mu,h}_{t+j}$, $\pi^{\mu,l}_{t+j}$ and x^{μ}_{t+j} are the inflation and output gap outcomes conditional on $p_{t+j} = \mu$. The terminology p^{μ}_{t+j} denotes the unconditional probability that belief p_{t+j} will be equal to μ at time *t+j*. $p^{s|\mu}_{t+j+1}$ is the conditional probability that belief p_{t+j+1} will be equal to s at time *t+j+1*, given that beliefs at time *t+j* are $p_{t+j} = \mu$. It is defined by equation (8) as $p^{s|\mu}_{t+j+1} = P(B(\mu, x^{\mu}_{t+j}, \varepsilon_{t+j}) = s)$.

$$
\begin{aligned}
&E_t \sum_{j=0}^{\infty} \beta^j \Big\{ \sum_{\mu} p^{\mu}_{t+j} [(\pi^{\mu}_{t+j})^2 + \lambda_x (x^{\mu}_{t+j})^2 + \lambda_r (r^{\mu}_{t+j})^2] \\
&+2\sum_{\mu} \phi^{\mu}_{1t+j} [x^{\mu}_{t+j} - \sum_{s} p^{s|\mu}_{t+j+1} x^{s}_{t+j+1} + \sigma^{-1}(r^{\mu}_{t+j} - r^{n}_{t+j} - \sum_{s} p^{s|\mu}_{t+j+1} \pi^{s}_{t+j+1})] \\
&+2\sum_{\mu} \phi^{\mu}_{2t+j} [\pi^{\mu,h}_{t+j} - \kappa_h x^{\mu}_{t+j} - \beta \sum_{s} p^{s|\mu}_{t+j+1} \pi^{s}_{t+j+1} - \varepsilon_{t+j}] \\
&+2\sum_{\mu} \phi^{\mu}_{3t+j} [\pi^{\mu,l}_{t+j} - \kappa_l x^{\mu}_{t+j} - \beta \sum_{s} p^{s|\mu}_{t+j+1} \pi^{s}_{t+j+1} - \varepsilon_{t+j}] \\
&+2\sum_{\mu} \phi^{\mu}_{4t+j} [\pi^{\mu}_{t+j} - \mu \pi^{\mu,h}_{t+j} - (1-\mu) \pi^{\mu,l}_{t+j}] \\
&+\beta \sum_{\mu} \phi^{\mu}_{5t+j+1} [\sum_{s} p^{s}_{t+j} P(B(s, x^{s}_{t+j}, \varepsilon_{t+j}) = \mu) - p^{\mu}_{t+j+1}] \Big\}
\end{aligned} \tag{A.1}
$$

The first line in the Lagrangian is the expected period-by-period loss function from equation (10). The next four lines in the Lagrangian relate to the set of constraints (12)-(15), with corresponding Lagrange multipliers $\phi^{\mu}_{1t+j} \cdots \phi^{\mu}_{4t+j}$. The forward-looking expectations in the IS and aggregate supply curves have been replaced by their conditional expectations. The final line of the Lagrangian is the learning constraint, with Lagrange multipliers $\beta\phi^{\mu}_{5t+j+1}$. The first derivatives of (A.1) with respect to $\pi^{\mu}_t, \pi^{\mu,h}_t, \pi^{\mu,l}_t, x^{\mu}_t, r^{\mu}_t$ and p^{μ}_t define the optimal active learning policy. They are shown in equations (A.2)-(A.7)

$$0 = E_t p^{\mu}_t \pi^{\mu}_t - (\beta\sigma)^{-1} \sum_s p^{\mu|s}_t \phi^{s}_{1t-1} - \sum_s p^{\mu|s}_t \phi^{s}_{2t-1} - \sum_s p^{\mu|s}_t \phi^{s}_{3t-1} + \phi^{\mu}_{4t} \quad \text{(A.2)}$$

$$0 = \phi^{\mu}_{2t} - \mu\phi^{\mu}_{4t} \quad \text{(A.3)}$$

$$0 = \phi^{\mu}_{3t} - (1-\mu)\phi^{\mu}_{4t} \quad \text{(A.4)}$$

$$\begin{aligned} 0 &= p^{\mu}_t \lambda_x x^{\mu}_t + \phi^{\mu}_{1t} - \beta^{-1} \sum_s p^{\mu|s}_t \phi^{s}_{1t-1} - \kappa_h \phi^{\mu}_{2t} - \kappa_l \phi^{\mu}_{3t} \\ &- E_t \sum_s [\phi^{\mu}_{1t} x^{s}_{t+1} + (\sigma^{-1}\phi^{\mu}_{1t} + \beta\phi^{\mu}_{2t} + \beta\phi^{\mu}_{3t})\pi^{s}_{t+1} - \tfrac{\beta}{2}\phi^{s}_{5t+1}] \frac{\partial P(\mathcal{B}(\mu, x^{\mu}_t, \varepsilon_t) = s)}{\partial x^{\mu}_t} \end{aligned} \quad \text{(A.5)}$$

$$0 = p^{\mu}_t \lambda_r r^{\mu}_t + \sigma^{-1}\phi^{\mu}_{1t} \quad \text{(A.6)}$$

$$0 = E_t(\pi^{\mu}_t)^2 + \lambda_x (x^{\mu}_t)^2 + \lambda_r (r^{\mu}_t)^2 + \beta E_t \sum_s \phi^{s}_{5t+1} P(\mathcal{B}(\mu, x^{\mu}_t, \varepsilon_t) = s) - \phi^{\mu}_{5t} \quad \text{(A.7)}$$

At time t, $p^{\mu}_t = p^{\mu|s}_t = 1$ for $\mu = p_t$ and $p^{\mu}_t = p^{\mu|s}_t = 0$ for $\mu \neq p_t$. Hence, $\phi^{\mu}_{1t} = \phi^{\mu}_{2t} = \phi^{\mu}_{3t} = \phi^{\mu}_{4t} = 0$ for $\mu \neq p_t$ and we can drop some of the subscripts on $\phi_{1t} \cdots \phi_{4t}, \pi_t, x_t$ and r_t. The system of equations simplifies to (A.8)-(A.13).

$$0 = E_t \pi_t - (\beta\sigma)^{-1}\phi_{1t-1} - \phi_{2t-1} - \phi_{3t-1} + \phi_{4t} \tag{A.8}$$

$$0 = \phi_{2t} - p_t \phi_{4t} \tag{A.9}$$

$$0 = \phi_{3t} - (1 - p_t)\phi_{4t} \tag{A.10}$$

$$0 = \lambda_x x_t + \phi_{1t} - \beta^{-1}\phi_{1t-1} - \kappa_h \phi_{2t} - \kappa_l \phi_{3t}$$
$$-E_t \sum_s [\phi_{1t} x^s_{t+1} + (\sigma^{-1}\phi_{1t} + \beta\phi_{2t} + \beta\phi_{3t})\pi^s_{t+1} - \tfrac{\beta}{2}\phi^s_{5t+1}] \frac{\partial P(\mathcal{B}(\mu, x^\mu_t, \varepsilon_t) = s)}{\partial x^\mu_t} \tag{A.11}$$

$$0 = \lambda_r r_t + \sigma^{-1}\phi_{1t}$$

$$0 = E_t(\pi^\mu_t)^2 + \lambda_x (x^\mu_t)^2 + \lambda_r (r^\mu_t)^2 + \beta E_t \sum_s \phi^s_{5t+1} P(\mathcal{B}(\mu, x^\mu_t, \varepsilon_t) = s) - \phi^\mu_{5t} \tag{A.13}$$

The first order conditions (A.8)-(A.13) define the active learning policy. In the case of passive learning, the central bank ignores the effects of changes in x_t on the expected future distribution of beliefs. The first order conditions for the passive learning policy are then given by equations (A.14)-(A.18).

$$0 = E_t \pi_t - (\beta\sigma)^{-1}\phi_{1t-1} - \phi_{2t-1} - \phi_{3t-1} + \phi_{4t} \tag{A.14}$$

$$0 = \phi_{2t} - p_t \phi_{4t} \tag{A.15}$$

$$0 = \phi_{3t} - (1 - p_t)\phi_{4t} \tag{A.16}$$

$$0 = \lambda_x x_t + \phi_{1t} - \beta^{-1}\phi_{1t-1} - \kappa_h \phi_{2t} - \kappa_l \phi_{3t} \tag{A.17}$$

$$0 = \lambda_r r_t + \sigma^{-1}\phi_{1t} \tag{A.18}$$