

The Learning Cost of Interest Rate Reversals

Martin Ellison^{a,b†}

^a University of Warwick, Coventry, CV4 7AL, UK; ^b CEPR, London EC1V 7RR, UK

Received 13 November 2003; Received in Revised Form Date; Accepted Date

Abstract

Many central banks in many time periods have sought to avoid interest rate reversals, but at present there is no good explanation of this phenomenon. Our analysis identifies a new learning cost associated with reversing the interest rate. In a standard monetary model with forward-looking expectations, data uncertainty and parameter uncertainty, a policy that frequently reverses the interest rate makes learning the key parameters of the model more difficult. Optimal monetary policy internalises this learning cost and therefore has a lower number of interest rate reversals.

Keywords: Interest Rate Smoothing, Learning, Monetary Policy

JEL classification: D83, E52, E58

*I would like to thank John Carlson, Neil Rankin, Jouko Vilmunen and the editor of this journal, Robert King, for helpful comments and suggestions. Financial support is acknowledged from an ESRC Research Fellowship “Improving Monetary Policy for Macroeconomic Stability in the 21st Century” (RES-000-27-0126). Giacomo Carboni provided excellent research assistance.

[†]Tel.: +44 2476 523472; fax +44 2476 523032; *E-mail address:* martin.ellison@warwick.ac.uk

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 (2003a), data uncertainty, Orphanides (2003), and parameter uncertainty, Sack (2000).¹

The purpose of this paper is to suggest that learning provides an additional reason for central banks to avoid interest rate reversals. We show, in a standard monetary model with learning, that it is more difficult for the central bank and private agents to learn a key feature of the monetary transmission mechanism if there are frequent reversals in monetary policy. Optimal monetary policy under learning is 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 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 slope of the Phillips curve in the canonical New-Keynesian model (the 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 Honkapo-

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

hja (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 present the first order conditions for optimal policy under passive and active learning.² Section 3 discusses our calibration and solution method. The results are presented in Section 4, which examines the nature of optimal monetary policy under different assumptions about 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 (2003a), 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 one 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.*).

²The derivation of the first order conditions is available from the author on request.

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, i_t . The short-term nominal interest rate is assumed to be the instrument of monetary policy.

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

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

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.

$$r_{t+1}^n = \rho r_t^n + \nu_{t+1} \quad (3)$$

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. We interpret time-variation in κ_t as representing changes in the link between the output gap and inflation.³ 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.⁴ However, since κ_t is unobservable, the central bank and private agents

³The motivation for introducing parameter uncertainty in this way is that learning becomes policy-dependent. If κ_t were to be interpreted as reflecting structural variation in the output gap then a separation principle holds and learning is independent of policy, as in the framework of Svensson and Woodford (2003).

⁴The assumption of a persistent two-state Markov process can be considered as a stylised representation of other persistent processes such as an AR(1).

have to infer its current value of the basis of observations of the output gap and inflation. The inference problem is complicated by the presence of an unobserved i.i.d. measurement error in inflation, ε_t , which corresponds to data uncertainty and is normally distributed with mean zero and standard deviation σ_ε .⁵

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 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 avoid simultaneity problems due to time t dated expectations in the IS and Phillips curves, we assume that inflation is only observed at the end of each period. Expectations and policy are set before this, on the basis of the information contained in $\{r_t^n, i_t, x_t, p_t\}$ and lagged values of all variables. To learn at the end of the period, the central bank and private agents must assess whether the observed inflation is 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 ,

⁵Following Clarida, Gali and Gertler (1999) and Woodford (2003b), it is conventional to give serial correlation to the shock in the Phillips curve. Allowing for serially correlated measurement errors would make the signal extraction problem for learning more complicated, but would not overturn our central results.

the current output gap and forward-looking inflation expectations. The central bank and private agents have to infer whether observed inflation is most likely to have come from distribution (4) or (5).

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

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

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 the new information available at the end of each period. 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.

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

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.

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

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 (2003a), we assume that the central bank minimises a loss function of the form (9). The discount factor satisfies $0 < \beta < 1$.

$$W = E_0 \left(\sum_{t=0}^{\infty} \beta^t [\pi_t^2 + \lambda_x x_t^2 + \lambda_i i_t^2] \right) \tag{9}$$

The objective function quadratically penalises the central bank for any deviations in inflation, the output gap and the short-term nominal interest rate from their long-run levels. All long-run levels are normalised to zero and λ_x, λ_i represent the weights placed on output gap and short-term 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 Woodford (2003b) shows, minimising a loss function such as (9) 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.⁶ 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.⁷

⁶The loss function (9) is not completely micro-founded since it ignores the fact that stochastic variation in κ_t should be reflected in the weight λ_x placed on output gap deviations. In this sense, our model only partially satisfies the cross-equation restrictions implied by the underlying theory.

⁷See Woodford (2003b) 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 rates 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 (10)-(13). The central bank minimises the expected loss defined by equation (9), subject to the IS curve (1) and the aggregate supply curve (2).

$$\begin{aligned} \min E_0 \left(\sum_{t=0}^{\infty} \beta^t [\pi_t^2 + \lambda_x x_t^2 + \lambda_i i_t^2] \right) \\ s.t. \end{aligned} \quad (10)$$

$$x_t = E_t x_{t+1} - \sigma(i_t - E_t \pi_{t+1} - r_t^n) \quad (11)$$

$$\pi_t = \kappa_t x_t + \beta E_t \pi_{t+1} + \varepsilon_t \quad (12)$$

$$p_{t+1} = \mathcal{B}(p_t, x_t, \varepsilon_t) \quad (13)$$

The final constraint (13) 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 (13) and the problem reduces to minimising (10) subject to (11)-(12). 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 (13) and the problem of the central bank is to minimise (10) subject to (11)-(13). 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.⁸ 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 satisfies the set of first order conditions (14)-(16), obtained by optimising with respect to π_t, x_t and i_t . ϕ_{1t} and ϕ_{2t} are the Lagrange multipliers associated with the IS curve and the aggregate supply curve 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’.

$$0 = E_t \pi_t - \beta^{-1} \sigma \phi_{1t-1} + \phi_{2t} - \phi_{2t-1} \quad (14)$$

$$0 = \lambda_x x_t + \phi_{1t} - \beta^{-1} \phi_{1t-1} - E_t \kappa_t \phi_{2t} \quad (15)$$

$$0 = \lambda_i i_t + \sigma \phi_{1t} \quad (16)$$

The first order conditions (14)-(16) are equivalent to those derived by Woodford (2003a), except for additional expectations due to uncertainty surrounding the structural parameter κ_t and inflation π_t . If $p_t = 1$ or $p_t = 0$ then κ_t is known with certainty and our first order conditions coincide. We therefore refer readers to Woodford (*ibid.*) for more discussion on the economic intuition of the system of equations (14)-(16). The passive learning policy that is defined by our first order conditions arises in the Woodford (*ibid.*) model if there is parameter uncertainty.

⁸Full details are available from the author on request.

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. Optimising with respect to π_t, x_t, i_t and p_{t+1} gives first order conditions (14), (17), (16) and (18). Condition (17) is a modification of (15) to allow for active learning, whereas (18) is the new first order condition and ϕ_{3t} 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 the first order condition (17) for the output gap. The extra terms are superscripted to indicate that they are 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 .

$$0 = \lambda_x x_t + \phi_{1t} - \beta^{-1} \phi_{1t-1} - E_t \kappa_t \phi_{2t} + E_t \int_0^1 [\phi_{1t} x_{t+1}^s + (\sigma \phi_{1t} + \beta \phi_{2t}) \pi_{t+1}^s + .5 \beta \phi_{3t+1}^s] \frac{\partial P(\mathcal{B}(p_t, x_t, \varepsilon_t) = s)}{\partial x_t} ds \quad (17)$$

$$0 = E_t \pi_t^2 + \lambda_x x_t^2 + \lambda_i i_t^2 + \beta E_t \int_0^1 \phi_{3t+1}^s P(\mathcal{B}(p_t, x_t, \varepsilon_t) = s) ds - \phi_{3t}^\mu \quad (18)$$

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 (19). Written in this way, the Lagrange multiplier ϕ_{3t} has an intuitive recursive structure of the type identified by Marcet and Marimon (1998).

$$\phi_{3t} = E_t \pi_t^2 + \lambda_x x_t^2 + \lambda_i i_t^2 + \beta E_t \int_0^1 \phi_{3t+1}^s P(\mathcal{B}(p_t, x_t, \varepsilon_t) = s) ds \quad (19)$$

ϕ_{3t} is equal to the sum of the expected period-by-period loss and the discounted value of ϕ_{3t+1} expected for the next period. 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 p_t at time t .

The analogy to a value function helps explain the presence of the extra terms in the first

order condition (17) 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.

3. Calibration and solution

We use the parameter values of Woodford (2003a) as a baseline quarterly 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 (holding expectations fixed) 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 κ_t remains unchanged is $1/0.025 = 40$ quarters. Some evidence on suitable values for σ_ε is given by the Gali and Gertler (1999) literature on estimation of the New Keynesian Phillips curve. Research in this area does not typically report specific estimates of σ_ε , but (unpublished) calculations based on Sbordone (2002) and Kurmann (2005) suggest values of 0.20 and 0.26. Table 1 reports the full calibration of our model. With respect to the persistence of shocks to the Wicksellian natural rate of interest, we report our main results with $\rho = 0.35$ but also consider other values in sensitivity analysis.

[TABLE 1 ABOUT HERE]

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.⁹ 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. As a *no learning* benchmark, we also simulate optimal policy with the slope of the Phillips curve held constant at its unconditional mean $(\kappa_h + \kappa_l)/2$. In all cases, 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, active and no 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.¹⁰ For the passive and no learning policies, the short-term nominal interest rate shows considerable persistence, with an autocorrelation coefficient around 0.7 for all beliefs. This is partly due to persistence in the natural rate of interest itself ($\rho=0.35$ in the simulations), but also because the presence of forward-

⁹A standard iterative approach such as the one suggested by Wieland (2000) cannot be used because our commitment policies are not time-consistent. The Bellman optimality principle is satisfied in our model, but only in the recursive sense defined by Marcet and Marimon (1998). A GAUSS program to solve our model and test its accuracy is available from the author on request.

¹⁰The estimates in Figure 1 are obtained by applying a standard Nadaraya-Watson kernel estimator of conditional autocorrelation to simulations of the model under passive, active and no learning policies. The results for the passive and no learning policies are represented by the same dashed line as they are almost identical.

looking expectations in the model creates an incentive for inertia in policy, as emphasised by Woodford (2003a).

[FIGURE 1 ABOUT HERE]

Turning to the active learning policy, we observe that the short-term nominal interest rate is more persistent than it was under the passive and no learning policies 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 the 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 unit positive innovation in the natural rate, when uncertainty is initially at its highest.¹¹

[FIGURE 2 ABOUT HERE]

The response of the short-term nominal interest rate under the active learning policy is initially muted when compared to the passive or no learning policy. However, the response is

¹¹We use the same kernel estimator as in Figure 1 to estimate the impulse response function conditional on initial beliefs being close to 0.5. The results under passive and no learning are very similar, so are represented by the same dashed line.

more persistent, remaining above baseline for longer under active than passive or no learning.

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 (20), the output gap x_t is a function of the sum of current and future expected deviations of the real short-term interest rate, $i_{t+j} - \pi_{t+j+1}$, from the natural rate, r_{t+j}^n . 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.

$$x_t = -\sigma E_t \sum_{j=0}^{\infty} (i_{t+j} - \pi_{t+j+1} - r_{t+j}^n) \quad (20)$$

The problem facing the central bank is how best to influence the output gap (which depends on long-term interest rates) when the policy instrument is a short-term interest rate. One idea would be to allow large fluctuations in the short-term nominal interest rate, but this would lead to problems with interest rates being distortionary or constrained by their zero bound, as discussed in Section 2.4. A better idea is for policy to incorporate a degree of persistence in the short-term nominal interest rate. Persistence is beneficial since then the current short-term nominal interest rate acts as a signal for future short-term rates. A given increase in the short-term nominal interest rate signals higher short-term rates in the future, so long-term rates and the output gap react more. This is the subject of Woodford (2003a).

Our active learning policy exhibits more persistence in short-term nominal interest rates than either the passive or no learning policy. The reason for the additional persistence is that learning creates an incentive for policy to induce extra volatility in the output gap. Increased volatility in the output gap promotes learning by making 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).¹² In an environment such as ours where volatility of the short-term nominal interest rate causes problems, the most efficient way of inducing extra volatility in the output gap is by increasing the persistence of the short-term nominal interest rate. 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.

[TABLE 2 ABOUT HERE]

The active learning policy induces greater volatility in the output gap because it is beneficial for learning. In our simulations, beliefs therefore converge faster to their true (extreme) values under active learning than passive learning. This explains the higher correlation between beliefs and the true value of the Phillips curve slope for the active learning policy in Table 2.

[FIGURE 3 ABOUT HERE]

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.

¹²This 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.

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-term nominal interest rate does translate into less interest rate reversals under active learning than under passive or no 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 all 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 with any of our assumptions about learning.

4.5. Sensitivity analysis

Our result is robust to a range of alternative (sensible) calibrations for the parameters of the model. The only occasion we find qualitatively different results is when 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 intuition behind this result is that, when shocks are persistent, the strategic interplay between the central bank and private agents becomes more important. The results of Cripps (1991)

explain that learning can be detrimental to welfare in such cases, and the active learning policy is designed to retard rather than promote learning.

The penalty on interest rate deviations in the central bank objective function is as important in generating our results as it is in Woodford (2003a). If the weight on interest rate deviations is reduced to zero, then policy can ensure that $x_t = 0$ and $\pi_t = \varepsilon_t$ for $\forall t$ without any knowledge of the Phillips curve slope.¹³ In this case, there is no incentive for the central bank to adjust policy to influence learning: the no learning, passive learning and active learning policies are identical. The short-term nominal interest rate always inherits the persistence properties of the Wicksellian natural rate of interest, so the active learning policy has just as many reversals as the passive and no learning policies.

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 or no learning policy). The active learning policy has less interest rate reversals, since then movements in the short-term nominal interest rate lead to larger changes in long-term rates, which translate into 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 (2003a), Orphanides (2003) and Sack (2000), for models with forward-looking expectations, data uncertainty and parameter uncertainty.

¹³Setting the short-term nominal interest rate according to $i_t = r_t^n + E_t\pi_{t+1} + \sigma^{-1}E_tx_{t+1}$ achieves this outcome independently of beliefs about κ_t .

References

- Balvers, R.J. and T.F. Cosimano, 1994, Inflation Variability and Gradualist Monetary Policy, *Review of Economic Studies* 61, 721-738
- Bertocchi, G. and M. Spagat, 1993, Learning, Experimentation and Monetary Policy, *Journal of Monetary Economics* 32, 169-183
- Clarida, R., Gali, J. and M. Gertler, 1999, The Science of Monetary Policy: A New Keynesian Perspective, *Journal of Economic Literature* 37, 1661-1707
- Cripps, M, 1991, Learning rational expectations in a policy game, *Journal of Economic Dynamics and Control* 15, 297-315
- 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
- den Haan, W.J. and A. Marcet, 1994, Accuracy in Simulations, *Review of Economic Studies* 61, 3-17
- Evans, G. and S. Honkapohja, 2001, *Learning and Expectations in Macroeconomics*, Princeton University Press, New Jersey
- Gali, J. and M. Gertler, 1999, Inflation Dynamics, A Structural Econometric Analysis, *Journal of Monetary Economics* 44, 195-222
- Giannoni, M. and M. Woodford, 2002, Optimal Interest Rate Rules I: General Theory, NBER Working Paper No. 9419
- Kurmann, A., 2005, VAR-based Estimation of Euler Equations with an Application to New Keynesian Pricing, mimeo

- Marcet, A. and R. Marimon, 1998, Recursive Contracts, mimeo
- 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
- Orphanides, A., 2003, Monetary Policy Evaluation with Noisy Information, *Journal of Monetary Economics* 50, 605-631
- Rotemberg, J.J. and M. Woodford, 1997, An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy, *NBER macroeconomics annual* 1997, 297-346
- Sack, B., 2000, Does the Fed Act Gradually? A VAR Analysis, *Journal of Monetary Economics* 46, 229-56
- 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
- Sbordone, A.M., 2002, Price and Unit Labour Costs: Testing Models of Pricing, *Journal of Monetary Economics* 29, 265-292
- Svensson, L.E.O. and M. Woodford, 2003, Optimal Policy with Partial Information in a Forward-Looking Model: Certainty Equivalence Redux, *NBER Working Paper* No. 9430
- 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.
- Woodford, M., 2003a, Optimal Interest-Rate Smoothing, *Review of Economic Studies* 70, 861-886

Woodford, M., 2003b, Interest and Prices: Foundations of a Theory of Monetary Policy
(Princeton University Press, New Jersey)

Structural parameters	β	0.99
	σ^{-1}	0.157
	κ_h	0.036
	κ_l	0.024
Shock processes	γ	0.975
	ρ	0.35
	$sd(r^n)$	3.72
	σ_ε	0.25
Loss function	λ_x	0.048
	λ_i	0.236

Table 1: Calibrated parameter values

Learning	σ_i	σ_x	$\rho_{p,\kappa}$
<i>passive</i>	1.40	12.67	0.60
<i>active</i>	1.34	13.21	0.62

Table 2: Properties of the short-term nominal interest rate, output gap and beliefs under passive and active learning




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




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




Figure 3: Distribution of beliefs.