

Learning uncertainty and central bank activism in an economy with strategic interactions

Martin Ellison^{a, *} and Matia Vallà^a

^aDepartment of Economics

European University Institute

Via dei Roccellini 9

50016 San Domenico di Fiesole (FI)

Italy

July 28, 2000

^{*}Corresponding author; telephone + 39 055 4851, fax + 39 055 599887, e-mail: ellison@iue.it. We are grateful to Michael Artis, Giuseppe Bertola, Michael Ehmann, Ramon Marimon, Mark Salmon, Andrew Scott and two anonymous referees for helpful comments and suggestions. We also thank participants at the ECB-CFS "Monetary Policy-Making Under Uncertainty" conference in Frankfurt, December 1999; the "Theories and Methods in Macroeconomics" conference in Paris, May 2000; CEF 2000; and workshops at the Bank of England, European University Institute and OECD.

Abstract

In this paper we examine the optimal level of central bank activism in a standard model of monetary policy with uncertainty, learning and strategic interactions. We calibrate the model using US data and find that the presence of strategic interactions between the central bank and private agents implies that optimality unambiguously recommends caution in monetary policy. An active policy designed to help learning and reduce future uncertainty creates extra volatility in inflation expectations and is detrimental to welfare.

JEL classification: D81, D83, D84, E52, E58

Keywords: activism, learning monetary policy, strategic interactions, uncertainty

1 Introduction

Should a central bank be cautious or activist in its monetary policy? Central bankers on both sides of the Atlantic such as B. Linder (1998) and Issing (1999) have recently stressed the practical importance of this question. Academic economists have responded to these concerns both by drawing on established work and by innovative state-of-the-art research. Over thirty years ago, Brainard (1967) recommended caution if there is uncertainty about the effects of monetary policy, a result which was later challenged by Bertocchi and Spagat (1993) with the suggestion that policy should be more activist since we learn more about the key parameters of the economy that way. Recent studies by Wieland (1998, 2000b) have synthesised and developed these ideas.

In this paper we develop arguments suggesting that caution may indeed be the optimal policy. We argue that an activist policy translates into more volatile inflation expectations, which cause problems for a central bank attempting to keep inflation low and smooth output fluctuations. By being more cautious, the central bank is able to dampen the volatility in inflation expectations and so create a more favourable environment for the conduct of monetary policy. To establish our result we use a standard monetary model, essentially that of Barro and Gordon (1983), in a dynamic setting. Rules for uncertainty and learning are created by assuming persistent but unobservable regimes in which monetary policy has different effects.

Our results depend on the strategic interactions inherent in the model,

which create a link between the activism of the central bank and the volatility of inflation expectations. The volatility of expectations reacts to the activism of a central bank because an activist policy produces more information, helping private agents to learn and adjust their expectations faster. Existing frameworks, by not taking strategic interactions into account, do not adequately specify the costs and benefits of an activist policy.

The paper is structured as follows. Our model is described in detail in Section 2 and then calibrated in Section 3 using empirical estimates of asymmetric regimes in the G7 economies. Section 4 solves the model for different assumptions concerning central bank behaviour. We focus on how policies differ depending on how the central bank takes learning issues into account. Section 5 concludes.

2 The model

2.1 Structure of the economy

The economy is characterised by an expectations-augmented Phillips curve relationship (1) between inflation surprises $\pi_t | \pi_t^e$ and output y_t , defined as deviation from trend. Inflation π_t is assumed to be completely under the control of the central bank and is the instrument of monetary policy.¹ ϵ_t is an i.i.d. output shock. It is assumed to be normally distributed with mean

¹Allowing the central bank complete control over the inflation rate abstracts from uncertainties related to the monetary transmission mechanism. These could be incorporated into the model but would not affect our conclusions.

zero and variance $\frac{1}{4}$. The shock itself is unobservable to the central bank and private agents but a signal is observable, giving information about the current output shock subject to noise. We define the activism of the central bank as the extent to which it reacts to the observed signal.

$$y_t = \bar{\alpha}_{s_t} (\frac{1}{4} \epsilon_t + \frac{1}{4} \epsilon_t^e) + \epsilon_t \quad (1)$$

It is immediately clear from equation (1) that volatile inflation expectations are problematic for the central bank. If inflation is kept close to a target then the volatility in inflation expectations is transmitted into volatile inflation surprises, and consequently volatile output. An alternative policy in which inflation is adjusted to meet expectations would reduce the problem of output fluctuations but create an exactly opposite problem in terms of inflation volatility. To avoid these problems the central bank prefers to conduct monetary policy in an environment where inflation expectations are less volatile.

To introduce learning issues we assume that the economy can be in either one of two unobservable regimes, $s_t = H$ or L , corresponding to high and low monetary policy effectiveness. The regime dependent parameter $\bar{\alpha}_{s_t}$ takes the value $\bar{\alpha}_H$ in the effective and $\bar{\alpha}_L$ in the ineffective regime. Since this parameter differs across regimes there will be asymmetry in the effects of monetary policy on output, depending on which is the current regime. The regimes are assumed to follow a hidden two state Markov switching process, in other words the economy switches between periods of high and low

monetary policy effectiveness. The conditional probabilities of not switching regime, i.e. $\frac{1}{2}_H = P(s_{t+1} = H | s_t = H)$ and $\frac{1}{2}_L = P(s_{t+1} = L | s_t = L)$ are assumed exogenous although not necessarily symmetric. The higher the probability of not switching the longer the regime is expected to last.




Figure 1: Timing of the model

The timing of the model is shown in Figure 1. The signal can be observed by the central bank before it makes its inflation choice, but by this time private agents have already set their inflation expectations for the period. Since only the central bank is able to react to the signal it creates asymmetries and a basis for stabilisation actions. If the signal indicates a large positive shock the central bank is able to tighten monetary policy accordingly to minimise losses. The activism of the central bank is reflected in the degree to which it reacts to the signal and is therefore measured by the difference between inflation π_t and inflation expectations π_t^e .

The signal z_t is assumed to be equal to the real output shock ϵ_t plus a classical measurement error ν_t as defined by equation (2). ν_t is assumed to be normally distributed with mean zero and variance σ_ν^2 . All variances are assumed to be known and independent of the regime so the signal extraction

problem of the central bank is to make a best estimate of the actual output shock, given the signal received. The problem reduces to finding the conditional distribution of ϵ_t given z_t . Equation (3) shows its solution, where the signal extraction parameter, \hat{A} , is defined as $\hat{A} = \frac{\sigma_{\epsilon}}{\sigma_{\epsilon} + \sigma_{\eta}}$.²

$$z_t = \epsilon_t + \eta_t \quad (2)$$

$$\epsilon_t | z_t \sim N[\hat{A}z_t; \hat{A}\sigma_{\epsilon}] \quad (3)$$

2.2 Central bank loss function

The central bank is assumed to have a per-period quadratic loss function (4) in the absolute level of inflation and deviations of output from a regime invariant target level y^r , as in the original Barro-Gordon (1983) model. The parameter \hat{A} reflects the weight placed by the policy maker on inflation versus output deviations from target. A non-inflationary bias arises in the model because at zero inflation the central bank has an incentive to exploit the Phillips-curve relationship (1) and reflate the economy through a positive inflation surprise. Rational private agents fully understand this incentive and so expect inflation to be higher than zero. In rational expectations equilibrium, the central bank no longer has an incentive to reflate the economy but there is a bias due to the non-zero inflation.

$$L_t = (\pi_t - y^r)^2 + \hat{A}\epsilon_t^2 \quad (4)$$

²Since ϵ_t and z_t have a bivariate normal distribution, the solution is given by standard conditioning results, see Illioud, Graybill and Boes (1974), pp. 167-8.

A central bank loss function of this type is not popular with central bankers. It is increasingly agreed that the solution to the problem of inflationary bias lies in institutional arrangements which prevent the central bank from targeting output above its natural rate, see inter alia Svensson (1999). We choose to retain the inflation bias because, if the central bank targets output at its natural rate, the only rational expectation for inflation in the model is the inflation target itself, in our case zero. Inflation expectations are constant and always equal to the inflation target, irrespective of the effectiveness of monetary policy. To give a role to inflation expectations we consequently relax the assumption of targeting output at its natural rate and assume the existence of an inflation bias.

2.3 Beliefs

At any point in time the central bank and private agents have beliefs about whether the economy is in the effective or ineffective monetary policy regime. Since the information available to the central bank and private agents is identical there is no scope for asymmetry in beliefs and the central bank and private agents always agree upon the probability of being in a particular regime. The beliefs of the central bank when it makes its inflation choice are the same as those of private agents when they set inflation expectations because, even though the signal is observed in the meantime, the signal on its own says nothing about the current regime. It is only when the signal is combined with other information, notably the inflation choice and realised

output, that it becomes useful in inferring the current state of the economy. The thus symmetric beliefs can conveniently be summarised by a single variable $p_t = P(s_t = H)$, which is the belief at time t that the economy is currently in the effective regime. If $p_t = 1$ then there is complete certainty that the economy is in the effective regime. Similarly, $p_t = 0$ implies that the ineffective regime is current.

2.4 Learning

Beliefs are not static in this model. They evolve over time as the central bank and private agents learn about which is the most likely current regime. The extent to which learning is possible depends critically on the informativeness of the actions taken by the central bank. In general, an activist policy involves taking actions which are more informative. Under the activist policy, the central bank reacts to the new information available in its signal and so its inflation choice is likely to differ significantly from expectations. With a large inflation surprise in the Phillips curve (1), it is easier to infer the monetary policy effectiveness parameter and hence learn what is the current regime. It is this learning benefit of an activist policy that has led to calls for greater activism in central bank policy by Balvers and Cosimano (1994) and Bertocchi and Spagat (1993).

To see how an activist policy translates into faster learning in the model, consider the distribution of output y_t conditional on prior information $I_t(\frac{1}{2}; Z_t; \frac{1}{2})$ and the state s_t . Equations (5) and (6) show the distribution of output

conditional on being in the high or low monetary policy effectiveness state respectively

$$y_t | I_{t:S_t=H} \gg N \left[\bar{\mu}_H (\mu_t | \mu_t^e) + \bar{\mu}_Z z_t; \bar{\mu}_\sigma \right] \quad (5)$$

$$y_t | I_{t:S_t=L} \gg N \left[\bar{\mu}_L (\mu_t | \mu_t^e) + \bar{\mu}_Z z_t; \bar{\mu}_\sigma \right] \quad (6)$$

Under a cautious policy, the inflation choice of the central bank is close to that expected so the inflation surprise is very small, i.e. $\mu_t | \mu_t^e \approx 0$. The two distributions (5) and (6) are nearly identical; the realisation of output y_t is almost equally likely to be from either regime and very little useful information is produced. Under this policy it is difficult to learn about the effectiveness of monetary policy. In contrast, under an activist policy the central bank does surprise private agents and the two distributions have significantly different means. The expected value of output is different in each regime so if the central bank then makes an inflation surprise, it can compare realised output y_t with the expected values to infer the current state of the economy. For example, if output is still close to trend even after a large positive inflation surprise the inference would be that monetary policy must currently be ineffective.

The fact that there are only a discrete number of states in the economy and that switching between states is exogenous means that the formation of beliefs takes a particularly simple form. A simple application of Bayes rule describes how beliefs are updated on the basis of new information. Equation (7) shows how the initial beliefs p_t are updated to p_t^+ at the end of the period, after the realisation of y_t . Under such Bayesian learning p_t^+ depends on the

relative probability of observing the outcome y_t in the two regimes.

$$p_t^+ = \frac{p_t P(y_t |_{t:St=H})}{p_t P(y_t |_{t:St=H}) + (1 - p_t) P(y_t |_{t:St=L})} \quad (7)$$

p_t^+ is the optimal inference of the current monetary policy effectiveness regime given the current realisation of output and the output signal. The central bank is hence able to make a prediction p_{t+1} of which regime will apply in the next period by taking account of the probability that there will be a regime shift at the beginning of the next period. In equation (8) the prediction is calculated as a weighted average of the probability of remaining in the high effectiveness regime and the probability of switching out of the low effectiveness regime

$$p_{t+1} = p_t^+ \frac{1}{2} + (1 - p_t^+) (1 - \frac{1}{2}) \quad (8)$$

Equations (7) and (8), combined with the normal distributions (5) and (6) for y_t give a non-linear equation (9) for updating beliefs. Updated beliefs are a function of the current belief, the inflation surprise, the signal and realised output. $B(\Phi)$ represents the Bayesian operator modified to take account of Markov-switching effects.

$$p_{t+1} = \frac{\frac{1}{2} p_t e^{i \frac{1}{2} \frac{y_t - \frac{1}{2} (\frac{1}{2} \mu_t^H + \frac{1}{2} \mu_t^L) - A z_t}{A \frac{1}{2}}}}{p_t e^{i \frac{1}{2} \frac{y_t - \frac{1}{2} (\frac{1}{2} \mu_t^H + \frac{1}{2} \mu_t^L) - A z_t}{A \frac{1}{2}}} + (1 - p_t) e^{i \frac{1}{2} \frac{y_t - \frac{1}{2} (\frac{1}{2} \mu_t^H + \frac{1}{2} \mu_t^L) - A z_t}{A \frac{1}{2}}}} + (1 - \frac{1}{2}) (1 - p_t) e^{i \frac{1}{2} \frac{y_t - \frac{1}{2} (\frac{1}{2} \mu_t^H + \frac{1}{2} \mu_t^L) - A z_t}{A \frac{1}{2}}}} \quad (9)$$

$$= B(p_t; \frac{1}{2} \mu_t^H, \frac{1}{2} \mu_t^L, z_t; y_t)$$

2.5 Rational expectations equilibrium

Private agents are assumed to be fully rational when making their expectation of the inflation rate. According to the definition of rational expectations equilibrium, these expectations have to be consistent with the actual behaviour of the central bank. In Figure 1, the inflation expectations formed ex ante before the observation of the signal must be equal to the average inflation choice made ex post by the central bank, i.e. inflation expectations have to satisfy equation (10), where π_t^e is the inflation choice of the policy maker given p_t , π_t^e and z_t .

$$\pi_t^e = \int \pi_t^e(p_t, \pi_t^e; z_t) f(z_t) dz_t \quad (10)$$

3 Calibration

The model should be calibrated at a frequency that reflects how often monetary policy decisions are made. In reality, the stance of monetary policy is reviewed regularly by the central bank so a natural choice is to make the model monthly.³ Table 1 shows our baseline calibration.

³In reality, monthly (and even quarterly) inflation data are rather noisy and far from a choice variable. However, our conclusions still remain valid in a more complex model in which the central bank sets monthly interest rates and the monetary transmission mechanism is modelled explicitly.

Parameter	Value
$\bar{\pi}_H$	3
$\bar{\pi}_L$	0:5
$\frac{1}{2}\bar{\pi}_H$	0:975
$\frac{1}{2}\bar{\pi}_L$	0:975
$\frac{3}{4}\alpha_1$	0:01
$\frac{3}{4}\alpha_2$	0:0033
\hat{A}	0:75
γ^π	0:01
\hat{A}	9
\pm	0:997

Table 1: Baseline calibrated values

The ...rst ...ve parameters in Table 3 are chosen on the basis of empirical estimation of the model for the G7 countries for 1980:1 to 1998:2.⁴ The sample period was chosen to avoid identifying the 1970's as a common observable regime of volatile shocks across all countries. Calibrated values of $\bar{\pi}_H = 3$ and $\bar{\pi}_L = 0:5$ imply that a 0.1 percentage point inflation surprise

⁴Estimation results are presented in Appendix A. We employ a structural vector autoregression (SVAR) approach in a Markov-switching context. The $\bar{\pi}$'s reported are in the range of the estimated impact effects of inflation surprises across the G7 economies in periods of high and low monetary policy effectiveness. For more details of this approach see Ehrmann, Ellison and Valla (2000).

leads to monthly output being 0.3% or 0.05% above trend, depending on the current monetary policy effectiveness regime. In other words, the output effect of an inflation surprise is six times higher when monetary policy is effective. Regimes are calibrated to be symmetrically persistent: a persistence parameter of $\rho = 0.975$ means that the average duration of each regime is $1/(1 - 0.975) = 40$ months. The standard deviation of output shocks, σ_{ϵ} , is set in the range of estimated values.

The natural level parameters in Table 1 cannot be estimated directly from the data. In the baseline calibration we set y^* , the target for the level of output above trend, to be equal to one standard deviation of the output shock but perform sensitivity analysis for a set of values $y^* \in \{0.15; 0.1; 0.05\}$. Since y^* is logarithmic these values correspond to target levels for output above trend of 0.5%, 1% and 1.5% respectively. λ reflects the weight that the central bank places on inflation as opposed to output deviations from target. We have calibrated this parameter so that the average inflation choice of the passive learning central bank corresponds to a monthly inflation rate of 0.2%. The signal extraction parameter, $\hat{\lambda}$, is difficult to calibrate so we take a range of parameters $\hat{\lambda} \in \{0.5; 0.75\}$. The standard deviation of the measurement error, σ_{η} , is set to match the signal extraction parameter. The discount factor, β , implies a quarterly discount rate of 1%.

Sensitivity analysis on the parameters y^* and $\hat{\lambda}$ did not yield qualitatively different behaviour of the model. Henceforth we only report results for the baseline calibration.⁵

⁵ Details of other results are available from the authors on request.

4 Results

We analyse the model under two different assumptions about how the central bank takes learning issues into account. Under the passive learning policy the central bank learns but makes no conscious effort to adjust its policies to affect the speed at which it learns. This policy serves as a benchmark in which the central bank learns but does not take into account that its current actions affect future benefits and losses through learning. In contrast, with the active learning policy the central bank fully internalises the effects of its actions on learning. We refer to this latter as the optimal policy.⁶

We begin by defining a benchmark, the full information policy in which the central bank and private agents always know the true state of the economy. In this economy there is no role for beliefs or learning.

4.1 Full information benchmark

The full information benchmark would be followed by the central bank if the true state of the economy was always known with certainty. Under these conditions the central bank minimises the expected loss each period subject to the known Phillips curve.

⁶The policy is optimal in the absence of a commitment technology by which the central bank can solve the time inconsistency problem.

$$\begin{aligned} \min_{\pi_t} E_t(L_t \pi_t^e; Z_t, S_t) \\ \text{s.t. } y_t = -\frac{1}{s_t}(\pi_t - \pi_t^e) + \epsilon_t \end{aligned} \quad (11)$$

By substituting the Phillips curve into the loss function the problem can be written as equation (12).

$$\min_{\pi_t} E_t \left[\frac{1}{2} \left(y_t + \frac{1}{s_t}(\pi_t - \pi_t^e) \right)^2 + \frac{1}{2} \pi_t^2 \right] \quad (12)$$

The only stochastic variable in this expression is the output shock ϵ_t . Hence we can use the conditional distribution $\epsilon_t | Z_t \sim N[\bar{\epsilon}, \bar{\sigma}^2]$ from equation (3) and then take the first order condition for loss minimisation (13) to derive the optimal central bank policy (14) for given $\pi_t^e; Z_t$ and s_t .

$$\frac{d}{d\pi_t} E(L_t \pi_t^e; Z_t, S_t) = 0 \quad (13)$$

$$\pi_t = \frac{1}{-2 \frac{s_t}{s_t + \bar{\Delta}}} \left(y_t + \frac{1}{s_t} \bar{\epsilon} \right) + \frac{1}{-2 \frac{s_t}{s_t + \bar{\Delta}}} \pi_t^e \quad (14)$$

Private agents take expectations of (14) to form their rational expectations of inflation (15), which further implies that the rational expectations equilibrium level of inflation is given by equation (16).

$$\pi_t^e = \frac{1}{\bar{\Delta}} y_t = \pi_{St}^d \quad (15)$$

$$\pi_t = \pi_t^e + \frac{1}{-2 \frac{s_t}{s_t + \bar{\Delta}}} \bar{\Delta} Z_t \quad (16)$$

According to equation (15) the inflation expectations of private agents directly follow the state of the economy. Expectations are exactly equal to the inflationary bias associated with the current regime μ_H^d in the high and μ_L^d in the low effectiveness regime, where $\mu_H^d > \mu_L^d$.

Equation (16) shows that inflation consists of a systematic component captured by expectations and a linear reaction to the observed signal z_t . The extent to which the central bank reacts to the signal depends on the state s_t , the signal extraction parameter \hat{A} , and its distaste for inflation parameter $\hat{\lambda}$. The central bank is more active in responding to the signal if it is in the effective monetary policy regime, if the signal is a good predictor of the output shock, or if the distaste for inflation is low.

4.2 Passive learning policy

The passive learning policy optimally accounts for current uncertainty but fails to realise that expected future losses can be affected by adjusting current policy actions; learning is ignored. Since learning and the updating of beliefs are the only source of dynamics in the model, the problem of the passive learning central bank reduces to that of minimising the expected one period loss function each period, subject to the Phillips curve. It is shown in equation (17).

The problem of the passive learning central bank is identical to the full information problem except that the state s_t is unknown. Equation (18) shows that the expected one period loss is a weighted average of the expected

losses conditional on the true state of the economy.

$$\begin{aligned} \min_{\pi_t} E_t(L_t | z_t; p_t) \\ \text{s.t. } y_t = -s_t(\pi_t | \pi_t^e) + \pi_t \end{aligned} \quad (17)$$

$$\min_{\pi_t} \begin{aligned} & p_t E_t \left[\left(\gamma^H i^{-1} (\pi_t | \pi_t^e) i^{-1} \right)^2 + \hat{A} \pi_t^2 \right]_{\pi_t; z_t; p_t; S_t=H} \\ & + (1 - p_t) E_t \left[\left(\gamma^L i^{-1} (\pi_t | \pi_t^e) i^{-1} \right)^2 + \hat{A} \pi_t^2 \right]_{\pi_t; z_t; p_t; S_t=L} \end{aligned} \quad (18)$$

Solving the minimisation problem of the central bank gives the passive learning policy, in which inflation is chosen according to equation (19).

$$\pi_t = \frac{p_t^{-2} + (1 - p_t)^{-2}_L}{p_t^{-2} + (1 - p_t)^{-2}_L + \hat{A}} (\gamma^H i^{-1} \hat{A} z_t) + \frac{p_t^{-2} + (1 - p_t)^{-2}_L}{p_t^{-2} + (1 - p_t)^{-2}_L + \hat{A}} \pi_t^e \quad (19)$$

Equations (20) and (21) characterise rational expectations equilibrium under this policy.

$$\pi_t^e = p_t \pi_t^d + (1 - p_t) \pi_t^L \quad (20)$$

$$\pi_t = \pi_t^e i^{-1} \frac{p_t^{-2} + (1 - p_t)^{-2}_L}{p_t^{-2} + (1 - p_t)^{-2}_L + \hat{A}} \hat{A} z_t \quad (21)$$

Inflation expectations given by equation (20) are a weighted average of the inflation biases that prevailed under the full information policy. The greater the probability assigned to the effective regime the higher is expected inflation. The inflation choice of the central bank in equation (21) again comprises a systematic component and a linear reaction to the observed signal z_t . However, it now is the inferred regime probabilities, p_t and $1 - p_t$, that determine how actively the central bank responds to shocks.

Beliefs p_t about the current state of the economy have two effects on the behaviour of the economy: they affect expectations directly through equation (20) and have implications for the size of surprises since they appear in the reaction to the signal in equation (21).

Figure 2 shows expected inflation and the inflation choices made by a passive learning central bank in the calibrated economy, as a function of prior beliefs and the signal observed. The central line shows inflation expectations, which are equal to the inflation choice made by the central bank after observing a zero signal. The upper line is the response to a one standard deviation negative signal, $z = -\frac{1}{4}z$, and the lower line is for a one standard deviation positive signal, $z = +\frac{1}{4}z$. The vertical distance between the upper and lower lines measures the degree to which the central bank reacts to its signal for a given prior belief.

The figure reveals that, in the calibrated economy, changes in beliefs have a significant effect on both inflation expectations and the extent to which the central bank reacts to the observed signal. Increased belief in the effective monetary policy regime is associated with higher inflation expectations and a stronger reaction to the signal. This is true for a wide range of alternative calibrations.

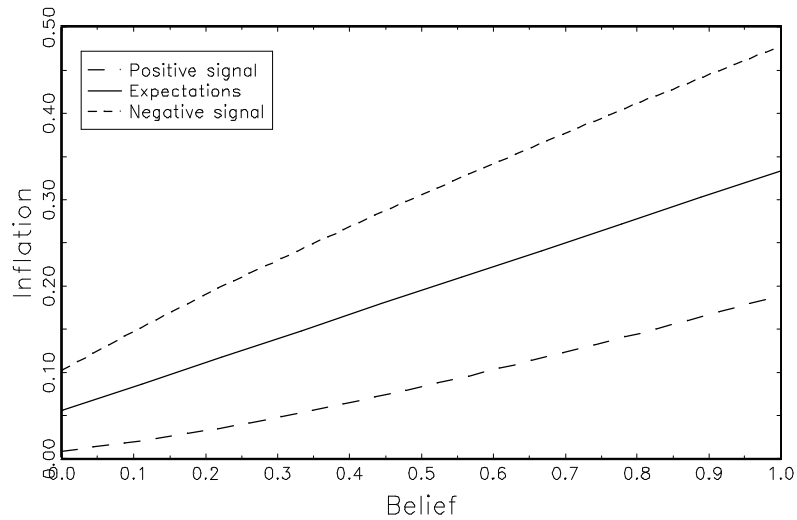


Figure 2: Inflation expectations and inflation choices under the passive learning policy

Although the passive learning policy ignores learning issues, this does not mean that there is no learning in the economy. Over time private agents and the central bank do receive information which helps them to infer the state of the economy. Private agents and the central bank learn and update their beliefs with considerable sophistication, as discussed in Section 2.3. Equation (9) summarising the symmetric Bayesian updating process is non-linear so simulations are needed to gain an insight into the dynamic behaviour of the economy. Table 2 shows some stylised facts calculated on the basis of 1000 simulations of the calibrated economy, each of 250 periods. In the table, $\frac{3}{4}$ represents standard deviation.

Passive learning policy		
Inflation	$\frac{3}{4} \pi_t$	0.131
Output	$\frac{3}{4} y_t$	0.844
Beliefs	$\frac{3}{4} p_t$	0.258
Expectations	$\frac{3}{4} \pi_t^e$	0.072
Surprises	$\frac{3}{4} \pi_t^{ue}$	0.109
Welfare loss	L	2.215

Table 2: Stylised facts of the passive learning policy

The dynamics simulations reveal that inflation is more volatile than either inflation expectations or inflation surprises. Indeed, since by definition inflation is equal to the sum of its expected and unexpected components, inflation volatility is due to volatility in both expectations and surprises and their covariance. A simple ANOVA analysis suggests that approximately 25% of the volatility in inflation can be attributed to volatility in inflation expectations and 75% to volatility in inflation surprises.⁷

The final row of Table 2 shows the average per-period welfare loss calculated according to equation (4).

⁷A analysis of variance (ANOVA) in this case is equivalent to calculating the goodness-of-fit measure R^2 in the linear regression model $\pi_t = a\pi_t^e + \epsilon_t$. Since $R^2 = 0.234$ we conclude that variations in inflation expectations π_t^e explain approximately 25% of the variation in inflation π_t .

4.3 Active learning policy

The work of Bertocchi and Spagat (1993) and Balvers and Cosimano (1994) suggests that the passive learning policy is suboptimal. They claim that a central bank should be more activist in its response to the observed signal because this provides valuable information about the state of the economy.⁸ By being more activist the central bank learns more quickly about the economy and so is better able to stabilise output shocks in the future. This argument was made recently by Wieland (1998, 2000b) in the context of a central bank learning the natural unemployment rate or learning whether the monetary policy transmission mechanism changed after German reunification.

To assess this argument in our model with strategic interactions we derive the active learning policy in which the central bank takes all learning into account. The active learning policy internalises the impact of central bank actions on learning by incorporating the learning constraints in the dynamic loss-minimisation problem (22). The central bank internalises both its own learning and that of private agents through equation (9) of Section 2.3. In addition, it internalises the consequences of learning by private agents for future inflation expectations, i.e. $\pi_{t+1}^e = \pi_{t+1}^e(p_{t+1})$, analogous to equation (20) in the passive learning case. The problem becomes intertemporal since future beliefs and inflation expectations both depend on current actions.

⁸ In the learning literature this behaviour is known as experimentation or probing

$$\min_{\mathbf{g}} \mathbb{E}_t \sum_{n=0}^{\infty} \beta^n L_{t+n}(y_{t+n}; \mathcal{W}_{t+n}^e) \quad \text{s.t.}$$

s.t

$$\begin{aligned} y_t &= -s_t(\mathcal{W}_{t+1} | \mathcal{W}_t^e) + \mathbf{1}_t \\ p_{t+1} &= B(p_t; \mathcal{W}_{t+1} | \mathcal{W}_t^e; z_t; y_t) \\ \mathcal{W}_{t+1}^e &= \mathcal{W}_{t+1}^e(p_{t+1}) \end{aligned} \quad (22)$$

This problem has a recursive nature so the active learning policy must satisfy the Bellman equation (23).

$$V(p_t; z_t; \mathcal{W}_t^e) = \min_{\mathcal{W}_t} \mathbb{E}_t L(p_t; \mathcal{W}_t^e; z_t; \mathcal{W}_t) + \beta V(p_{t+1}; z_{t+1}; \mathcal{W}_{t+1}^e) \quad (23)$$

No closed form analytical solution exists to this problem. However, Whited (2000a) shows how a standard dynamic programming algorithm can be used to obtain a numerical solution to the Bellman equation and an approximation to the active learning policy.⁹

⁹ Since Blackwell's sufficiency conditions are satisfied for this class of problem, see Kiefer and Nyerko (1989), it is possible to define a contraction mapping which converges to a unique fixed point. Hence, repeated iterations over the Bellman equation will converge to the stationary optimal policy and associated value function. Further details of the solution technique appear in Appendix A.

Figure 3 shows the active and passive learning policies. In the figure the inflation expectations associated with each belief are sufficiently close under the two policies that only those of the passive learning policy are shown. However, there is a marked difference in the degree of activism of the central bank. In contrast to the passive learning policy, the active learning inflation choices are closer to expectations. In other words, the central bank's reaction to the observed signal is muted. The active learning central banker is more cautious than its passive learning counterpart.

The largest deviation between the two policies occurs at $p_t = 0.5$, where uncertainty is at its greatest. For beliefs closer to certainty, the potential for learning by both the central bank and private agents is less, so the precise way the central bank takes learning into account becomes less important and the active learning policy approaches the passive learning one. In general, there is no reason to expect symmetry in the deviations between the two policies because the loss function is quadratic in the levels of inflation and output. However, in this calibration the deviations are almost symmetric so inflation expectations are almost identical under the two policies.




Figure 3: Inflation expectations and inflation choices under the active learning policy

The dynamic properties of the active and passive learning policies are compared in Table 3. Under the active learning policy the volatility of inflation surprises naturally falls as the central bank becomes more cautious. This decrease in volatility is generated by the central bank to dampen the undesirable volatility in inflation expectations of the private sector. Indeed, the volatility of inflation expectations also falls. The policy of an actively learning central bank therefore reduces the volatility of inflation both directly through the volatility of surprises and indirectly through the volatility of expectations.

		Passive learning policy	Active learning policy
Inflation	$\frac{3}{4} \pi$	0:131	0:095
Output	$\frac{3}{4} y$	0:844	0:887
Beliefs	$\frac{3}{4} p$	0:258	0:205
Expectations	$\frac{3}{4} \pi^e$	0:072	0:056
Surprises	$\frac{3}{4} \pi^u$	0:109	0:076
Welfare loss	L	2:215	2:213

Table 3: Stylised facts of the active learning policy

The decreased volatility of inflation has a positive effect on welfare at only the cost of a small increase in output fluctuations. The final row of Table 3 shows that the average welfare loss falls with caution, despite rising output volatility. The relatively small welfare gain of active learning reflects the fact that the costs and benefits of activism in this calibration almost exactly offset one another and the passive learning policy is approximately optimal.

5 Conclusions

We began this paper by asking whether a central bank should be cautious or active in its monetary policy. Our model suggests further reasons why caution may be the correct answer. The presence of strategic interactions

in the economy questions the claim that the central bank should follow an activist policy. On the contrary, it may be optimal for a central bank to do the opposite and become more cautious.

The problem with the activist policy is that it induces additional volatility into the inflation expectations of private agents. The rise in the volatility of inflation itself is therefore more pronounced; it rises due to more volatility in both surprises and expectations. From a learning perspective, volatile surprises are informative but volatile expectations give no benefits. This additional cost associated with the extra volatility in inflation expectations counteracts the call for greater activism made by authors such as Bertocchi and Spagat (1993), Balvers and Cosimano (1994) and Wieland (1998, 2000b). The results of our sensitivity analysis suggest that the additional cost actually dominates for a wide range of parameter values so activist policy is unlikely to be welfare improving.

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] Bлиндер, A. S., 1998, *Central banking in theory and practice* (MIT Press, Cambridge and London).
- [5] Brainard, W., 1967, Uncertainty and the Effectiveness of Policy, *American Economic Review Papers and Proceedings* 57, 411-425.
- [6] Ehrmann, M., M. Ellison and M. Valla, 2000, Regime dependent impulse response functions in a Markov-switching vector autoregression model, European University Institute mimeo
- [7] Issing O., 1999, The monetary policy of the ECB in a world of uncertainty, Contribution to the policy panel at the conference on "Monetary policy making under uncertainty" organised by the European Central Bank and the Center for Financial Studies, 3-4 December, Frankfurt am Main.

- [8] Kiefer, N.M. and Y. Nyarko, 1989, Optimal Control of an Unknown Linear Process with Learning, *International Economic Review* 30, 571-586
- [9] Mood, A.M., F.A. Graybill and D.C. Boes, 1974, *Introduction to the theory of statistics*, 3rd ed., (McGraw-Hill, New York).
- [10] O'Leary, D. and S.P. Vahey, 1995, Measuring Core Inflation, *Economic Journal* 105, 1130-1144.
- [11] Svensson, L.E.O., 1999, Price Level Targeting versus Inflation Targeting: A Free Lunch?, *Journal of Money, Credit and Banking* 31, 277-295.
- [12] Wieland, V., 1998, Monetary Policy and Uncertainty about the Natural Unemployment Rate, *Federal Reserve Board Finance and Economics Discussion Series* No. 22.
- [13] 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.
- [14] Wieland, V., 2000b, Monetary Policy, Parameter Uncertainty and Optimal Learning, *Journal of Monetary Economics*, forthcoming

A Estimation and calibration of the model

A.1 Estimation

We use the structural vector autoregression (SVAR) approach in a Markov-switching context. In a standard SVAR the structural form of the economy is recovered by imposing identifying restrictions on the moving average representation of an unrestricted vector autoregression. Our two stage procedure is to allow for Markov switching in the estimation of the unrestricted vector autoregression and then, for each regime separately, impose restrictions on the moving average representation calculated. We are thus able to trace out the impulse response functions corresponding to the effects of fundamental shocks in each regime.

In stage 1 we estimate the unrestricted vector autoregression in equation (A.1). The MAH(2)-VARX(p) specification shown was chosen to allow simultaneous Markov switches in the autoregressive parameters and the variance covariance matrix. A symmetry in the impact effects of monetary policy will be reflected in a switch in the variance-covariance matrix whereas dynamic asymmetry will switch the autoregressive parameters. The choice of a 2-regime model was made for consistency with our theoretical analysis.

$$i_{\frac{y}{t}}^{\Phi} = i_{v_1}^{\Phi} + A_1(s_t) i_{\frac{y}{t-1}}^{\Phi} + \dots + A_p(s_t) i_{\frac{y}{t-p}}^{\Phi} + i(s_t) t + i_{e_2 t}^{\Phi} \quad (A.1)$$

$$\text{Var} \begin{pmatrix} e_1 \\ e_2 \end{pmatrix} = \Sigma(s_t)$$

In stage 2, we use the Quah and Vahey (1995) procedure to identify inflation surprises by the restriction that, whilst they may have a long run

effect on inflation, they have no long run effect on output. In contrast, output shocks may have long run effects on both inflation and output.

For our estimation we used quarterly data from 1980:1 to 1998:2 for the G-7 countries, excluding Germany.¹⁰ Data were taken from the OECD Statistical Compendium, output being proxied by an index of domestic industrial production and inflation measured by the rate of change in the consumer price index. An exogenous trend was included in the vector autoregression to ensure stationarity of the dependent variables, particularly inflation. After some experimentation, a common lag length of one for output growth and two for inflation was accepted as giving clearly defined regimes in each country. All estimations were performed using the Matlab package for VAR(2). Results are reported in Table A.1.¹¹

	Persistence		Impact effect		Output shocks $\frac{3}{4}_1$
	$\frac{1}{2}_H$	$\frac{1}{2}_L$	$\frac{\partial y}{\partial \frac{1}{4} \pi^e}_H$	$\frac{\partial y}{\partial \frac{1}{4} \pi^e}_L$	
US	0.82	0.90	1.58	0.66	0.0055
UK	0.89	0.94	1.07	0.47	0.0084
France	0.88	0.87	0.85	0.78	0.0091
Italy	0.97	1.00	1.06	0.26	0.0184
Japan	0.90	0.64	0.94	0.13	0.0118
Canada	0.92	0.89	1.13	0.14	0.0124

¹⁰ Germany was excluded because of problems due to reunification occurring in the middle of the sample period.

¹¹ A full analysis of the estimation results is beyond the scope of this appendix.

Table A .1: Estimated regime persistence, impact effects and standard deviation of output shocks

The first two columns of Table A .1 show that the estimated regimes are highly persistent. An average quarterly probability of remaining in any regime of 0.9 translates into an expected regime duration of 10 quarters, i.e. two and a half years. Italy and to some extent Japan appear to be outliers in terms of persistence. The next two columns report numerical estimates of the impact effect of an inflation surprise. The final column of Table A .1 reports estimates of the standard errors of the identified output shocks.

Figure A .1 shows our identified impulse response functions.¹² Each country's panel shows the response of the level of output to a one percentage point inflation surprise. The panels reveal evidence of asymmetry in the effects of inflation surprises in all six countries. For all countries except the US the pattern is very similar: in the effective regime a one percentage point inflation surprise increases output immediately to 1% above trend. This effect peaks after one quarter and then slowly dies away. In the ineffective regime there is only a small positive impact effect on output.

¹² Estimates of the standard error bands are omitted for clarity. Paucity of data, a total of 74 observations meaning that there are only about 37 observations per regime, inevitably leads to wide standard error bands in any case.

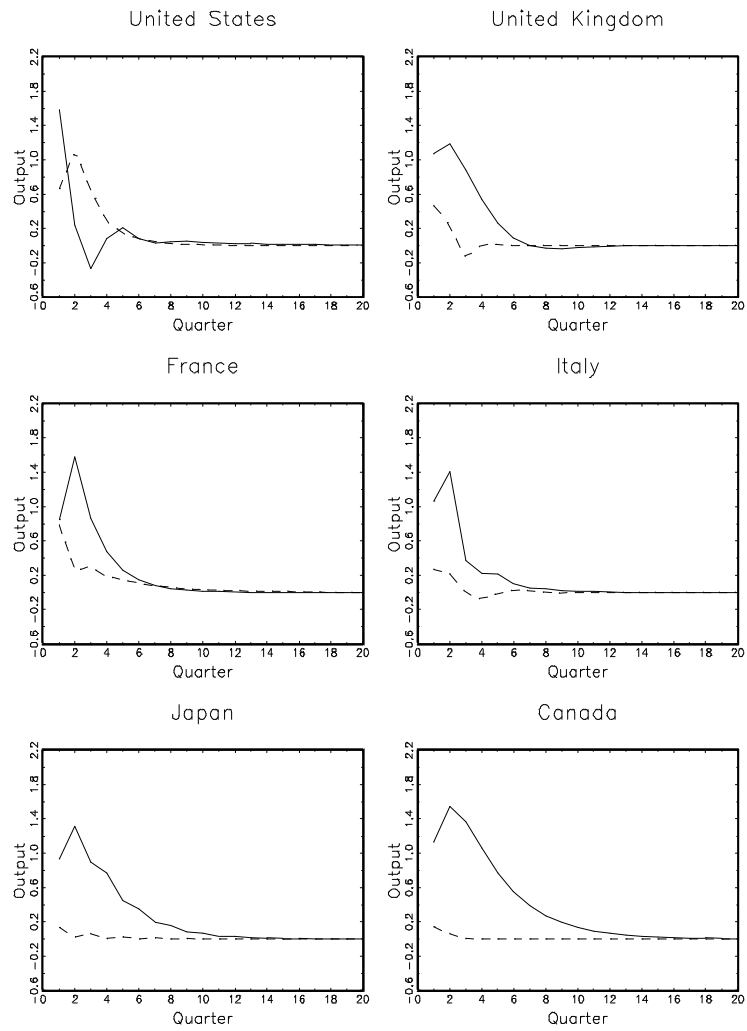


Figure A .1: Estimated response of output to a one percentage inflation rate surprise in each regime

A .2 Calibration

Table A .2 shows the calibrations implied by our estimation. The transform from our quarterly estimates to monthly parameter values is not trivial because y_t is a level defined as the logarithm of output deviation from trend, while π_t is the rate of change in prices.

Parameter	Quarterly calibration	Monthly calibration
\bar{h}	1	3
\bar{l}	0:167	0:5
$\frac{1}{2}h$	0:9	0:975
$\frac{1}{2}l$	0:9	0:975
$\frac{3}{4}a_1$	0:01	0:01
$\frac{3}{4}a_2$	0:0033	0:0033
\hat{A}	0:75	0:75
y^a	0:01	0:01
\hat{A}	1	9
\pm	0:99	0:997

Table A .2: Derivation of the calibrated values

B Numerical approximation of the active learning policy

An approximation of the active learning policy is calculated by solving the Bellman equation (23) numerically. To do this it is necessary to obtain expressions for the expected one period loss, $E_t L_t$, and the expected continuation value, $E_t V_{t+1}$ for a given inflation choice π_t . The former is defined by equation (4) and the latter can be written as equation (A.2), in which future beliefs, ρ_{t+1} , have been replaced by the non-linear updating equation (9).

$$E_t V_{t+1} = E_t V(B(\rho_t; \pi_t^e; z_t; \pi_t; y_t); z_{t+1}; \pi_{t+1}^e) \quad (\text{A.2})$$

The expectation in (A.2) has to be evaluated by the central bank before the realisation of both current output, y_t and the signal of next period's output shock, z_{t+1} . Hence, before setting current inflation, the central bank must calculate the double integral in equation (A.3).

$$E_t V_{t+1} = \int \int V(B(\rho_t; \pi_t^e; z_t; \pi_t; y_t); z_{t+1}; \pi_{t+1}^e) f(y_t | \rho_t; \pi_t^e; z_t; \pi_t) f(z_{t+1}) dy_t dz_{t+1} \quad (\text{A.3})$$

$f(z_{t+1})$ is the distribution of z_{t+1} and $f(y_t | \rho_t; \pi_t^e; z_t; \pi_t)$ is the predictive distribution of y_t . Their distributions are independent, a normal and a mix-

ture of normals respectively, as described by equations (A .4) and (A .5).

$$f(z_{t-1}) = N [0; \frac{1}{2}z] \quad (A .4)$$

$$f(y_t | p_t; \frac{1}{2}z_t; z_t; \frac{1}{2}z_t) = p_t N [-\frac{1}{2}H (\frac{1}{2}z_t | \frac{1}{2}z_t^e) + \hat{A}z_t; \frac{1}{2}z_t (1 | \hat{A})] \\ + (1 - p_t) N [-\frac{1}{2}L (\frac{1}{2}z_t | \frac{1}{2}z_t^e) + \hat{A}z_t; \frac{1}{2}z_t (1 | \hat{A})] \quad (A .5)$$

Our computational algorithm begins by defining a grid of points in the state space $(p_t; z_t)$. The grid points for beliefs, p_t , are placed uniformly over the interval $[0; 1]$, whereas grid points for the output shock signal, z_t , are bunched around the mean according to a cosine weighting function to increase accuracy. At each grid point we assign an inflation choice, $\frac{1}{2}z_t$, and a value for the value function, V_t , using as starting values equilibrium inflation choices of a passive learning central bank from equation (21) and the present discounted value of the associated expected one period loss (4).

One iteration of the Bellman equation is achieved by passing through the grid point-by-point. At each grid point the inflation choice is recalculated by minimising the right hand side of equation (23), using equation (4) for the expected one period loss and equations (A .2) to (A .5) for the expected continuation value. Numerical evaluation of the expected continuation value relies on Gaussian Quadrature methods to approximate the double integral in equation (A .3) and linear interpolation of adjacent grid points to evaluate the value function contained within the integral. Once a new inflation choice is calculated, a new $\frac{1}{2}z_t$ and V_t are assigned to the grid point. An iteration of the Bellman equation is completed when the inflation choice and value function have been updated for each grid point.

Repeating the above procedure to iterate the value function converges to the active learning policy. Over a grid of 10×10 points, we define the value function as converged when the values associated with each grid point change by less than 0.0001 over successive iterations. When updating the inflation choice we use a convergence tolerance of 0.00001. 32 ordinates are used in the Gaussian quadrature approximation of equation (A.3).

To calculate a rational expectations equilibrium we employ a simple iterative algorithm. Firstly, the inflation choices of the central bank $\pi_t(\rho_t; \pi_t^e; z_t)$ are calculated for given initial inflation expectations $\pi_t^e(\rho_t)$, according to the procedure described above. In the next stage, a new set of inflation expectations is calculated according to equation (10), i.e., $\pi_t^e = \int \pi_t(\rho_t; \pi_t^e; z_t) f(z_t) dz_t$. These expectations, $\pi_t^e(\rho_t)$, are used as the basis for calculating the new inflation choices $\pi_t(\rho_t; \pi_t^e; z_t)$. This procedure is iterated until convergence to a rational expectations equilibrium is achieved.

We use a uniform grid of 10 points over the state space ρ_t for inflation expectations and take starting values π_t^e from equation (20), inflation expectations under the passive learning policy. Calculating the new inflation expectations π_t^e requires a Gaussian quadrature approximation of the integral in equation (10). For each ordinate of the approximation the first expression inside the integral, the inflation choice $\pi_t(\cdot)$, is given by linear interpolation of adjacent inflation choices. The second term, z_t , has a normal distribution $f(z_{t+1}) = N[0; \frac{1}{3}z]$. Convergence is accepted when the change in inflation expectations between iterations is less than 0.000001 for each grid point. 32 ordinates are used in the Gaussian quadrature approximation.