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

Martin Ellison and Natacha Valla
Department of Economics
European University Institute

December 15, 1999

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 G7 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

*Paper presented at the "Monetary Policy-Making Under Uncertainty" conference organised by the European Central Bank and Center for Financial Studies, Frankfurt, December 1999. We thank Michael Artis, Giuseppe Bertola, Michael Ehrmann, Ramon Marimon, Andrew Scott and seminar participants at the Bank of England and the European University Institute for helpful comments and suggestions. *Address for correspondence:* Department of Economics, European University Institute, Via dei Roccettini 9, Badia Fiesolana, 50016 San Domenico di Fiesole (FI), Italy. e-mail: *ellison@iue.it* and *valla@iue.it*

1 Introduction

Should a central bank be cautious or active in its monetary policy? Central bankers think they have the answer to this problem: cautious. As Blinder (1998) puts it, "compute the direction and magnitude of the optimal policy move ... and then do less". Academic economists are not so sure. Brainard (1967) recommends caution if there is uncertainty about the effects of monetary policy whilst Bertocchi and Spagat (1993) suggest policy should be more active since we learn more about the key parameters of the economy that way. Recent studies by Wieland (1998a, 1998b) have revived interest in this debate.

In this paper we agree with the central bankers and find unambiguously that caution is the optimal policy. To establish our result we use a standard monetary model, essentially that of Barro and Gordon (1983), in a dynamic setting. Roles for uncertainty and learning are created by assuming persistent but unobservable regimes in which monetary policy has different effects. We argue that activism 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.

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 active policy produces more information, helping private agents to learn and adjust their expectations faster. Even the presence of small strategic interactions is sufficient to support our result. Existing frameworks, by not taking strategic interactions into account, do not adequately specify the costs and benefits of a more active 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 two different types of central bank behaviour: a myopic policy ignoring learning issues and optimality in which learning issues are taken 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.

$$y_t = \beta_{s_t}(\pi_t - \pi_t^e) + \mu_t \tag{1}$$

Inflation π_t is assumed to be completely under the control of the central bank and is the instrument of monetary policy. 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 strictly 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, $s = H$ or L, corresponding to high and low monetary policy effectiveness. The regime-dependent parameter β_{s_t} takes the value β_H in the effective and β_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, so the economy switches between periods of high and low monetary policy effectiveness. The conditional probabilities of not switching regime, i.e. $\rho_H = P(s_{t+1} = H\,|_{s_t=H})$ and $\rho_L = P(s_{t+1} = L\,|_{s_t=L})$ are assumed exogenous although not necessarily symmetric. The higher the probability of not switching the longer the regime is expected to last.

μ_t is an i.i.d. output shock. It is assumed to be normally distributed with mean zero and variance σ_μ. 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. As shown in the timing of the model 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 would be able to tighten monetary policy accordingly to minimise losses.

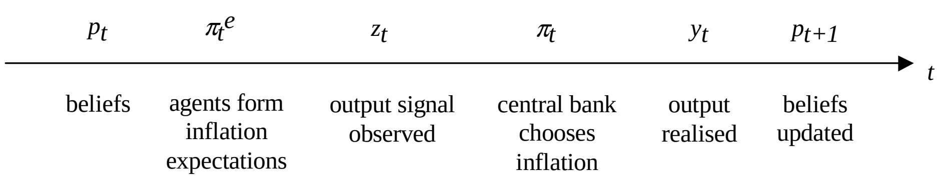

Figure 1: Timing of the model

The signal z_t is assumed to be equal to the real output shock μ_t plus a classical measurement error ε_t, as defined by equation (2). ε_t is assumed to be normally distributed with mean zero and variance σ_ε. All variances are assumed to be known so that the signal extraction problem of the central bank is to make a best estimate of the actual output shock, given the signal received. The solution of the signal extraction problem is shown in equation (3), which applies standard conditioning results with $\phi = \frac{\sigma_\mu}{\sigma_\varepsilon + \sigma_\mu}$.

$$z_t = \mu_t + \varepsilon_t \tag{2}$$

$$\mu_t \,|z_t \sim N\left[\phi z_t; \phi\sigma_\varepsilon\right] \tag{3}$$

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^*, as in the original Barro-Gordon (1983) model. The parameter χ reflects the weight placed by the policy maker on inflation versus output deviations from target.

$$\mathcal{L}_t = (y^* - y_t)^2 + \chi\pi_t^2 \tag{4}$$

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 (1996). We choose to retain the inflation bias for two reasons. Firstly, if the central bank targets output at its natural rate then the only rational expectation for inflation in the model is the inflation target itself, in our case zero. Inflation expectations are 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. Secondly, in the real world private agents may expect an inflationary bias even though none is actually present. Including an inflation bias means that our model is closer to that perceived by private agents, if not that perceived by the central bank. We prefer this because it is the perceptions of private agents that are the focus of our paper.

2.3 Learning

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 completely certainty that the economy is in the effective regime. Similarly $p_t = 0$ implies that the ineffective regime is current.

Beliefs are not static in this model. They evolve over time as information is produced which can be used to learn about which is the most likely current regime. Actions by the central bank lead to outcomes, which together can be used to infer the monetary policy effectiveness regime currently prevailing in the economy. In general, the more active the central bank is the more information is produced and the easier it is to infer the current regime. If the central bank is cautious then the surprises to private agents are small and the trade-off between unexpected inflation and output is not really used. Such a policy is not very informative; it is difficult to find out how monetary policy works if you never use it in the real world. It is only when the central bank is more active in exploiting the Phillips curve that information is produced that can be used to learn the current state of the economy. To see this in the model consider the distribution of y_t conditional on prior information $\mathcal{I}_t(\pi_t^e, z_t, \pi_t)$ and the state s_t given by equations (5) and (6).

$$y_t \left|_{\mathcal{I}_t, s_t=H}\right. \sim N\left[\beta_H(\pi_t - \pi_t^e) + \phi z_t; \phi\sigma_\varepsilon\right] \tag{5}$$

$$y_t \left|_{\mathcal{I}_t, s_t=L}\right. \sim N\left[\beta_L(\pi_t - \pi_t^e) + \phi z_t; \phi\sigma_\varepsilon\right] \tag{6}$$

If there is no inflation surprise, i.e. $\pi_t - \pi_t^e = 0$, then the two distributions are identical; the realisation of output y_t is equally likely to be from either regime and no useful information is produced. When the central bank does surprise private agents the means of the distributions differ and so the expected value of output is different in each regime. 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.

$$p_t^+ = \frac{p_t P(y_t \left|_{\mathcal{I}_t, s_t=H}\right.)}{p_t P(y_t \left|_{\mathcal{I}_t, s_t=H}\right.) + (1 - p_t) P(y_t \left|_{\mathcal{I}_t, s_t=L}\right.)} \tag{7}$$

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.

$$p_{t+1} = p_t^+ \rho_H + (1 - p_t^+)(1 - \rho_L) \tag{8}$$

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. $\mathcal{B}$ represents the Bayesian operator modified to take account of Markov-switching effects.

$$\begin{aligned} p_{t+1} &= \frac{\rho_H p_t e^{-\left(\frac{y_t - \beta_H(\pi_t - \pi_t^e) - \phi z_t}{\phi \sigma_\varepsilon}\right)^2} + (1-\rho_L)(1-p_t) e^{-\left(\frac{y_t - \beta_L(\pi_t - \pi_t^e) - \phi z_t}{\phi \sigma_\varepsilon}\right)^2}}{p_t e^{-\left(\frac{y_t - \beta_H(\pi_t - \pi_t^e) - \phi z_t}{\phi \sigma_\varepsilon}\right)^2} + (1-p_t) e^{-\left(\frac{y_t - \beta_L(\pi_t - \pi_t^e) - \phi z_t}{\phi \sigma_\varepsilon}\right)^2}} \\ &= \mathcal{B}(p_t, \pi_t - \pi_t^e, z_t, y_t) \end{aligned} \tag{9}$$

2.4 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 $\pi^*(.)$ is the inflation choice of the policy maker given p_t, π_t^e and z_t.

$$\pi_t^e = \int \pi^*(p_t, \pi_t^e, z_t) f(z_t) dz_t \tag{10}$$

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. Table 1 shows our baseline calibration.

Parameter	Value
β_H	3
β_L	0.5
ρ_H	0.975
ρ_L	0.975
σ_μ	0.01
σ_ϵ	0.0033
ϕ	0.75
y^*	0.01
χ	9
δ	0.997

Table 1: Baseline calibrated values

The first five 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.[1] Calibrated values of $\beta_H = 3$ and $\beta_L = 0.5$ imply that a 0.1 percentage point inflation surprise 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_μ, is set in the range of estimated values.

The final five 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 $y^* \in \{0.015, 0.01, 0.005\}$. Since y^* is logarithmic these values correspond to target levels for output above trend of 0.5%, 1% and 1.5% respectively. χ reflects the weight that the central bank places on inflation as opposed to output deviations from target.

[1] Estimation results are presented in Appendix A. We employ a structural vector autoregression (SVAR) approach in a Markov-switching context. The β'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 (1999).

We have calibrated this parameter so that the average inflation choice of the myopic central bank corresponds to a monthly inflation rate of 0.2%. The signal extraction parameter, ϕ, is difficult to calibrate so we take a range of parameters $\phi \in \{0.5, 0.75\}$. The standard deviation of the measurement error, σ_ε, is set to match the signal extraction parameter. The discount factor, δ, implies a quarterly discount rate of 1%.

Sensitivity analysis on the parameters y^* and ϕ did not yield quantitatively different behaviour of the model. Henceforth we only report results for the baseline calibration.[2]

4 Results

We analyse the model under two different assumptions about central bank behaviour. The first is that the central bank is myopic and minimises the expected one-period loss function each period. This central bank ignores learning issues by not taking into account that its current actions affect future beliefs and losses. The second is that the central bank minimises the expected present discounted value of all future losses. The policy under this assumption internalises learning issues and is therefore optimal.[3]

4.1 Myopic policy

A central bank following the myopic policy makes an inference about the current state of the economy and then minimises the expected one-period loss accordingly. It therefore optimally accounts for current uncertainty but fails to realise that it can affect expected future losses by adjusting its current policy actions; learning issues are ignored. The problem of the myopic central bank (11) is to minimise the expected one-period loss function each period subject to the Phillips curve.

$$\begin{aligned} &\min_{\pi_t} E_t \left[\mathcal{L}_t(y_t, \pi_t) \left|_{\pi_t^e, z_t, p_t} \right. \right] \\ &\text{s.t.} \quad y_t = \beta_{s_t}(\pi_t - \pi_t^e) + \mu_t \end{aligned} \tag{11}$$

[2] Details of other results are available from the authors on request.

[3] In the language of the learning literature, when the central bank follows the myopic policy it is a *passive learner* and when it follows the optimal policy it is an *active learner*.

By substituting the Phillips curve into the loss function the problem can be written as equation (12), in which the expected one-period loss is a weighted average of the expected losses conditional on the true state of the economy.

$$\min_{\pi_t} \left\{ \begin{array}{c} p_t E_t \left[(y^* - \beta_H(\pi_t - \pi_t^e) - \mu_t)^2 + \varkappa \pi_t^2 \left|_{\pi_t^e, z_t, p_t, s_t = H} \right. \right] \\ +(1-p_t) E_t \left[(y^* - \beta_L(\pi_t - \pi_t^e) - \mu_t)^2 + \varkappa \pi_t^2 \left|_{\pi_t^e, z_t, p_t, s_t = L} \right. \right] \end{array} \right\} \tag{12}$$

The only stochastic variable in this expression is the output shock μ_t. Hence we can use the conditional distribution $\mu_t \left| z_t \right. \sim N[\phi z_t; \phi \sigma_\varepsilon]$ from equation (3) and then apply the first order condition for loss minimisation, i.e. $\frac{d}{d\pi_t} E(\mathcal{L}_t \left|_{\pi_t^e, z_t, p_t} \right.) = 0$, to derive the optimal central bank policy (13) for given π_t^e, z_t and p_t.

$$\pi_t = \frac{p_t \beta_H + (1-p_t)\beta_L}{p_t \beta_H^2 + (1-p_t)\beta_L^2 + \chi}(y^* - \phi z_t) + \frac{p_t \beta_H^2 + (1-p_t)\beta_L^2}{p_t \beta_H^2 + (1-p_t)\beta_L^2 + \chi} \pi_t^e \tag{13}$$

Private agents take expectations of (13) to form their rational expectations of inflation (14), which further implies that the rational expectations equilibrium level of inflation is given by equation (15).

$$\pi_t^e = p_t \pi_H^d + (1-p_t)\pi_L^d \tag{14}$$

$$\pi_t = \pi_t^e - \frac{p_t \beta_H + (1-p_t)\beta_L}{p_t \beta_H^2 + (1-p_t)\beta_L^2 + \chi} \phi z_t \tag{15}$$

Inflation expectations given by equation (14) are a weighted average of the inflation biases that would prevail if the state was known with certainty, $\pi_H^d = \frac{\beta_H y^*}{\varkappa}$ and $\pi_L^d = \frac{\beta_L y^*}{\varkappa}$.[4] The greater the probability assigned to the effective regime the higher is expected inflation. Equation (15) 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 inferred regime probability p_t, the signal extraction parameter ϕ, and its distaste for inflation parameter χ. The central bank is more active in responding to the signal if the signal is a good predictor of the output shock or if the distaste for inflation is low.

[4] If the state of the economy is known then inflation expectations are exactly equal to the inflationary bias associated with the current regime: $\pi_t^e = \frac{\beta_H y^*}{\chi}$ in the high and $\pi_t^e = \frac{\beta_L y^*}{\chi}$ in the low effectiveness regime.

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 (14) and have implications for the size of surprises since they appear in the reaction to the signal in equation (15).

Figure 2 shows expected inflation and the inflation choices made by a myopic 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, $z = -\sigma_z$, and the lower line is for a one standard deviation positive signal, $z = +\sigma_z$. 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.

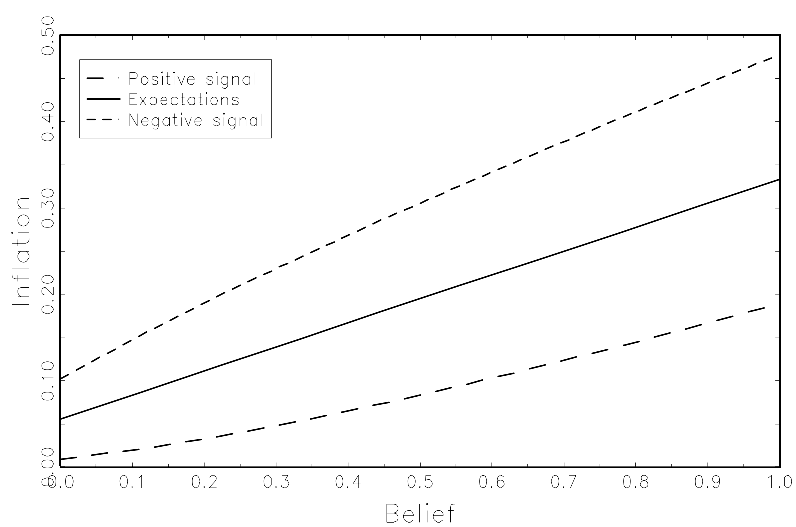

Figure 2: Inflation expectations and inflation choices under the myopic policy

Whilst the myopic policy ignores learning issues, it 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, σ represents standard deviation.

		Myopic policy
Inflation	σ_π	0.131
Output	σ_y	0.844
Beliefs	σ_p	0.258
Expectations	σ_{π^e}	0.072
Surprises	$\sigma_{\pi^{ue}}$	0.109
Welfare loss	$\mathcal{L}$	2.215

Table 2: Stylised facts of the myopic policy

The dynamic 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.

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

4.2 Optimal policy

The work of Bertocchi and Spagat (1993) and Balvers and Cosimano (1994) suggests that the myopic policy is suboptimal. They claim that a central bank should be more active in its response to the observed signal because this provides valuable information about the state of the economy. By being more active the central bank learns faster about the economy and so is better able to stabilise output shocks in the future.[5] This argument has been made recently by Wieland (1998a, 1998b) 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 test this argument we derive the fully optimal policy in which the central bank takes learning into account. Under the optimal policy the central bank minimises the present discounted value of expected current and future losses subject to the Phillips curve and learning considerations. 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. $\pi^e_{t+1} = \pi^e_{t+1}(p_{t+1})$, analogous to equation (14) in the myopic case. The problem becomes intertemporal since future beliefs and inflation expectations both depend on current actions.

$$\min_{\{\pi_t\}} E_t \sum_{n=0}^{\infty} \delta^n \left[\mathcal{L}_{t+n}(y_{t+n}, \pi_{t+n}) \,\middle|_{\pi^e_t, z_t, p_t} \right]$$

s.t.

$$\begin{aligned} y_t &= \beta_{s_t}(\pi_t - \pi^e_t) + \mu_t \\ p_{t+1} &= \mathcal{B}(p_t, \pi_t - \pi^e_t, z_t, y_t) \\ \pi^e_{t+1} &= \pi^e_{t+1}(p_{t+1}) \end{aligned} \tag{16}$$

This problem has a recursive nature so the optimal policy must satisfy the Bellman equation (17).

$$V(p_t, z_t, \pi^e_t) = \min_{\pi_t} E_t \left[\mathcal{L}(p_t, \pi^e_t, z_t, \pi_t) + \delta V(p_{t+1}, z_{t+1}, \pi^e_{t+1}) \right] \tag{17}$$

No closed-form analytical solution exists to this problem. However, Wieland (1999) shows how a standard dynamic programming algorithm can be used

[5] In the learning literature this behaviour is known as *experimentation*.

to obtain a numerical solution to the Bellman equation and an approximation to the optimal policy.[6]

Figure 3 shows the optimal and myopic policies. In the figure the inflation expectations associated with each belief are almost identical under the two policies but there is a marked difference in the degree of activism of the central bank. In contrast to the myopic policy, the optimal inflation choices are closer to expectations, implying that the central bank's reaction to the observed signal is muted. The optimal central banker is more cautious than its myopic counterpart.

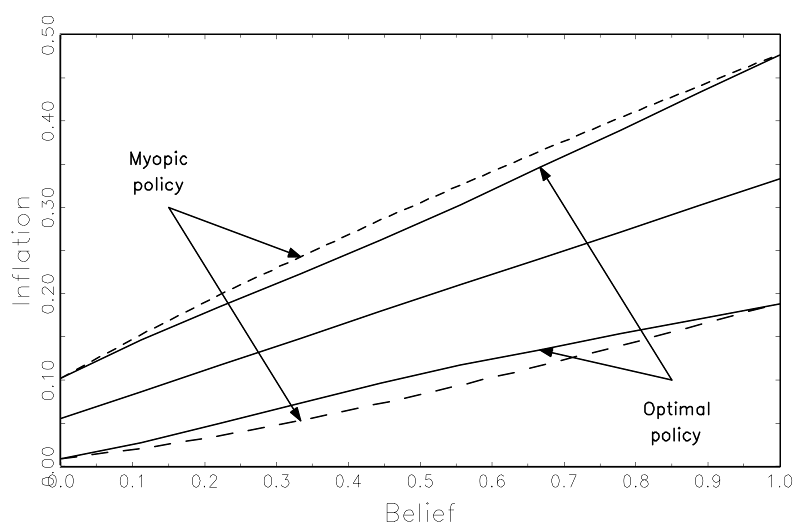

Figure 3: Inflation expectations and inflation choices under the optimal policy

The dynamic properties of the optimal and myopic policies are compared in Table 3. Under the optimal policy the volatility of inflation surprises naturally falls as the central bank becomes more cautious. This decrease in

[6]Since Blackwell's sufficiency conditions are satisfied for this class of problem, see Kiefer and Nyarko (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.

volatility is generated by the central bank for the sole purpose of dampening the inflation expectations of the private sector. Indeed, the volatility of inflation expectations also falls.

An optimal central bank therefore reduces volatility not only in inflation surprises but also in inflation expectations. Since inflation volatility depends on both surprises and expectations we therefore find that inflation volatility falls by more than the volatility of surprises.

		Myopic policy	Optimal policy
Inflation	σ_π	0.131	0.095
Output	σ_y	0.844	0.887
Beliefs	σ_p	0.258	0.205
Expectations	σ_{π^e}	0.072	0.056
Surprises	$\sigma_{\pi^{ue}}$	0.109	0.076
Welfare loss	$\mathcal{L}$	2.215	2.213

Table 3: Stylised facts of the optimal policy

The decreased volatility of inflation has a positive effect on welfare which is only partially offset by the costs of greater output fluctuations. The final row of Table 3 shows that the average welfare loss falls with caution, despite rising output volatility.

5 Conclusions

We began this paper by asking whether a central bank should be cautious or active in its monetary policy. Our model gives an unambiguous answer: cautious. The presence of strategic interactions in the economy repudiates the claim that the central bank should follow an active policy. On the contrary, a central bank should do the opposite and become more cautious.

The problem with the active 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 is sufficient to overturn the call for greater activism made by authors such as Bertocchi and Spagat (1993), Balvers and Cosimano (1994) and Wieland (1998a, 1998b). The results of our sensitivity analysis suggest this is true for a wide range of parameter values. Only if the inflation bias is very low or the asymmetry in monetary policy effectiveness is small will the active policy be welfare-improving.

On the call for caution, Blinder (1998) writes "My intuition tells me that this finding is more general - or at least more wise - in the real world than the [simple] mathematics will support. And I certainly hope it is ...". Our results suggest one way in which it is.

References

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

[2] Barro, R.J. and Gordon, D.B. (1983) "Rules, discretion and reputation in a model of monetary policy", *Journal of Monetary Economics*, 12:101-21

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

[4] Blinder, A. S., (1998) *Central Banking in Theory and Practice*, The MIT Press, Cambridge, Massachusetts

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

[6] Ehrmann, M., Ellison, M. and Valla, N. (1999) "Conditional impulse response functions in a Markov-switching model", *mimeo*, European University Institute

[7] Kiefer, N. and Nyarko, Y. (1989) "Optimal Control of an Unknown Linear Process with Learning", *Journal of Economic Dynamics and Control*, 30:571-586

[8] Quah, D. and Vahey, S.P. (1995) "Measuring Core Inflation", *Economic Journal*, 105:1130-1144

[9] Svensson, L.E.O. (1996) "Price Level Targeting vs. Inflation Targeting: A Free Lunch?", NBER Working Paper, No. 5719

[10] Wieland, V. (1998a) "Monetary Policy, Parameter Uncertainty and Optimal Learning", *mimeo*, Federal Reserve Board

[11] Wieland, V. (1998b) "Monetary Policy and Uncertainty about the Natural Unemployment Rate", *Federal Reserve Board Finance and Economics Discussion Series*, No. 22

[12] Wieland, V. (1999) "Learning by Doing and the Value of Optimal Experimentation", *Journal of Economic Dynamics and Control*, 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 MSAH(2)-VARX(p) specification shown was chosen to allow simultaneous Markov switches in the autoregressive parameters and the variance-covariance matrix. Asymmetry 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.

$$\begin{gathered}\binom{\Delta y_t}{\pi_t} = \binom{v_1}{v_2} + \mathbf{A}_1(s_t)\binom{\Delta y_{t-1}}{\pi_{t-1}} + ... + \mathbf{A}_p(s_t)\binom{\Delta y_{t-1}}{\pi_{t-1}} + \Gamma(s_t)t + \binom{e_{1t}}{e_{2t}} \\ Var\binom{e_1}{e_2} = \Omega(s_t)\end{gathered} \tag{A.1}$$

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 G7 countries, excluding Germany.[7] 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

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

and two for inflation was accepted as giving clearly-defined regimes in each country. All estimations were performed using the Msvar package for Ox 2.10. Results are reported in Table A.1.

	Persistence		Impact effect		Output shocks
	ρ_H	ρ_L	$\left.\frac{\partial y}{\partial \pi^{ue}}\right\vert_H$	$\left.\frac{\partial y}{\partial \pi^{ue}}\right\vert_L$	σ_μ
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

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 reports 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.[8] 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.

[8]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.

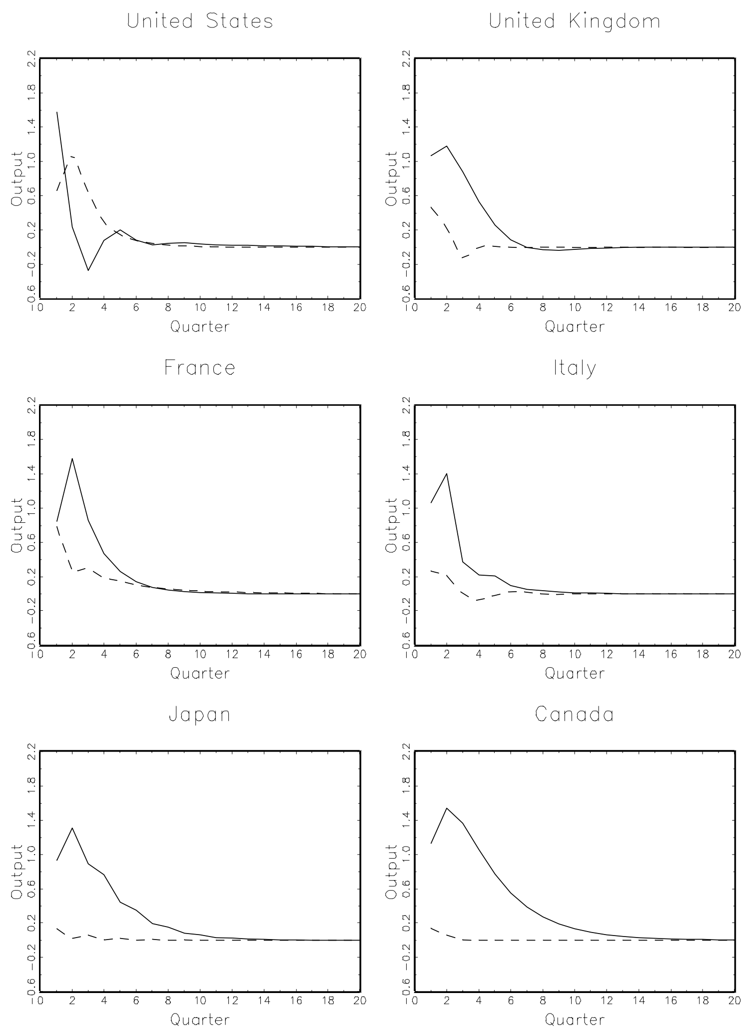

Figure A.1: Estimated response of output to a one percentage inflation rate surprise in each regime

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
β_H	1	3
β_L	0.167	0.5
ρ_H	0.9	0.975
ρ_L	0.9	0.975
σ_μ	0.01	0.01
σ_ϵ	0.0033	0.0033
ϕ	0.75	0.75
y^*	0.01	0.01
χ	1	9
δ	0.99	0.997

Table A.2: Derivation of the calibrated values

B Numerical approximation of the optimal policy

An approximation of the optimal policy is calculated by solving the Bellman equation (17) numerically. To do this it is necessary to obtain expressions for the expected one-period loss, $E_t\mathcal{L}_t$, and the expected continuation value, E_tV_{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, p_{t+1}, have been replaced by the non-linear updating equation (9).

$$E_tV_{t+1} = E_tV(\mathcal{B}(p_t, \pi_t^e, z_t, \pi_t, y_t), z_{t+1}, \pi_{t+1}^e) \tag{A.2}$$

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).

$$E_tV_{t+1} = \iint V(\mathcal{B}(p_t, \pi_t^e, z_t, \pi_t, y_t), z_{t+1}, \pi_{t+1}^e) f(y_t \left|_{p_t,\pi_t^e,z_t,\pi_t}\right.) f(z_{t+1}) dy_t dz_{t+1} \tag{A.3}$$

$f(z_{t+1})$ is the distribution of z_{t+1} and $f(y_t \left| p_t, \pi_t^e, z_t, \pi_t \right.)$ is the predictive distribution of y_t. Their distributions are independent, a normal and a mixture of normals respectively, as described by equations (A.4) and (A.5).

$$\begin{aligned} f(z_{t+1}) &= N[0; \sigma_z] && \text{(A.4)} \\ f(y_t \left|_{p_t,\pi_t^e,z_t,\pi_t}\right.) &= p_t N[\beta_H(\pi_t - \pi_t^e) + \phi z_t; \sigma_\mu(1-\phi)] \\ &\quad +(1-p_t)N[\beta_L(\pi_t - \pi_t^e) + \phi z_t; \sigma_\mu(1-\phi)] && \text{(A.5)} \end{aligned}$$

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, π_t, and a value for the value function, V_t, using as starting values equilibrium inflation choices of a myopic central bank from equation (15) 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 gridpoint the optimal inflation choice is recalculated by minimising the right hand side of equation (17), 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 optimal inflation choice is calculated, a new π_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 optimal 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 optimal 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 optimal inflation choices of the central bank $\pi(p_t, \pi^e_{t_0}, z_t)$ are calculated for given initial inflation expectations $\pi^e_{t_0}(p_t)$, according to the procedure described above. In the next stage, a new set of inflation expectations is calculated according to equation (10), i.e., $\pi^e_{t_1} = \int \pi(p_t, \pi^e_{t_0}, z_t) f(z_t) dz_t$. These expectations, $\pi^e_{t_1}(p_t)$, are used as the basis for calculating the new optimal inflation choices $\pi(p_t, \pi^e_{t_1}, z_t)$. 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 p_t for inflation expectations and take starting values $\pi^e_{t_0}$ from equation (14), inflation expectations under the myopic policy. Calculating the new inflation expectations $\pi^e_{t_1}$ requires a Gaussian Quadrature approximation of the integral in equation (10). For each ordinate of the approximation the first expression inside the integral, the optimal inflation choice $\pi(.)$, is given by linear interpolation of adjacent inflation choices. The second term, z_t, has a normal distribution $f(z_{t+1}) = N[0; \sigma_z]$. 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.