

Structural Uncertainty and Central Bank
Conservatism: The Ignorant Should Shut Their
Eyes

Joao Mauricio Rosal
Department of Economics
Royal Holloway College
University of London
Egham, Surrey, TW20 0EX
UK
j.m.rosal@rhul.ac.uk

Michael Spagat ¹
Department of Economics
Royal Holloway College
University of London
Egham, Surrey, TW20 0EX
UK
M.Spagat@rhul.ac.uk

22nd August 2002

¹Send editorial correspondence to: Professor Michael Spagat, Department of Economics, Royal Holloway College, University of London, Egham, Surrey TW20 OEX, United Kingdom, M.Spagat@rhul.ac.uk

Abstract

We study the problem of a central bank whose policy actions simultaneously affect the information flow about its expectations-augmented Phillips curve and its reputation for toughness in fighting inflation. In an environment with an unknown relationship between inflation surprises and output, big inflation surprises yield big short-term output gains and a strong information flow. Yet optimal policy is very conservative because inflation surprises yield information that increases the volatility of both future inflationary expectations and inflation itself. In fact, the more there is that can be learned about the Phillips curve the less does optimal policy aim towards learning.

JEL Numbers: E5, E58, D8, D83

Key Words: Monetary Policy, Learning, Experimentation, Reputation, Central Bank Conservatism

1 Introduction

Monetary authorities routinely conduct policy without full knowledge of the underlying structure of their economy. Some uncertain factors are truly unknowable at the time policymakers set their instruments. Other structural characteristics can be learned through experience. The possibility of learning about structural parameters creates potential for experimentation, i.e., policy partially aimed at generating useful information to improve the quality of future policy. For example, some have argued that US monetary policy in the middle to late 1990s was concerned precisely with learning how far interest rates could be lowered without kindling inflation, taking seriously the possibility that the structural relationship between interest rates and inflation might have shifted for the better [e.g. Bean (1999)].

Papers on experimentation and monetary policy include Bertocchi and Spagat (1993), Balvers and Cosimano (1994), and Wieland (2000).¹ Bertocchi and Spagat (1993) provide a simple framework showing how the desire to learn can bias monetary choices away from those yielding the highest present payoff and in a direction that generates more information. Balvers and Cosimano (1994) also consider active learning strategies for a central bank but with the aim of helping the public to improve its inflation forecasts. Wieland (2000) uses computational methods to analyze post-unification German monetary policy, arguing that the country paid dearly from the failure of the Bundesbank to experiment with looser monetary policy. These papers fill an important gap in the literature by considering learning and experimentation by central banks but are incomplete in the post-rational expectations era because they leave the public inactive.

A key to activating the public is to give it rational inflationary expectations. This can mean, in part, that the public observes the monetary

¹Early work on experimentation outside the context of monetary policy includes Prescott (1972) and Grossman, Kihlstron and Mirman (1977).

authority, learns about its toughness toward inflation and incorporates this information into its inflationary expectations as in Backus and Driffill (1984) and Barro (1986).² In this literature the public holds Bayesian beliefs about central banker toughness. Weak bankers have an incentive to maintain an illusion of toughness because once the public learns the truth it becomes difficult to conduct effective monetary policy.

In the present paper we combine all of the above elements.³ There is structural uncertainty about the relationship between inflation surprises and output.⁴ In general, different inflation rates will yield different information about this relationship. In addition, the public will form inflationary expectations rationally and this calculation will involve Bayesian updating of beliefs about the toughness of the central bank.

Our results are opposite to those in the present experimentation literature. In the standard literature, e.g., Bertocchi and Spagat (1993) structural uncertainty drives policy away from myopic optimality and in the direction of generating more information. However, activating the public reverses this conclusion, driving optimal policy to the other side of myopic optimality. In other words, equilibrium policy in our framework tends to generate less information than the myopically optimal policy would. Moreover, the more there is to learn, i.e., the greater is the structural uncertainty, the more conservative will be the central bank's behavior, shunning experimentation increasingly as its ignorance grows.

²Kydland and Prescott (1977) is the early key reference in this literature but it does not share our Bayesian approach.

³Caplin and Leahy (1996) combine a learning central bank with an active public in a different context. Their bank gropes for an unknown interest rate that will cause enough investment to pull the economy out of recession without generating so much as to overheat the economy. The problem is, e.g., that investors will wait to invest if they believe the bank will grope by gradually lowering interest rates.

⁴Throughout the paper we use the term structural uncertainty to mean potentially learnable uncertainty about parameters of the economic model.

Our work is loosely related to a large literature on learning and expectations in macroeconomics surveyed in Evans and Honkapohja (2001) that has developed largely independently of the above learning literature. Most of this work focuses on learning by the public but recent work considers central bank learning. Evans and Honkapohja (2002) and Honkapohja and Mitra (2002) study the issue of convergence to rational expectations equilibria when both the public and the central bank are learning. This whole research line differs from the present paper in two substantial ways. First, it uses “statistical” or “econometric” learning rather than the Bayesian approach of our paper. Second, it is interested primarily in long-run issues, mainly convergence or non-convergence to rational expectations equilibria, while we study inflation bias in a reputation and learning environment.⁵

The structure of the paper is as follows. In section 2 we present a model based on an expectations augmented Phillips curve with an unknown parameter governing the relationship between the inflation surprise and output. There are two types of policymaker with one caring only about inflation and the other caring about both inflation and output. The public holds Bayesian beliefs about the relative likelihood of these two types as well as about a structural parameter of the Phillips curve. The public’s inflationary expectations are rational given its information. In section 3 we solve the optimization problem for the weak policymaker, showing that he will actively shun new information about the structural parameter because such information is unusable and, in fact, positively harmful because it introduces unwanted variability into the inflation rate. A nice intuition underpins this result. The information of an inflation choice is increasing in the distance of the inflation rate from the level the public expects, i.e., the bigger the inflationary surprise the greater is the information. But a policy that aims

⁵Sargent (1999), on the other hand, does try to explain America’s postwar inflation experience based on full learning dynamics for the Fed.

at big inflationary surprises generates high volatility both in future inflationary expectations and in future inflation that are, in equilibrium, mutually reinforcing. Volatility is harmful so the central bank will strive to avoid big surprises and will, consequently, learn little about the unknown structural parameter.

In section 4 we study the equilibrium distinguishing between the case where the weak policymaker can commit to not using new information from the case where he cannot make such a commitment. There are two types of equilibria. In one the weak policymaker acts tough and pools with the strong one by choosing zero inflation. In the other equilibrium the weak policymaker exploits the output-enhancing potential of the Phillips curve while revealing his type to the detriment of future policy. The policymaker is more likely to pool in the no-commitment case than he is in the commitment case. We also show that the negative effect of new information puts pressure on the weak policymaker to act tough, that is, structural uncertainty argues for conservative monetary policy.

In section 5 we study in depth the tractable case where shocks to the Phillips curve have a uniform distribution. We show that there is a monotonic relationship between structural uncertainty and equilibria, i.e., for parameter values related to low structural uncertainty the policymaker is more likely to separate, and for those related to high structural uncertainty the policymaker is more likely to pool. There is also an intermediate region where equilibrium indeterminacy prevails, as the equilibrium can be either pooling or separating depending on the public's inflationary expectations. We draw conclusions in section 6.

2 The Model

The model is a simple aggregate supply and demand framework augmented by a monetary authority (the “policymaker”) with an unknown willingness

to fight inflation. The novelty of this work is that, in addition to reputation issues related to the policymaker, both the public and the policymaker have incomplete information about one of the underlying structural parameters of the economy.

2.1 The economic environment

The time horizon has two periods, $t = 1, 2$, and the following supply and demand curves describe the economic environment:

$$y_t = b(\pi_t - \pi_t^e) + \epsilon_t \quad (\text{Supply}) \quad (1)$$

$$\pi_t = m_t \quad (\text{Demand}) \quad (2)$$

where y_t is output growth, π_t is actual inflation, π_t^e is expected inflation and m_t is money growth, all in period t . The variable ϵ_t is taken to be a period- t random shock, and is distributed according to a cumulative distribution $F(\cdot)$, with a differentiable density $f(\cdot)$ and with full support on $[-\bar{\epsilon}_t, \bar{\epsilon}_t]$, where this support may well be the entire real line. For simplicity, it is assumed that $E(\epsilon_t) = 0$.

We assume that the supply curve's slope, given by the parameter b , is unknown by all agents and can take the values \underline{b} and \bar{b} , $0 < \underline{b} < \bar{b}$. For this lack of information, we use the term "structural uncertainty". At the outset, both the public and the policymaker attach a subjective probability θ , $\theta \in (0, 1)$, to the event that $b = \underline{b}$.

2.2 Policymakers

There are two types of policymakers, weak (ω) and strong (τ), differentiated by their period- t payoff. These types are policymakers' private information, and the public attaches an initial probability p , $0 < p < 1$, to the event of the policymaker being of the weak type.

Policymakers' pay-offs are given by the following expression:

$$\left\{ \begin{array}{ll} \text{Type Weak } (\omega) & w_t^\omega = y_t - a \frac{\pi_t^2}{2} \quad \text{with probability} = p \\ \text{Type Tough } (\tau) & w_t^\tau = -a \frac{\pi_t^2}{2} \quad \text{with probability} = 1 - p \end{array} \right.$$

where w_t^i stands for the policymaker type i 's payoff in period t , with $i = \omega, \tau$.

We assume that $a > 0$.

The specification above aims to formalize the idea that a weak policymaker is the one who may consider exploring the short-run trade off between inflation and output described by the supply curve. The tough policymaker, however, cares only about inflation. This specification also formalizes the fact that the policymakers dislike inflation variability, as both are risk averse towards inflation. This feature is particularly important in our analysis.

2.3 Strategies and Equilibrium Definition

The simplicity of the model allows us to make relatively simple definitions of strategies and equilibrium. Policymakers choose money growth in each period, however, from equation 2 this is equivalent to choosing inflation rates in both periods. The public rationally adjusts its expectations in both periods. The equilibrium is defined as follows.

Consider a time-and-type-indexed profile of policymakers' inflation strategies $\{\pi_{it}\}$, where $i = \omega, \tau$ and $t = 1, 2$. Also, let $\{\pi_1^e, \pi_2^e(\pi_1, y_1), p, \alpha(\pi_1, y_1), \theta, \mu(\pi_1, y_1)\}$ be the profile of expectations and beliefs held by agents, such that $\{\pi_1^e, \pi_2^e(\pi_1, y_1)\}$ are the public's inflation expectations at $t = 1, 2$; $\{p, \alpha(\pi_1, y_1)\}$ are respectively the public's period-1 and period-2 beliefs about the policymaker being weak; and $\{\theta, \mu(\pi_1, y_1)\}$ are the period-1 and period-2 beliefs held by the public and policymakers about b being \underline{b} . Then the equilibrium requirements are:

- i) Given all beliefs, inflation strategies π_{it}^* , $i = \omega, \tau$ and $t = 1, 2$, are sequentially rational;

- ii) Whenever possible, for history (π_1, y_1) period-2 beliefs $\alpha(\pi_1, y_1)$ and $\mu(\pi_1, y_1)$ are updated in a Bayesian fashion. Moreover, whenever $\pi_1 \neq 0$ we set $\alpha = 0$; ⁶
- iii) Inflation expectations $\pi_2^e(\pi_1, y_1)$ are equal to the expected inflation strategies over both policymakers' types, where expectations are taken using the public's beliefs in the relative probabilities of these types $\alpha(\pi_1, y_1)$.

Given the previous -somewhat informal- equilibrium definition, we proceed by solving the model backwards.

3 Optimal Choice for Policymakers

The main result of this section is that the weak policymaker tends towards inflation levels that generate only small surprises for the public. This result derives from the fact that the policymaker can only improve his information about b by increasing the degree to which he upsets public expectations at $t = 1$ while the bigger the surprise at time 1 the larger will be the variability of inflation at $t = 2$.

3.1 The Tough Policymaker

The tough policymaker cares only about inflation, so his choice is trivial:

$$\pi_{\tau 1} = \pi_{\tau 2} = 0 \tag{3}$$

These choices allow us to simplify the notation by dropping the subscript i for type and referring to the weak policymaker's inflation choices as simply π_1 and π_2 . Recall from the equilibrium definition that we assume that if

⁶This last requirement encapsulates the usual equilibrium requirement that agents attach zero probability in their beliefs to the event that the tough policymaker may play his dominated strategy $\pi_{\tau 1} \neq 0$.

the public observes π_1 different than zero it infers that the policy maker is weak. Therefore, we can once more simplify the notation and redefine the public's beliefs over types at $t = 2$ solely as function of π_1 , i.e., $\alpha(\pi_1)$, as this variable becomes a sufficient statistic for the policymaker's type.

3.2 The Weak Policymaker

3.2.1 Optimal Choice in Period 2

The public and the policymaker fully observe both output and inflation and update their beliefs about b in a Bayesian fashion, the expression for this belief being:

$$\mu(y_1, \pi_1; \pi_1^e) = \frac{\theta f_b}{\theta f_b + (1 - \theta) f_{\bar{b}}} \quad (4)$$

where, recalling that $f(\cdot)$ is the density function of ϵ_t , $f_b(y_1, \pi_1; \pi_1^e) = f(y_1 - b(\pi_1 - \pi_1^e))$. Note that $f_b(\cdot, \pi_1, \pi_1^e)$ is the density function of y_1 conditional on b and π_1 . As this density depends on π_1^e , this variable is also included in the expression of $\mu(\cdot)$. In fact, because the weak policymaker's inflation strategy at $t = 2$ will depend on $\mu(\cdot)$, and hence on π_1^e , we make use of similar notation for that variable, that is, the $t = 2$ inflation strategy is now denoted by $\pi_2(\pi_1, y_1; \pi_1^e)$.

The period-2 optimal choice for the weak policymaker solves:⁷

$$Max_{\pi_2} \int_{\underline{\epsilon}}^{\bar{\epsilon}} [\mu \underline{b}(\pi_2 - \pi_2^e) + (1 - \mu) \bar{b}(\pi_2 - \pi_2^e) - a \frac{\pi_2^2}{2}] f(\epsilon) d\epsilon$$

Intuitively, this means that the weak policymaker optimally weighs the constant marginal effect of unexpected inflation on output against the increasing marginal intensity of his dislike of inflation. The solution of this problem is:

$$\pi_2^*(y_1, \pi_1; \pi_1^e) = \frac{\mu \underline{b} + (1 - \mu) \bar{b}}{a} = \frac{E(b|y_1, \pi_1, \pi_1^e)}{a} \quad (5)$$

Note that while π_2^e appears in the policymaker's optimization problem, the solution to the problem does not depend on it. Nevertheless, period-2

⁷Notice that we have used the fact that $E(\epsilon_t) = 0$.

inflation expectations are still relevant because they affect the policymaker's payoff in period 2 and, hence, his behavior in period 1.

3.2.2 Inflation Expectations in Period 2

There are only two possible optimal choices for the weak policymaker in period 1. Either he poses as the strong policymaker and sets $\pi_1 = 0$ or he separates himself by choosing some $\pi_1 \neq 0$. Define q , $0 \leq q \leq 1$, as the probability the weak policymaker sets $\pi_1 = 0$. Also, recalling that $\alpha(\pi_1)$ is the public's posterior belief about policymaker's type being the weak one, note that $\alpha(0) = \frac{qp}{qp+(1-p)}$, with $0 < \alpha(0) < p$, and that $\alpha(\pi_1) = 1$ if $\pi_1 \neq 0$. Hence, the public's period-2 inflation expectation is:

$$\pi_2^e = \alpha\pi_2^*. \quad (6)$$

Therefore, using equations 5 and 6, one can show that the weak policymaker's payoff in period 2 is:

$$W_2(y_1, \pi_1, \pi_1^e) = a[(1 - \alpha) - \frac{1}{2}](\pi_2^*)^2 \quad (7)$$

3.2.3 The Learning Process

Figure 1 illustrates the process of learning about b when ϵ_t has a compact support and π_1^e is fixed. The dashed region represents the distribution of y_1 conditional on \bar{b} being the truth for each possible choice of π_1 . The light shaded region gives the same distribution conditional on \underline{b} . For sufficiently high or sufficiently low π_1 these regions do not overlap and all agents learn the truth with probability 1. For π_1 close enough to π_1^e the two support regions intersect in the dark shaded region. When output falls into this overlapping area there is only partial learning; when it falls outside there is full learning. In the case where ϵ_t has a support over the real line, the dark shaded region becomes the whole (y_1, π_1) plane.




Figure 1: Learning Regions.

As one can see from equation 4, the likelihood ratio $LR(y_1, \pi_1, \pi_1^e) \equiv \frac{f_{\bar{b}}}{f_b}$ plays a key role in the learning process, indicating how agents update beliefs after observing y_1 . We introduce the following properties for the likelihood function associated with the density function $f(\cdot)$ that apply throughout the paper:

- P.1) (Crossing property) There is a non-empty interval I such that $LR(y_1, \pi_1, \pi_1^e) = 1$ if $y_1 \in I$ and $LR(y_1, \pi_1, \pi_1^e) \neq 1$ if $y_1 \notin I$.
- P.2) (Monotonicity) If $\pi_1 \geq \pi_1^e$ the function $LR(y_1, \pi_1, \pi_1^e) = 1$ is non-increasing in y_1 ; if $\pi_1 \leq \pi_1^e$ it is non-decreasing.

Notice that in case where ϵ_t has compact support, P1 restrains the analysis to the case where the two supply curves always overlap, which also restrains the feasible choices of π_1 . Notice also that this property generalizes

the analysis to shock distributions like the uniform, with regions where $f(\cdot)$ is flat. Property *P.2* means that in the case $\pi_1 \geq \pi_1^e$ higher output growth is associated with higher probability on \bar{b} , and likewise, if $\pi_1 \leq \pi_1^e$ higher output growth is associated with higher probability on \underline{b} . This property is illustrated in figure 1.⁸

For simplification of the exposition, with the exception of section 5, properties *P.1* and *P.2* will be considered to hold *strictly*. That is, the interval I will be given by a single point, and the likelihood function $LR(y_1, \pi_1, \pi_1^e)$ will be taken to be strictly monotonic. Moreover, the support distribution of the shock ϵ_t will spread through the entire real line. As is illustrated during the analysis of the uniform case (see section 5), these restrictions should represent no loss of generality.

3.2.4 Optimal Choice in Period 1

We can express the full payoff function for the weak policymaker as:

$$W_1(\pi_1, \pi_1^e) = [\theta \underline{b} + (1 - \theta) \bar{b}] (\pi_1 - \pi_1^e) - a \frac{(\pi_1)^2}{2} + \beta a [(1 - \alpha) - \frac{1}{2}] E_{y_1} [(\pi_2^*)^2] \quad (8)$$

where $E_{y_1}(\cdot)$ stands for the expectation operator over all possible values of y_1 , and $\beta \in (0, 1)$ is a discount factor. Then, the policymaker will choose π_1 to maximize this payoff function.

Suppose the weak policymaker decides to reveal his weakness in period 1, that is, $\pi_1 \neq 0$ and $\alpha(\pi_1) = 1$. The next result shows that his choice will be biased towards inflation levels that generate small surprises for the public. In other words, in period 2 the policymaker will be penalized to the extent he has surprised the public in period 1.

Proposition 1 *Consider the case where $\pi_1 \neq 0$. If $\pi_1 \geq \pi_1^e$, the function $W_2(y_1, \pi_1, \pi_1^e)$ is strictly decreasing in π_1 . If $\pi_1 < \pi_1^e$ the function*

⁸These properties hold for most standard distributions such as the uniform distribution and the normal distribution.

$W_2(y_1, \pi_1, \pi_1^e)$ is strictly increasing in π_1 .⁹ $E_{y_1}[W_2(y_1, \pi_1, \pi_1^e)]$, treated as a function of π_1 , is single peaked and achieves its maximum at $\pi_1 = \pi_1^e$.

The proof is in the appendix. The key is to show that the closer the inflation level is to π_1^e the lower is the volatility of period-2 inflation in the sense of second order stochastic dominance. Since the central bank has no information advantage over the public, the former will be unable to explore the output-inflation trade-off on an improved basis in period 2 based on knowledge of b gleaned from inflation policies very different than π_1^e in period 1. In fact, a better information flow about b will only increase the volatility of both expected inflation and inflation itself. As the policy maker is risk averse toward inflation, once he reached period 2 he will prefer period-1 inflation levels that limit inflation volatility in period 2. In fact, the proof basically relies on showing that from a period-1 perspective inflation levels closer to π_1^e dominates those further away from it.

Interestingly, proposition 1 implies that when the weak policymaker chooses a separating strategy at $t = 1$, this policy will depend on expected inflation at that time. Hence, the optimal separating policy will solve the following trade-off. On the one hand the policymaker may try to take advantage of the immediate gains of high inflation, but on other hand he has to minimize the effects of upsetting the public expectations over $t = 2$ expected pay off.

4 Equilibrium Analysis

The equilibrium of the model is given by the strategies (π_1^*, π_2^*) , beliefs $\mu(\cdot)$ and expectations $\alpha(\cdot)$ as described in the previous section. In addition, the condition $\pi_1^e = qp\pi_1^*$ must hold, where q is the probability that the weak policymaker chooses a separating policy $\pi_1^* \neq 0$, and p is the prior probability

⁹Later we will show that π_1 will be either zero or strictly positive.

that the policymaker is weak. For simplicity, we focus on the pure strategy case, that is, either $q = 1$ (separating) or $q = 0$ (pooling).

The following result simplifies the analysis. Define $\pi^\theta \equiv \frac{(\theta)\underline{b} + (1-\theta)\bar{b}}{a}$. This is the inflation level that maximizes the first-period payoff in equation 8.

Lemma 1 *Any equilibrium has to be such that $0 \leq \pi_1^* < \pi^\theta$.*

Proof:

In a pooling equilibrium $\pi_1^* = 0 < \pi^\theta$, so the lemma is true. Hence we need to consider just the case of separating equilibrium.

Notice that in any separating equilibrium $\alpha = 1$, and so the effect of second-period expected payoff on equation 8 is unambiguously negative. It follows from that and from proposition 1 that for $\pi_1 > \pi_1^e$ the period-2 expected payoff is decreasing in π_1 , and for $\pi_1 < \pi_1^e$ it is increasing on π_1 .

Assume first that $\pi_1^* < 0$. In this case, the weak policymaker strictly prefers choosing $\pi_1 = \pi_1^e$ to any $\pi_1 < \pi_1^e$, since both the expected period-2 payoff and the period-1 payoff are increasing in π_1 for all $\pi_1 < \pi_1^e$ when $\pi_1 < 0$. However, since $\pi_1^e = p\pi_1$, such a π_1^* cannot be an equilibrium, and so for any potential separating equilibrium, it must be that $\pi_1^* \geq 0$.

Assume now $\pi_1^* > \pi^\theta$. In equilibrium $\pi_1^* > p\pi_1^* = \pi_1^e$ so the second-period payoff is strictly decreasing in π_1 for all $\pi_1 > \pi_1^e$. It is also true that the first-period payoff is strictly decreasing in π_1 for all $\pi_1 > \pi^\theta$. Hence, there is some π_1 , $\pi_1 \leq \text{Max}[\pi^\theta, \pi_1^e] < \pi_1^*$, such that the policymaker is strictly better off by choosing it rather than π_1^* . Hence, $\pi_1^* \leq \pi^\theta$. Moreover, if $\pi_1 = \pi^\theta$ the derivative of the first-period payoff with respect to π_1 is zero. The second-period payoff is strictly decreasing in π_1 beginning at $\pi_1 = \pi_1^e$, therefore if it is differentiable at π^θ , the derivative must be strictly negative at this point, so $\pi_1^* < \pi^\theta$. But $E_{y_1}[W_2(y_1, \pi_1, \pi_1^e)]$ is monotonic for $\pi_1 \geq \pi_1^e$ so it must be differentiable almost everywhere and so the proof is finished.

□

As we pointed out in the previous section information is harmful to the policymaker. Therefore, it is plausible to imagine that he would like to either blind himself to new information or else commit himself to ignore it. This distinction is useful for studying the effect that learning has on the choice of π_1 . We proceed by comparing the optimal choice for the policymaker of period-1 inflation when he cannot credibly commit to an information-insensitive rule with what his optimal choice would be if he were allowed to ignore or not respond to new information. We will refer to this latter, case as the “commitment case” and the former case as the “no-commitment case”.

Proposition 2 *In the commitment case the weak policymaker is more willing to reveal himself as weak than is the case in the no-commitment case, i.e., in the former case $\pi_1 \neq 0$ for a wider set of parameter values than in the latter case.*

Proof:

In any separating equilibrium, $\alpha = 1$, hence the effect of expected period-2 on equation 8 is unambiguously negative. In the no-commitment case, any period-2 optimal strategy $\pi_2^*(\cdot)$ is such that the following equality holds:

$$E(\pi_2^*(y_1, \pi_1; \pi_1^e)) = \frac{E(E(b|y_1, \pi_1; \pi_1^e))}{a} = \frac{E(b)}{a} = \frac{\theta \underline{b} + (1 - \theta) \bar{b}}{a}$$

In the commitment situation, we have $\pi_2^* = \pi^\theta \equiv \frac{\theta \underline{b} + (1 - \theta) \bar{b}}{a}$. Then, from Jensen’s inequality, it follows that $(\pi_2^*)^2 = (\pi^\theta)^2 = [E(\pi_2^*(y_1, \pi_1; \pi_1^e))]^2 \leq E[(\pi_2^*(y_1, \pi_1; \pi_1^e))^2]$, and the period-2 expected payoff in the commitment case is always larger than that in the no-commitment case.¹⁰ As the period-1 payoff is the same in both situations, for any separating inflation rate

¹⁰Note that the commitment period-2 strategy takes no argument. This was done on purpose, as it should not depend on either y_1 , π_1 nor π_1^e . For more details, see the discussion following proposition 3.

$0 < \pi_1 \leq \pi^\theta$ the weak policy maker will enjoy a larger overall payoff in the commitment case than in the non-commitment case.

The result then follows from the observation that the pooling payoff for the weak policy maker is the same regardless of the possibility of commitment. This is because in a pooling equilibrium $\pi^e = \pi_1 = 0$, implying that $LR(\cdot) = 1$ and, hence (4), $\mu(\cdot) = \theta$ for every y_1 . Since nothing about b is learned in a pooling equilibrium the possibility of committing to ignore new information is irrelevant in this case.

□

The next result gives a necessary and sufficient condition applicable to the commitment case ensuring that the only equilibrium is a pooling one. Note that, given result 2, this condition is also a sufficient condition for the equilibrium to be a pooling one in the no-commitment case.

Proposition 3 *In the commitment case, the weak policymaker chooses the pooling strategy $\pi_1 = 0$ if and only if $(1-p)\beta > \frac{1}{2}$. Moreover, if this condition is satisfied the weak policy maker will never choose a separating strategy in the no-commitment case.*

Proof:

In the commitment situation the weak policymaker always sets $\pi_2^* = \pi^\theta \equiv \frac{\theta(b) + (1-\theta)(\bar{b})}{a}$. Hence, if he separates in period 1, $W_1(\pi^\theta, \pi_1^e) = -a\pi_1^e\pi^\theta + a(\frac{1}{2} - \frac{\beta}{2})(\pi^\theta)^2$, and if instead he pools, $W_1(0, \pi_1^e) = -a\pi_1^e\pi^\theta + a[(1-p) - \frac{1}{2}]\beta(\pi^\theta)^2$. Thus, the pooling strategy dominates the separating one if and only if:

$$W_1(\pi^\theta, \pi^e) < W_1(0, \pi^e) \iff \frac{1}{2} < (1-p)\beta$$

□

As the proofs of the propositions 2 and 3 suggest, there is a important distinction between the commitment situation and the no-commitment

situation. As in many pure reputation model, in the commitment situation the weak policymaker's strategy does not depend on period-1 inflationary expectations. In particular, the optimum separating strategy will be the one that maximizes the period-1 pay off, $\pi_1^* = \pi^\theta$, and will not depend on the period-2 pay off. However, as noted in the previous section, this is in marked contrast to the no-commitment case, where the expected $t = 1$ inflation π_1^e plays a key role in case the weak policy maker decides to separate, as this variable determines how much will be learned, and hence how much period-2 policy will fluctuate.

To sum up this section we note that forces of conservatism work on the weak policymaker in two distinct ways. First, they push him toward of zero-inflation pooling equilibrium. Second, even if he separates, creating a positive inflation surprise, he still chooses less inflation than the myopically optimal level, π^θ . This last result runs against the standard one in the experimentation literature that suggest the bank should adjust from π^θ in the direction of more information, i.e., up.

5 The Uniform Case

To gain precision and clarify the exposition, we now consider a case that yields closed-form solutions. The analysis illustrates the effects of the structural uncertainty over the supply curve's slope on the equilibrium outcome. In particular, we identify a region in the parameter space representing structural uncertainty where the unique equilibrium is a pooling one, another region with a unique separating equilibrium and third region where there are multiple equilibria, with the pooling and separating equilibria co-existing.

Let ϵ_t be distributed uniformly, that is, $f(\epsilon_t) = \frac{1}{2\bar{\epsilon}}$ if $\epsilon_t \in [-\bar{\epsilon}, \bar{\epsilon}]$, and 0 otherwise. Thus, the likelihood ratio is such that $LR(y_1, \pi_1, \pi_1^e) = 1$ if y_1 is in the partial learning region (the dark shaded region in figure 1), and $LR(y_1, \pi_1, \pi_1^e)$ is either 0 or ∞ , depending on whether y_1 falls respectively in

the \underline{b} or \bar{b} full-learning zones (the dashed and light shaded regions in figure 1).¹¹ Equation 4 for the posterior probabilities becomes:

$$\mu(\pi_1, y_1; \pi_1^e) = \begin{cases} \theta & \text{if non-learning region} \\ 1 & \text{if } \underline{b} \text{ full learning region} \\ 0 & \text{if } \bar{b} \text{ full learning region} \end{cases}$$

Equation 5 for period-2 inflation $\pi_2^*(.)$ specializes to:

$$\pi_2^*(\pi_1, y_1; \pi_1^e) = \begin{cases} \pi_2^\theta \equiv \frac{\theta b + (1-\theta)\bar{b}}{a} & \text{non-learning region} \\ \pi_2^{\underline{b}} \equiv \frac{\underline{b}}{a} & \text{if } \underline{b} \text{ full learning region} \\ \pi_2^{\bar{b}} \equiv \frac{\bar{b}}{a} & \text{if } \bar{b} \text{ full learning region} \end{cases}$$

It is also useful to calculate the probability the economy falls in the non-learning region, $P(\pi_1, \pi_1^e)$. The expression for this is:

$$P(\pi_1, \pi_1^e) = \begin{cases} \text{Max}\{1 - \frac{(\Delta b)(\pi_1 - \pi_1^e)}{2\bar{\epsilon}}, 0\} & \text{if } \pi_1 \geq \pi_1^e \\ \text{Max}\{1 + \frac{(\Delta b)(\pi_1 - \pi_1^e)}{2\bar{\epsilon}}, 0\} & \text{if } \pi_1 < \pi_1^e \end{cases}$$

where $\Delta b \equiv \bar{b} - \underline{b}$. Note that if $\pi_1 \leq \underline{\pi}_1 \equiv \pi_1^e - \frac{2\bar{\epsilon}}{\Delta b}$ or if $\pi_1 \geq \bar{\pi}_1 \equiv \pi_1^e + \frac{2\bar{\epsilon}}{\Delta b}$ then there is complete learning with probability 1.¹² It follows that $P(\pi_1, \pi_1^e)$ has a triangular shape, where the upper vertex is at π_1^e , the side vertices are $\underline{\pi}_1$ and $\bar{\pi}_1$, and its support is the interval $[\underline{\pi}_1, \bar{\pi}_1]$

Given the expressions for $\mu(.)$, and $\pi_2^*(.)$ and $P(\pi_1, \pi_1^e)$, equation 8 becomes:

$$\begin{aligned} W_1(\pi_1, \pi_1^e) &= a\pi^\theta(\pi_1 - \pi_1^e) - a\frac{\pi_1^2}{2} + \beta a[(1 - \alpha) - \frac{1}{2}] \cdot \\ &\quad \cdot \{[\theta(\frac{\underline{b}}{\pi_2^e})^2 + (1 - \theta)(\frac{\bar{b}}{\pi_2^e})^2] - P(\pi_1, \pi_1^e) \underbrace{[\theta(\frac{\underline{b}}{\pi_2^e})^2 + (1 - \theta)(\frac{\bar{b}}{\pi_2^e})^2 - (\frac{\theta}{\pi_2^e})^2]}_{\gamma}\} \end{aligned} \tag{9}$$

Note that the constant γ in 9 is the unconditional variance of period-2 inflation policy, and that $a^2\gamma$ is the unconditional variance of b . As one can

¹¹The partial learning zone implies no learning at all in this case precisely because the likelihood ratio is constant in this region.

¹²In figure 1, these thresholds are the right and left extremes of the dark shaded region.

readily see, the appearance of γ in that expression - along with the triangular shape $P(\pi_1, \pi_1^e)$ - is just a version of result 1 to the current specialization of the model. That is, if the policymaker separates and $\alpha = 1$, it follows from $\gamma > 0$ that π_1 's closer to π_1^e will be preferred to those farther away.

Suppose the weak policymaker separates and sets $\pi_1 > 0$. He will then maximize equation 9 with respect to π_1 , the first order condition being:

$$\frac{dW_1}{d\pi_1} = (\pi^\theta - \pi_1) + \frac{dP(\pi_1, \pi_1^e)}{d\pi_1} \gamma \frac{\beta}{2} \geq 0. \quad (10)$$

Note that this function is not continuous at the points $\pi_1 = \pi_1^e$, $\underline{\pi}_1$ and $\overline{\pi}_1$, and hence the equality does not necessarily hold. More precisely, note that

$$\frac{dP(\pi_1, \pi_1^e)}{d\pi_1} = \begin{cases} \frac{\Delta b}{2\epsilon} \geq 0 & \text{if } \pi_1 \in [\underline{\pi}_1, \pi_1^e] \\ \frac{-\Delta b}{2\epsilon} \leq 0 & \text{if } \pi_1 \in [\pi_1^e, \overline{\pi}_1] \\ 0 & \text{if } \pi_1 \notin (\underline{\pi}_1, \overline{\pi}_1) \end{cases}$$

The discontinuity stems from the compactness of ϵ_t and appears in figure 1 in the form of kinks on the edges of the dark shaded region when π_1 is either, π_1^e , $\underline{\pi}_1$ or $\overline{\pi}_1$.¹³

Define $\phi \equiv \gamma\beta\frac{\Delta b}{4\epsilon} > 0$. There are three possible maximizing points for equation 9, depending on the parameters π^θ and ϕ , and on π_1^e . If $\pi^\theta - \phi > \pi_1^e$, the possible solutions are $\pi^\theta - \phi$, if this policy is on the support of $P(\pi_1, \pi_1^e)$, and π^θ , if this policy falls off the support of the same function. This is so because in this case these two policies satisfy the first order condition of the separating problem.

If $\pi^\theta - \phi \leq \pi_1^e$, the possible solutions are π_1^e and π^θ . In fact, if this inequality holds, it follows from equation 10 that $\frac{dW_1}{d\pi_1} > 0$ at $\pi^\theta - \phi$, so the maximum of $W_1(\pi_1, \pi_1^e)$ within the support of $P(\pi_1, \pi_1^e)$ is achieved at π_1^e .

¹³It may appear that at π_1^e , $\underline{\pi}_1$ and $\overline{\pi}_1$ the derivative $\frac{dP(\pi_1, \pi_1^e)}{d\pi_1}$ is ill defined, as it may assume different values depending on the direction of the approximations. However, one should note that at these points we are slightly abusing the definition of the derivative and extending it to the case where the left hand and right hand derivatives are different.

Moreover, for the same reason as in the previous paragraph, if π^θ falls outside the support of $P(\pi_1, \pi_1^e)$ it is also a possible solution for the separating problem.

For simplification of the analysis, we will restrict the range of the study to the case where $\pi^\theta \leq \frac{2\bar{c}}{\Delta b}$, that is, we require the demand shock to be large relative to the difference of the possible supply curve slopes. Given this assumption, we have that π^θ is always on the support of $P(\pi_1, \pi_1^e)$ in any equilibrium. This is so, because by this assumption together with lemma 1 imply that $\bar{\pi}_1 = \pi_1^e + \frac{2\bar{c}}{\Delta b} > \frac{2\bar{c}}{\Delta b} \geq \pi^\theta$ as $\pi_1^e \geq 0$. Note also that $\pi_1 = \pi^\theta - \phi$ will also be on the support of $P(\pi_1, \pi_1^e)$, as $\pi^\theta - \phi < \pi^\theta$. Hence, under the stated assumption the first order condition of the separating problem will be satisfied with equality only for $\pi^\theta + \phi > \pi_1^e$, and π^θ is never a policymaker's best action.

It is worth noting that the assumption $\pi^\theta \leq \frac{2\bar{c}}{\Delta b}$ creates a similar environment to the one we would have if the support of ϵ_t were unbounded, that is, a case where the interval I in the definition of property P.2 is always non-empty. In these environments any separating policy implies some learning, driving the policymaker to chose inflation lower than π^θ , the optimal separating policy in the commitment case. In contrast, if $\pi^\theta > \frac{2\bar{c}}{\Delta}$ then there will always be a threshold $\bar{\pi}_1$ above which the policymaker receives no additional information about b , making it possible that π^θ could be optimal.

We now consider the effect of mean-preserving risk spreads on the prior distribution of b on the equilibrium. That is, we gradually increase γ while keeping the expected value of b constant. These mean-preserving spreads on b will increase the value of ϕ , as both γ and Δb will increase, forcing the weak policy maker to chose lower levels of π_1 in case of separation as $\pi^\theta - \phi$ decreases. Moreover, mean-preserving spreads also decreases the willingness of the weak policymaker to separate himself from the tough one.¹⁴

¹⁴Referring to figure 1, this exercise implies “squeezing” the non-informative region,

The following proposition fully characterizes the equilibrium within the range of study.

Proposition 4 *Suppose that $(1-p)\beta < \frac{1}{2}$ and that $\pi^\theta - \phi \geq 0$. There are three equilibrium regions in (p, γ) -space, where changes in γ correspond to mean-preserving spreads on b :*

- i) A unique equilibrium region with a separating equilibrium $\pi_1^* = \pi^\theta - \phi$ and $\pi_1^e = p(\pi^\theta - \phi)$.*
- ii) A unique equilibrium region with a pooling equilibrium $\pi_1^* = 0$ and $\pi^e = 0$.*
- iii) A multiple equilibria region, with a separating equilibrium as in i) and a pooling equilibrium as in ii).*

We leave the details of the proof to the appendix but give here its main argument. As pointed out above, there are two policies to compare, a pooling one with $\pi_1 = 0$, and separating one with $\pi_1 = \pi^\theta - \phi$. The proof consists just in comparing $W_1(\pi^\theta - \phi, p(\pi^\theta - \phi))$ with $W_1(0, p(\pi^\theta - \phi))$ when $\pi_1 = \pi^\theta - \theta$ and $\pi_1^e = p(\pi^\theta - \theta)$ is the candidate equilibrium, and likewise, $W_1(0, 0)$ with $W_1(\pi^\theta - \phi, 0)$, when $\pi_1 = \pi^e = 0$ is the candidate equilibrium.

We summarize the main features of the equilibrium in figure 2. The figure describes the equilibrium for different mean preserving spreads on b and for different prior probabilities p of the policy maker being of weak type. We know from previous discussion that the larger is γ the less willing the policymaker will be to separate. In addition, and in common with standard reputation models, the larger is p the more willing to separate is the weak policy maker. Hence, the figure captures the basic tension between these two forces in the model.

increasing the information content of larger π_1 's as it moves away from π_1^e .

Consider first the case where the policy maker can commit to ignore any new information. For this situation, the relevant line is the vertical line at $p = 1 - \frac{1}{2\beta}$. From result 2, we know that to the left of this line there is a unique pooling equilibrium, and to the right a unique separating equilibrium. That is, the possible deterrent effect of high γ 's have no bite in the commitment case and the policymaker is only concerned with his reputation.

Matters change considerably when the policy maker is unable to commit against any new information. Define γ_{π^θ} as the value of γ such that $\pi^\theta = \phi$. This line divides the commitment separating area into two regions. For the region above this γ_{π^θ} , lemma 1 implies the pooling equilibrium is the unique equilibrium, which in the present context accounts to say that the deterrence effect of future information discovery is so strong that no separating equilibrium arises, regardless of p . For the region under this horizontal line, the result 4 holds, according to which this region can be separated into three subregions. The region at the bottom (dark shaded) where the separating equilibrium $(\pi^\theta - \phi, p(\pi^\theta - \phi))$ is the unique equilibrium; the region at the top, where the pooling equilibrium is the unique one; and the region in the middle of these two, where the two equilibria co-exist.

The main properties of the regions in between the straight lines $p = 1 - \frac{1}{2\beta}$ and $\gamma = \gamma_{\pi^\theta}$ (proven in the appendix) are as follows. Define $B^p(p, \gamma)$ as the set of (p, γ) such that $W_1(0, 0) = W_1(\pi^\theta - \phi(\gamma), 0)$, and analogously $B^s(p, \gamma)$ as set of (p, γ) such that $W_1(\pi^\theta - \phi(\gamma), p(\pi^\theta - \phi(\gamma))) = W(0, p(\pi^\theta - \phi(\gamma)))$. We show in the appendix that the pooling equilibrium region lies above $B^p(p, \gamma)$ and that the points $(1 - \frac{1}{2\beta}, 0)$ and $(1, \gamma_{\pi^\theta})$ belong to $B^p(p, \gamma)$.¹⁵ We also show that the separating region lies below $B^s(p, \gamma)$ and that the points $(1 - \frac{1}{2\beta}, 0)$ and $(1, \gamma_{\pi^\theta})$ belong to this region as well. Finally, we have

¹⁵It is also show in the appendix that $B^p(p, \gamma)$ implicitly defines an increasing relationship between γ and p .




Figure 2: Slope Mean-Preserving Spread ($\pi^\theta = \text{constant}$)

the property that the line $B^p(p, \gamma)$ lies below the line $B^s(p, \gamma)$, implying that the multiple equilibrium region is made up of the intersection of the separating equilibrium region and the pooling equilibrium region.

In light of these properties, the intuition for figure 2 is clear. As in pure reputation models, the higher is p the more likely is the weak policy maker to separate. In the present context, if $\gamma = 0$, the same result is replicated here. However, as γ grows larger, the learning effect grows in importance, shrinking the set of p for which separation occurs. In fact, by proposition 4, for values of γ high enough, the learning effect is so strong that separation occurs only when the policymaker is quite likely to be weak.

There does exist a region with multiple equilibria in between the separating and the pooling regions, where the combinations of γ and p are such that neither the learning effect through γ nor the reputation effect through

p are strong enough to dominate. If it were not for the dependence of the learning effect on expected inflation, as shown in proposition 1, we would observe a sudden change from a separating to a pooling equilibrium as γ increases, rather than having the two equilibria co-exist. That is, we would have $B^p(p, \gamma) = B^s(p, \gamma)$. But, as argued above, this is not case, as the curve $B^p(p, \gamma)$ lies above the curve $B^s(p, \gamma)$. Within this multiple equilibrium region the prevailing equilibrium may well depend on exogenous factors which can coordinate agents inflationary expectations either to $\pi_1^e = 0$ or $\pi_1^e = p(\pi^\theta - \phi)$. However, in either case the expectations will be such that the weak policy maker will be forced to make a choice such that those expectations will necessarily be fulfilled.

Finally, consider mean-preserving spreads on the shock ϵ_1 . The definition $\phi \equiv \gamma\beta\frac{\Delta b}{4\bar{\epsilon}}$ and the proof of proposition 4 make clear that these will make policy less conservative in two senses. First, there will be separating equilibria for a larger set of parameter values and pooling equilibria for a smaller set. Second, even in separating equilibria mean-preserving spreads push π_1^* closer to π^θ . In other words, increasing the variability of the shocks has the opposite effect of increasing the spread on the b 's.

6 Conclusion

We can envision various extensions of this work that could change or modify some of the results. For example, giving the central bank more limited control over monetary policy would be more realistic and might give it some scope for experimenting without completely losing its reputation. Another change that might lead to more experimentation would be to give the bank some scope for maintaining an information advantage over the public.¹⁶

¹⁶In this context some recent literature of the possible value of having central banks that are not completely transparent to the public could be relevant [Jensen (2001), Gersbach (2002)].

Despite these caveats, it is still interesting to take the model in its present form seriously and think about what it might imply about policy. In this spirit we emphasize our result that high structural uncertainty is grounds for conservatism. This idea sounds similar to that of Brainard (1967), recently reemphasized in Blinder (1998) who offers the following advice for central bankers: "...Estimate how much you need to tighten or loosen monetary policy to 'get it right.' Then do less." (Blinder[1998], p. 17]. However, while our conclusion is similar to Brainard's our reasoning is different. Brainard's parameter uncertainty is fundamentally unlearnable while ours is. Thus, our model suggests that central bankers who care about output should be conservative out of fear of learning too much about how the economy functions, and thereby introducing excessive instability. The less a central bank knows about economic structure the more conservative it should be *ceteris paribus*. Thus, new central banks, such as the European Central Bank and those for economies in transition from communism, should be among the most conservative whereas well-established banks such as the US Federal Reserve Bank can afford to be somewhat more experimental in their policies.

References

- [1] Backus, D. and J. Driffill, 1984, Inflation and reputation, *American Economic Review* 75, 530-538.
- [2] Barro, R., 1986, Reputation in a model of monetary policy with incomplete information, *Journal of Monetary Economics* 17, 1-20.
- [3] Balvers, R. and T. Cosimano, 1994, Inflation variability and gradualist monetary policy, *Review of Economic Studies* 61, 721-738.
- [4] Bean, C., 1999, Discussion of professor goodhart's lecture: 'Central bankers and uncertainty', *Bank of England Quarterly Bulletin* 39, 115-116.

- [5] Bertocchi, G. and M. Spagat, 1993, Learning, experimentation, and monetary policy, *Journal of Monetary Economics* 32, 169-183.
- [6] Blinder, A., 1998, *Central banking in theory and practice*, MIT Press, Cambridge MA.
- [7] Brainard, W., 1967, Uncertainty and the effectiveness of policy, *American Economic Review* 57, 411-425.
- [8] Caplin, A. and J. Leahy, 1996, Monetary policy as a process of search, *American Economic Review* 86, 689-702.
- [9] Evans, G. and S. Honkapohja, 2001, *Learning and expectations in macroeconomics*, Princeton University Press, Princeton NJ.
- [10] Evans, G. and S. Honkapohja, 2002, Expectations and the stability problem for optimal monetary policies, forthcoming in *Review of Economic Studies*.
- [11] Gersbach, H., 2002, On the negative social value of central banks' knowledge transparency, unpublished manuscript, University of Heidelberg.
- [12] Grossman, S., R. Kihlstrom, and L. Mirman, 1977, A Bayesian approach to the production of information and learning by doing, *Review of Economic Studies* 44, 533-547.
- [13] Honkapohja, S. and K. Mitra, 2002, Performance of monetary policy with internal central bank forecasting, working paper 127, European Central Bank.
- [14] Kydland, F. and E. Prescott, 1977, Rules rather than discretion: the inconsistency of optimal plans, *Journal of Political Economy*, 85, 473-491.

- [15] Jensen, Henrik, 2001, Optimal degrees of transparency in monetary policy-making, CEPR Discussion Paper, 2689.
- [16] Laffont, Jean-Jacques, 2002, The economics of uncertainty and information, MIT Press, Cambridge MA.
- [17] Prescott, E., 1972, The multi-period control problem under uncertainty, *Econometrica* 40, 1043-1058.
- [18] Sargent, T., 1999, The conquest of American inflation, Princeton University Press, Princeton NJ.
- [19] Wieland, V, 2000, Monetary policy, parameter uncertainty and optimal learning, *Journal of Monetary Economics*, 46, 199-228.

A Appendix

A.1 Proof of Proposition 1

Consider the case where $\pi_1 \geq \pi_1^e$. The opposite case being similar. Take $\pi_1'' > \pi_1' \geq \pi_1^e$. Since the period-2 payoff is given by a concave function in π_2 when $\alpha = 1$, it is enough to prove that $\pi_2^*(y_1, \pi_1''; \pi_1^e)$ is a mean-preserving spread of $\pi_2^*(y_1, \pi_1'; \pi_1^e)$.

Consider the function $g(\pi_1)$ as defined below:

$$g(\pi_1) = \int_{-\infty}^{\bar{w}} P(\pi_2^*(y_1, \pi_1'; \pi_1^e) \leq w) dw \quad (\text{A.1})$$

for any \bar{w} . To prove that $\pi_2^*(y_1, \pi_1'')$ is a mean-preserving spread of $\pi_2^*(y_1, \pi_1')$ it is sufficient to prove:

$$\int_{-\infty}^{\bar{w}} [P(\pi_2^*(y_1, \pi_1'; \pi_1^e) \leq w) - P(\pi_2^*(y_1, \pi_1''; \pi_1^e) \leq w)] dw = g(\pi_1') - g(\pi_1'') < 0$$

since $E(\pi_2^*(y_1, \pi_1; \pi_1^e)) = \frac{\theta \bar{b} + (1-\theta)\bar{b}}{a}$ for every π_1 (see Laffont [1989], pp. 25-26). Moreover, since $\pi_1'' > \pi_1'$, it is sufficient to prove $\frac{dg(\pi_1)}{d\pi_1} > 0$ for every $\pi_1 \geq \pi_1^e$.

Consider first a few useful expressions. Recall that the expression for the likelihood function is $L(y_1, \pi_1; \pi_1^e) = \frac{f_b(\cdot)}{f_{\bar{b}}(\cdot)}$, where $f_b(\cdot) \equiv f(y_1 - b(\pi_1 - \pi_1^e))$.

We calculate the derivatives:

$$\frac{\partial LR(y_1, \pi_1, \pi_1^e)}{\partial y_1} = [f_{\bar{b}}(\cdot) \frac{\partial f_b(\cdot)}{\partial y_1} - f_b(\cdot) \frac{\partial f_{\bar{b}}(\cdot)}{\partial y_1}] / (f_{\bar{b}}(\cdot))^2 \quad (\text{A.2})$$

$$\frac{\partial LR(y_1, \pi_1, \pi_1^e)}{\partial \pi_1} = -[b f_{\bar{b}}(\cdot) \frac{\partial f_b}{\partial y_1} - \bar{b} f_b(\cdot) \frac{\partial f_{\bar{b}}(\cdot)}{\partial y_1}] / (f_{\bar{b}}(\cdot))^2 \quad (\text{A.3})$$

Note that $\frac{\partial LR(y_1, \pi_1, \pi_1^e)}{\partial y_1} < 0$ if and only if $\psi \equiv [f_{\bar{b}}(\cdot) \frac{\partial f_b(\cdot)}{\partial y_1} - f_b(\cdot) \frac{\partial f_{\bar{b}}(\cdot)}{\partial y_1}] < 0$. Recall that $\pi_2^*(\cdot)$ is decreasing in $LR(\cdot)$ (equations 4 and 5) and the $LR(\cdot)$ is strictly decreasing in y_1 by property P.2. Thus, $\pi_2^*(y_1, \pi_1; \pi_2^*)$ is strictly increasing in y_1 and strictly decreasing in π_1 .

Define $h(w; \pi_1; \pi_1^e) = \{y_1; \pi_2^*(y_1, \pi_1; \pi_1^e) = w\}$. Then we have that this function is well defined for w within the range of possible period-2 inflation

rates because of property *P.2*. To calculate $\frac{\partial h(w, \pi_1)}{\partial \pi_1}$ precisely, first re-write $\pi_2^*(.) = [LR(.)\underline{b}\theta + \bar{b}(1 - \theta)]/[LR(.)\theta + (1 - \theta)]$ and the use this expression to obtain the following derivatives:

$$\frac{\partial \pi_2^*}{\partial \pi_1} = \frac{1}{\rho} \left\{ \frac{\partial LR(.)}{\partial \pi_1} \underline{b}\theta [LR(.)\theta + (1 - \theta)] - \frac{\partial LR(.)}{\partial \pi_1} \theta [LR(.)\underline{b}\theta + \bar{b}(1 - \theta)] \right\} \quad (\text{A.4})$$

$$\frac{\partial \pi_2^*}{\partial y_1} = \frac{1}{\rho} \left\{ \frac{\partial LR(.)}{\partial y_1} \underline{b}\theta [LR(.)\theta + (1 - \theta)] - \frac{\partial LR(.)}{\partial y_1} \theta [LR(.)\underline{b}\theta + \bar{b}(1 - \theta)] \right\} \quad (\text{A.5})$$

where $\rho = [LR\theta + (1 - \theta)]^2$. It follows from using equations A.2, A.3 and A.4, A.5 and that:

$$\frac{\partial h(w, \pi_1)}{\partial \pi_1} = \frac{dy_1}{d\pi_1} = -\frac{\frac{\partial \pi_2^*}{\partial \pi_1}}{\frac{\partial \pi_2^*}{\partial y_1}} = -\frac{\frac{\partial LR(.)}{\partial \pi_1}}{\frac{\partial LR(.)}{\partial y_1}} = \underbrace{\frac{[\underline{b}f_{\bar{b}}(.)\frac{\partial f_{\underline{b}}}{\partial y_1} - \bar{b}f_{\underline{b}}(.)\frac{\partial f_{\bar{b}}}{\partial y_1}]}{[f_{\bar{b}}(.)\frac{\partial f_{\underline{b}}}{\partial y_1} - f_{\underline{b}}(.)\frac{\partial f_{\bar{b}}}{\partial y_1}]}}_{\psi} \quad (\text{A.6})$$

Given the function $h(w, \pi_1) = y_1$, equation A.1 can be written as follows:

$$g(\pi_1) = \int_{-\infty}^{\bar{w}} P(y_1 \leq h(w, \pi_1)) dw$$

and from that follows:

$$\frac{dg(\pi_1)}{d\pi_1} = \int_{-\infty}^{\bar{w}} \frac{\partial P(y_1 \leq h(w, \pi_1))}{\partial \pi_1} dw \quad (\text{A.7})$$

Notice that the expression for $P(y_1 \leq h(w, \pi_1))$ is as below:

$$P(y_1 \leq h(w, \pi_1)) = \int_{-\infty}^{h(w, \pi_1)} \theta f(y_1 - \underline{b}(\pi_1 - \pi_1^e)) + (1 - \theta) f(y_1 - \bar{b}(\pi_1 - \pi_1^e)) dy_1$$

Hence, by applying Leibniz's rule:

$$\begin{aligned} \frac{\partial P(y_1 \leq h(w_1, \pi_1))}{\partial \pi_1} &= \\ &\{ \theta f(h(w, \pi_1) - \underline{b}(\pi_1 - \pi_1^e)) + (1 - \theta) f(h(w, \pi_1) - \bar{b}(\pi_1 - \pi_1^e)) \} \frac{\partial h(w, \pi_1)}{\partial \pi_1} \\ &- \left\{ \int_{-\infty}^{h(w, \pi_1)} \underline{b}\theta f'(y_1 - \underline{b}(\pi_1 - \pi_1^e)) + \bar{b}(1 - \theta) f'(y_1 - \bar{b}(\pi_1 - \pi_1^e)) dy_1 \right\} \\ &\{ \theta f(h(w_1, \pi_1) - \underline{b}(\pi_1 - \pi_1^e)) + (1 - \theta) f(h(w, \pi_1) - \bar{b}(\pi_1 - \pi_1^e)) \} \frac{\partial h(w, \pi_1)}{\partial \pi_1} \end{aligned}$$

$$\begin{aligned}
& -[\underline{b}\theta f(h(w, \pi_1) - \underline{b}(\pi_1 - \pi_1^e)) + \bar{b}(1 - \theta)f(h(w, \pi_1) - \bar{b}(\pi_1 - \pi_1^e))] = \\
& \theta\left(\frac{\partial h(w, \pi_1)}{\partial \pi_1} - \underline{b}\right)f_{\underline{b}}(h(\cdot), \pi_1; \pi_1^e) + (1 - \theta)\left(\frac{\partial h(w, \pi_1)}{\partial \pi_1} - \bar{b}\right)f_{\bar{b}}(h(\cdot), \pi_1; \pi_1^e)
\end{aligned} \tag{A.8}$$

Using equation A.6, we arrive at expressions for the coefficients of the equation A.8:

$$\begin{aligned}
\frac{\partial h(w, \pi_1)}{\partial \pi_1} - \underline{b} &= -[(\Delta b)f_{\underline{b}}(\cdot)\frac{\partial f_{\bar{b}}(\cdot)}{\partial y_1}]/\psi \\
\frac{\partial h(w, \pi_1)}{\partial \pi_1} - \bar{b} &= -[(\Delta b)f_{\bar{b}}(\cdot)\frac{\partial f_{\underline{b}}(\cdot)}{\partial y_1}]/\psi
\end{aligned}$$

where $\Delta b = \bar{b} - \underline{b} > 0$, $\psi < 0$ are as defined above. These expressions cannot be signed immediately because of the $\frac{\partial f_{\bar{b}}(\cdot)}{\partial y_1}$ and $\frac{\partial f_{\underline{b}}(\cdot)}{\partial y_1}$ terms. However, using the definition of ψ , we can re-write these expressions as:

$$\frac{\partial h(w, \pi_1)}{\partial \pi_1} - \underline{b} = \frac{-\Delta b}{[(f_{\bar{b}}\frac{\partial f_{\underline{b}}}{\partial y_1})/(f_{\underline{b}}\frac{\partial f_{\bar{b}}}{\partial y_1})] - 1} \tag{A.9}$$

$$\frac{\partial h(w, \pi_1)}{\partial \pi_1} - \bar{b} = \frac{-\Delta b}{[1 - (f_{\underline{b}}\frac{\partial f_{\bar{b}}}{\partial y_1})/(f_{\bar{b}}\frac{\partial f_{\underline{b}}}{\partial y_1})]} \tag{A.10}$$

Now notice that:

$$\psi \equiv [f_{\bar{b}}(\cdot)\frac{\partial f_{\underline{b}}(\cdot)}{\partial y_1} - f_{\underline{b}}(\cdot)\frac{\partial f_{\bar{b}}(\cdot)}{\partial y_1}] < 0 \iff (f_{\underline{b}}\frac{\partial f_{\bar{b}}}{\partial y_1})/(f_{\bar{b}}\frac{\partial f_{\underline{b}}}{\partial y_1}) > 1$$

Therefore both equations A.9 and A.10 are positive. It follows that equation A.8 is positive and so is equation A.7.

A.2 Proof of Proposition 4

We start by restating the value function for the weak policy maker, since the proof heavily depends on it:

$$\begin{aligned}
W_1(\pi_1, \pi_1^e) &= a\pi^\theta(\pi_1 - \pi_1^e) - a\frac{\pi_1^2}{2} + \beta a[(1 - \alpha) - \frac{1}{2}] \cdot \\
&\quad \cdot \underbrace{\{[\theta(\pi_2^b)^2 + (1 - \theta)(\pi_2^{\bar{b}})^2] - P(\pi_1, \pi_1^e)[\theta(\pi_2^b)^2 + (1 - \theta)(\pi_2^{\bar{b}})^2 - (\pi_2^\theta)^2]\}}_{\gamma}
\end{aligned} \tag{A.11}$$

where

$$P(\pi_1, \pi_1^e) = \begin{cases} \text{Max}\{1 - \frac{(\Delta b)(\pi_1 - \pi_1^e)}{2\bar{\epsilon}}, 0\} & \text{if } \pi_1 \geq \pi_1^e \\ \text{Max}\{1 + \frac{(\Delta b)(\pi_1 - \pi_1^e)}{2\bar{\epsilon}}, 0\} & \text{if } \pi_1 < \pi_1^e \end{cases}.$$

Define as in section 0.4 $\phi = \gamma \frac{\Delta b}{4\bar{\epsilon}}$ and note that for mean-preserving spreads on b , $\frac{d\phi(\gamma)}{d\gamma} > 0$.

Item i):

For this region it must be that $W_1(\pi^\theta - \phi, p(\pi^\theta - \phi)) \geq W_1(0, p(\pi^\theta - \phi))$.

This relationship holds if and only if the following set of inequalities hold:

$$\begin{aligned} & \frac{(\pi^\theta)^2}{2} - \frac{(\phi)^2}{2} - \frac{\beta}{2}[\theta(\pi^b)^2 + (1 - \theta)(\pi^{\bar{b}})^2] - P(\pi^\theta - \phi, p(\pi^\theta - \phi))\gamma \\ \geq & \beta[(1 - p) - \frac{1}{2}][\theta(\pi^b)^2 + (1 - \theta)(\pi^{\bar{b}})^2] - P(0, p(\pi^\theta - \phi))\gamma \iff \\ & \frac{(\pi^\theta)^2}{2} - \frac{(\phi)^2}{2} - \frac{\beta}{2}[P(0, p(\pi^\theta - \phi)) - P(\pi^\theta - \phi, p(\pi^\theta - \phi))]\gamma \\ \geq & \beta(1 - p)[1 - [P(0, p(\pi^\theta - \phi))]\gamma + (\pi^\theta)^2] \end{aligned}$$

Note that from the expression for $P(\pi_1, \pi_1^e)$, $\frac{\beta}{2}\gamma[P(0, p(\pi^\theta - \phi)) - P(\pi^\theta - \phi, p(\pi^\theta - \phi))] = \phi[(\pi^\theta - \phi)(1 - 2p)]$ and $\gamma\beta[1 - P(0, p(\pi^\theta - \phi))] = 2p\phi(\pi^\theta - \phi)$. Substituting these equalities into the last inequality above, we have a new set of relationships:

$$\begin{aligned} \frac{(\pi^\theta)^2}{2} - \frac{(\phi)^2}{2} - \phi(1 - 2p)(\pi^\theta - \phi) & \geq 2(1 - p)p\phi(\pi^\theta - \phi) + \beta(1 - p)(\pi^\theta)^2 \\ \iff & \frac{(\pi^\theta - \phi)^2}{2} + 2p^2(\pi^\theta - \phi)\phi - \beta(1 - p)(\pi^\theta)^2 \geq 0 \end{aligned} \tag{A.12}$$

Define $B^s(p, \phi)$ as when inequality A.12 holds as an equality. Since ϕ is a (increasing) function of γ , define as well $B^s(p, \gamma) \equiv B^s(p, \phi(\gamma))$. Then the following properties follow:

1. The inequality A.12 holds for $\gamma = 0$ and every $p \geq 1 - \frac{1}{2\beta}$.

To see this, take $\phi = 0$. Then, it is easy to see that $p = 1 - \frac{1}{2\beta}$ is the

solution of the equation $B^s(p, 0)$, and for every $p \geq 1 - \frac{1}{\beta}$ inequality A.12 holds. But $\gamma = 0$ implies $\phi = 0$, thus $\gamma = 0$ implies that A.12 holds for all $p \geq 1 - \frac{1}{2\beta}$.

2. Define γ_{π^θ} as the value of γ such that $\pi^\theta = \phi$. Then if $p = 1$, the inequality A.12 holds for every $\gamma \leq \gamma_{\pi^\theta}$.

Take $p = 1$ and note that in this case $\phi = \pi^\theta$ is the solution of the equation $B^s(1, \phi)$, and that inequality A.12 holds for every $\phi \leq \pi^\theta$. As $\frac{d\phi}{d\gamma} > 0$, for $p = 1$ there is a number γ_{π^θ} such that A.12 holds for every $\gamma \leq \gamma_{\pi^\theta}$, where γ_{π^θ} is as defined above.

Notice that the definitions for $B^s(p, \gamma)$ made here and the one in section 3 are exactly the same. Hence, from the properties just demonstrated, it follows that the curve $B^s(p, \gamma)$ is as in figure 2, and so the separation region is as depicted there.

Item ii):

In this case, we have to analyze the values of p and γ in the (p, γ) space such that $W_1(0, 0) \geq W_1(\pi^\theta - \phi, 0)$ holds. Using the value function's expression, we have that the following set of inequalities must hold:

$$\begin{aligned} \beta[(1-p) - \frac{1}{2}](\pi^\theta)^2 &\geq \frac{(\pi^\theta)^2}{2} - \frac{(\phi)^2}{2} - \frac{\beta}{2}[\theta(\pi^b)^2 + (1-\theta)(\pi^{\bar{b}})^2 - P(\pi^\theta - \phi, 0)] \\ &\iff \beta[(1-p)](\pi^\theta)^2 \geq \frac{(\pi^\theta)^2}{2} - \frac{(\phi)^2}{2} - \frac{\beta}{2}[1 - P(\pi^\theta - \phi, 0)]\gamma \end{aligned}$$

As $\gamma \frac{\beta}{2}[1 - P(\pi^\theta - \phi, 0)] = \phi(\pi^\theta - \phi)$, the inequality above can be written as below:

$$\beta[(1-p)](\pi^\theta)^2 - \frac{(\pi^\theta - \phi)^2}{2} \geq 0 \quad (\text{A.13})$$

As in the previous item, define $B^p(p, \phi)$ as when inequality A.13 holds as an equality, and set $B^p(p, \gamma) \equiv B^p(p, \phi(\gamma))$. Again, following steps analogous as those in *item1*, the following properties can be shown:

1. If $\gamma = \gamma_{\pi^\theta}$ (that is, $\pi^\theta = \phi$), the inequality A.13 holds for all $0 \leq p \leq 1$.

2. If $\gamma = 0$ ($\phi = 0$), inequality A.13 holds for all $0 \leq p \leq 1 - \frac{1}{2\beta}$.
3. The curve $B^p(p, \gamma)$ defines a function implicitly, $p(\gamma)$, such that $\frac{dp(\gamma)}{d\gamma} > 0$. In fact, from $B^p(p, \phi)$ we can find that $\frac{dp}{d\phi} = \frac{\pi^\theta - \phi}{\beta(\pi^\theta)^2} > 0$. As $\frac{d\phi}{d\gamma} > 0$, then $\frac{dp}{d\gamma} > 0$.

Hence, the pooling equilibrium region is as depicted in figure 2 and we verify that $B^p(p, \gamma)$ as defined in section 5 is just $B^p(p, \gamma) = B^p(p, \phi(\gamma))$.

Item iii):

Take any point $(\hat{p}, \hat{\gamma})$ that satisfies the equation $B^p(\hat{p}, \phi(\hat{\gamma}))$. As $\phi(\theta - \phi) > 0$, it follows that inequality A.12 is satisfied for $(\hat{p}, \phi(\hat{\gamma}))$. Thus, the set of (p, γ) that satisfies $B^p(p, \gamma)$ lies below the one that satisfies $B^s(p, \gamma)$. Moreover, the multiple equilibria region is non-empty and lies on the intersection of the pooling and the separating regions, as indicated in figure 2.