The Optimal Monetary Policy Rule: the Role of Asymmetries and Reputation*

Filip Rozsypal†

October 16, 2011

Abstract

I study an optimal policy rule of a central bank under an asymmetric information reputation building setting. I find that the optimal reaction function is nonlinear, despite a standard quadratic loss function and linear Phillips curve. The asymmetry arises from the signal extraction problem of the agents. The central bank is modelled as having standard quadratic preferences with its bliss point at output higher than the natural rate. The central bank maximizes the welfare by optimizing over possible announcements of the strategy it plans to follow.

1 Introduction

The central bank is undoubtedly a very important economic policy institution which significantly affects macroeconomic conditions for all agents in the economy. In order make decisions and act optimally, agents therefore need to have a good sense of the central bank's behavior. Learning about the central bank true objectives from economic outcomes is not straightforward as the transmission mechanism involves lags and noise. Consequently, estimation of monetary policy rules and loss functions is challenging task with is suppose to help economic agents to understand and predict the central banks' behavior.

The recent revival of interest in estimation of policy rules brought about more advanced econometric techniques designed to extend our knowledge of the monetary policy conduct, in particular, new methods have been used to detect any form of asymmetry, either in policy rules themselves, or in the underlying preferences of the monetary authority. These techniques have potential to better describe how central banks have acted in the past which in turn improves understanding of current and future monetary policy. While documenting the nonlinearities in the policy rule is a standard econometric exercise, the preferences of the central bank are not directly observed and hence not directly measurable. The standard approach is to model some preferences of the central bank, derive the optimal behavior and then test if these implications are supported by the data. The working hypothesis seems to be that it is the asymmetric loss function that drives the asymmetric policy function. In this paper, I argue that this solution raises some methodological concerns and I propose an alternative explanation for the origin of the nonlinear policy functions.

*I would like to thank to Chryssi Giannitsarou and Petra Geraats for helpful suggestions.

†Faculty of Economics, University of Cambridge, email: fr282@cam.ac.uk

In the tradition of the Lucas critique, a sound model should be based on microfoundations rather than relying on *ad hoc* fixes, which might be well mimicking the observed behavior in the past, but might fail completely if the underlying conditions change. In the present context, this means that the actions of a central bank should be a result of its own intertemporal optimization problem, rather than result of a creative alteration of central bank's preferences.

To illustrate this point, consider the concept of inflationary bias, a situation of unproductively high inflation with no movement towards the socially optimal level of output (Kydland and Prescott, 1977). Some authors in the literature implicitly dismissed this problem by arguing that because the central bankers know what the outcome would be, they would not try to follow such policies and hence the preferences considered by Kydland and Prescott do not describe the real behavior of a modern central bank (DeLong, 1995; Blinder, 1998). Indeed, in the Kydland-Prescott setting, the inflationary bias arises only if the central bank targets higher then natural level of output.

This approach is a prime example of an ad hoc change in the preferences of the central bank. While it is probably true that economists working in central banks do know the literature on dynamic inconsistency and inflationary bias, this fact does not imply that their preferences have changed to accommodate this possibility. Should there be a magic policy which somehow eliminated the problem, the central bank should and would use such a policy to attain the socially optimal level of output. In another words, if there is a view that, for whatever reason, the output is lower than what would be the social optimum, then the central bank preferences should have bliss point at this socially optimal level of output, in spite of the fact that there is no policy which would deliver such an allocation.

In line with a micro-founded approach to the behavior of the central bank, in this paper I consider a model where the central bank has a standard quadratic loss function targeting the socially optimal level of output. The Kydland-Prescott time inconsistency is addressed by a reputation building repeated game between the central bank and a representative agent. The central bank announces the strategy it plans to follow and the representative agent sets her inflation expectations as a function of the model fundamentals and trustworthiness of the central bank. The fact that the representative agent observes only a noisy signal of the allocation chosen by the central bank complicates the reputation building mechanism and hence the representative agent uses maximum likelihood to extract the signal from the noisy observed allocation.

The optimal behavior of the central bank turns out to be a nonlinear, despite the fact that the preferences of the central bank are modelled as quadratic and the Phillips curve is linear. This is outcome of the fact that private agents have imperfect information about the central bank's behavior. The central bank finds such a deviation from the policy rule which does not hamper the rationality of the inflation expectations of the representative agent while it decreases the expected loss.

The research program of estimating central banks' policy and loss functions has witnessed a vigorous revival in early 2000's. If the loss function is indeed asymmetric, as Nobay and Peel (2003) and Cukierman and Gerlach (2003) have shown, a new type of inflation bias would arise, even if the central bank targets the natural level of output. The inflation bias is clearly suboptimal and policies should be designed so the bias is minimized.

However, such policies would be welfare detrimental, if the asymmetric behavior is driven by something else then the asymmetric preferences. Here, I argue that observed nonlinearities and existence of inflation bias can be addressed without altering the standard preferences. While there is no reason why a symmetric quadratic loss function is necessarily the right description of central banks preferences, it is important to understand it properly and see if and under what circumstances asymmetries can arise even without assuming some asymmetric loss function.

2 Literature review

2.1 Inflationary bias - the beginning

My model builds on the simplest possible setting with the inflationary bias. This literature was initiated by Kydland and Prescott (1977), who first noticed that even a perfectly benevolent policy maker with the same preferences as the rest of the society might not deliver an optimal policy. The reason for this failure is that the optimal monetary policy plan is not sub-game perfect. In the setting where a central bank takes inflation expectations as given when it chooses the inflation-output gap allocation. After the inflation expectations has been fixed, the central bank might find it optimal to increase the inflation above the expected level in order to boost output and thus deviation from the *ex ante* optimal policy. Perfectly rational agents foresee this lack of commitment and set their inflation expectations to a value where central bank does not find it profitable to push it even further. The outcome is positive inflation while the output at its natural level. This situation can be avoided if the central bank can credibly commit itself to following some rule. Without commitment, the inflationary bias arises.

The insights of Kydland and Prescott led to voluminous subsequent research.[1] Barro and Gordon (1983) showed that the inflationary bias can be eliminated even without any commitment, if the model is viewed as an infinite repetitive game. The "cooperation" is enforced by a reputation building repeated game.

Rogoff (1985) suggested that the bias can be reduced if the central bank is more hawkish then the rest of the society, or in another words, if the central bank is more penalized for inflation deviations from its target then general public. There are important reservations to the solution proposed by Rogoff. First, within the model, the change in the preference parameter generates a wedge between the policy which would be chosen normally and the one which would be chosen by the conservative central banker. The bigger is the supply shock (in absolute terms), the bigger is the difference in allocations is and hence the bigger is the cost paid for lowering the inflationary bias is. In another words, the conservative central banker implements suboptimal trade-off between inflation and output stabilization. Second, Rogoff did not consider the credibility of the arrangement, which might be an issue if the central bank is not perfectly independent. Finally, from a methodological perspective, introducing a conservative central banker is a technical short-cut in the sense that given the preferences of the society, the optimal monetary policy should be modelled as a optimal solution to the central bank problem.

Flood and Isard (1988) observed that the the costs of installing a Rogoff style conservative central banker are increasing in the magnitude of the supply shock. As a solution they suggested an exit clause which would allow a policymaker (a politician appointing the central banker) to tem-

[1] For general overview and discussion of inflationary bias literature see Walsh (2001, chapter 8).

porarily remove the conservative central banker from power at some fixed cost for the policymaker. For low magnitude shocks, the conservative central banker successfully implements hawkish monetary policy, extreme shocks are accommodated by the policy maker and the credibility of the central bank is preserved.

More recently, Lohmann (1992) extends the model of Flood and Isard by modeling the interplay between the policymaker and the central banker. The central banker would prefer setting a hawkish allocation, but if such policy is chosen, it would be overridden by the policymaker. However, as the central banker knows the (fixed) cost of removing her from office, she can set her policy such that marginal benefits of invoking the exit clause are zero. Such policy is still more preferable for the central banker than the allocation that would be chosen by the policymaker alone. In equilibrium, the central banker is never removed from the office. Lohmann shows that this setting generates a nonlinear policy rule. Lohmann's contribution is significant to the problem outlined in this paper, because it shows how an asymmetric reaction function can arise even in the setting with quadratic loss function and a linear Phillips curve. However, the solution proposed by Lohmann still relies no the short-cut notion of the conservative central banker, whereas my model is fully micro founded. Furthermore, Lohmann's solution does not completely remove the inflationary bias neither it is true that the trade-off between the inflation and output is optimal, whereas these points are addressed in my setting by allowing for any possible trade-off and any possible output target announcement.

Walsh (1995) suggested that a suitably written contract with the governor of the central bank might credibly prevent the central bank from discretionary behavior. The optimal contract can be difficult to implement from both theoretical (it might depend on potentially unobservable characteristic of the central banker) or practical (the performance bonus is low to any outside option the governor is very high already so any change in the monetary incentive of plausible magnitude is not going to affect the behavior by much).

Finally, if the central bank targets the natural rate of output, then the inflationary bias does not arise, because the ex ante policy is no longer sub-game imperfect. This approach is sometime called *target* conservative (as opposed to the term *weight* conservative that refers to changing the weights in the loss function proposed by Rogoff (1985)). The implication of this stream of literature was to increase the independence of the central bank and hence insulate the monetary policy from political influence. Blinder (1998) argued that because the central bankers know the literature on inflationary bias, they do not attempt to target socially optimal level of output and hence the discretionary policy does not induce the inflationary bias, because there is no incentive to deviate and boost the economy. While it is true that the central bank now implements the correct trade-off between inflation and output stabilization, it is still the case that the Blinder's suggestion is merely a shortcut delivering a reasonable behavior rather a truly micro-founded solution. As such, it in some sense resembles calvo pricing; it provides a computationally simple setting which captures the main behavioral aspects at cost of not modeling the microfoundations. The fact, that observed behavior is close to the one which would be implemented by Blinder's central bankers, does not necessarily constitute that the central banker in charge is has preferences with the bliss point at natural level of unemployment/output.

There have been attempts to estimate the existence and importance of the inflationary bias empirically. Ireland (1999) uses a version of Kydland-Prescott model and tests its predictions on inflation and unemployment data of the US since 60's. He assumes that the natural rate of unemployment follows a unit root process. The result of the central bank optimization implies that the inflation is cointegrated with the natural rate of unemployment. Ireland then tests the restrictions generated by the cointegration relation and finds that the long run predictions of the inflationary bias model are supported by the data. Indeed, when looking on smoothed inflation time series in the US, there is a clear difference between the early post-war period of 1960's to early 1980, where the smoothed time series were trending up, and from late 1980's onwards, where the trend was reversed and inflation rates started to decline.

However, the predictions of the model for the short run are rejected. Ireland argues that this not surprising and this result should be used as a evidence against the models as such: "*Given the large number of restrictions... it comes as no surprise that statistical tests reject the constrained model. After all, the underlying theory is described by three simple equations... it would be unreasonable to expect this simple structure to account for all of the dynamics that can be found in the data.*"(Ireland, 1999, page 289). The original models of Kydland and Prescott and Barro and Gordon were not meant to explain short term fluctuations in inflation. Furthermore, the way they capture the gap between the target and the natural rate is constructed to be a simple representation. The fact that the inflationary bias depends on the absolute value of the unemployment which is modeled as time varying, is thus rather an artefact of the adapted mathematical formalism, rather than a deliberate intention of the original authors.

2.2 Models of reputation with asymmetric information

The second main ingredient of my model is the informational asymmetry between the representative agent and the central bank. I model the central bank as observing the true shocks and the public as observing only a noisy signal of the allocation chosen by the central bank. Romer and Romer (2000) support this view by documenting that the forecast of the FED were much better then the forecast of all other entities. I abstract from intertemporal aspects introduced by imperfect control which necessitates forward looking. However, in the spirit of Romer and Romer results, the central bank is still better informed about the economic fundamentals then the representative agent.

The literature in the past focused typically on settings, where the central bank does not observe the economic fundamentals perfectly. Such models hence try to capture different aspect of reality then my model. However, many insights about the mechanism of reputation building for repeated monetary policy games with asymmetric information are present in my model as well so it is worth summarizing and contrasting the main contributions. Furthermore, reviewing the literature of imperfect information sheds light on why asymmetry arises in my setting while it does not typically arise in the other models with asymmetric information.

Canzoneri (1985) commented on results of Barro and Gordon by observing that the reputation building is complicated in the presence of imperfect information. In his model, the goal of the central bank is to accommodate shocks in the money demand by printing money. However, the central bank observes only a noisy signal about future money demand. The private agents observe the final allocation and decide whether they believe that the outcome potentially suggesting a deviation is a deliberate attempt to inflate the economy, or only a result of a mistaken forecast on a part of the central bank.

Canzoneri constructs the equilibrium in a similar fashion as in Barro and Gordon (1983). However, he shows that in order to enforce the cooperative equilibrium, the agents will set up a threshold level of inflation and they would not interpret an excessive value as an attempt of cheating unless it higher than the threshold level. In another words, the threshold governs how bad forecast the agents are willing to tolerate before starting to suspect that the central bank has deviated. Interestingly, in equilibrium it is never optimal for the central bank to try to cheat, yet because the forecast error is stochastic, there are some periods where agents trigger the punishment.

Given the context of nonlinear loss and policy functions, it is worth noting that there cannot be any asymmetry in the reaction function in Canzonery's setting simply because the bank sets its policy before the uncertainty is realized. Although closely related, my paper differs from that of Canzoneri in one crucial aspect. Canzoneri assumes that the central bank has imperfect information about future demand. In my model, the central bank has complete knowledge of fundamentals and it is the agents who observe only a noisy signal about the allocation chosen by the central bank (the noise shock is realized only after the central bank chooses the allocation).

Backus and Driffill (1985) explore the implications of uncertainty of the agents about the preferences of the central bank. In their model, there are two types of central banks which differ in terms of preferences towards inflation. The agents update their beliefs based on the observed outcomes. Backus and Driffill show that even a central bank of an inflationist type might find it profitable to pretend that its type is hawkish. The model also predicts that the disinflation policies would be more costly for a bank with relatively worse reputation.

Cukierman and Meltzer (1986) effectively combine the both information asymmetries of Canzonery then Backus-Driffill models. They explore a setting where the agents do observe neither the direct action of the central bank nor the time-varying preferences of the central bank. The public observes the outcomes, updates its beliefs about the state of the central bank's preferences and adjusts its inflation preferences accordingly. Under this setting, Cukierman and Meltzer's show that it using the most precise instrument possible is not necessarily the instrument minimizing the loss of the central bank.

There are two ingredients to this result. First, the more precise the instrument is, the faster the public learns about the preferences of the central bank by observing its actions. Second, the stochastic and persistent preferences of the central bank make the relative benefits of inflating and boosting the economy time varying. With these two combined, when the central bank is hit by a preference shock which increases its desire to inflate the economy, a relatively less precise instrument means that public will need more time to learn about the shift in the central bank's preferences and will ultimately adjust the inflation expectations less and later. A less precise instrument thus slows down the reaction of the inflation expectations making the expansionary policy more efficient. Cukierman and Meltzer can be viewed as one of the first paper studying the role of transparency in monetary policy. This stream of literature became very influential as more central banks adopted inflation targeting as a framework for conducting monetary policy. For an overview of the literature on transparency, see Geraats (2002).

In a review chapter, Stokey (2003) takes Canzoneri's insight and analyzes the optimal instrument choice. If there are two instruments available, it is generally better to use the one which has higher correlation with the target variable. However, if the other one is much more observable and hence the private agents can use it in order to verify central banks actions, Stokey argues that it might be better to use the less correlated, but more transparent instrument.

2.3 Asymmetric loss function and new theory of inflationary bias

The new literature on the inflationary bias shows that bias can arise in the setting where the central bank is targeting the natural rates, but the loss function is asymmetric. My paper thus offers an alternative to the papers discussed in this section.

Nobay and Peel (2003) explore the implication of having an asymmetric loss function and study the situation where the loss function has a linex form.[2]. They note that with asymmetric preferences, it is not necessarily the case that targeting the natural rate is welfare maximizing, depending on the skewness of the loss function (affecting the direction) and the variance of the supply shock.

Cukierman and Gerlach (2003) and Cukierman and Muscatelli (2008) provide a deeper economic intuition to the argument of Nobay and Peel argument. Their explanation of the bias is as follows: Consider a central bank which targets the natural rate of unemployment, but has an asymmetric loss function, so unemployment above the target is relatively more harmful than a deviation of the same magnitude in the opposite direction. Second, assume that the central bank chooses the policy before observing the supply shock. Given the skewed expected losses, the central bank prefers to err on the negative side of unemployment gap and hence chooses an accommodative policy. The public recognizes this point and incorporate this effect into inflation expectations, making the inflationary bias even bigger. The model has testable implications, because it predicts that the bigger the variance of supply shocks is, the bigger is "precautionary" action taken by the central bank and hence the bigger is the resulting inflationary bias is.

Cukierman and Muscatelli (2008) estimate a nonlinear policy rule using smooth transition regression with hyperbolic tangent function, documenting a concave Taylor rule for the UK prior to adoption of inflation targeting in 1992. Cukierman and Muscatelli also report interesting switch in the asymmetry with respect to inflation in the US, where McChesney-Martin tenure at the FED could be characterized as having asymmetrically high losses for high inflation, whereas under Greenspan's chairmanship the FED could be characterized as having asymmetrically higher loss from output gap.

Ruge-Murcia (2003) test whether the inflationary bias is better explained by the standard setting where central banks targets positive output gap, or rather where the inflationary bias is caused by the asymmetric preferences of the central bank. The linex loss functions nests the quadratic preferences as a special case, so it is possible to discriminate the two based on testing restrictions on the parameters in the functional form. The identification relies on the fact, that past unemployment is informative for future inflation. However, this relation is linear for the standard inflationary bias model, whereas it is nonlinear for the Cukierman model. Ruge-Murcia finds that the standard inflationary bias is rejected by the data, whereas the Cukierman model is not. Ruge-Murcia (2004) extends the analysis Canada, France, Italy, Japan, UK and the US and obtains mixed results. However, as pointed Ireland (1999), the Barro-Gordon model is a very

[2] Linex form: $g(x) = \frac{\exp(\gamma x) - \gamma x - 1}{\gamma^2}$, where $x = u - ku^*$, u^* is the natural rate of unemployment

simple model and hence it would be too ambitious to require it to explain fully short term swings in inflation.

Surico (2007) explores the asymmetric preferences of the FED. He starts with a general loss function where the quadratic terms in unemployment and inflation are exchanged by a linex functions. He then tests whether and which part of the loss function is asymmetric and during which time period. He finds the FED's loss function to be asymmetric with respect to the output from 60's to early 80's, but there is no asymmetry detected with respect to inflation in the whole sample.Surico thus finds that the Cukierman's style inflationary bias is present only in the pre-Volcker period and it caused an increase of inflation by 1.48 percentage points. Surico (2008) confirms these results using slightly different methodology, finding that the implicit inflation target declined from 3.81% to about 2%, the inflationary bias was about 1 percentage point in the period from 1960's to early 1980's and the bias disapeared in the period starting in the late 1980's.

Dolado et al. (2004) analyze the optimal monetary policy rule for a setting where the central bank has a linex loss function loss function and the Phillips curve is also nonlinear. The estimates of the model on the US data suggest that the Phillips curve is linear, but there is nonlinearity in the loss function of the FED after 1983. Dolado et al. (2005) extend the analysis to data from France, Germany, Spain and the Eurozone as whole and the US. The main finding is that the central banks behave nonlinearly since 1980's, reacting more strongly to deviation above the target then to deviations below the targets in inflation and output.

These contradicting empirical results (Surico and Dolado at al.) suggest that there is no clear consensus regarding the asymmetry in the central banks' loss function. Generally, there is a tendency to estimate the policy functions separately for periods os different chairmen of the FED which conflict with GMM's notoriously bad small sample properties. Furthermore, even with unlimited amounts of data, the estimation of forward looking policy rules suffers from the weak instrument problem and the identification issues (Mavroeidis, 2010). Another issue is the fact that there is not limit on the possible functional form that the nonlinear policy rule can take and this introduces another degree of freedom in the estimation. This concern is addressed by Osborn et al. (2005), who apply Hamilton (2001)'s framework which allows to relax strict parametric assumptions about the shape of the estimated nonlinear function. They find that the interaction between inflation and output gap should be included in the specification (Dolado et al., 2005) of the policy rule.

Over all, it is fair to say that the existing empirical results do not provide a fully convincing evidence for or against asymmetric loss functions. Therefore there is scope in exploring other possible sources of asymmetry in the policy functions. This paper offers such an alternative.

3 Model

3.1 Basic settings

There are two players in the economy: a central bank and a representative agent. The economic environment is characterized by the inflation expectations augmented Phillips curve,

$$y_t = \alpha(\pi_t - \pi_{t|t-1}) + \varepsilon_t, \tag{1}$$

where y_t and π_t denote the output gap and the inflation at the period t, $\pi_{t|t-1}$ represents the expectations formed by the representative agent at period $t-1$ about the inflation in the period t. Both the central bank and the representative agent have the same loss function:

$$L_t = (y_t - y^*_{high})^2 + \beta_{low}\pi_t^2, \tag{2}$$

where y^*_{high} is the socially optimal level of output. The output target is at the socially optimal level of output y^*_{high}. This value is higher then the natural level of output (which is normalized to zero). y^*_{high} might be positive for example due to the monopolistic competition on the markets leading to higher prices and lower output. The preference parameter β_{low} is common for both the central bank and the rest of the society.

The subscripts of β and y^* indicate the solutions to the inflationary bias in the literature, either empower a conservative central banker whose preference parameter $\beta_{high} > \beta_{low}$, or give up aiming for the socially optimal level of output $y^*_{low} < y^*_{high}$

In this paper, the central bank cannot decide to change its preferences parameters y^*_{high} and β_{low}, as conjectured by Rogoff or Blinder. However, the central bank *announces* that it plans to behave according to some β_{high} and y^*_{low}. The representative agent understands that the true preferences are described by $(\beta_{low}, y^*_{high})$ and hence she tests every period if the central bank behaves according to its announcement rather then the true preference parameters.

The central bank minimizes the discounted expected loss (assuming the economy starts at $t = 0$) using a discount factor δ:

$$L = \mathsf{E}\sum_{t=0}^{\infty}\delta^t L_t$$

The representative agent does observe neither the supply shock faced by the central bank, nor the chosen allocation (as one would be perfectly informative about the other), but she observes only a noisy signal about the allocation chosen by the central bank. The reputation building mechanism works in the following way: the central banks announces a policy $(\beta_{high}, y^*_{low})$ and the representative agent trusts the central bank as long as they do not observe an outcome which would suggest that the central bank has deviated (from behavior which would be optimal given $(\beta_{high}, y^*_{low})$). The details on the deviation detection in section 3.4. The exact timing is following:

1. supply shock ϵ_t is realized and observed by the central bank
2. central bank chooses an allocation $y_t = y(\varepsilon_t), \pi_t = \pi(\varepsilon_t)$
3. noise shocks $\xi_{\pi,t}, \xi_{y,t}$ are realized and the representative agent observes

$$\tilde{y}_t = y_t + \xi_{y,t} \tag{3}$$

$$\tilde{\pi}_t = \pi_t + \xi_{\pi,t} \tag{4}$$

4. the representative agent faces a signal extraction problem and infers the likelihood of the central bank *deviating*, denoted by $l_D(\tilde{y}_t, \tilde{\pi}_t)$, or *not deviating*, denoted by $l_{ND}(\tilde{y}_t, \tilde{\pi}_t)$, and revise the inflation expectations $\pi_{t+1|t}$ for the next period accordingly:

- if they believe that the central bank has not deviated, they set their expectations of future inflation as $\pi_{t+1|t} = \pi(\beta_{high}, y^*_{low})$
- if they believe that the central bank has deviated, they set their expectations of future inflation as $\pi_{t+1|t} = \pi(\beta_{low}, y^*_{high})$

I will call choosing $\pi_{t+1|t} = \pi(\beta_{low}, y^*_{high})$ as *triggering punishment*, because it effectively punishes the central bank from deviating in the current period by increasing the loss in the next period (higher inflation expectations worsen the output gap - inflation trade-off) The punishment is either forever, or only for a finite period. The finite punishment can be interpreted as if the governor looses his job after the public looses the confidence in the central bank.

For the sake of exposition, I will consider that there is only one possible deviation, i.e. the central bank can behave either according to $(\beta_{high}, y^*_{low})$ or deviating by behaving according to $(\beta_{low}, y^*_{high})$, but this assumption is released lated and numerical results without this assumption are presented later. The representative agents always tests the hypothesis of no deviation versus the hypothesis of complete deviation. The detailed description of the signal extraction problem can be found in the section 3.4.

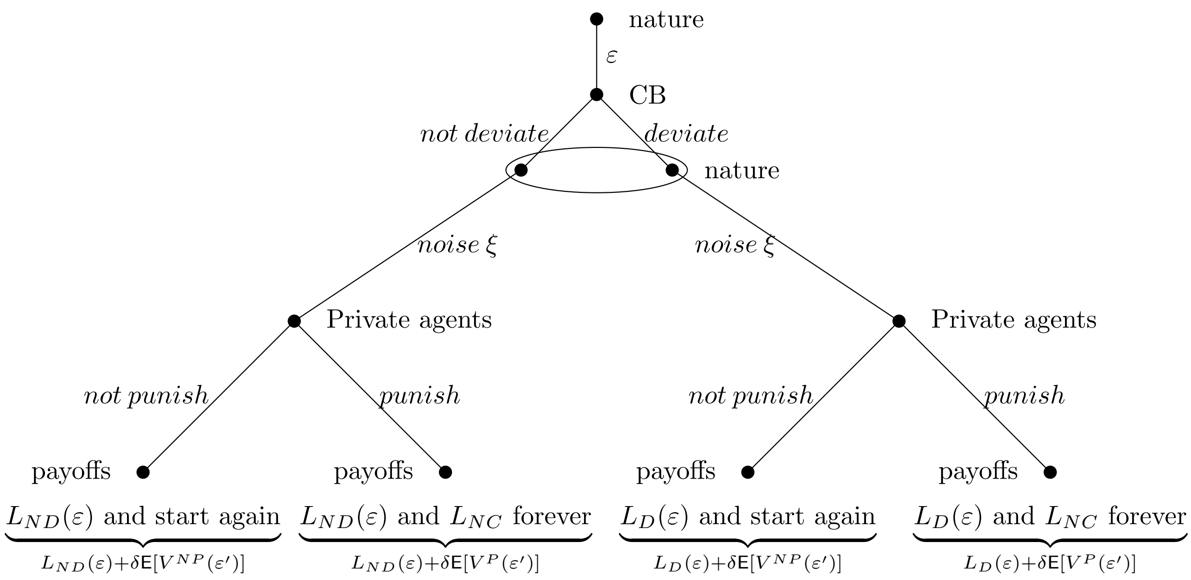

Figure 1: Game scheme: 0-1 deviation setting

The situation with only two possible actions by the central bank and eternal punishment is depicted in figure 1. The figure 1 uses the following notation: $L_{ND}(\varepsilon_t)$ denotes the value of the loss function when the central bank does not deviate while facing the supply shock ε_t. Similarly, L_D is the value of the loss function if the central bank deviates. V^P and V^{NP} are the value function of being in ether *punished* or *not punished* state.

3.2 Basic Algebra of the model

It is convenient to derive the basic results of the textbook setting of the inflationary bias. Doing so helps to develop the intuition and to build the tools which are used later.

3.2.1 Basic setting

Let's start with a generic loss function:

$$L_t = \frac{1}{2}\left(y_t - y^*\right)^2 + \frac{\beta}{2}\pi_t^2, \tag{5}$$

where y is output gap, $y^* \geq 0$ allows for positive output gap target. π is inflation (inflation target normalized to 0) and β is coefficient of inflation aversion, the inflation expectations augmented Phillips curve is:

$$y_t = \alpha(\pi_t - \pi_{t|t-1}) + \varepsilon_t,$$

where $\pi_{t|t-1}$ is the expected inflation, ε_t denotes the supply shock. I assume that the central bank can choose any combination of output gap and inflation consistent with the Phillips curve and abstract from modelling the policy instrument.

The regime in which the central bank cannot credibly pre-commit to follow some policy rule is called discretion. The inflations expectations are thus not function of the policy/preferences of the central bank, but they are preset and the central bank acts only after the inflation expectations are realized. For any given inflation expectations, using the first order conditions of the central bank problem, the optimal reaction to the supply shock ε_t can be found to be (the superscript d stands for discretion):

$$\pi_t^d = \frac{\alpha}{\alpha^2+\beta}y^* + \frac{\alpha^2}{\alpha^2+\beta}\pi_{t|t-1} - \frac{\alpha}{\alpha^2+\beta}\varepsilon_t, \tag{6}$$

$$y_t^d = \frac{\alpha^2}{\alpha^2+\beta}y^* - \frac{\alpha\beta}{\alpha^2+\beta}\pi_{t|t-1} + \frac{\beta}{\alpha^2+\beta}\varepsilon_t. \tag{7}$$

Let's find the equilibrium. Under the rational expectations, the expectation of inflation must be consistent with the true data generation process:

$$\pi_{t|t-1} = \mathsf{E}[\pi_t] \tag{8}$$

then

$$\begin{aligned}
\pi_{t|t-1} = \mathsf{E}[\pi_t^d] = \mathsf{E}\left[\frac{\alpha^2}{\alpha^2+\beta}\left(\pi_{t|t-1} + \frac{1}{\alpha}y^*\right) - \frac{\alpha}{\alpha^2+\beta}\varepsilon_t\right] \\
\pi_{t|t-1} = \frac{\alpha^2}{\alpha^2+\beta}\left(\pi_{t|t-1} + \frac{1}{\alpha}y^*\right) \\
\pi_{t|t-1} = \frac{\alpha}{\beta}y^*.
\end{aligned} \tag{9}$$

In order to obtain the optimal choice of inflation given any value of the supply shock shock ε, let's plug the equation (9) into (6):

$$\begin{aligned}
\pi_t^{ed} &= \frac{\alpha^2}{\alpha^2+\beta}\left(\frac{\alpha}{\beta}y^* + \frac{1}{\alpha}y^*\right) - \frac{\alpha}{\alpha^2+\beta}\varepsilon_t, \\
\pi_t^{ed} &= \frac{\alpha}{\beta}y^* - \frac{\alpha}{\alpha^2+\beta}\varepsilon_t
\end{aligned} \tag{10}$$

and we can notice that the expectation of inflation in equation (10) is indeed the inflation expectation $\pi_{t|t-1}$. From here we can observe that $\frac{\partial \pi_t^D}{\partial \varepsilon_t} < 0$, so a very inflation averse central bank ($\beta \to \infty$ in the loss function) does keep inflation at zero all the time which implies zero volatility in inflation as all shocks are accommodated by movements in the output gap. On the other hand, a central bank which does not care about inflation stabilization at all would be modelled as having $\beta \to 0$.

If we plug this result into the Phillips curve (1), we obtain the optimal choice for the output:

$$\begin{aligned} y_t^{ed} &= \alpha \left(\frac{\alpha}{\beta} y^* - \frac{\alpha}{\alpha^2 + \beta} \varepsilon_t - \frac{\alpha}{\beta} y^* \right) + \varepsilon_t, \\ y_t^{ed} &= \frac{\beta}{\alpha^2 + \beta} \varepsilon_t. \end{aligned} \tag{11}$$

The superscript *ed* captures the fact that these are the optimal allocations in equilibrium, by which I mean that the inflation expectations of the representative agent are correct in the sense that the equation (8) is satisfied.

3.2.2 Expansion path parametrization

The optimal choices of the central bank lie on a straight line. This fact is useful to understand the mechanics of the signal extraction problem of the representative agent. To see that this is the case, let's start with the equation (6), which can be rewritten as

$$\varepsilon_t = y^* + \alpha \pi_{t|t-1} - \frac{\alpha^2 + \beta}{\alpha} \pi_t^D, \tag{12}$$

and plug this into the PC, equation (1). The inflation expectations cancel out and the result is

$$y_t = y^* - \frac{\beta}{\alpha} \pi_t. \tag{13}$$

Note, that the parametrization result does not depend on $\pi_{t|t-1}$, so it holds no matter if the inflations expectations capture correctly the preference parameters $(\beta_{high}, y_{low}^*)$ in the central bank loss function or not.

3.2.3 Introducing conservative central banker

The for a generic preferences (β, y^*), the equilibrium is described by equations (9), (10) and (11), where the generic values are replaced by the correct values of the preference parameters, i.e. for the not deviating central bank by $(\beta_{high}, y_{low}^*)$ and $(\beta_{low}, y_{high}^*)$ for the central bank which is being punished.

Imagine for the moment that nay deviation is immediately detected. Still, the inflation expectations are fixed for the current period and the punishment starts only from the next period onwards. Hence, there are three possible situations:

1. the central bank is not deviating
2. the central bank is deviating for the first time
3. the central bank has deviated in the past and it is now being punished

In order to correctly set up the problem of the central bank if and how much to deviate, it is important to understand the ramifications of the decision to deviate. In particular, the knowledge of expected value of the loss function in the not punished and in the punished state are important in order to decide whether the immediate gain from deviation is worth the risk of moving to the punished state. Due to the noise, the chance of detection is not perfect and at the same time, even if the central bank is not deviating, there is still a chance that the agents might still punish it, if the noise realization is particularly bad.

Not Deviating equilibrium Assume that the central bank with preferences $(\beta_{low}, y^*_{high})$ acts as if it has preferences $(\beta_{high}, y^*_{low})$. In equilibrium, the representative agent sets her inflation expectations according to the behavior $(\beta_{low}, y^*_{high})$ (and not the actual preferences $(\beta_{high}, y^*_{low})$), so the central banker implements an equilibrium allocation:

$$\pi_t^{ND}(\beta_{high}, y^*_{low}) = \frac{\alpha}{\beta_{high}} y^*_{low} - \frac{\alpha}{\alpha^2 + \beta_{high}} \varepsilon_t, \tag{14}$$

$$y_t^{ND}(\beta_{high}, y^*_{low}) = \frac{\beta_{high}}{\alpha^2 + \beta_{high}} \varepsilon_t. \tag{15}$$

Let's denote the value of the loss by L_{ND}:

$$\begin{aligned}
L_{ND}(\varepsilon_t) &\equiv \frac{1}{2}\Big(y_t^{ND}(\beta_{high}, y^*_{low}) - y^*_{high}\Big)^2 + \frac{\beta_{low}}{2}\left(\pi_t^{ND}(\beta_{high}, y^*_{low})\right)^2 \\
&= \frac{1}{2}\Big(\frac{\beta_{high}}{\alpha^2 + \beta_{high}}\varepsilon_t - y^*_{high}\Big)^2 + \frac{\beta_{low}}{2}\Big(\frac{\alpha}{\beta_{high}} y^*_{low} - \frac{\alpha}{\alpha^2 + \beta_{high}}\varepsilon_t\Big)^2 \\
&= \frac{1}{2}\left(\frac{\beta_{high}}{\alpha^2 + \beta_{high}}\right)^2 \varepsilon_t^2 - \frac{\beta_{high}}{\alpha^2 + \beta_{high}} y^*_{high}\varepsilon_t + \frac{1}{2}\left(y^*_{high}\right)^2 \\
&+ \frac{\beta_{low}}{2}\left(\frac{\alpha}{\beta_{high}}\right)^2 (y^*_{low})^2 - \beta_{low}\frac{\alpha^2}{\beta_{high}(\alpha^2 + \beta_{high})} y^*_{low}\varepsilon_t + \frac{\beta_{low}}{2}\left(\frac{\alpha}{\alpha^2 + \beta_{high}}\right)^2 \varepsilon_t^2 \\
&= \frac{1}{2}\varepsilon_t^2\left(\left(\frac{\beta_{high}}{\alpha^2 + \beta_{high}}\right)^2 + \beta_{low}\left(\frac{\alpha}{\alpha^2 + \beta_{high}}\right)^2\right) \\
&- \varepsilon_t\left(\frac{\beta_{high}}{\alpha^2 + \beta_{high}} y^*_{high} + \beta_{low}\frac{\alpha^2}{\beta_h(\alpha^2 + \beta_{high})} y^*_{low}\right) \\
&+ \frac{1}{2}\left((y^*_{high})^2 + \beta_{low}\left(\frac{\alpha}{\beta_h}\right)^2 (y^*_{low})^2\right)
\end{aligned}$$

and the expected value of not deviating is thus

$$\mathsf{E}L^{ND} \equiv \mathsf{E}[L_{ND}(\varepsilon_t)] = \frac{1}{2}\sigma_\varepsilon^2 \frac{\alpha^2\beta_{low} + \beta_{high}^2}{(\alpha^2 + \beta_{high})^2} + \frac{1}{2}\frac{\alpha^2\beta_{low}(y^*_{low})^2 + \beta_{high}^2 (y^*_{high})^2}{\beta_{high}^2}. \tag{16}$$

As a general rule and if possible, subscripts denote actions and superscripts denote states.

Equilibrium in the punished state In this state, the central bank has no longer any incentive to deviate and the representative agent has correct expectations set according to $(\beta_{low}, y^*_{high})$. The superscript P captures the fact that this is the equilibrium in the punished state. The optimal allocation is

$$\pi_t^P(\beta_{low}, y_{high}^*) = \frac{\alpha}{\beta_{low}} y_{high}^* - \frac{\alpha}{\alpha^2 + \beta_{low}} \varepsilon_t \tag{17}$$

$$y_t^P(\beta_{low}, y_{low}^*) = \frac{\beta_{low}}{\alpha^2 + \beta_{low}} \varepsilon_t, \tag{18}$$

so the value of the loss function is

$$\begin{aligned} L^P(\varepsilon_t) &\equiv \frac{1}{2}\Big(y_t^D(\beta_{high}, y_{low}^*) - y_{high}^*\Big)^2 + \frac{\beta_{low}}{2}\Big(\pi_t^D(\beta_{high}, y_{low}^*)\Big)^2 \\ &= \frac{1}{2}\left(\left(\frac{\beta_{low}}{\alpha^2 + \beta_{low}}\right)^2 \varepsilon_t^2 - 2\frac{\beta_{low}}{\alpha^2 + \beta_{low}} y_{high}^* \varepsilon_t + (y_{high}^*)^2\right) \\ &\frac{\beta_{low}}{2}\left(\left(\frac{\alpha}{\beta_{low}}\right)^2 (y_{high}^*)^2 - 2\frac{\alpha}{\beta_{low}}\frac{\alpha}{\alpha^2 + \beta_{low}} y_{high}^* \varepsilon_t + \left(\frac{\alpha}{\alpha^2 + \beta_{low}}\right)^2 \varepsilon_t^2\right), \end{aligned}$$

which simplifies to

$$\begin{aligned} &= \frac{1}{2}\varepsilon_t^2\left(\left(\frac{\beta_{low}}{\alpha^2 + \beta_{low}}\right)^2 + \beta_{low}\left(\frac{\alpha}{\alpha^2 + \beta_{low}}\right)^2\right) \\ &- \varepsilon_t\left(\frac{\beta_{low}}{\alpha^2 + \beta_{low}} y_{high}^* + \beta_{low}\frac{\alpha}{\beta_{low}}\frac{\alpha}{\alpha^2 + \beta_{low}} y_{high}^*\right) \\ &+ \frac{1}{2}(y_{high}^*)^2\left(1 + \beta_{low}\left(\frac{\alpha}{\beta_{low}}\right)^2\right). \end{aligned}$$

The expected loss is then

$$\mathsf{E}L^P \equiv \mathsf{E}[L^P(\varepsilon_t)] = \frac{1}{2}\sigma_\varepsilon^2 \frac{\beta_{low}}{\alpha^2 + \beta_{low}} + \frac{1}{2} y_{high}^* \frac{\alpha^2 + \beta_{low}}{\beta_{low}}.$$

Loss when deviating When the central bank is deviating, it is facing the representative agent expecting $\pi_{t+1|t} = \frac{\alpha}{\beta_{high}} y_{low}^*$, so this is not an equilibrium in the sense defined above. Given these inflation expectations, the deviating central banker implements inflation

$$\begin{aligned} \pi_t^D &= \frac{\alpha}{\alpha^2 + \beta_{low}} y_{high}^* + \frac{\alpha^2}{\alpha^2 + \beta_{low}} \pi_{t|t-1} - \frac{\alpha}{\alpha^2 + \beta_{low}} \varepsilon_t \\ &= \frac{\alpha}{\alpha^2 + \beta_{low}} y_{high}^* + \frac{\alpha^2}{\alpha^2 + \beta_{low}} \frac{\alpha}{\beta_{high}} y_{low}^* - \frac{\alpha}{\alpha^2 + \beta_{low}} \varepsilon_t \\ &= \frac{\alpha}{\alpha^2 + \beta_{low}} \left(y_{high}^* + \frac{\alpha^2}{\beta_{high}} y_{low}^*\right) - \frac{\alpha}{\alpha^2 + \beta_{low}} \varepsilon_t \end{aligned}$$

$$\pi_t^D = \frac{\alpha}{\alpha^2 + \beta_{low}} \frac{1}{\beta_{high}} \left(\beta_{high} y_{high}^* + \alpha^2 y_{low}^* - \beta_{high} \varepsilon_t\right) \tag{19}$$

and output

$$
\begin{aligned}
y_t^D &= \frac{\alpha^2}{\alpha^2+\beta_{low}} y^*_{high} - \frac{\alpha\beta_{low}}{\alpha^2+\beta_{low}} \pi_{t|t-1} + \frac{\beta_{low}}{\alpha^2+\beta_{low}} \varepsilon_t \\
&= \frac{\alpha^2}{\alpha^2+\beta_{low}} y^*_{high} - \frac{\alpha\beta_{low}}{\alpha^2+\beta_{low}} \frac{\alpha}{\beta_{high}} y^*_{low} + \frac{\beta_{low}}{\alpha^2+\beta_{low}} \varepsilon_t \\
&= \frac{1}{\alpha^2+\beta_{low}} \left(\alpha^2 y^*_{high} - \alpha\beta_{low} \frac{\alpha}{\beta_{high}} y^*_{low} + \beta_{low}\varepsilon_t \right) \\
y_t^D &= \frac{1}{\alpha^2+\beta_{low}} \frac{1}{\beta_{high}} \left(\alpha^2 \beta_{high} y^*_{high} - \alpha^2 \beta_{low} y^*_{low} + \beta_{low}\beta_{high}\varepsilon_t \right) \qquad (20)
\end{aligned}
$$

Loss comparisons It can be shown that $L_P(\varepsilon_t) > L_{ND}(\varepsilon_t) > L_D(\varepsilon_t)$. This situation is depicted in figure 2. The points ND, D and P depict "No Deviation", "Deviation" and "Punishment" outcomes, corresponding circle captures the value of the central bank's loss function for $\varepsilon_t = 0$. The inflation expectations π^e_{low} correspond to $\pi_{t|t-1}(\beta_{high}, y^*_{low})$ and $\pi^e_{low} = \pi_{t|t-1}(\beta_{low}, y^*_{high})$.

As was pointed out by the time inconsistency literature, not deviating is not *one period* equilibrium, which manifests itself on the figure by showing that the Phillips curve is not tangent to the iso-loss circles. Furthermore, the loss is smallest under cheating D, but in that case inflation expectations are not consistent with the outcome, so neither this point is an equilibrium.

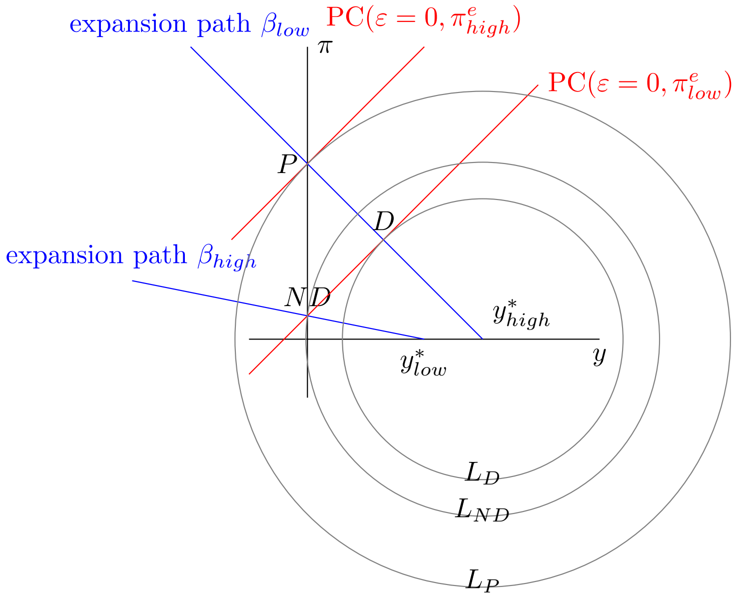

Figure 2: The central bank follows $(\beta_{high}, y^*_{low})$ instead its true preferences $(\beta_{low}, y^*_{high})$. β_{low} is set to 1, so the iso-loss curves are circles.

As an useful shortcut, when I say that central bank behaves according to $(\beta_{high}, y^*_{low})$, I mean that the central banks chooses an allocation defined by equations (10) and (11) with $\beta = \beta_{high}$ and $y^* = y^*_{low}$. To the contrary, when I say that the central bank is punished and behaves according to $(\beta_{low}, y^*_{high})$, I mean that it chooses an allocation given by (10) and (11) with β_{low} and y^*_{high}.

3.3 Value functions

3.3.1 basic definitions

Because of the recursive nature of the problem, the problem is solved by formulating the Bellman equations for the problem the central bank faces at the beginning of the game and finding the optimal policy function by implementing the value function iteration algorithm. In the problem of the central bank, there are two state variables: continuous supply shock ε_t and state of the beliefs of the agents. There is only one choice variable, the deviation (which is modelled as either discrete or continuous, section 3.5).

Let's denote the state of the beliefs by $\pi_{t|t-1}$. This is a slight abuse on notation, as the mapping between the inflation expectation $\pi_{t|t-1}$ and the preference parameters (β, y^*) is not a bijection, but since there is only one announcement, there is no risk of confusion. In general, the Bellman equation has form of

$$V(\varepsilon_t, \pi_{t|t-1}) = \min L(\varepsilon_t, \pi_{t|t-1}) + \delta \mathsf{E} V(\varepsilon_{t+1}, \pi_{t+1|t}), \tag{21}$$

where $V(\varepsilon_t, \pi_{t|t-1})$ is the value function. The problem is to find the policy function PF : $(\varepsilon_t, \pi_{t|t-1}) \to (y_t, \pi_t)$, which maps the states into allocations, in my case it tells what inflation $\pi(\varepsilon_t, \pi_{t|t-1})$ and output $y(\varepsilon_t, \pi_{t|t-1})$ the central bank chooses given the supply shock and the state of the beliefs. As there are only two values of possible beliefs, and since the behavior in the punished state was already solved for, the problem is to solve for the policy function in the not punished state, so the only state variable is in fact ε_t.

In my case, the supply shock ε is exogenous and iid, so out of control of the central bank. However, the beliefs are under control (albeit imperfect), because the decision to stick with $\pi^e_{low} = \pi_{t|t-1}(\beta_{high}, y^*_{low})$ or punish by $\pi^e_{high} = \pi_{t|t-1}(\beta_{low}, y^*_{high})$ is a function of observed $\tilde{y}_t, \tilde{\pi}_t$, which are (noisy) signal about the allocation actually chosen by the central bank. Because of the binary nature of the beliefs, I will denote $V(\varepsilon, \hat{\beta})$ as either $V^P(\varepsilon)$ (for $\hat{\beta} = \beta_{low}$) or $V^{NP}(\varepsilon)$ (for $\hat{\beta} = \beta_{high}$).

The policy function can be viewed as generated by *a deviation function*,

$$d(\varepsilon_t, (\beta, y^*)) \to \mathbb{R}, \tag{22}$$

which assign a real number to each point in the state space, capturing how much is optimal to deviate from the allocation given by $(\beta_{high}, y^*_{low})$:

$$\pi_t(\beta_{high}, y^*_{low}, d) = \frac{\alpha}{\beta_{high}} y^*_{low} - \frac{\alpha}{\alpha^2 + \beta_{high}} \varepsilon_t + \frac{1}{\alpha} d, \tag{23}$$

$$y_t(\beta_{high}, y^*_{low}, d) = \frac{\beta_{high}}{\alpha^2 + \beta_{high}} \varepsilon_t + d. \tag{24}$$

In detail, if the central bank deviates by d, it chooses following allocation by moving along the Phillips curve. Furthermore, if the central bank implements the $0-1$ deviation, then that does not mean that $\forall \varepsilon, d(\varepsilon) \in \{0, 1\}$, but the $d(\varepsilon)$ is a linear function of the supply shock, as can be seen from the figure 2, because the slopes of the two expansion paths are different.

Because I always start from the situation where the central bank announces $(\beta_{high}, y^*_{low}, d)$ and the announcement is believed, I will write $\pi_t(\varepsilon_t, d)$ and $y_t(\varepsilon_t, d)$ as a short cut for $\pi_t(\beta_{high}, y^*_{low}, d)$

and $y_t(\beta_{high}, y^*_{low}, d)$ in situations where this shortcut cannot lead to confusion. Also, in the same way, I sometimes use PF as map between the states and the chosen outcomes $\pi_t(\varepsilon_t, d)$ and $y_t(\varepsilon_t, d)$, meaning PF : $(\varepsilon_t, \hat{\beta}_t) \to (\pi_t(\varepsilon_t, d), y_t(\varepsilon_t, d))$.

3.3.2 Value function depending of the state of beliefs

Value in the punished state using the notation defined above, $V^P(\varepsilon)$ denotes the value function when the representative agent no longer believes that $(\beta_{high}, y^*_{low})$ is correct description of CB actions (=being Punished). The inflation expectations are then set according to $(\beta_{low}, y^*_{high})$, which gives $\pi_{t|t-1} = \frac{\alpha}{\beta_{low}} y^*_{high}$, as derived in equation (9).

In the setting where the punishment period is infinite, the expected value of the loss function is just the discounted value of the expected losses at the punished state, because there will never by any change in the beliefs. This means that

$$\mathsf{E}[V^P(\varepsilon)] = \mathsf{E}\sum_{t=0}^{\infty} L_t(\beta_{low}) = \frac{1}{1-\delta}\mathsf{E}L^P. \tag{25}$$

On the other hand, if the punishment is for one period only, then

$$\mathsf{E}[V^P(\varepsilon)] = \mathsf{E}L^P + \delta \mathsf{E}V^{NP}(\varepsilon_t). \tag{26}$$

Value in the not punished state $V^{NP}(\varepsilon)$ labels the value of not being punished facing the supply shock ε, this is the problem which the central bank faces at the beginning of the game, and which is being analyzed here.

For the sake of the exposition, let's assume that the deviation can be either none or complete, denoted by 0-1 values, where 0 means no deviation, and 1 means the full deviation to $(\beta_{low}, y^*_{high})$ outcome. The central bank chooses the action minimizes its loss and hence the value of the problem is the smaller value of either case:

$$V^{NP}(\varepsilon) = \min\{V^{NP}_D(\varepsilon), V^{NP}_{ND}(\varepsilon)\}, \tag{27}$$

where

$$\begin{aligned} V^{NP}_D(\varepsilon) &= L_D(\varepsilon) + \delta\left[\mathsf{P}(\varepsilon, D)\mathsf{E}[V^P(\varepsilon')] + (1-\mathsf{P}(\varepsilon, D))\mathsf{E}[V^{NP}(\varepsilon')]\right] \\ V^{NP}_{ND}(\varepsilon) &= L_{ND}(\varepsilon) + \delta\left[\mathsf{P}(\varepsilon, ND)\mathsf{E}[V^P(\varepsilon')] + (1-\mathsf{P}(\varepsilon, ND))\mathsf{E}[V^{NP}(\varepsilon')]\right], \end{aligned}$$

where $\mathsf{P}(\varepsilon, X)$ denotes the probability of punishment when facing supply shock ε and choosing action $X \in \{D, ND\}$ and δ is the central bank discount factor. The results reported in the dedicated section of this paper are computed using $X = d(\varepsilon) \in \mathbb{R}$.

3.3.3 Probability of Punishment

$\mathsf{P}(\varepsilon, X)$ is the probability of the punishment. Given ε_t, the central bank can either decide to stick to the promised strategy and set (y_t, π_t) allocation according to $(\beta_{high}, y^*_{low})$, or it can deviate and behave according to $(\beta_{low}, y^*_{high})$.

Let's informally define the *Rejection Region* $\mathcal{R}$ as an area in output x inflation space, such that if $(\tilde{y}, \tilde{\pi}) \in \mathcal{R}$ is observed, the representative agent will decide to punish the central bank and revise their expectations for the next period 3.4.2. Let's denote the parameter which governs how strict the representative agent is during the signal extraction by κ, formally defined in equation (38) section 3.4.2.

If $\mathcal{R} = \mathcal{R}(\kappa)$ is the *Rejection Region*,[3] then $\mathsf{P}(\varepsilon, X)$ can be written formally as

$$\mathsf{P}(\varepsilon, X) = \mathsf{E}\boldsymbol{I}_{\{(y(\varepsilon,X)+\xi_y, \pi(\varepsilon,X)+\xi_y) \in \mathcal{R}\}},$$

where $x = (y(\varepsilon, X), \pi(\varepsilon, X))$ is the allocation chosen by the central bank and $\tilde{x} = (y(\varepsilon, X) + \xi_y, \pi(\varepsilon, X) + \xi_y) = x + \xi$ is the allocation observed by the private agents.[4]

I obtain $\mathsf{P}(\varepsilon, X)$ by simulation: for each point in the [output x inflation] grid (representing the possible allocation chosen by the central bank) I draw pairs of random numbers representing the noise shocks and for each realization of the shock I let the private agents compute the signal extraction problem, obtain the estimated shocks and compare the likelihoods. This way the probability of revision $\mathsf{P}(y, \pi)$ in the (y, π) space is computed and then this matrix is used to interpolate the probability $\mathsf{P}(\varepsilon, d)$.

The asymmetry is the result of following consideration: even when the central bank never deviates, the likelihood of punishment is increasing with the size of the shock. If it could, the central bank would like to choose to marginally deviate and make the policy tighter.[5] The intuition is that punishment leads to higher inflation expectations, hence worse inflation-output trade-off and ultimately higher loss for the central bank. The central bank thus wants to avoid punishment, or at least decrease it's likelihood.

So far I have explained the model as if the only two possible actions for the central bank were to deviate (choose allocation according to $(\beta_{low}, y^*_{high})$ or not to deviate $(\beta_{high}, y^*_{low})$). This was done for purely exposition purposes and I allow for a continuous deviation.

3.4 Agent's problem

Here I analyze the problem faced by the representative agent: after observing $\tilde{y}$ and $\tilde{\pi}$, she needs to assess the credibility of the central bank in order to set her inflation expectations for the next period. In another words, she has to decided on her beliefs about the conservativeness of the central bank. In particular, whether the central bank really follows the announced $(\beta_{high}, y^*_{low})$, or whether it has deviated and behaved according to $(\beta_{low}, y^*_{high})$.

In order to do so, the agent has to solve a signal extraction problem. The agent knows that there is a noise in the signal and she knows the variances of all shocks, supply and noise. The agent knows all the results about the optimal choices of the central bank given its preferences which were derived in the section 3.2.

The representative agent decides only between two possibilities: central bank behaves linearly

[3] *No rejection region* $\mathcal{N}$ is the complement region to the *Rejection Region*, i.e. a region such that if an allocation $x \in \mathcal{N}$ is observed, then the agents do not revise their beliefs, i.e. do not trigger punishment strategy. The region depends on the shape of likelihood functions and a cushion parameter κ.

[4] To complete the notation, $\hat{x}$ is the inferred allocation, i.e. the solution to the signal extraction problem of the private agents.

[5] Note that this deviation goes against deviation due to the time inconsistency.

as if it parameter was $(\beta_{high}, y^*_{low})$ or the bank deviated and it behaved exactly how discretionary central bank with $(\beta_{low}, y^*_{high})$ without any dynamic consideration would behave. In the earlier version of this paper, I implemented a continuous signal extraction, not a binary one which is presented here. The results were similar, but the computation time was significantly higher. The main reason is that there are closed form solutions for the binary problem, whereas additional numerical method has to applied in order to solve for the continuous case. This makes the computations of matrix $\mathsf{P}(\varepsilon, X)$ much more demanding, as the signal extraction problem has to be solved for each point in the grid in the (y, π) and for each of the simulation of the noise shock (ξ_y, ξ_π). While this is indeed a simplification, an extra constraint is imposed on the central bank behavior so the assumption about the behavior of the representative agent can be viewed as a reasonable simplification, rather then unrealistic assumption driving the results of this paper. This constraint is discussed in detail in section 3.5.

3.4.1 Agent's signal extraction problem

Assumption (1): Shocks ε_t, $\xi_{\pi,t}$ and $\xi_{y,t}$ are independently normally distributed with known variances σ^2_ε, σ^2_π and σ^2_y. ($\mathbf{\Sigma} = diag(\sigma^2_\varepsilon, \sigma^2_\pi, \sigma^2_y)$)

Here I analyze two cases, deviating and not deviating from the perspective of the representative agent. The advantage of imposing the binary beliefs structure is that the two cases can be solved using the same approach. Let's recapitulate the basic equations of the model, the optimal allocation chosen by the central bank described by some generic (β, y^*), facing some inflation expectations $\pi_{t|t-1}$ and supply shock ε_t, finding the derivative is then easy:

$$\begin{aligned}
\pi_t &= \frac{\alpha}{\alpha^2+\beta}y^* + \frac{\alpha^2}{\alpha^2+\beta}\pi_{t|t-1} - \frac{\alpha}{\alpha^2+\beta}\varepsilon_t \quad &\Rightarrow \quad \frac{\partial \pi_t}{\partial \varepsilon_t} = -\frac{\alpha}{\alpha^2+\beta} \\
y_t &= \frac{\alpha^2}{\alpha^2+\beta}y^* - \frac{\alpha\beta}{\alpha^2+\beta}\pi_{t|t-1} + \frac{\beta}{\alpha^2+\beta}\varepsilon_t \quad &\Rightarrow \quad \frac{\partial y_t}{\partial \varepsilon_t} = \frac{\beta}{\alpha^2+\beta}
\end{aligned}$$

The representative agent is aware of the optimal choices described by the two previous equations. The signal extraction is an exercise to obtain some estimates of $\hat{y}(\tilde{y}, \tilde{\pi})$ and $\hat{\pi}(\tilde{y}, \tilde{\pi})$. From equations (4) and (3)

$$\hat{\xi}_{\pi,t}(\tilde{y}, \tilde{\pi}) = \tilde{\pi}_t - \hat{\pi}_t(\tilde{y}, \tilde{\pi}) \quad \Rightarrow \quad \frac{\partial \hat{\xi}_{\pi,t}}{\partial \hat{\varepsilon}_t} = -\frac{\partial \hat{\pi}_t}{\partial \hat{\varepsilon}_t} = \frac{\alpha}{\alpha^2+\beta} \tag{28}$$

$$\hat{\xi}_{y,t} = \tilde{y}_t - \hat{y}_t(\tilde{y}, \tilde{\pi}) \quad \Rightarrow \quad \frac{\partial \hat{\xi}_{y,t}}{\partial \hat{\varepsilon}_t} = -\frac{\partial \hat{y}_t}{\partial \hat{\varepsilon}_t} = -\frac{\beta}{\alpha^2+\beta} \tag{29}$$

Now let's use these results for the signal extraction. Using the independence of the shocks, the log-likelihood of a triplet $(\varepsilon, \xi_y, \xi_\pi)$ for any generic central bank described by (β, y^*) can be written as

$$l(\varepsilon_t, \xi_{t,y}, \xi_{t,\pi}, \beta, y^*) = -\frac{1}{2}\log(2\pi\sigma^2_\varepsilon) - \frac{1}{2}\log(2\pi\sigma^2_\pi) - \frac{1}{2}\log(2\pi\sigma^2_y) - \frac{\varepsilon^2_t}{2\sigma^2_\varepsilon} - \frac{\xi^2_{\pi,t}}{2\sigma^2_\pi} - \frac{\xi^2_{y,t}}{2\sigma^2_y}. \tag{30}$$

Given the inflation expectations $\pi_{t|t-1}$, there is one to one mapping $\varepsilon \to (y(\varepsilon), \pi(\varepsilon))$ defined in the equations (7) and (6), and there is one to one mapping $\big((y,\pi),(\tilde{y},\tilde{\pi})\big) \rightarrow (\xi_y,\xi_\pi)$. Therefore the noise can be written as $(\xi_y,\xi_\pi) = (\xi_y(\varepsilon),\xi_\pi(\varepsilon))$, and in this sense

$$l(\varepsilon_t,\xi_{t,y},\xi_{t,\pi},\beta,y^*) = l(\varepsilon_t,\beta,y^*,\tilde{y},\tilde{\pi}). \tag{31}$$

The agents solve the signal extraction problem by finding the most likely allocation given the observed allocation $(\tilde{y},\tilde{\pi}_t)$ and using the fact that the noise is the difference between the allocation chosen by the central bank and the observed noisy signal:

$$\begin{aligned}
\frac{\partial l(\varepsilon_t,\beta,y^*,\tilde{y},\tilde{\pi})}{\partial \varepsilon} &= -\frac{\varepsilon_t}{\sigma_\varepsilon^2} - \frac{\xi_{\pi,t}}{\sigma_\pi^2}\frac{\partial \xi_{\pi,t}}{\partial \varepsilon_t} - \frac{\xi_{y,t}}{\sigma_y^2}\frac{\partial \xi_{y,t}}{\partial \varepsilon_t} \\
&= -\frac{\varepsilon_t}{\sigma_\varepsilon^2} - \frac{\xi_{\pi,t}}{\sigma_\pi^2}\frac{\alpha}{\alpha^2+\beta} + \frac{\xi_{y,t}}{\sigma_y^2}\frac{\beta}{\alpha^2+\beta} \\
&= -\frac{\varepsilon_t}{\sigma_\varepsilon^2} - \frac{\tilde{\pi}_t-\pi_t}{\sigma_\pi^2}\frac{\alpha}{\alpha^2+\beta} + \frac{\tilde{y}_t-y_t}{\sigma_y^2}\frac{\beta}{\alpha^2+\beta} \\
0 &= -\frac{\varepsilon_t}{\sigma_\varepsilon^2} - \frac{\tilde{\pi}_t - \left(\frac{\alpha}{\alpha^2+\beta}y^* + \frac{\alpha^2}{\alpha^2+\beta}\pi_{t|t-1} - \frac{\alpha}{\alpha^2+\beta}\varepsilon_t\right)}{\sigma_\pi^2}\frac{\alpha}{\alpha^2+\beta} \\
&\quad + \frac{\tilde{y}_t - \left(\frac{\alpha^2}{\alpha^2+\beta}y^* - \frac{\alpha\beta}{\alpha^2+\beta}\pi_{t|t-1} + \frac{\beta}{\alpha^2+\beta}\varepsilon_t\right)}{\sigma_y^2}\frac{\beta}{\alpha^2+\beta}
\end{aligned}$$

$$\begin{aligned}
0 &= -\varepsilon_t\left(\frac{1}{\sigma_\varepsilon^2} + \frac{1}{\sigma_\pi^2}\left(\frac{\alpha}{\alpha^2+\beta}\right)^2 + \frac{1}{\sigma_y^2}\left(\frac{\beta}{\alpha^2+\beta}\right)^2\right) \\
&\quad + \frac{1}{\alpha^2+\beta}\left(\frac{\beta}{\sigma_y^2}\tilde{y}_t - \frac{\alpha}{\sigma_\pi^2}\tilde{\pi}_t\right) \\
&\quad + y^*\left(\frac{\alpha}{\alpha^2+\beta}\right)^2\left(\frac{1}{\sigma_\pi^2} - \beta\frac{1}{\sigma_y^2}\right) \\
&\quad + \pi_{t|t-1}\frac{\alpha}{(\alpha^2+\beta)^2}\left(\alpha^2\frac{1}{\sigma_\pi^2} - \beta^2\frac{1}{\sigma_y^2}\right),
\end{aligned}$$

so the resulting inferred shock $\hat{\varepsilon}_t$ has to be

$$\hat{\varepsilon}(\beta,y^*,\tilde{y}_t,\tilde{\pi}_t) = \frac{\frac{1}{\alpha^2+\beta}\left(\frac{\beta}{\sigma_y^2}\tilde{y}_t - \frac{\alpha}{\sigma_\pi^2}\tilde{\pi}_t\right) + y^*\left(\frac{\alpha}{\alpha^2+\beta}\right)^2\left(\frac{1}{\sigma_\pi^2} - \beta\frac{1}{\sigma_y^2}\right) + \pi_{t|t-1}\frac{\alpha}{(\alpha^2+\beta)^2}\left(\alpha^2\frac{1}{\sigma_\pi^2} - \beta^2\frac{1}{\sigma_y^2}\right)}{\frac{1}{\sigma_\varepsilon^2} + \frac{1}{\sigma_\pi^2}\left(\frac{\alpha}{\alpha^2+\beta}\right)^2 + \frac{1}{\sigma_y^2}\left(\frac{\beta}{\alpha^2+\beta}\right)^2},$$

where the inflation expectations satisfy $\pi_{t|t-1} = \frac{\alpha}{\beta_{high}}y^*_{low}$, because agents are assumed to believe the announcement at the start of the game.

Now, given $\hat{\varepsilon}(\beta,y^*,\tilde{y}_t,\tilde{\pi}_t)$, the inferred allocation chosen by a central bank (following (β,y^*))

is

$$\hat{\pi}_t(\beta, y^*, \tilde{y}_t, \tilde{\pi}_t) = \frac{\alpha}{\alpha^2+\beta} y^* + \frac{\alpha^2}{\alpha^2+\beta}\pi_{t|t-1} - \frac{\alpha}{\alpha^2+\beta}\hat{\varepsilon}(\beta, y^*, \tilde{y}_t, \tilde{\pi}_t) \tag{32}$$

$$\hat{y}_t(\beta, y^*, \tilde{y}_t, \tilde{\pi}_t) = \frac{\alpha^2}{\alpha^2+\beta} y^* - \frac{\alpha\beta}{\alpha^2+\beta}\pi_{t|t-1} + \frac{\beta}{\alpha^2+\beta}\hat{\varepsilon}(\beta, y^*, \tilde{y}_t, \tilde{\pi}_t) \tag{33}$$

and the noise shock are then

$$\hat{\xi}_\pi(\beta, y^*, \tilde{y}_t, \tilde{\pi}_t) = \tilde{\pi}_t - \hat{\pi}_t(\beta, y^*, \tilde{y}_t, \tilde{\pi}_t) \tag{34}$$

$$\hat{\xi}_y(\beta, y^*, \tilde{y}_t, \tilde{\pi}_t) = \tilde{y}_t - \hat{y}_t(\beta, y^*, \tilde{y}_t, \tilde{\pi}_t) \tag{35}$$

as a shortcut, let's denote the inferred the values inferred using $(\beta_{high}, y^*_{low})$ by subscript ND (for example $\hat{\varepsilon}_{ND}(\tilde{y}_t, \tilde{\pi}_t) \equiv \hat{\varepsilon}(\beta_{high}, y^*_{low}, \tilde{y}_t, \tilde{\pi}_t)$) and the values inferred using $(\beta_{low}, y^*_{high})$ by subscript D ($\hat{\varepsilon}_D(\tilde{y}_t, \tilde{\pi}_t) \equiv \hat{\varepsilon}(\beta_{low}, y^*_{high}, \tilde{y}_t, \tilde{\pi}_t)$), and similarly for all inferred variables.

Let's denote the log-likelihood that the observed allocation $(\tilde{y}_t, \tilde{\pi}_t)$ is an outcome of the central bank either *not deviating* or *deviating* by

$$l_{ND}(\tilde{y}_t, \tilde{\pi}_t) = l\left(\hat{\varepsilon}_{ND}(\tilde{y}_t, \tilde{\pi}_t), \hat{\xi}_{\pi,ND}(\tilde{y}_t, \tilde{\pi}_t), \hat{\xi}_{y,ND}(\tilde{y}_t, \tilde{\pi}_t)\right) \tag{36}$$

$$l_{D}(\tilde{y}_t, \tilde{\pi}_t) = l\left(\hat{\varepsilon}_{D}(\tilde{y}_t, \tilde{\pi}_t), \hat{\xi}_{\pi,D}(\tilde{y}_t, \tilde{\pi}_t), \hat{\xi}_{y,D}(\tilde{y}_t, \tilde{\pi}_t)\right) \tag{37}$$

3.4.2 Comparing log-likelihoods and Revision Region $\mathcal{R}$

The representative agent revises her beliefs about central bank trustworthiness if the log-likelihood given $(\beta_{high}, y^*_{low})$, $l_{ND}(\tilde{y}_t, \tilde{\pi}_t)$ is too low compared to $l_D(\tilde{y}_t, \tilde{\pi}_t)$. However, it is not necessary welfare maximizing to have the revise the beliefs when the one likelihood just bigger, so the representative agent might want to have some cushion in decision making, which is captured by the parameter κ. Let's formally define the Rejection region $\mathcal{R}$:

$$(\tilde{y}, \tilde{\pi}) \in \mathcal{R} \Leftrightarrow \frac{l_D(\tilde{y}_t, \tilde{\pi}_t)}{l_{ND}(\tilde{y}_t, \tilde{\pi}_t)} < \kappa \tag{38}$$

The parameter κ affects how stringent the agent is when evaluating the deviations. She chooses κ to minimize her own loss. However, because the loss function of the representative agent is the same as the loss of the central bank (disregarding the noise), we can solve for κ as a minimizer of the value function, given that the constraint on the behavior of the central bank are satisfied.

3.5 Additional Constraints on Central Bank's Behavior $\mathbf{E}[d(\varepsilon)] = 0$

During the signal extraction exercise, the representative agent considers only two possibilities: the central bank behaves either according to $(\beta_{high}, y^*_{low})$ or $(\beta_{low}, y^*_{high})$. This assumption was imposed mainly due to computational aspects. While it is true that this limits the rationality of the representative agent, I am going to argue that it less strict assumption then it looks like. In this section I discuss the constraint I put on the behavior of the central bank so the mental process of private agents can be considered plausible.

First, in general, the representative private agents tests against each other two hypotheses about the slope of two linear expansion paths. I believe this is reasonable setting as even today, most of macroeconomic models feature a linear policy rule.

Second, given the constraint defined later in this section, the linear expansion path given by $(\beta_{high}, y^*_{low})$ is the best linear description of the behavior of the central bank and it is exactly this expansion path which would be fitted by the representative agent if she wanted to estimate a linear regression between inflation and output.

The nonlinearity then arises, because the information asymmetries makes perfect immediate detection impossible, however, the central bank cannot possibly deviate in a linear fashion, because such behavior would be detectable in the long run, because, the condition of the rational expectations of inflation $\mathsf{E}\pi = \pi_{t|t-1}$ would not be satisfied.

The constraint $\mathsf{E}[d(\varepsilon)] = 0$ hence seems as a natural way to enforce a deviation function such that the inflation expectations of the representative agent given the announcement $(\beta_{high}, y^*_{low})$ correspond with the mean of the distribution of the realized inflation.

Furthermore, if $\mathsf{E}[d(\varepsilon)] = 0$ is satisfied, then the (linear) expansion path given by β_{high} cannot be rejected against any other linear policy rule. Hence, the restriction, together with the fact that there is no persistence in the supply shock, justify the assumption that the agents do not use any past observations for inference about the current action of the central bank.

3.5.1 Implementing the Value function iteration algorithm with a constraint

The standard value function algorithm converges point-wise, by which I mean the fact that in each step, the choice of optimal policy at any point of the state space is not in any way affected by choice at different point of the state space.

However, imposing $\mathsf{E}[d(\varepsilon)] = 0$ requires an adjustment in the value function iteration algorithm, because choosing $d(\varepsilon_1)$ marginally higher, requires, that on (appropriately weighted) average $d(\varepsilon)$ has to be lower for all other ε in the state space.

I proceeded in following way. The variable being minimized is not the loss at each point of the state space, but the *expected value* of the value function. Originally, policy function is an infinite dimensional object, it maps each point in the state space to particular value of the deviation. Here, policy function is approximated and the optimization runs over much fewer coefficients of this approximation. I use two possible approximations techniques: chebyshev polynomials and piece-wise linear approximation.

Formally, let's recall that the policy function PF maps shocks ε into d, deviations from expansion path defined by β_{high}. Assume that $PF(\varepsilon)$ can be approximated by some function of parameters $\boldsymbol{\theta}$, $\mathcal{P}_{\boldsymbol{\theta}}(\varepsilon)$. Lets recall, that the value function at the not-punished state is (defining $d_t \equiv d(\varepsilon_t)$)

$$V^{NP}(\varepsilon_t) = \min_{d_t} \left\{ L(\varepsilon_t, d_t) + \delta \left[P(\varepsilon_t, d_t)\mathsf{E}V^{NP} + (1 - P(\varepsilon_t, d_t))\mathsf{E}V^P \right] \right\}. \tag{39}$$

The problem is that $d(\varepsilon_t)$ is in theory infinite dimensional object and even with discretization, the grid is still too fine to solve for it. However, if we use the approximation $\mathcal{P}_n(\varepsilon) \approx d(\varepsilon_t)$, we get

$$V^{NP}(\varepsilon_t) = \min_{\boldsymbol{\theta}} \left\{ L(\varepsilon_t, \mathcal{P}_{\boldsymbol{\theta}}(\varepsilon)) + \delta \left[P(\varepsilon_t, \mathcal{P}_{\boldsymbol{\theta}}(\varepsilon))\mathsf{E}V^{NP} + (1 - P(\varepsilon_t, \mathcal{P}_{\boldsymbol{\theta}}(\varepsilon)))\mathsf{E}V^P \right] \right\}, \tag{40}$$

where $\boldsymbol{\theta} = \{\theta_i\}_{i=0}^{n-1}$ might be the parameters defining an approximation of the policy function by a polynomial $\mathcal{P}_{\boldsymbol{\theta}}(\varepsilon) = \sum_{i=0}^{n-1} \theta_i \varepsilon_t^i$, or any other approximation of the policy function. $\boldsymbol{\theta}$ is an order of magnitude smaller then the grid for the state variable ε.

Furthermore, some mechanism which which would link the choices for all values of the state variable is needed. The natural choice is to minimize *the expected value* of the value function, $\mathsf{E}V^{NP}(\varepsilon_t)$. The expected value is approximated by numerically by using a quadrature. This setting can be directly implemented using Matlab's nonlinear minimization procedure `fmincon`, which incorporates the nonlinear constrain.

3.5.2 Implementation summary

To sum up, in each step of the value function iteration, the problem is to find parameters $\boldsymbol{\theta} = \{\theta_i\}_{i=0}^{n-1}$ of some choice of approximation function, which minimizes the expected value of the value function. Formally,

$$\text{find} \quad \boldsymbol{\theta} \quad \text{to minimize}$$

$$\mathsf{E}\Big[V^{NP}(\varepsilon_t)\Big] = \mathsf{E}\Bigg[L(\varepsilon_t, \mathcal{P}_{\boldsymbol{\theta}}(\varepsilon)) + \delta\Big[P(\varepsilon_t, \mathcal{P}_{\boldsymbol{\theta}}(\varepsilon))\mathsf{E}V^{NP} + \Big(1 - P(\varepsilon_t, \mathcal{P}_{\boldsymbol{\theta}}(\varepsilon))\Big)\mathsf{E}V^{P}\Big]\Bigg]$$

$$\text{such that} \quad \mathsf{E}[\mathcal{P}_{\boldsymbol{\theta}}(\varepsilon)] = 0.$$

Assuming that there are initial guesses for $\mathsf{E}V^{NP}$ and $\mathsf{E}V^{P}$, the algorithm is following:

1. run the minimization sub-routine:
   (a) construct the probability of rejection $P(\varepsilon, d) = P(\varepsilon_t, \mathcal{P}_{\boldsymbol{\theta}}(\varepsilon))$ by interpolation of $P(y, \pi)$ matrix as a function of ε and $\boldsymbol{\theta}$
   (b) construct the value function as a function of ε and θ using the interpolated matrix $P(\varepsilon_t, \mathcal{P}_{\boldsymbol{\theta}}(\varepsilon))$
   (c) construct the expected value of the value function as a function of θ using some numerical integration method, $\mathsf{E}V(\varepsilon, \theta) = f(\boldsymbol{\theta})$
   (d) construct $\mathsf{E}[d(\varepsilon)] = g(\boldsymbol{\theta})$ as a function of coefficients in the policy function approximation
   (e) run `fmincon` on $f(\boldsymbol{\theta})$ using $g(\boldsymbol{\theta}) = 0$ as a constraint, obtaining a new value function V^{NP}
2. update $\mathsf{E}V^{NP}$ and then $\mathsf{E}V^{P} = \mathsf{E}L^{P} + \delta\mathsf{E}V^{NP}$
3. check if the difference in $\mathsf{E}V^{NP}$ is small, if so, stop, if not start again

The results presented the dedicated section are computed using Simpson numerical integration method. As a robustness check, I also used Gaussian quadrature. As the approximation function, I used Chebyshev polynomials and linear approximation. The number of points in the ε grid is much higher then the number of nodes in the Gauss-Hermite quadrature (of 12th order). Using 12th order guaranties the precise value of the integral if the function being integrated can be written as a Hermite polynomial of up to 23rd order. In theory this should be sufficient for any reasonably smooth function, however, it seems that the policy function can have jumps and hence is be well approximated by polynomials of any order. Therefore the Simpson method seems to be more reliable.

4 Parameter values

Let's recall the parameters of the model:

- Phillips curve slope α: $y_t = \alpha(\pi_t - \pi_{t|t-1}) + \varepsilon_t$
- preferences of the central bank captured by the loss function: $L_t = \frac{1}{2}(y_t - y^*_{high})^2 + \frac{\beta_{low}}{2}\pi_t^2$
  - socially optimal level of output y^*_{high}
  - $(\beta_{high}, y^*_{low})$ the values the central bank announces to act upon
  - the true preference of the central bank $(\beta_{low}, y^*_{high})$. This is also the value of the parameter of all private agents (rest of the society)
  - central bank discount factor δ
- variances of the shocks
  - variance of the supply shock σ^2_ε
  - variance of the noise in the output σ^2_y
  - variance of the noise in the inflation σ^2_π
- punishment cushion parameter κ: agents trigger punishment if log-likelihoods satisfy

$$(\tilde{y}, \tilde{\pi}) \in \mathcal{R} \Leftrightarrow \frac{l_D(\tilde{y}_t, \tilde{\pi}_t)}{l_{ND}(\tilde{y}_t, \tilde{\pi}_t)} < \kappa$$

The values of $(\beta_{high}, y^*_{low})$ and κ can be found by a grid search conditional on the rest of the parameters as the values which minimizes the expected value of the central bank's (and hence also the social) value of the problem. δ is assumed to be 0.99. α by obtained by estimating the Phillips curve from inflation and HP filtered gap measure of output. y^*_{high} is obtained from using a textbook new keynesian approach. The relative variance of the noise shocks is obtained by exploring the difference between the first release of data and the data revised after one year. Finally, empirical literature is consulted to get an insight what the ration between β_{high} and β_{low} in Rogoff setting.

4.1 Slope of the Phillips curve α

The Phillips curve in the model is defined as

$$y_t - y_t^{POT} = \alpha(\pi_t - \pi_{t|t-1}) + \varepsilon_t, \tag{41}$$

where $\pi_{t|t-1}$ is the expectation about period t inflation formed at time $t-1$, the potential output y_t^{POT} is normalized to zero. This PC is different to standard NK (hybrid) PC, as investigated for example by (Galí and Gertler, 1999), because of the timing of the expectations. The results from the standard literature are therefore not directly applicable and I estimate the parameter myself.

4.1.1 Estimation

To capture $\pi_{t|t-1}$ I use the inflation expectation time series compiled by University of Michigan obtained from the database FRED,[6] to obtain output gap y I apply the HP filter with smoothing parameter $\lambda = 1600$ on logarithm of real GDP. Finally, for inflation π I use the year-year percentage change.[7]

The equation 41 cannot be estimated as it is. The reason is that in my model, there is no persistence in the variables, whereas the data is strongly persistent. In order to eliminate this issue, I estimate following equation:

$$y_t = \alpha(\pi_t - \pi_{t|t-1}) + \beta_1 y_{t-1} + \beta_2 y_{t-2} + \varepsilon_t.$$

Two lags of y on the right hand side of the equation are needed in order not to reject H_0 of no autocorrelation among the fitted residuals. However, due to the strong persistence in the forecast error $(\pi_t - \pi_{t|t-1})$, which is assumed away in my model, I add lagged inflation expectation error and estimate:

$$y_t = \alpha_1(\pi_t - \pi_{t|t-1}) + \alpha_2(\pi_{t-1} - \pi_{t-1|t-2}) + \beta_1 y_{t-1} + \beta_2 y_{t-2} + \varepsilon_t \tag{42}$$

The output gap parameters are not affected, but the lagged inflation forecast error is significant.

parameter	Coefficient	Std. Error	t-Statistic	Prob.
β_1	1.152350	0.090676	12.70842	0.0000
β_2	-0.355578	0.090902	-3.91167	0.0002
α_1	0.904448	0.286005	3.16235	0.0020
α_2	-0.658455	0.274285	-2.40063	0.0180

Table 1: Estimation of the slope of the Phillips curve II

The contemporaneous effect is then measured to be $\alpha \approx 0.9$. Hence, in my model I use $\alpha = 0.9$. Some other estimated might suggest a lower values, for example Ireland (1999, page 289) estimated ($\alpha = 0.15$), but results is not directly applicable either, as in his setting the central bank acts before observing the shock, not after as in my model, and he considers the unemployment, not the output gap.

4.2 Preferences of the central bank

4.2.1 Socially optimal level of output y^*_{high}

Galí (2008, page 48) shows that the natural level of output of standard NK model with Calvo

[6]Alternative would be to use the survey of professional forecasters, http://www.philadelphiafed.org/research-and-data/real-time-center/survey-of-professional-forecasters/, Mankiw et al. (2003) document significant differences between the two time series.

[7]inflation expectation: http://research.stlouisfed.org/fred2/series/MICH,
GDP: http://research.stlouisfed.org/fred2/series/GDP,
CPI: http://research.stlouisfed.org/fred2/series/CPIAUCSL

pricing can be derived to be

$$y_t^n = \frac{1+\phi}{\sigma(1-\alpha)+\phi+\alpha} a_t - \frac{(1-\alpha)(\mu - \log(1-\alpha))}{\sigma(1-\alpha)+\phi+\alpha}, \tag{43}$$

where σ is the risk aversion, ϕ is the labour supply elasticity, α is the production function parameter, μ is the logarithm of the flexible price mark-up, equal to $\frac{\varepsilon}{\epsilon-1}$, where ε is the price elasticity of demand. a_t is percentage deviation of the technology from its steady state.

The simplest possible parametrization is following: $\alpha = 0.36$, $\sigma = \phi = 1$. A conservative value for the desired mark-up is about 20%, which gives $\mu = 0.2$ (Ireland, 2001, page 10).

The equilibrium under perfect competition would be higher by $\frac{(1-\alpha)\mu}{\sigma(1-\alpha)+\phi+\alpha}$ percent, which is about 6.4% given the parameter values discussed in the previous paragraph. However, the true risk aversion is probably larger then one and also ϕ might be higher in reality, resulting in higher denominator. Furthermore, the socially optimal level of output might not go as far.

Hence, I assume that the social optimal 4% above the potential level of output. Because the potential is normalized to zero, this means that $y^* = 4\%$.

4.2.2 Rogoff case: $\beta_{high} > \beta_{low}, y^*_{low} = y^*_{high}$

Here I try to find values for β_{high} and β_{low} which are consistent with each other in the Rogoff setting of conservative central banker, so $y^*_{high} = y^*_{low} = y^*$. First, I will look for empirical evidence on β_{high} and then given the choice of β_{high} I will try to find a plausible values for β_{low}. This parametrization is then used to find the optimal policy and show that this policy is nonlinear even in the setting conjectured by Rogoff.

Choosing β_{high} The standard approach to estimate the preferences of a central bank is to set up a system of equations and estimate them simultaneously using ML or as restriction by GMM. The empirical results in the literature are inconclusive. Some authors report that weight on inflation stabilization is about 8-10 times higher then the one of the output, as for example Givens (2009) antwoFavero and Rovelli (2003). However, there also a paper arguing that the ration is about 1.5 (Ozlale, 2003). The evidence on the central bank inflation-output gap trade-off preference parameter is reported in the table 3 on page 36.

I set $\beta_{high} = 8$, i.e. the central bank puts 8 times more weight on stabilizing deviations of inflation from its steady state then the stabilization of output gap.

Choosing β_{low} Now, given β_{high}, what is a conform value of β_{low}? Higher value of β decreases the inflation bias, but at the same time the inflation-output gap trade off becomes less favorable. The optimal value of can be found using the first order condition on the expected loss $L_t(\beta_1, \beta_2, \varepsilon_t)$, a loss of a central bank with preferences described by β_1, which is acting as it has preferences β_2

facing shock ε_t:

$$L_t(\beta_{low}, \beta_{high}, \varepsilon_t) = \frac{1}{2}\left(\frac{\alpha}{\beta_{high}}y^* - \frac{\alpha}{\alpha^2+\beta_{high}}\varepsilon_t - y^*\right)^2 + \frac{\beta_{low}}{2}\left(\frac{\beta_{high}}{\alpha^2+\beta_{high}}\varepsilon_t\right)^2$$
$$= \frac{1}{2}\left[\left(\frac{\beta_{high}}{\alpha^2+\beta_{high}}\right)^2\varepsilon_t^2 - 2y^*\frac{\beta_{high}}{\alpha^2+\beta_{high}}\varepsilon_t + (y^*)^2\right]$$
$$+ \frac{\beta_{low}}{2}\left[\left(\frac{\alpha}{\beta_{high}}y^*\right)^2 - 2\frac{\alpha}{\beta_{high}}y^*\frac{\alpha}{\alpha^2+\beta_{high}}\varepsilon_t + \left(\frac{\alpha}{\alpha^2+\beta_{high}}\right)^2\varepsilon_t^2\right] \quad (44)$$

$$\mathsf{E}L_t(\beta_{low}, \beta_{high}) = \frac{1}{2}\left[\left(\frac{\beta_{high}}{\alpha^2+\beta_{high}}\right)^2\sigma_\varepsilon^2 + (y^*)^2\right] + \left[\left(\frac{\alpha}{\beta_{high}}y^*\right)^2 + \left(\frac{\alpha}{\alpha^2+\beta_{high}}\right)^2\sigma_\varepsilon^2\right]$$
$$= \frac{1}{2}\left(\beta_{high}^2 + \alpha^2\beta_{low}\right)\left(\frac{\sigma_\varepsilon^2}{(\alpha^2+\beta_{high})^2} + \frac{(y^*)^2}{\beta_{high}^2}\right) \quad (45)$$

$$\frac{\partial \mathsf{E}L_t(\beta_{low}, \beta_{high})}{\partial \beta_{high}} = \frac{1}{2}\left(\frac{\sigma_\varepsilon^2}{(\alpha^2+\beta_{high})^2}\right)2\beta_{high} + \frac{1}{2}\left(\beta_{high}^2+\alpha^2\beta_{low}\right)\left(-2\frac{\sigma_\varepsilon^2}{\alpha^2+\beta_{high}} - 2\frac{(y^*)^2}{\beta_{high}^2}\right)$$

$$\frac{\partial \mathsf{E}L_t(\beta_{low}, \beta_{high})}{\partial \beta_{high}} = \alpha^2\left(\frac{\sigma_\varepsilon^2(\beta_{high}-\beta_{low})}{(\alpha^2+\beta_{high})^3} - \frac{(y^*)^2}{\beta_{high}^3}\beta_{low}\right) \quad (46)$$

For given β_{low} (social preferences), this nonlinear equation can be solved numerically to obtain the optimal β_{high} (optimal "pretended" conservativeness of the central banker). However, the social preferences are not observables, because the only observable variables are realized inflation and output, which were chosen by the conservative central banker. Fortunately, I have already found the optimal degree of conservativeness (β_{high}), I can invert the equation to see, what β_{high} implies for possible values of β_{low}. I am using the values of parameters from above. The results are depicted at figure 3. In particular, if is the variance of the supply shock is equal to 2, we can see that β_{low} being between 1 and 2 implies $\beta_{high} \approx 8 - 12$. This is very close to some of the empirical results mentioned above (furthermore, the slope of the expansion path is the inverse of the parameter value, which makes the differences even smaller) and I hence I claim that after calibrating the volatility of the supply shock to $\sigma_\varepsilon = 2.5$, the other parameter choices are reasonable.

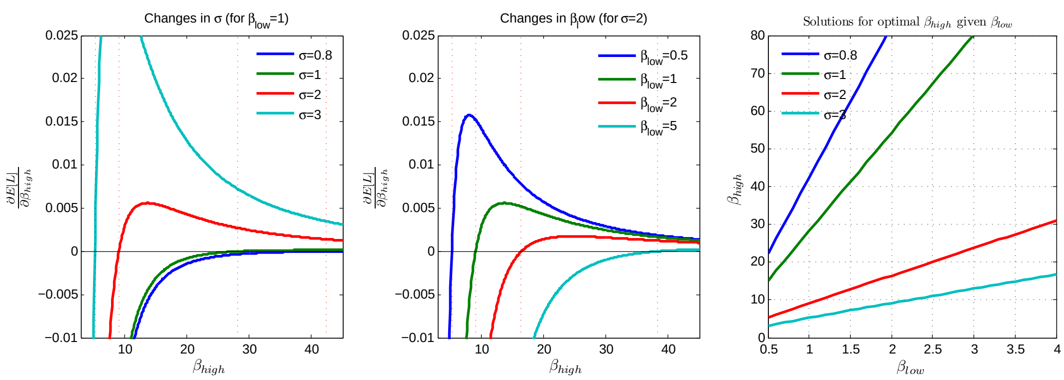

(a) changing σ_ε (b) changing β_{low} (c) β_{high} as a function of β_{low}

Figure 3: The relation between β_{high} and β_{low} in Rogoff's conservative banker setting.

Taylor (1979) argued that while there is no long run trade-off between inflation and output, there is *a second order trade-off*, where the long run volatility of output can be decreased only at cost of increasing the volatility of inflation. Given the results reported in table 2, the long run variance of inflation is about 4 times lower then the variance of output. Considering that the slope of the Phillips curve is about 1, this would suggest that β which the central bank actually follows should be be around 4 as well. This fact might suggest that, in the Rogoff case, setting $\beta_{high} = 8$ with $\beta_{low} = 1$ might be at the more conservative end of the interval for the plausible specification.

4.3 Signal to noise ration

Signal to noise ratio is important factor for the signal extraction problem. Te get an ideas about the relative magnitudes, I take data for CPI and GDP from FRED database, get year on year growth and inflation and compare different vintage time series. I define noise at time t as the difference between the just released number and the revised number one year afterwards.

As can be seen from figure 2, the GDP data is much more noisy then CPI data. The estimated standard deviations are listed in table 2.

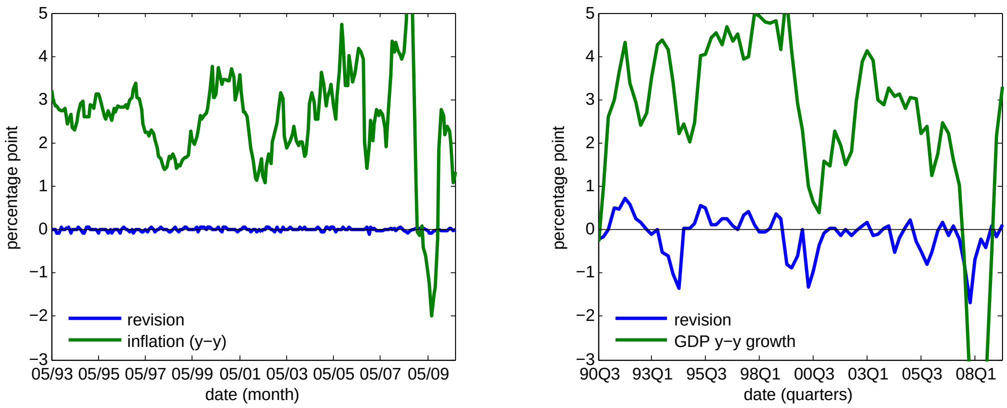

(a) CPI (b) GDP

Figure 4: Noise and signal in the data

	GDP	CPI
The std of signal	2.09	1.14
The std of the noise	0.46	0.04
Signal to noise ratio	4.57	29.16

Table 2: Signal and Noise for CPI and GDP

4.4 Summary

All the numerical results to be presented in the next sections are computed using following parameter values:

parameter	δ	β_{low}	β_{high}	α	y^*	σ_ε	σ_π	σ_y	κ
value	0.99	1	8	0.9	4	2.5	0.2	1	0.8

5 Results

In order to see the effects of the information asymmetries introduced in the model, I analyze first the optimal policy if the central bank announces the Rogoff setting ($\beta_{high} > \beta_{low}, y^*_{high} = y^*_{low}$) and Blinder setting ($\beta_{high} = \beta_{low}, y^*_{low} = 0$).

The results are presented by a showing a two-panel figure. The panel (a) shows the policy functions, the panel (b) captures the implied allocations of output and inflation. Each panel contains policy functions or allocations corresponding to 5 different regimes:

1. the *announced policy* captures what the behavior would be if there was no deviation from the announced policy (captured in black dashed line),
2. *true preferences* line shows what would be the allocation if central bank chooses the allocation according to its true preferences and hence fully deviating by doing so (magenta dashed line),
3. *no_con* depicts the optimal policy given the informational structure but without implementing the constraint $\mathsf{E}[d(\varepsilon) = 0]$ (blue line),
4. the optimal policy approximated by 12th order chebyshev polynomials and linear approximation using the Simpson formula for numerical integration (green and red line).

Furthermore, each panel contains information about the area between 2.5th and 97.5th percentile of ε; the panels (a) contain vertical dashed lines and the allocation corresponding to either quantile is depicted by a dot in the panels (b). Finally, the allocations panels (b) also capture the no rejection region by a green area.

Let's first some general observations. Ignoring the $\mathsf{E}[d(\varepsilon)] = 0$ constraint for a moment and hence focusing on the blue line of the figures, the central bank tries to deviate towards the deviating outcome. However, getting closer to the edge of the Rejection region increases probability of beliefs revision and punishment. Given the relative variance of the noise shocks, the central bank can choose an allocation much closer to the edge along the inflation dimension then the output dimension. This can be observed from the (b) panel of figure 6.

Secondly, let's consider I, the point of intersection of the two expansion paths given given by the announcement and the true preferences. At this point the effective deviation switches. To see this, let's assume a standard case where $\beta_{high} > \beta_{low}$ and $y^*_{low} \leq y^*_{high}$. In such a situation, the two expansion paths intersect at some positive value of y and negative value of π. The deviation is defined as movement along the Phillips curve and in particular, a positive values of d moves in the north-east direction. To the left of the intersection point I, positive values of d move the allocation towards the fully deviation allocation, however, to the right of the intersection point, it is the negative values of d which moves the allocation towards the full deviation.

There are two different robustness checks. First, the deviation function is approximated either by 12th order chebyshev polynomial or by a piece-wise linear function. Second, numerical integration is done either by 12th order Gauss-Hermite quadrature or by using Simpson formula. The results reported are computed using the latter.

5.1 Blinder proposition

The Blinder's suggestion was that the central bank in fact target the natural level of output rather then the socially optimal. In my setting, this translates to the announcement of $(\beta_{high}, y^*_{low}) = (1, 0)$.

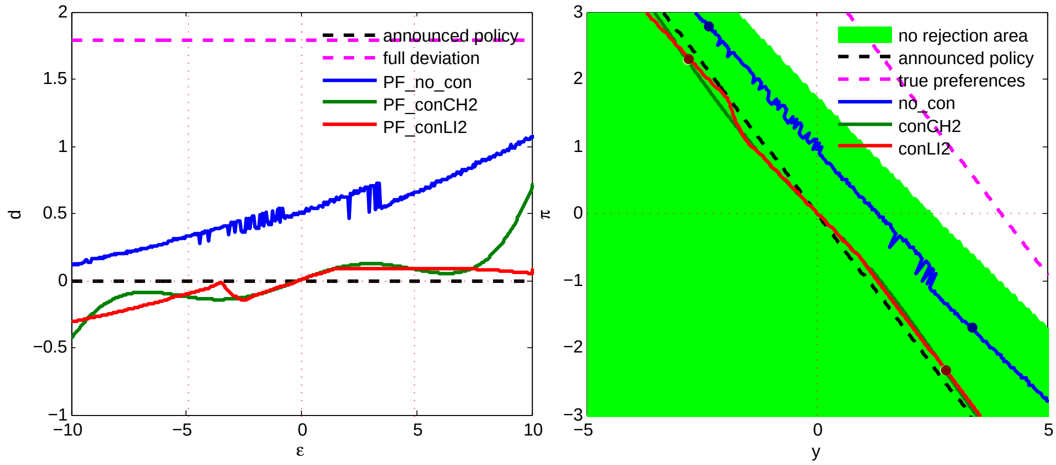

(a) Policy function

(b) allocations

Figure 5: Rogoff proposition: $(\beta_{high} > \beta_{low}, y^*_{low} = y^*_{high})$

The optimal policy in this setting is more accommodative for positive shocks and less so for negative shocks. This comes form the shape of the no revision region. To the contrary to what one might suspect, in the situation where $\beta_{low} = \beta_{high}$ the border of No revision region is not a straight line. To see this, consider the signal extraction problem of the representative agent. The agent analyzes 2 possible actions which could have led to the observed outcome. Assume that the observed allocation is $(\tilde{y}, \tilde{\pi})$ and that it can be explained in the no deviation case by $\hat{\varepsilon}_{ND}$ and in deviation case by $\hat{\varepsilon}_{ND}$. If the central bank deviates, it is increasing both inflation and output. The higher values of the supply shock increase output and decrease inflation and $\hat{\varepsilon}_D < \hat{\varepsilon}_{ND}$ and this affects the likelihoods. The interplay of this effect with the effect of deviation on inflation ultimately causes the nonlinear shape of the border.

5.2 Rogoff proposition

Rogoff (1985) suggested that the inflationary bias can be reduced by introducing a conservative central banker. In my setting, this corresponds to the central bank announcing $(\beta_{high}, y^*_{low}) = (8, 4)$ where the actual numbers were discussed in the previous section. Any $\beta_{high} \in (4, 16)$ would deliver very similar policy functions. The results can be seen in the figure 6.

While there are still some differences between the two methods of approximation, the general pattern is common: the central bank deviates positively for negative values of shocks and to fulfill the constraint $\mathsf{E}[d(\varepsilon)] = 0$, it deviates negatively. The motivation for such a behavior steps from the quadratic loss function; it is *relatively* more beneficial to deviate when facing a very negative shock, then when facing a positive shock, so the central bank finds it optimal to be more

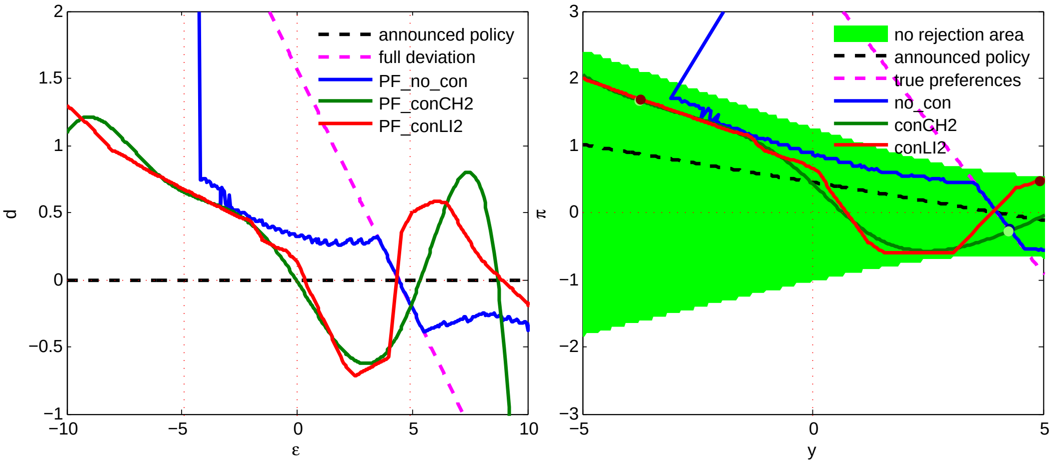

(a) Policy function (b) allocations

Figure 6: Rogoff proposition: $(\beta_{high} > \beta_{low}, y^*_{low} = y^*_{high})$

accommodative in recessions and pay the price of being more hawkish in the booms.

5.3 Optimal announcement

The optimal announcement is such announcement $(\beta_{high}, y^*_{low})$ which minimizes the expected loss of the central bank given the value of κ, Because the central bank has the same loss function as the representative agent, this is the socially optimal policy.

The policy is found by looking at a grid of parameters $(\beta_{high}, y^*_{low}, \kappa)$

Figure 7 shows the expected value of the value function, i.e. the expected value of the loss function given the optimal policy response to ϵ for different values of parameters β_{high}, y^*_{low} and κ, starting from the case when the central bank targets the natural level of output and keeps the original trade-off parameter, $(\beta_{high}, y^*_{low}) = (1, 0)$.

As pointed out already by Nobay and Peel (2003), the base scenario $(\beta_{high}, y^*_{low}) = (1, 0)$ is not welfare maximizing. In my setting, the central bank can decrease the expected loss by targeting even smaller output (panel (a) of figure 7), or become less conservative by choosing β_{high} actually below 1 $\beta_{high} < 1 = \beta_{low}$, (panel (b)). There seem to be no optimal choice for κ in the interval that was searched. The full grid search is to be done. Obviously, such a point might not exists, because as κ get smaller, the agents are less likely to punish the central bank, so the central bank can deviate more and by doing to increase the social welfare. However, the constraint $\mathsf{E}d = 0$ should prevent the central bank from deviating further and further and so it seems reasonable to assume that at some κ_0 there is no benefit of decreasing it below κ_0.

6 Conclusion

The aim of this paper was to show how information asymmetry can give rise to a nonlinear reaction function of the central bank characterized by a standard (perfectly symmetric) quadratic

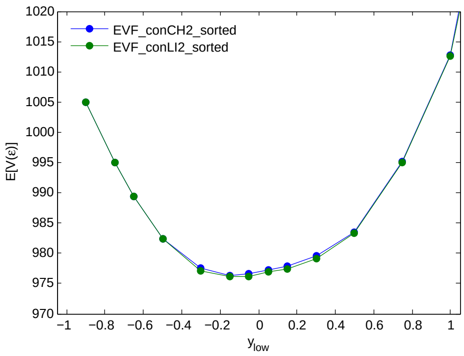

(a) keeping $(\beta_{high}, \kappa) = (1, 0.8)$, changing y^*_{low}

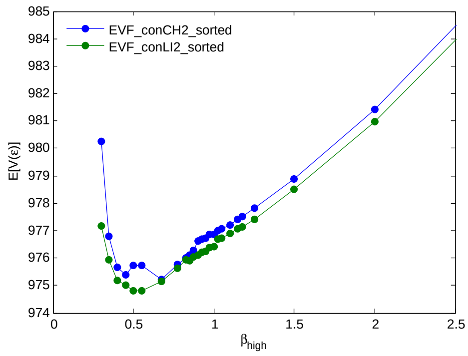

(b) keeping $(\kappa, y^*_{low}) = (0.8, 0)$, changing β_{high}

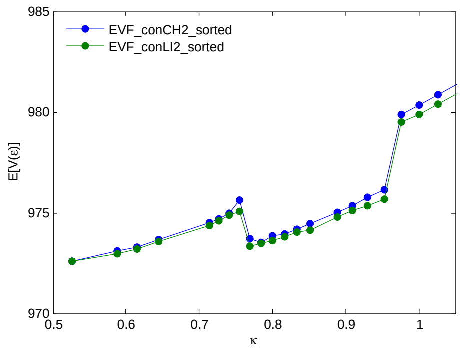

(c) keeping $(\beta_{high}, y^*_{low}) = (1, 0)$, changing κ

Figure 7: Value of the $\mathsf{E}[V]$ for different values of parameters, the Blinder case (baseline: $\beta_{high} = \beta_{low}, y^*_{low} = 0$)

loss function and by showing so, contribute to the discussion about the nonlinear monetary policy rules and asymmetric preferences of the central bank.

In my model, the central bank announces a policy in order to decrease inflationary bias. However, the central bank cannot change its preferences, as was implicitly assumed by Rogoff and Blinder in their contributions to this literature. The decrease in the inflationary bias is then a result of a reputation building repeated game. Reputation building is nontrivial as the representative agent observes only a noisy signal about the allocation chosen by the central bank. Observing the noisy signal, the representative agent tests the announced policy against the deviation policy and chooses to either believe the central bank's announcement or revise her expectations.

I analyzed the two famous answers to the inflationary bias from the literature, Rogoff's conservative central bank and Blinder's targeting the natural level of output. In my setting, the optimal monetary policy is asymmetric in both cases. Furthermore, starting from the Blinder case where the central bank targets the natural rate of output and uses the same output-inflation trade-off parameter β as the rest of the society, the central bank can deliver lower expected social loss by decreasing the output target below zero or by decreasing the trade-off parameter β (full grid search is yet to be completed).

As persistence of the economic fundamentals was assumed away, the representative agent does not use past observations to help him to see if the central bank deviated or not. While this is an obvious simplification, the agents still has perfectly rational expectations in the sense that her inflation expectations correspond with the mean of inflation. Introducing persistence would make the model richer, but only at the cost of increasing the state space, which would significantly increase computational difficulties.

Finally, the optimal nonlinear policy in my model is derived as a result of numerical optimization and as such it might be difficult to implement as opposed to some simple rule. While simple rules are undeniably a useful descriptive and communication tool, reality is more complicated and all the ways to improve social welfare should be exploited.

References

Backus, D. and Driffill, J. (1985). Inflation and reputation. *American Economic Review*, 75(3):530–38.

Barro, R. J. and Gordon, D. B. (1983). Rules, discretion and reputation in a model of monetary policy. *Journal of Monetary Economics*, 12(1):101–121.

Blinder, A. S. (1998). *Central Banking in Theory and Practice.* MIT Press.

Canzoneri, M. B. (1985). Monetary policy games and the role of private information. *American Economic Review*, 75(5):1056–70.

Cukierman, A. and Gerlach, S. (2003). The inflation bias revisited: theory and some international evidence. *Manchester School*, 71(5):541–565.

Cukierman, A. and Meltzer, A. H. (1986). A theory of ambiguity, credibility, and inflation under discretion and asymmetric information. *Econometrica*, 54(5):1099–1128.

Cukierman, A. and Muscatelli, A. (2008). Nonlinear taylor rules and asymmetric preferences in central banking: Evidence from the united kingdom and the united states. *The B.E. Journal of Macroeconomics*, 8(1):7.

DeLong, J. B. (1995). America's peacetime inflation: The 1970s. J. Bradford De Long's Working Papers 104, University of California at Berkeley, Economics Department.

Dennis, R. (2006). The policy preferences of the us federal reserve. *Journal of Applied Econometrics*, 21(1):55–77.

Dolado, J., Maria-Dolores, R., and Ruge-Murcia, F. J. (2004). Nonlinear monetary policy rules: Some new evidence for the u.s. *Studies in Nonlinear Dynamics & Econometrics*, 8(3).

Dolado, J. J., Maria-Dolores, R., and Naveira, M. (2005). Are monetary-policy reaction functions asymmetric?: The role of nonlinearity in the phillips curve. *European Economic Review*, 49(2):485–503.

Favero, C. A. and Rovelli, R. (2003). Macroeconomic stability and the preferences of the fed: A formal analysis, 1961-98. *Journal of Money, Credit and Banking*, 35(4):545–56.

Flood, R. P. and Isard, P. (1988). Monetary policy strategies. NBER Working Papers 2770, National Bureau of Economic Research, Inc.

Galí, J. (2008). *Monetary policy, Inflation, and the Business Cycle.* Princeton University Press.

Galí J. and Gertler, M. (1999). Inflation dynamics: A structural econometric analysis. *Journal of Monetary Economics*, 44(2):195–222.

Geraats, P. M. (2002). Central bank transparency. *Economic Journal*, 112(483):532–565.

Givens, G. E. (2009). Estimating central bank preferences under commitment and discretion. Working Papers 200905, Middle Tennessee State University, Department of Economics and Finance.

Hamilton, J. D. (2001). A parametric approach to flexible nonlinear inference. *Econometrica*, 69(3):537–73.

Ireland, P. N. (1999). Does the time-consistency problem explain the behavior of inflation in the united states? *Journal of Monetary Economics*, 44(2):279–291.

Ireland, P. N. (2001). Sticky-price models of the business cycle: Specification and stability. *Journal of Monetary Economics*, 47(1):3–18.

Kydland, F. E. and Prescott, E. C. (1977). Rules rather than discretion: The inconsistency of optimal plans. *Journal of Political Economy*, 85(3):473–91.

Lohmann, S. (1992). Optimal commitment in monetary policy: Credibility versus flexibility. *American Economic Review*, 82(1):273–86.

Mankiw, N. G., Reis, R., and Wolfers, J. (2003). Disagreement about inflation expectations. Working Paper 9796, National Bureau of Economic Research.

Mavroeidis, S. (2010). Monetary policy rules and macroeconomic stability: Some new evidence. *American Economic Review*, 100(1):491–503.

Nobay, A. R. and Peel, D. A. (2003). Optimal discretionary monetary policy in a model of asymmetric central bank preferences. *Economic Journal*, 113(489):657–665.

Osborn, D. R., Kim, D. H., and Sensier, M. (2005). Nonlinearity in the fed's monetary policy rule. *Journal of Applied Econometrics*, 20(5):621–639.

Ozlale, U. (2003). Price stability vs. output stability: tales of federal reserve administrations. *Journal of Economic Dynamics and Control*, 27(9):1595–1610.

Rogoff, K. (1985). The optimal degree of commitment to an intermediate monetary target. *The Quarterly Journal of Economics*, 100(4):1169–89.

Romer, C. D. and Romer, D. H. (2000). Federal reserve information and the behavior of interest rates. *American Economic Review*, 90(3):429–457.

Ruge-Murcia, F. J. (2003). Does the barro-gordon model explain the behavior of us inflation? a reexamination of the empirical evidence. *Journal of Monetary Economics*, 50(6):1375–1390.

Ruge-Murcia, F. J. (2004). The inflation bias when the central bank targets the natural rate of unemployment. *European Economic Review*, 48(1):91–107.

Söderström, U., Söderlind, P., and Vredin, A. (2005). New-keynesian models and monetary policy: A re-examination of the stylized facts. *Scandinavian Journal of Economics*, 107(3):521–546.

Stokey, N. L. (2003). Rules vs. discretion; after twenty-five years. In *NBER Macroeconomics Annual 2002, Volume 17*, NBER Chapters, pages 9–64. National Bureau of Economic Research, Inc.

Surico, P. (2007). The fed's monetary policy rule and u.s. inflation: The case of asymmetric preferences. *Journal of Economic Dynamics and Control*, 31(1):305–324.

Surico, P. (2008). Measuring the time inconsistency of us monetary policy. *Economica*, 75(297):22–38.

Taylor, J. B. (1979). Estimation and control of a macroeconomic model with rational expectations. *Econometrica*, 47(5):1267–86.

Walsh, C. E. (1995). Optimal contracts for central bankers. *American Economic Review*, 85(1):150–67.

Walsh, C. E. (2001). *Monetary Theory and Policy.* MIT Press.

paper	estimated equation	estimate(SE)	method	period
Söderström et al. (2005)	$\min_i \text{var}(\pi) + \lambda\text{var}(x_t) + \nu\text{var}(\Delta i_t)$	$\lambda = 0(0)$	matching sim. moments	1984Q4-1999Q4
Dennis (2006)	$E_t \sum_{j=0}^{\infty} \beta^j \left[(\pi_{t+j}^a - \pi^*)^2 + \lambda x_{t+j}^2 + \nu(\Delta i_{t+j})^2\right]$	2.94(5.685)	Quasi-FIML	1982Q1-2000Q2
Ozlale (2003)	$(1-\beta)E_t \sum_{j=0}^{\infty} \beta^j \left[\lambda_\pi(\bar{\pi}_{t+j})^2 + \lambda_x x_{t+j}^2 + \lambda_i(\Delta i_{t+j})^2\right]$ such that $\lambda_\pi + \lambda_x + \lambda_i = 1$	$\lambda_\pi = 0.39(0.016)$ $\lambda_x = 0.26(0.012)$ $\lambda_i = 0.35(0.017)$	ML	1970Q1-1999Q1
Favero and Rovelli (2003)	$E_t \sum_{j=0}^{\infty} \beta^j \left[(\pi_{t+j} - \pi^*)^2 + \lambda x_{t+j}^2 + \nu(\Delta i_{t+j})^2\right]$	$\lambda = 0.00125$ (0.0002)	GMM	1980Q3-1998Q3
Givens (2009)	$E_t(1-\beta) \sum_{j=0}^{\infty} \beta^j \left[\pi_{t+j}^2 + \lambda x_{t+j}^2 + \nu(\Delta i_{t+j})^2\right]$	$\lambda_{commit} = 0.135(0.058)$ $\lambda_{disc} = 0.099(0.043)$	ML	1982Q1-2008Q4

Table 3: Central bank preference parameter estimates