21 May 2006

Optimal Monetary Policy under Adaptive Learning

Vitor Gaspar, Frank Smets and David Vestin[1]

Abstract

This paper investigates the implications of private sector adaptive learning for the conduct of monetary policy. The first contribution is to analyze the optimal monetary policy response to shocks and the associated macro-economic outcomes, when the central bank minimizes an explicit loss function. We assume the admittedly extreme case that the central bank has full information about the structure of the economy (a standard assumption under rational expectations) including the precise mechanism generating private sector's expectations. The focus on optimal policy allows us to investigate to what extent a relatively small change in the assumption of how agents form their inflation expectations affects the principles of optimal monetary policy. Moreover, this analysis serves as a benchmark for the analysis of two simple policy rules that would be optimal under rational expectations with and without central bank commitment respectively. The second contribution is to show that a simple representation of the optimal commitment policy under rational expectations that exhibits history dependence is surprisingly robust to changes in the way inflation expectations are formed. It leads to macro-economic outcomes that are very close to those under the optimal policy. By responding persistently to cost-push shocks, the simple commitment rule is able to significantly lower the degree of inflation persistence estimated by the private agents and thereby stabilize inflation and inflation expectations.

Key words: Optimal policy, adaptive learning, rational expectations, policy rules.

JEL Classification System: E52.

---

[1] Correspondence: Vitor_R_Gaspar@bportugal.pt; frank.smets@ecb.int; david.vestin@ecb.int.
The views expressed are the authors' own and do not necessarily reflect those of the European Central Bank, the Banco de Portugal or the Eurosystem. We thank George Evans, our discussant at the 2006 American Economic Association meetings for insightful and constructive comments.

1. Introduction

The assumption of rational expectations (Muth, 1961) has become standard in modern macroeconomics. For example, in the context of a micro-founded New Keynesian model Woodford (2003) systematically explores the implications of rational expectations for the optimal conduct of monetary policy. However, rational expectations assume economic agents who are extremely knowledgeable (Evans and Honkapohja, 2001). A reasonable alternative is to assume adaptive learning. In this case, agents have limited knowledge of the precise working of the economy, but as time goes by, and available data changes, they update their knowledge and the associated forecasting rule. Adaptive learning may be seen as a minimal departure from rational expectations in an environment of pervasive structural change. Moreover, some authors, for example Orphanides and Williams (2004) and Milani (2005), have found that adaptive learning models are able to reproduce important features of empirically observed inflation expectations.

This paper looks at the implications of private sector adaptive learning for the conduct of monetary policy. Our first contribution is to analyze the optimal monetary policy response to shocks and the associated macro-economic outcomes, when the central bank minimizes an explicit loss function.[2] In this analysis we assume the admittedly extreme case that the central bank has full information about the structure of the economy (a standard assumption under rational expectations), including the precise mechanism generating private sector's expectations. To emphasize that we assume the central bank knows everything about the expectations' formation mechanism, we have labeled this

[2] In doing so, we build on Svensson's (2003) distinction between "instrument rules" and "targeting rules". An instrument rule expresses the central bank's policy-controlled instrument, typically a short-term interest rate, as a function of observable variables in the central bank's information set. A targeting rule, in contrast, expresses it implicitly as the solution to a minimization problem of a loss function. Svensson stresses the importance of looking at optimal policy and targeting rules in order to understand modern central banking.

extreme case "sophisticated" central banking in Gaspar, Smets and Vestin (2006). The focus on optimal policy has two objectives. It allows us to investigate to what extent a relatively small change in the assumption of how agents form their inflation expectations affects the principles of optimal monetary policy. Second, it serves as a benchmark for the analysis of two simple policy rules that would be optimal under rational expectations with and without central bank commitment respectively. Here, the objective is to investigate how robust these policy rules are to changes in the way inflation expectations are formed.

Our paper is closely related to the work by Orphanides and Williams (2005). They have shown that, for the case of linear feedback rules, inflation persistence increases when adaptive learning is substituted for rational expectations. They also show that a stronger response to inflation helps limiting the increase in inflation persistence and that, in such a context, a strategy of stricter inflation control helps to reduce both inflation and output gap volatility. In a predecessor to this paper (Gaspar, Smets and Vestin, 2005), we have found that under adaptive learning optimal policy responds persistently to cost-push shocks. Such a persistent response to shocks allows central banks to stabilize inflation expectations, reduce inflation persistence and inflation variance at little cost in terms of output gap volatility. Persistent policy responses and well-anchored inflation expectations resemble optimal monetary policy under commitment and rational expectations. However, the mechanisms are very different. In the case of rational expectations, it operates through expectations of *future* policy. In the case of adaptive learning, it operates through a reduction in inflation persistence, as perceived by economic agents, given the past history determined by shocks and policy responses. Of course; there is no dichotomy between the two mechanisms anchoring inflation expectations. On the contrary, the central bank's ability to influence expectations about the future course of policy rates and its track record in preserving stability are complements.

In this paper, we build on this work by characterizing more fully optimal policy under adaptive learning and contrasting it with the simple rules that correspond to optimal monetary policy under rational expectations. In line with Orphanides and Williams (2005) and Woodford (2005), we show that lessons for the conduct of monetary policy under model-consistent expectations are strengthened, when policy takes modest departures from rational expectations into account. Woodford's (2005) expressed concern that giving the central bank superior knowledge about the expectation formation process might, by exploring systematic forecasting mistakes of the right kind, lead to outcomes superior, to those possible under rational expectations. Such disturbing possibility does not apply to the cases we will discuss in this paper. On the contrary, as stressed above, the main intuition is closely related to Orphanides and Williams (2005) and Woodford (2005), in that departures from rational expectations increase the potential for instability in the economy, thereby strengthening the importance of managing (anchoring) inflation expectations. We also find that the simple commitment rule under rational expectations is robust when expectations are formed in line with adaptive learning. As a matter of fact, for our baseline calibration, macroeconomic outcomes, under the simple commitment rule, are surprisingly close to those under full optimal policy.

The paper is organized as follows. In section 2, we introduce a simple New Keynesian model with adaptive learning and present our benchmark calibration assumptions. In section 3, we present the macro-economic outcomes under different policy regimes and characterize the optimal policy to state variables, especially to cost-push shocks, lagged inflation and perceived inflation persistence. Section 4 contains some robustness analysis with respect to different assumptions regarding the calibration. In section 5, we conclude.

2. New Keynesian model with adaptive learning.

2.1. A simple New Keynesian model of inflation dynamics under rational expectations

Throughout the paper, we use the following standard New Keynesian model of inflation dynamics, which under rational expectations, as extensively discussed in Woodford (2003), can be derived from a consistent set of microeconomic assumptions:

(1) $$\pi_t - \gamma\pi_{t-1} = \beta\left(\mathrm{E}_t\pi_{t+1} - \gamma\pi_t\right) + \kappa x_t + u_t ,$$

where π_t is inflation, x_t is the output gap and u_t is a cost-push shock (assumed i.i.d.). Furthermore, β is the discount rate, κ is a function of the underlying structural parameters including the degree of Calvo price stickiness, α, and γ captures the degree of intrinsic inflation persistence due to partial indexation in the goods market. Galï and Gertler (1999) and Galí, Gertler and Lopez-Salido (2001) have shown that such a hybrid New Keynesian Phillips curve fits the actual inflation process in the United States and the euro area quite well.

In addition, we assume that the central bank uses the following loss function to guide its policy decisions:

(2) $$L_t = (\pi_t - \gamma\pi_{t-1})^2 + \lambda x_t^2 \ .$$

Woodford (2003) has shown that, under rational expectations and the assumed microeconomic assumptions, such a loss function can be derived as a quadratic approximation of the (negative of the) period social welfare function, where $\lambda=\kappa/\theta$ measures the relative weight on output gap stabilization and θ is the elasticity of substitution between the differentiated goods. We implicitly assume that the inflation target is zero. To keep the model simple, we also abstract from any explicit representation of the transmission mechanism of monetary policy and simply assume that the central bank controls the output gap directly.

As discussed in the introduction, we consider two assumptions regarding the formation of inflation expectations in equation (1): rational expectations and adaptive learning. Moreover, we assume that with the exception of the expectations operator, equations (1) and (2) are invariant to these assumptions.[3]

In this subsection, we first solve for optimal policy under rational expectations with and without commitment by the central bank. This will serve as a benchmark for the analysis of optimal policy under adaptive learning. A simple representation of the optimal policy behaviour under rational expectations will also serve to investigate the robustness of those policies under adaptive learning.

Defining $z_t = \pi_t - \gamma\pi_{t-1}$, equations (1) and (2) can be rewritten as:

---

[3] It is clear that in general both the inflation equation (1) and the welfare function (2) may be different when adaptive learning rather than rational expectations are introduced at the micro level (Preston, 2005). In this paper, we follow the convention in the adaptive learning literature and assume that the structural relations (besides the expectations operator) remain identical when moving from rational expectations to adaptive learning.

(1’) $z_t = \beta E_t z_{t+1} + \kappa x_t + u_t$

(2’) $L_t = z_t^2 + \lambda x_t^2$.

Optimal policy under discretion.

If the central bank can not commit to its future policy actions, it will not be able to influence expectations of future inflation. In this case, there are no endogenous state variables and since the shocks are iid, the rational expectations solution (which coincides with the standard forward-looking model) must have the property $E_t z_{t+1} = 0$. Thus:

(1’’) $z_t = \kappa x_t + u_t$

Hence, the problem reduces to a static optimization problem. Substituting (1’’) into (2’) and minimizing the result with respect to the output gap, implies the following policy rule:

(3) $x_t = -\dfrac{\kappa}{\kappa^2 + \lambda} u_t$.

Under the optimal discretionary policy, the output gap only responds to the current cost-push shock. In particular, following a positive cost-push shock to inflation, monetary policy is tightened and the output gap falls. The strength of the response depends on the slope of the New Keynesian Phillips curve, κ, and the weight on output gap stabilization in the loss function, λ.[4]

Using (3) to substitute for x_t in (1’’) and the definition of z_t implies:

(4) $\pi_t = \gamma \pi_{t-1} + \dfrac{\lambda}{\kappa^2 + \lambda} u_t$.

Note that in this case inflation follows an AR(1) process.

Or, expressing inflation directly as a function of the output gap:

(5) $\pi_t = \gamma \pi_{t-1} - \dfrac{\lambda}{\kappa} x_t$.

[4] The reaction function in (3) contrasts with the one derived in Clarida, Gali and Gertler (1999). They assume that the loss function is quadratic in inflation (instead of the quasi-difference of inflation, z_t) and the output gap. They find that, in this case, lagged inflation appears in the expression for the reaction function, corresponding to optimal policy under discretion.

This equation expresses the usual tradeoff between inflation and output gap stability in the presence of cost-push shocks. In the standard forward-looking model (corresponding to γ=0), there should be an appropriate balance between inflation and the output gap. The higher the λ, the higher is inflation in proportion to (the negative of) the output gap, because it is more costly to move the output gap. When κ increases, inflation falls relative to the output gap. When γ>0, it is the balance between the quasi difference of inflation and the output gap that matters. If last periods inflation was high, there is a tendency that current inflation *should* also be high. The reason is that price dispersion drives the welfare criterion. When prices are partially indexed to lagged inflation, other prices must rise in proportion to this indexation in order to avoid price dispersion.

Optimal monetary policy under commitment

As shown above, under discretion optimal monetary policy only responds to the exogenous shock and there is no inertia in policy behaviour. In contrast, as discussed extensively in Woodford (2003), if the central bank is able to credibly commit to future policy actions, optimal policy will feature a persistent "history dependent" response. In particular, Woodford (2003) shows that optimal policy will now be characterized by the following equation:

(6) $$z_t = -\frac{\lambda}{\kappa}(x_t - x_{t-1}) .$$

In this case, the expressions for the output gap and inflation can be written as:

(7) $$x_t = \partial x_{t-1} - \frac{\kappa}{\lambda} u_t \text{, and}$$

(8) $$\pi_t = \gamma\pi_{t-1} + \frac{\lambda(1-\partial)}{\kappa} x_{t-1} + \partial u_t ,$$

where $\partial = \left(\tau - \sqrt{\tau^2 - 4\beta}\right)/2\beta$ and $\tau = 1 + \beta + k^2/\lambda$ (see Clarida, Gali and Gertler, 1999). Comparing equation (3) and (7), it is clear that under commitment optimal monetary policy is characterized by history dependence in spite of the fact that the shock is temporary. The intuitive reason for this is that under commitment perceptions of future policy actions help stabilize current inflation, through their effect on expectations. By ensuring that, under rational expectations, a decline in inflation expectations is associated with a positive cost-push shock, optimal policy manages to spread the impact of the shock over time.

2.2. Inflation expectations under adaptive learning.

In this section we specify the model under adaptive learning. As shown in equation (4), under rational expectations and discretionary monetary policy, the equilibrium dynamics of inflation will follow a first-order autoregressive process:

(4') $$\pi_t = \rho\pi_{t-1} + \tilde{u}_t$$

Under adaptive learning, we assume that the private sector believes the inflation process is well approximated by such an AR(1) process. However, as the private agents do not know the underlying parameters, they estimate the equation recursively, using a "constant-gain" least squares algorithm, implying perpetual learning. Thus, the agents estimate the following reduced-form equation for inflation,[5]

(9) $$\pi_t = c_t\pi_{t-1} + \varepsilon_t .$$[6]

Agents are bounded rational because they do not take into account the fact that the parameter c varies over time. The c parameter captures the estimated, or perceived, inflation persistence.

The following equations describe the recursive updating of the parameters estimated by the private sector.

(10) $$c_t = c_{t-1} + \phi R_t^{-1}\pi_{t-1}(\pi_t - \pi_{t-1}c_{t-1})$$

(11) $$\mathrm{R}_t = \mathrm{R}_{t-1} + \phi(\pi_{t-1}^2 - \mathrm{R}_{t-1}),$$

[5] We assume that the private sector knows the inflation target (equal to zero). In future research, we intend to explore the implications of learning about the inflation target.

[6] Alternatively, we could assume that the private sector assumes that also lagged output gap affects inflation as in the case of commitment (equation 8). However, this would introduce three additional state variables in the non-linear optimal control problem and thereby make it computationally infeasible to numerically solve the model. In this paper, we therefore stick to the simpler univariate AR(1) case.

where ϕ is the gain. Note that due to the learning dynamics the number of state variables is expanded to four: u_t, π_{t-1}, c_{t-1}, R_t, The last two variables are predetermined and known by the central bank at the time they set policy at time t.

A further consideration regarding the updating process concerns the information the private sector uses when updating its estimates and forming its forecast for next period's inflation. We assume that agents use current inflation when they forecast future inflation, but not in updating the parameters. This implies that inflation expectations, in period t, for period t+1 may be written simply as:

(12) $$\mathrm{E}_t \pi_{t+1} = c_{t-1} \pi_t$$

Generally, there is a double simultaneity problem in forward-looking models with learning. In (1), current inflation is determined, in part, by future expected inflation. However, according to (12), expected future inflation is not determined until current inflation is determined. Moreover, in the general case also the estimated parameter, c, will depend on current inflation. The literature has taken (at least) three approaches to this problem. The first is to lag the information set such that agents use only t-1 inflation when forecasting inflation at t+1, which was the assumption used in Gaspar and Smets (2002). A different and more common route is to look for the fixed point that reconciles both the forecast and actual inflation, but not to allow agents to update the coefficients using current information (i.e. just substitute (12) into (1) and solve for inflation). This has the benefit that it keeps the deviation from the standard model as small as possible (also the rational expectations equilibrium changes if one lags the information set), while keeping the fixed-point problem relatively simple. At an intuitive level, it can also be justified by the assumption that it takes more time to re-estimate a forecasting model than to apply an existing model. Finally, a third approach is to also let the coefficients be updated with current information. This results in a more complicated fixed-point problem.[7]

Substituting equation (12) into the New-Keynesian Phillips curve (1) we obtain:

[7] It is possible to solve this problem in the current setting. However, we leave this for future research.

(13) $$\pi_t = \frac{1}{1+\beta(\gamma - c_{t-1})}(\gamma\pi_{t-1} + \kappa x_t + u_t).$$

2.3. Solution method for optimal monetary policy under adaptive learning

Under adaptive learning we want to distinguish between the case where the central bank follows a simple rule (specifically the rules given in equation (3) and (7)) and fully optimal policy under the loss function (2). In the first case, the simple rule (3 or 7), the Phillips curve (1) and equations (10), (11) and (12) determine the dynamics of the system. Standard questions, in the adaptive learning literature, are whether a given equilibrium is learnable and which policy rules lead to convergence to rational expectations equilibrium (Evans and Honkapohja , 2001). By focusing on optimal policy, we aim at a different question. Namely: suppose the central bank knows fully the structure of the model including that agents behave in line with adaptive learning, what is the optimal policy response? And, how will the economy behave? In this case, the central banker is well aware that policy actions influence expectations formation and thereby inflation dynamics. To emphasize that we assume the central bank knows everything about the expectations' formation mechanism, we have labeled this extreme case "sophisticated" central banking in Gaspar, Smets and Vestin (2006). "Sophisticated" central banking implies solving the full dynamic optimization problem, where the parameters associated with the estimation process are also state variables.

Specifically, in this case the central bank solves the following dynamic programming problem:

(14) $$V(u_t, \pi_{t-1}, c_{t-1}, R_t) = \max_{x_t}\left\{-\frac{(\pi_{t,} - \gamma\pi_{t-1})^2 + \lambda x_t^2}{2} + \beta \mathrm{E}_t V(u_{t+1}, \pi_t, R_{t+1}, c_t)\right\},$$

subject to equation (13) and the recursive parameter updating equations (10) and (11).[8]

The solution characterizes optimal policy as a function of the states and parameters in the model, which may be written simply as:

(15) $$x_t = \psi(u_t, \pi_{t-1}, c_{t-1}, R_t) \, .$$

See Appendix A.2 for further details on optimal policy under adaptive learning. The presence of learning instead of fully rational agents introduces three modifications relative to the standard framework under rational expectations. First, the agents simply run their regression and make their forecast, so that actual inflation is not the outcome of a game between the central bank and the private sector (as is the case under discretion and rational expectations). Second, promises of future policy play no role as agents look only at inflation outcomes. Hence, there is no scope for the type of commitment gains discussed in the rational expectations literature. Third, we leave the linear-quadratic world, as the learning algorithm makes the model non-linear.

From a technical perspective, the first two aspects simplify finding the optimal policy whereas the third is a complication. The value function will not be linear-quadratic in the states and hence we employ the collocation-methods described in Judd (1998) and Miranda and Fackler (2002) to solve the model numerically. This amounts to approximating the value function with a combination of cubic splines and translates in a root finding exercise. Further information on numerical simulation procedures is outlined in appendix A3.

[8] The value function is defined as $V(.) = \max_{\{x_j\}} \left\{ -\sum_j \beta^j \left[(\pi_j - \gamma\pi_j)^2 + \lambda x_j^2 \right] s.t.(1),(8),(9) and (10) \right\}$, that is as maximizing the negative of the loss. It is important to bear this in mind when interpreting first order conditions.

2.4. Calibration of the model

In order to study the dynamics of inflation under adaptive learning we need to make specific assumptions about the key parameters in the model. In the simulations, we use the set of parameters shown in Table 1 as a benchmark.

[Insert Table 1]

Coupled with additional assumptions on the intertemporal elasticity of substitution of consumption and the elasticity of labor supply these structural parameters imply that κ=0.019; λ= 0.002. [9] γ is chosen such that there is some inflation persistence in the benchmark calibration. A value of 0.5 for γ is frequently found in empirically estimated new Keynesian Phillips curves (see, for example, Smets (2002) or Gali and Gertler (1999)). θ=10 corresponds to a mark-up of about 10%. 1-α measures the proportion of firms allowed to change prices optimally each period. α is chosen such that the average duration of prices is three quarters; which is consistent with US evidence. The constant gain, ϕ, is calibrated at 0.02. Orphanides and Williams (2004) found that a value in the range 0.01 to 0.04 is needed to match up the resulting model-based inflation expectations with the Survey of Professional Forecasters. A value of 0.02 corresponds to an average sample length of about 25 years. [10] In the limiting case, when the gain approaches zero, the influence of policy on the estimated inflation persistence goes to zero and hence plays no role in the policy problem.

[9] Here we follow the discussion in Woodford (2003). See especially pages 187 and 214-15.
[10] See Orphanides and Williams (2004). Similarly, Milani (2005) estimates the gain parameter to be 0.03 using a Bayesian estimation methodology.

3. Optimal monetary policy under adaptive learning

In this section, we first discuss the macro-economic performance under adaptive learning. We compare the outcomes under rational and adaptive expectations for both optimal monetary policy and the simple policy rules given by equations (3) and (7) above. Next, we characterize optimal monetary policy by looking at the shape of the policy function and mean dynamic impulse responses following a cost-push shock.

3.1. Optimal monetary policy, persistence and macroeconomic performance

Table 2 compares, for our benchmark calibration, five cases: two under rational expectations and three under adaptive learning. Under rational expectations we compare the discretionary and commitment policy; under adaptive learning we compare the optimal policy with the discretion and commitment rules (equation (3) and (7) respectively) that would be optimal under rational expectations.

It is instructive to first start walking a well-trodden path comparing the outcomes under commitment and discretion, under rational expectations. For such a case we have shown (see also Clarida, Gali and Gertler (1999) and Woodford (2003)) that commitment implies a long-lasting response to cost-push shocks persisting well after the shock has vanished from the economy. As already stated above, the intuition is that generating expectations of a reduction in the price level, in the face of a positive cost-push shock, optimal policy reduces the immediate impact of the shock, spreading it over time. With optimal policy under commitment, inflation expectations operate as automatic stabilizers in the face of cost-push shocks. Such intuition is clearly present in the results presented in Table 2. Clearly, the output gap is not persistent under the simple rule, (under the assumption that cost-push shocks are i.i.d.). In contrast, under commitment the output gap becomes very persistent with autocorrelation of 0.66. The reverse is true for inflation. Inflation persistence, under discretion, is equal to the assumed intrinsic persistence parameter at 0.5. Under commitment it comes down to less than half of that: 0.24. The inflation variance is about 85 % higher under discretion and the variance of the quasi-difference of inflation is about 37% higher. At the same time, output gap volatility is only about 5 % lower. The reduction in output gap volatility illustrates the stabilization bias under optimal discretionary monetary policy. Overall, the loss is about 28 % higher under discretion.

[Insert Table 2]

Following Orphanides and Williams (2002), it is also useful to compare the outcomes under rational expectations and adaptive learning for the case of the discretion and commitment rules (comparing the first and second columns with the third and fourth in Table 2). The comparison confirms the findings of Orphanides and Williams (2002). Clearly, the autocorrelation and the volatility of the output gap remain unchanged in both cases, under the simple rules the output gap only responds to the exogenous cost-push shock and (in the commitment case) its own lag. Nevertheless, under adaptive learning, the autocorrelation of inflation increases from 0.5 to about 0.56 in the discretion case and from 0.24 to 0.34 in the commitment case. As a result, the loss increases by about 8 pp under discretion and 11 pp under commitment. The intuition is that, under adaptive learning, inflation expectations operate as an additional channel magnifying the immediate impact of cost-push shocks and contributing to the persistence of their propagation in the economy. The increase in persistence and volatility are intertwined with dynamics induced by the learning process.

How does optimal monetary policy perform under adaptive learning (last column of Table 2). As expected, it is able to improve macro-economic performance relative to the simple linear rules that were optimal under rational expectations. Interestingly, it leads to similar outcomes as the commitment cases. Optimal policy induces considerable persistence in the output gap and thereby reduces sharply the persistence of inflation to about 0.34 (the same as under the commitment rule). As before, this is linked with a significant decline in inflation volatility relative to the discretionary outcomes. Inflation variance declines by 95 percentage points to about only 23% more than in case of commitment under rational expectations. The variance of the quasi-difference of inflation also falls by about 38 percentage points. At the same time, the volatility of the output gap is slightly higher than under the discretion rules. On balance, the expected welfare loss falls significantly, by about 28 percentage points, when optimal policy replaces the simple discretionary rule.

Overall, it appears that optimal policy under adaptive learning brings the loss close to the one under commitment and rational expectations, as we can see from a comparison between the second and the last column in Table 2. Moreover, in both cases the output gap exhibits significant persistence and inflation is much less persistent than under the discretion rule.. Nevertheless, it is still the case that even under optimal policy, adaptive learning makes inflation more persistent and the economy less stable than under rational expectations and the commitment rule. A second important conclusion to highlight is that the simple commitment rule, in which the output gap only responds to the cost-push shock and its own lag, does surprisingly well under adaptive learning. It delivers results very close to full optimal policy. In Gaspar, Smets and Vestin (2005), we argue that optimal policy under adaptive learning operates by responding persistently to cost push shocks, together with optimal response to the other state variables. The remarkable performance of the simple commitment rule under adaptive learning suggests that the ability of the central bank to adapt its response to cost-push shocks, depending on the state of the economy (e.g. lagged inflation and the perceived inflation persistence) is only second order importance relative, to its ability to bring the perceived persistence of the inflation process down, through a persistent response to cost-push shocks.

Figure 1 provides some additional detail concerning the distribution of the endogenous variables, i.e. the estimated persistence, output gap, inflation, quasi-difference of inflation, and the moment matrix, under optimal policy and the simple rules. First, panel (a) shows not only that the average of the estimated persistence parameter is significantly lower under the optimal policy and the simple commitment rule, but also that the distribution is more concentrated around the mean. It is important to note that, under optimal policy, the perceived inflation parameter does never go close to one, contrary to what happens under the simple discretion rule. In fact, the combination of the simple discretion rule and private sector's perpetual learning at times gives rise to explosive dynamics, when perceived inflation persistence goes to (or above) unity[11]. In order to portray the long run distributions, we have excluded explosive paths by assuming (following Orphanides and Williams, 2004) that when perceived inflation reaches unity the updating stops, until the updating pushes the estimated parameter downwards. Naturally, this assumption leads to underestimating the risks of instability under the discretion rule. In Gaspar, Smets and Vestin (2006) we looked at the transition from an economy, regulated by the discretion rule, taking off on an explosive path to the anchoring of inflation through optimal policy. Optimal monetary policy under adaptive learning succeeds in excluding such explosive dynamics.

Second, panels (b), (c) and (d) confirm the results reported in Table 2. Under the optimal policy and the simple commitment rule, the distributions of inflation (panel c) and of the quasi-difference of inflation (panel d) become more concentrated. At the same time, the distributions of the output gap, in panel (d), are very similar confirming the result that the variances of the output gap under the two regimes are identical. Finally, the distribution of the R matrix also shifts to the left and becomes more concentrated under optimal policy, reflecting the fact that the variance of inflation falls relative to the simple discretion rule.

[Insert Figure 1]

Overall, optimal monetary policy under adaptive learning shares some of the features of optimal monetary policy under commitment. To repeat, in both cases persistent responses to cost-push shocks induce a significant positive autocorrelation in the output gap, leading to lower inflation persistence and volatility, through stable inflation expectations. Nevertheless, the details of the mechanism, leading to these outcomes must be substantially different. As we have seen, under rational expectations commitment works through the impact of *future* policy actions on current outcomes. Under adaptive learning, the announcement of *future* policy moves is, by assumption, not relevant. We devote the rest of the section to characterizing optimal monetary policy under adaptive learning and how it works.

[11] Similar results, for the case of a Taylor rule, are reported by Orphanides and Williams (2004).

3.2. Optimal monetary policy under adaptive learning: how does it work?

As we have discussed before optimal policy may be characterized as a function of the four state variables in the model: $(u_t, \pi_{t-1}, c_{t-1}, R_t)$. In the appendix A. 2 we show that equation (15) can implicitly be written as:

$$(16) \qquad x_t = -\frac{\kappa\delta_t}{\kappa^2\delta_t + \lambda\chi_t^2}u_t + \frac{\kappa\gamma(\chi_t - \delta_t) + \beta\kappa\chi_t\phi R_t^{-1}E_t V_c}{\kappa^2\delta_t + \lambda\chi_t^2}\pi_{t-1} + \beta\frac{\kappa\chi_t}{\kappa^2\delta_t + \lambda\chi_t^2}E_t V_\pi$$

where $\delta_t = 1 - 2\beta\phi E_t V_R$, $\chi_t = 1 + \beta(\gamma - c_{t-1})$ and V_c, V_π and V_R denote the partial derivatives of the value function with respect to the variables indicated in the subscript. When interpreting equation (16) there are two important points to bear in mind. First, the partial derivatives V_c, V_π and V_R depend on the vector of states $(u_{t+1}, \pi_t, c_t, R_{t+1})$. The last three states, in turn, depend on the history of shocks and policy responses. Second, the value function is defined in terms of a maximization problem. In such a case, a positive partial derivative means that an increase in the state contributes favorably to our criterion. Or, more explicitly, that it contributes to a *reduction* in the loss.

In order to discuss some of the intuition behind the optimal policy reaction function, it is useful to consider a number of special cases. In particular, in the discussion that follows, we assume that $E_t V_R$ is zero, so that the expected marginal impact of changes in the moment matrix on the value function is zero. Such assumption provides a reasonable starting point for the discussion for reasons, which we make clear in the appendix A.4. If $E_t V_R$ is zero, then $\delta_t = 1$, making equation (16) much simpler.

The intra-temporal trade-off ($\pi_{t-1} = 0$)

If lagged inflation is equal to zero, $\pi_{t-1}=0$, the optimal monetary policy reaction (16) can be reduced to a simple response to the current cost-push shock:

(17) $$x_t = -\frac{\kappa}{\kappa^2 + \lambda\chi_t^2} u_t.$$

This is the case because clearly the second term on the right hand side of equation (16) is zero and moreover, it can be shown that for $\pi_{t-1} = 0$, $E_t V_\pi$ is zero.

If, in addition, $c_{t-1} = \gamma$ and as a result $\chi_t^2 = \chi_t = 1$, equation (17) reduces to the simple rule derived under rational expectations and discretion given by equation (3). In other words, when lagged inflation is zero and the estimated inflation persistence is equal to the degree of intrinsic persistence, the immediate optimal monetary policy response to a shock under adaptive learning coincides with the optimal response under discretion and rational expectations[12].

The reason for this finding is quite simple. From equation (10), it is clear that, when lagged inflation is zero, the estimated persistence parameter is not going to change irrespective of current policy actions. As a result, no benefit can possibly materialize from trying to affect the perceived persistence parameter. The same intuition holds true to explain why when the constant gain parameter is zero (ϕ=0) the solution under fully optimal policy coincides with (3), meaning that the simple discretion rule would lead to full optimal policy.

In this case, only the intra-temporal trade-off between output and inflaton stabilization plays a role. However, different from the discretionary policy under rational expectations, the optimal response under adaptive learning will in general depend on the perceived degree of inflation persistence. For example, when the estimated persistence is lower than the degree of intrinsic persistence, $\gamma > c_{t-1}$, the immediate response to a cost-push shock will be less, $\frac{\kappa}{\kappa^2 + \lambda\chi_t^2} < \frac{\kappa}{\kappa^2 + \lambda}$, than under the simple discretion rule. The reason is again intuitive. As shown in equation (13), the smaller the degree of perceived inflation persistence, the smaller the impact of a given cost-push shock on inflation, all other

12 However, it is clear from figure 2 that the policy response under optimal policy will persist contrary to the simple discretion rule. See further discussion below.

things constant. As a result, when balancing inflation and output gap stabilization, it is optimal for the central bank to mute its immediate response to the cost-push shock. This clearly illustrates the first-order benefits of anchoring inflation expectations. Conversely, when perceived inflation persistence is relatively high, the response of optimal policy to cost push shocks becomes stronger even on impact than under the simple rule.

In Figure 2, we illustrate this response by showing the mean dynamics response of the output gap, inflation and estimated persistence to a one-standard deviation (positive) cost push shock, taking lagged inflation to be initially zero, for different initial levels of perceived (or estimated) inflation persistence on the side of the private sector. Panel a) confirms the finding discussed above that as estimated persistence increases so does the output gap response (in absolute value). The stronger policy reaction helps mitigating the inflation response, although it is still the case (from panel b) that inflation increases by more when estimated inflation persistence is higher. This illustrates the worse trade-off the central bank is facing when estimated persistence is higher. Finally, from panel c) it is apparent that the estimated persistent parameter adjusts gradually to its equilibrium value, which is lower than the degree of intrinsic persistence.

[Insert Figure 2]

The intertemporal trade-off ($u_t = 0$)

Returning to equation (16) and departing from the assumption that $\pi_{t-1}=0$, we can discuss the second term, on the right hand side, which captures part of the optimal response to lagged inflation.

$$x_t = +\frac{\kappa\gamma(\chi_t - \delta_t) + \beta\kappa\chi_t\phi R_t^{-1} E_t V_c}{\kappa^2\delta_t + \lambda\chi_t^2}\pi_{t-1} + \ldots$$

Note that the first term in the numerator is zero when $\gamma = c_{t-1}$ (still using the simplifying assumption that $\delta_t = 1$). In such a case, inflation expectations adjust to past inflation just in line with the partial adjustment of inflation due to its intrinsic persistence (equation

11). Given the loss function this is a desirable outcome. In the absence of any further shock, inflation will move exactly enough so that the quasi-difference of inflation will be zero. Note that when $\gamma > c_{t-1}$ or $\chi_t > 1$ the response of the output gap to past inflation, according to this effect, is positive. Hence, past inflation justifies expansionary policy. At first sight, this is counter-intuitive. However, the reason is clear, when estimated persistence is below intrinsic persistence, past inflation does not feed enough into inflation expectations, to stabilize the quasi-difference of inflation. In order to approach such a situation an expansionary policy must be followed. This factor is important because it shows that, in the context of our model, there is a cost associated with pushing the estimated persistence parameter too low.

However, the important point to make is a different one. In general, the second term in the numerator of the reaction coefficient will be negative and dominate the first term ensuring a negative response of the output gap to inflation. This term reflects the intertemporal trade-off the central bank is facing between stabilizing the output gap and steering the perceived degree of inflation persistence by inducing forecast errors. In our simulations it turns out that the expected marginal cost (the marginal impact on the expected present discounted value of all future losses) of letting estimated inflation persistence increase is always positive, i.e. $V_c < 0$ and large. The intuition is that, as discussed above, a lower degree of perceived persistence will lead to a much smaller impact of future cost-push shocks on inflation, which tends to stabilize inflation, its quasi-difference and the output gap. As a result, under optimal policy the central bank will try to lower the perceived degree of inflation persistence. As is clear from the private sector's updating equation (10), it can do so by engineering unexpectedly low inflation when past inflation is positive and conversely by unexpectedly reducing the degree of deflation when past inflation is negative. In other words, in order to reap the future benefits of lowering the degree of perceived inflation persistence, monetary policy will tighten if past inflation is positive and will ease if past inflation is negative. Overall, this effect justifies a counter-veiling response to lagged inflation, certainly in the case of $\gamma = c_{t-1}$, when the first term in the numerator is zero.

Finally, the third term in equation (13) is also interesting. We have already noticed that when π_{t-1}=0, $E_t(V_\pi)$=0 and this term plays no role. Now, if π_{t-1}>0, and u_t=0 then $E_t(V_\pi)$<0 and this will reinforce the negative effect of inflation on the output gap discussed above. More explicitly, if lagged inflation is positive, this term will contribute to a negative output gap – tight monetary policy - even in the absence of a contemporary shock. This effect will contribute to stabilizing inflation close to zero. In the case π_{t-1}<0, and u_t=0, in contrast $E_t(V_\pi)$>0. Thus, when lag inflation is negative, this term will contribute to a positive output gap – loose monetary policy – even in the absence of a contemporary shock. Again this effect will contribute to stabilizing inflation close to zero.

Perceived persistence and symmetry.

Figures 3a and 3b summarize some of the important features of the shape of the policy function (15) in the calibrated model. Figure 3a plots the output gap (on the vertical axis) as a function of lagged inflation and the perceived degree of inflation persistence for a zero cost-push shock and assuming that the moment matrix R equals its average for a particular realization of c. A number of features are worth repeating. First, when lagged inflation and the cost-push shock are zero, the output gap is also zero irrespective of the estimated degree of inflation persistence. Second, when the shock is zero, the response to inflation and deflation is symmetric. Third, as the estimated persistence of inflation increases, the output gap response to inflation (and deflation) rises. It is then interesting to see how the output gap response differs when a positive cost-push shock hits the economy. This is shown in Figure 3b, which plots the differences in output gap response to a positive one-standard deviation cost-push shock and zero cost-push shock as a function of lagged inflation and the perceived persistence parameter. The output gap response is always negative and increases with the estimated degree of inflation persistence. The figure also shows the non-linear interaction with lagged inflation. In particular, the output gap response becomes stronger when inflation is already positive.

[Insert Figure 3]

Finally, it is also interesting to ask whether the symmetric response of optimal policy to inflation and deflation is more general. More formally, does the following equality hold?

(18) $\psi(u_{t,} - \pi_{t-1,} c_{t-1,} R_t) = -\psi(-u_{t,} \pi_{t-1,} c_{t-1,} R_t)$

The answer is yes, as we illustrate in figure 4, for the case of a positive (negative) cost-push shock when lag inflation is negative (positive). The policy response, apparent in panel (b), is fully symmetric. Moreover, from the panel (a) of Figure 4 it is clear that the adjustment of inflation is also symmetric. Finally, panel (d) shows that the adjustment of estimated persistence is the same in both cases (the small discrepancy in the figure is due to the numerical accuracy of our numerical procedure). The same would be true of the moment matrix (not shown).

[Insert Figure 4]

4. Some sensitivity analysis

In this section we analyze how some of the results depend on the calibrated parameters. First, we investigate how the results change with a different gain and a different degree of price stickiness. Second, we look at the impact of increasing the weight on output gap stabilization in the central bank's loss function.

[Insert Figure 5]

Figure 5 plots the realization of the average perceived inflation persistence in economies with different gains and two different degrees of price stickiness (α=0.66, corresponding to our baseline calibration and a higher degree of price stickiness, α=0.75). Remember that (1-α) measures the proportion of firms changing prices optimally each period. The other parameters are as in the calibration reported in Table 1. We focus on the perceived degree of persistence because this gives an idea about how the trade-off between lowering inflation persistence and stabilizing the output gap changes as those parameters change. As discussed above, when the gain is zero, the optimal policy converges to the simple discretion rule and the estimated degree of persistence equals the degree of intrinsic persistence in the economy (0.5 in the benchmark case). In this case, the central bank can no longer steer inflation expectations and the resulting equilibrium outcome is the same as under rational expectations. Figure 5 shows that an increasing gain leads to a fall in the average perceived degree of inflation persistence. With a higher gain, agents update their estimates more strongly in response to unexpected inflation developments. As a result, the monetary authority can more easily affect the degree of perceived persistence, which affects the trade-off in favor of lower inflation persistence. Figure 5 also shows that a higher degree of price stickiness increases the degree of inflation persistence. Again the intuition is straightforward. With higher price stickiness, it is more costly in terms of variation in the output gap to affect the degree of inflation persistence through unexpected inflation.

Finally, we look at the impact of increasing the weight on output gap stabilization in the central bank's loss function. Figure 6 shows that increasing the weight λ from 0.002 to 0.012 shifts the distribution of the estimated degree of inflation persistence to the right. The mean increases from 0.33 to 0.45. A higher weight on output gap stabilization makes it more costly to affect the private sector's estimation of the degree of inflation persistence and therefore leads to a higher average degree of inflation persistence.

[Insert Figure 6]

5. Conclusions

In this paper we look at optimal monetary policy when private sector expectations are determined in accordance to adaptive learning. As in Orphanides and Williams (2005) and Woodford (2005) our main conclusion is that the fundamental policy prescriptions under model consistent expectations continue to hold, or are even strengthened, by limited departures from rational expectations. Specifically, when expectations are formed in accordance with adaptive learning, the gains from anchoring inflation and inflation expectations increase significantly. Optimal policy under adaptive learning stabilizes inflation and inflation expectations mainly through persistent responses to cost-push shocks. The previous remark explains why, in our numerical examples, the simple commitment rule performs well under adaptive learning. By responding persistently to cost push shocks, the simple commitment rule is able to significantly lower the degree of estimated inflation persistence relative to the simple discretion rule. It is worthwhile stressing that the simple commitment rule is able to approximate quite closely the outcomes that could be obtained under full optimal policy.

In our set-up, monetary policy actions have intra-temporal and intertemporal effects. For example, we have seen that monetary policy responds relatively strongly to lag inflation and to inflation shocks, when the estimated persistence parameter is high. In such a case the central bank, facing positive inflation, will push down estimated persistence, by generating unexpectedly low inflation (in the case of deflation by generating unexpectedly high inflation). In our model simulations the intertemporal, long-term considerations, dominate optimal policy when trade-offs between intra-temporal and inter-temporal considerations arise. The importance of inter-temporal considerations helps to explain why optimal policy under adaptive learning pushes down the estimated persistence parameter to values well below intrinsic inflation persistence and the equilibrium value under the simple rule. By behaving in this way, optimal monetary policy provides an anchor for inflation and inflation expectations, thus contributing to the overall stability of the economy and to better macroeconomic outcomes, as evaluated by the social loss function. We view optimal monetary policy under adaptive learning as illustrating (once more) why medium term price stability and anchoring inflation expectations is key in environments characterized by endogenous inflation expectations.

We have also found that, even in the context of an over-simple model, the characterization of optimal policy becomes very involved. It is easy to imagine how much more difficult such a characterization would become if we would try to reckon the complexity of actual policy choices and the prevalence of economic change. Such considerations clearly limit the possibility of using our framework in a *prescriptive* way. However, we have shown in the paper, that Woodford's (2003) case for emphasizing central banking as management of expectations comes out even stronger when adaptive learning substitutes for model consistent expectations.

Table 1: Relevant parameters for the benchmark case.

β	γ	λ	θ	α	ϕ	κ	σ
0.99	0.5	0.002	10	0.66	0.02	0.019	0.004

Table 2: Summary of macro-economic outcomes

	Rational Expectations		Adaptive Learning		
	Discretion	Commitment	Discretion Rule	Commitment Rule	Optimal
$Corr(x_t, x_{t-1})$	0	0.66	0	0.66	0.54
$Corr(\pi_t, \pi_{t-1})$	0.50	0.24	0.56	0.34	0.34
$Var(x_t)$	0.95	1	0.95	1	1.02
$Var(\pi_t)$	1.85	1	2.18	1.27	1.23
$Var(\pi_t - \gamma\pi_{\tau-1})$	1.38	1	1.49	1.14	1.11
$E[L_t]$	1.29	1	1.37	1.11	1.09

Notes: $Var(x_t)$, $Var(\pi_t - \gamma\pi_{\tau-1})$ and $E[L_t]$ are measured as ratios relative to commitment

Figure 1: The distribution of the estimated inflation persistence (a), output gap (b), inflation (c), quasi-difference of inflation (d) and the moment matrix (e).

(a) Estimated inflation persistence

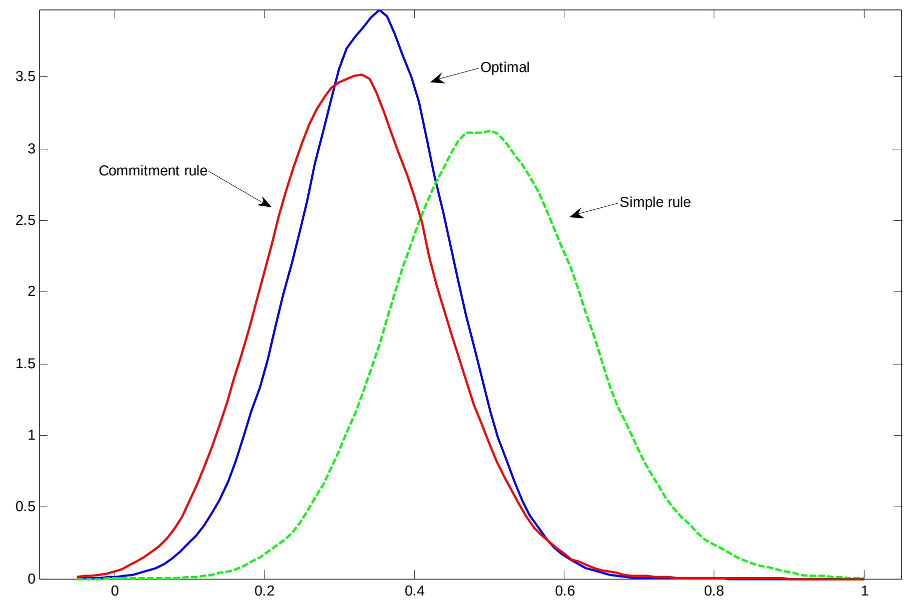

(b) Output gap

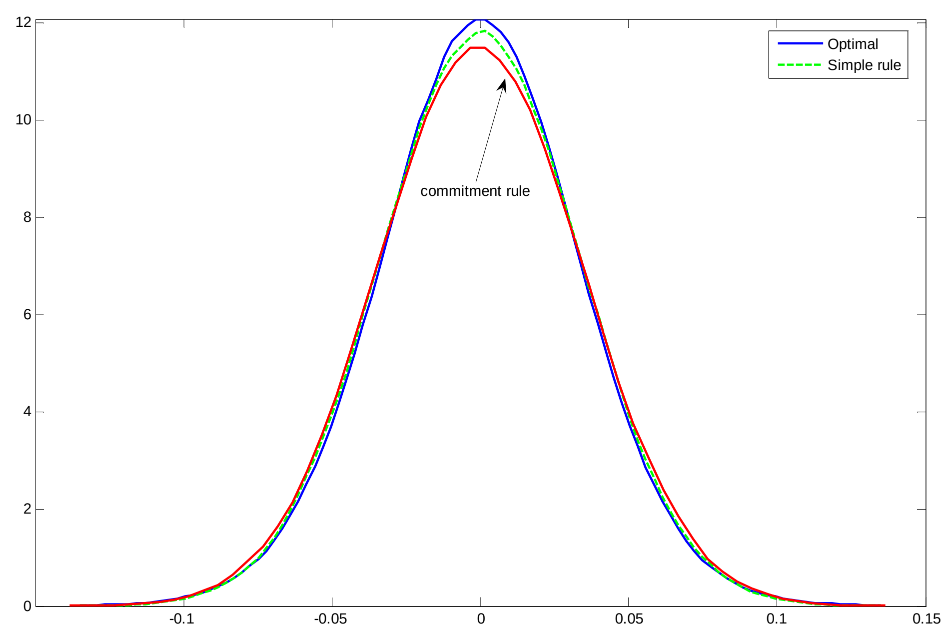

(c) Inflation

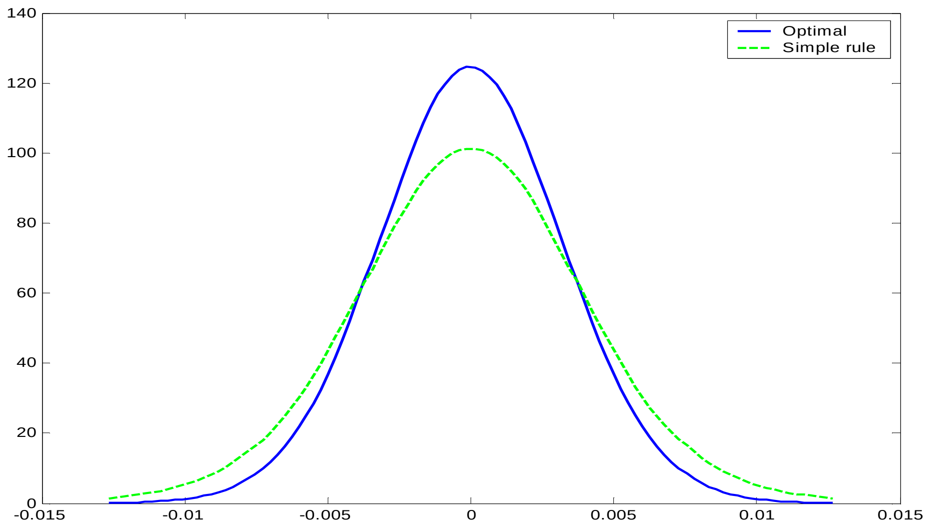

(d) Quasi-difference of inflation

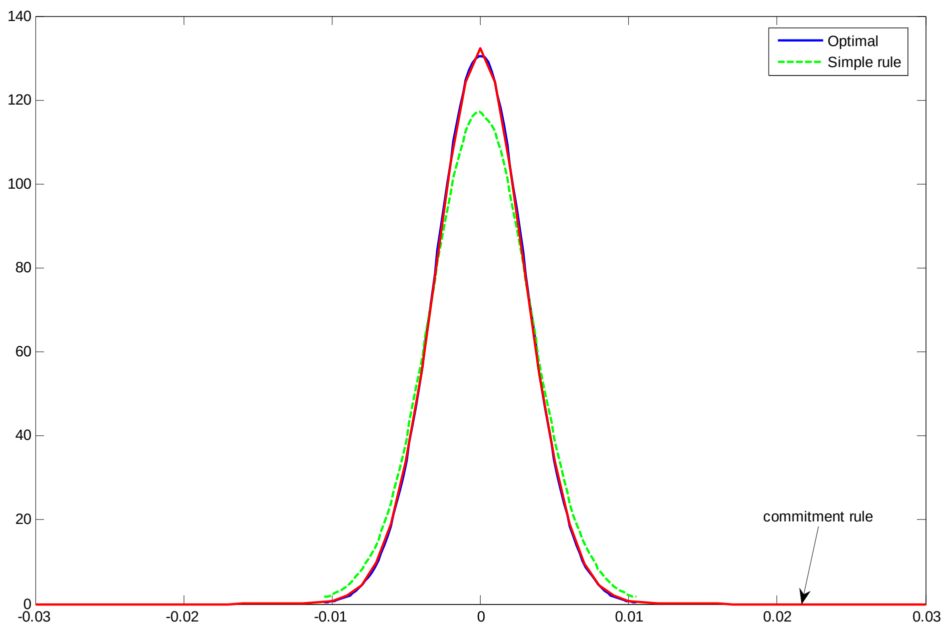

(e) R – moment matrix

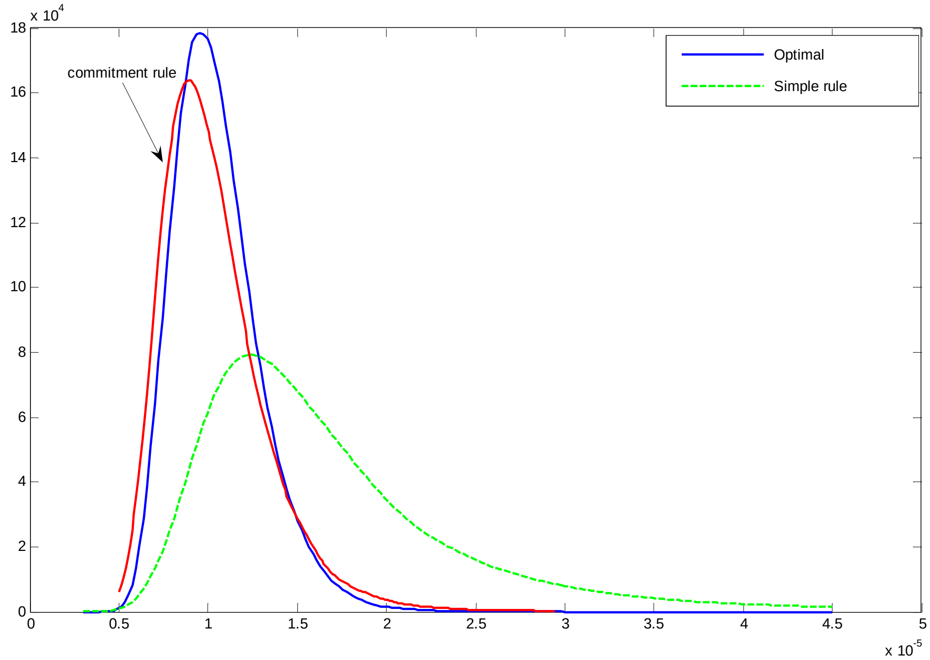

Figure 2: The mean dynamics of the output gap, inflation and the estimated inflation persistence following a one-standard deviation cost-push shock

(a) Output gap

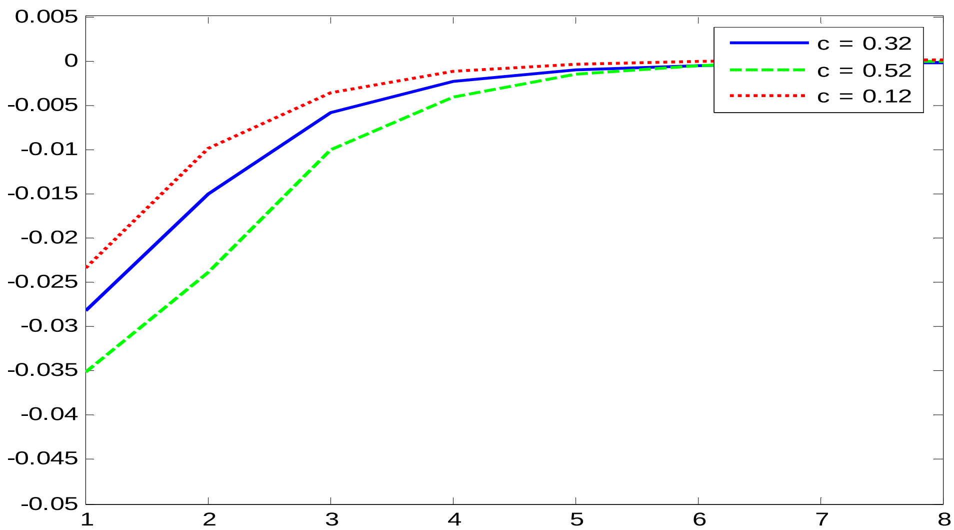

(b) Inflation

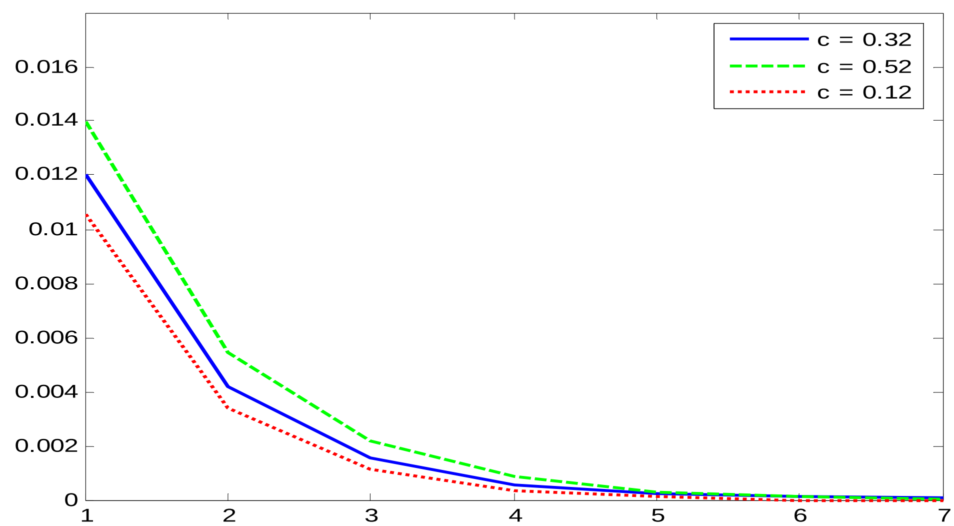

(c) Estimated inflation persistence

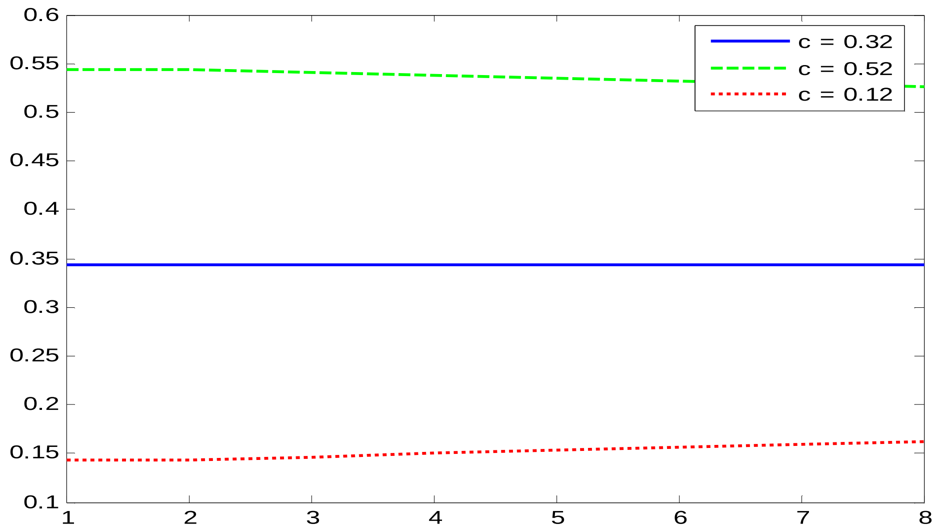

Figure 3: The policy function output gap as a function of lagged inflation and the estimated degree of inflation persistence.

(a)

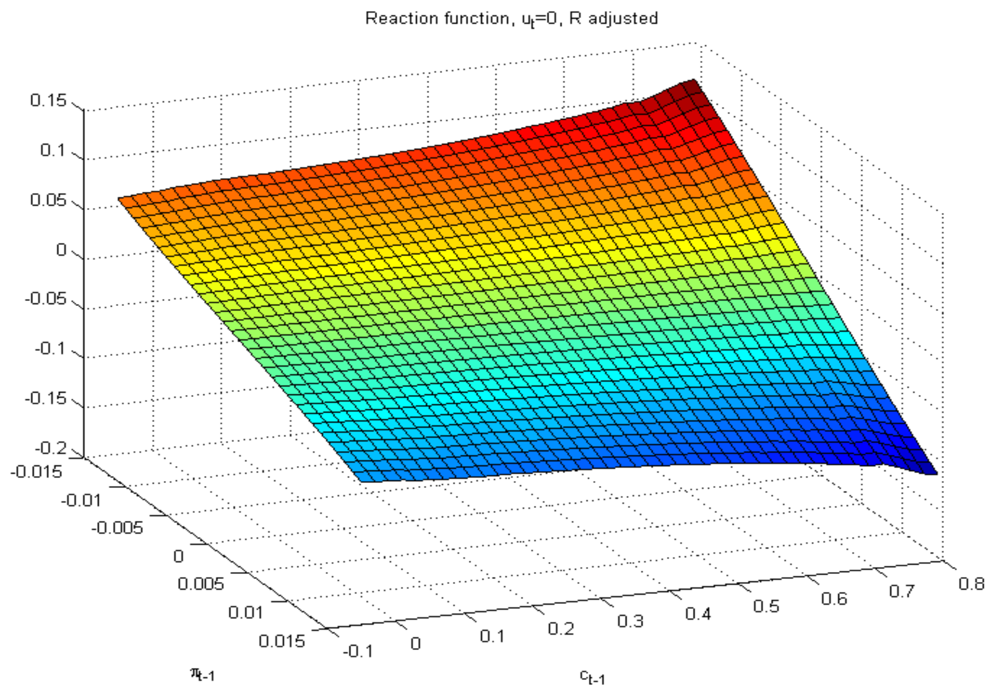

(b) Difference x(sigm,.)-x(0,.)

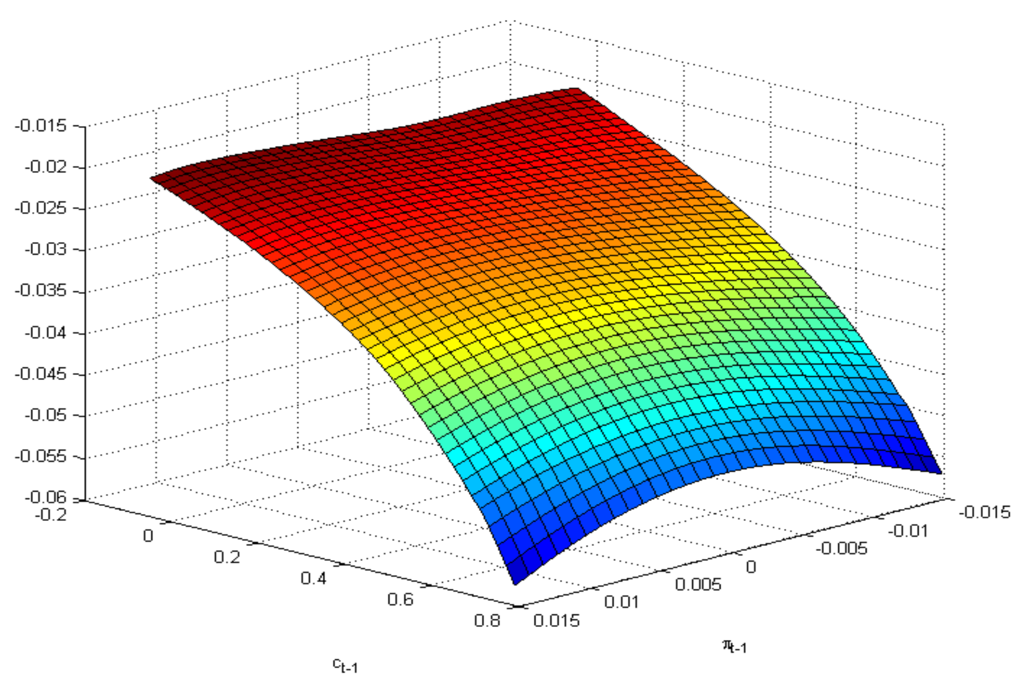

Figure 4: Illustration of symmetry in the response of policy (output gap) to inflation and deflation

(a) Inflation annualized

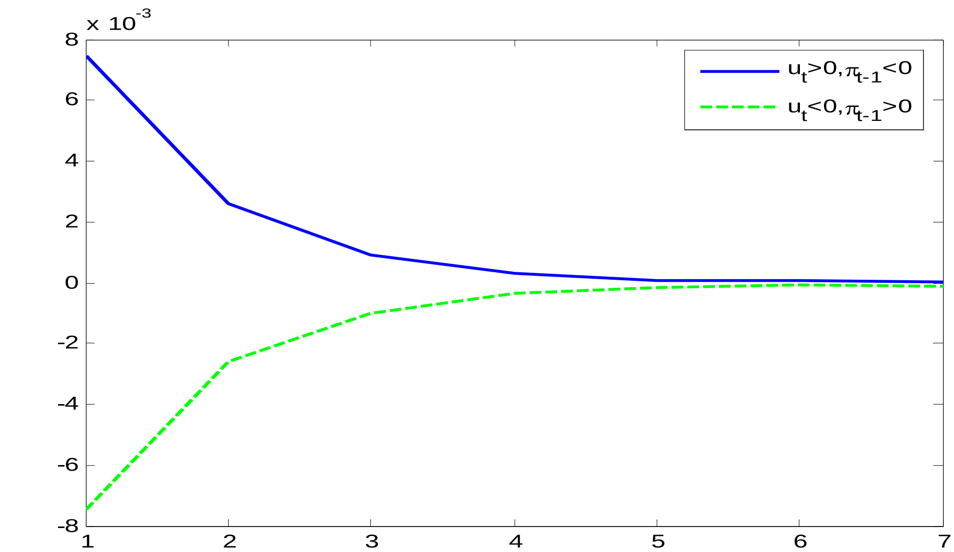

(b) Output Gap

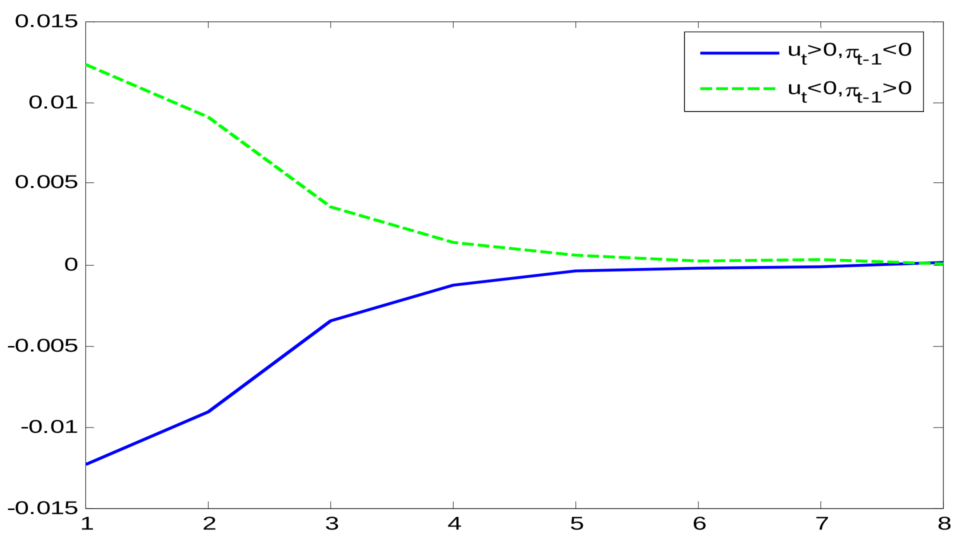

(c) Estimated c

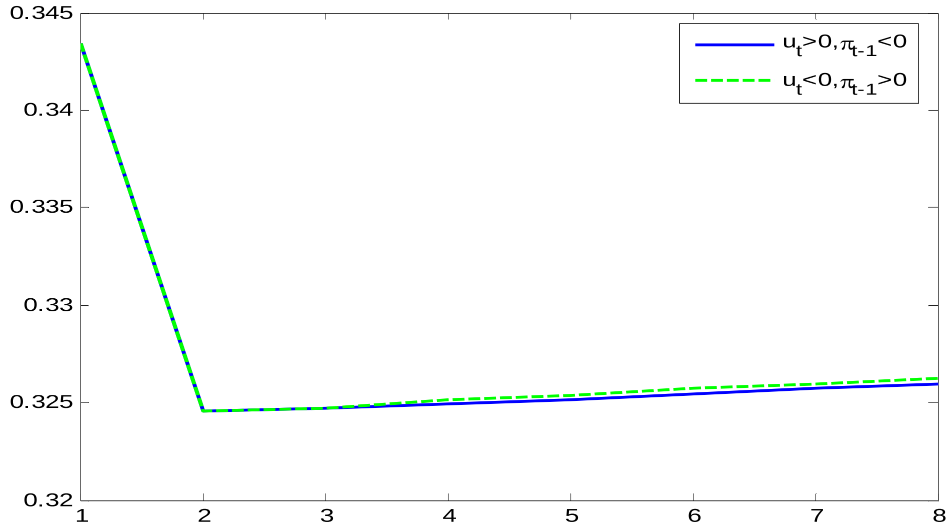

Figure 5: Sensitivity analysis: Average estimated persistence in function of the gain and the degree of price stickiness.

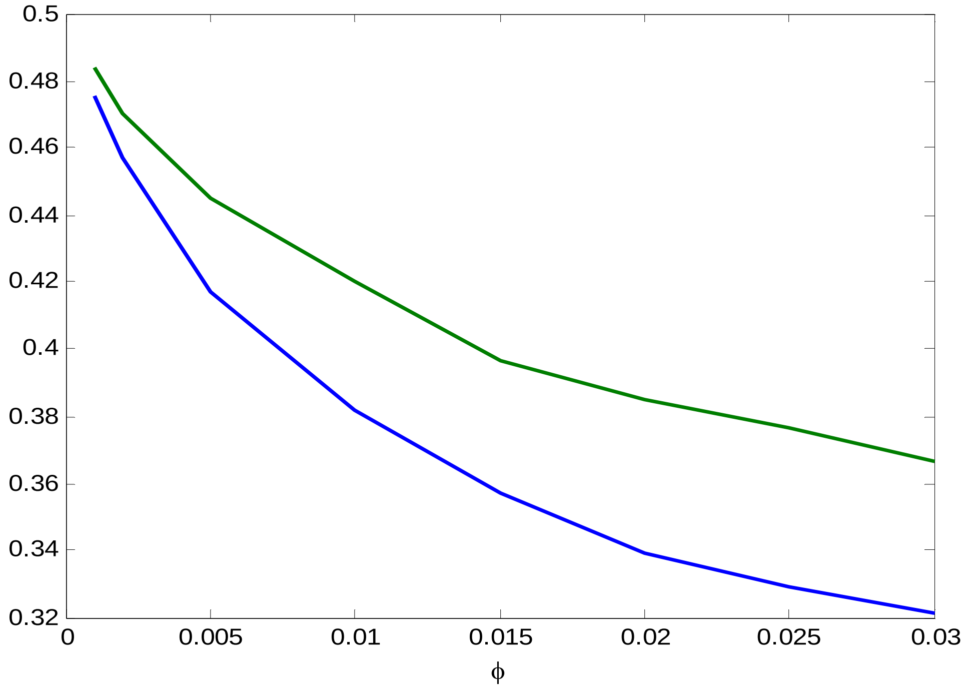

Figure 6: Distribution of estimated inflation persistence as a function of the weight on output gap stabilization.

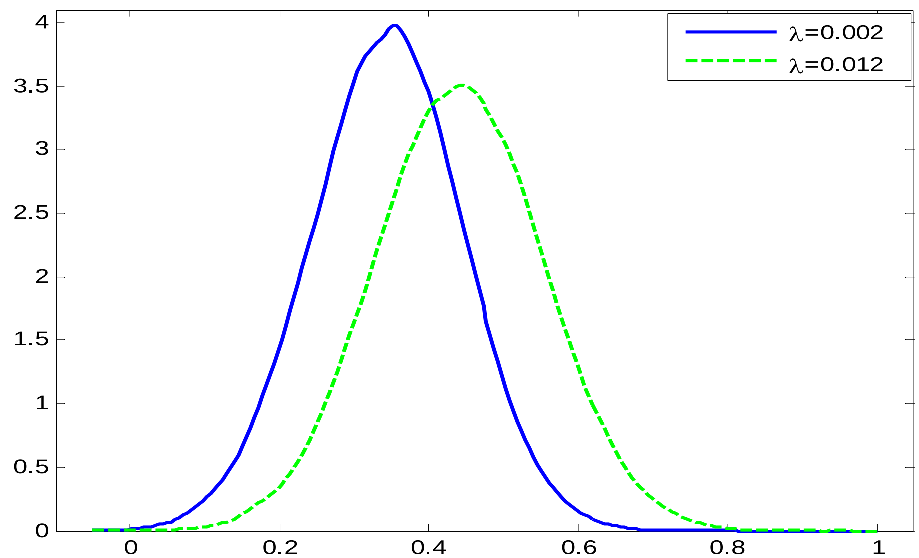

REFERENCES

Clarida, Richard, Jordi Galí and Mark Gertler (1999) The Science of Monetary Policy: a New Keynesian Perspective. *Journal of Economic Literature* 37, 1661-1707.

Dixit, Avinash K. and Joseph E. Stiglitz (1977) Monopolistic Competition and Optimum Product Diversity. *American Economic Review* 67, 297-308.

Evans, George. W. and Seppo Honkapohja (1998) Economic Dynamics with Learning: New Stability Results. *Review of Economic Studies 6, 23-44.*

Evans, George W. and Seppo Honkapohja (2001) *Learning and Expectations in Macroeconomics.* Princeton: Princeton University Press.

Gali, J. and M. Gertler, 1999, Inflatio Dynamics: a Structural Econometric Analysis, *Journal of Monetary Economics* 44 (2), 195-222.

Galí, J., M. Gertler and D. López-Salido (2001). European inflation dynamics. *European Economic Review*, 45(7), 1237-1270.

Gaspar, Vitor and Frank Smets (2002) Monetary Policy, Price Stability and Output Gap Stabilisation. *International Finance* 5, 193-202.

Gaspar, Vitor, Frank Smets and David Vestin (2005), Adaptive Learning, Persistence and Optimal Monetary Policy, paper presented at the 2005 Meetings of the European Economic Association, Amsterdam, 2005.

Gaspar, Vitor, Frank Smets and David Vestin (2006), Monetary Policy over Time, *Macroeconomic Dynamics,* 1-23..

Judd, Kenneth L. (1998) *Numerical Methods in Economics.* Cambridge: MIT Press.

Miranda, Mario J. and Paul Fackler (2002) *Applied Computational Economics and Finance.* Cambridge: MIT Press.

Milani, F., (2005), Expectations, Learning and Macroeconomic Persistence, Unpublished, Princeton University.

Muth, John (1961) Rational Expectations and the Theory of Price Movements. *Econometrica*, 29, 315-335.

Orphanides, Athanasios and John C. Williams (2002) Robust Monetary Policy Rules with Uknown Natural Rates. *Brookings Papers on Economic Activity*, 63-118.

Orphanides, Athanasios and John C. Williams, 2004, Robust Monetary Policy with Imperfect Knowledge, draft presented at the ECB Conference Monetary Policy and Imperfect Knowledge, Wurzburg, October 14-15, 2004.

Orphanides, Athanasios and John C. Williams (2005) 'Imperfect Knowledge, Inflation Expectations and Monetary Policy,' in *Inflation Targeting,* Ben Bernanke and Michael Woodford (eds), University of Chicago Press.

Preston, B. (2005), Learning About Monetary Policy Rules When Long Horizon Expectations Matter, *International Journal of Central Banking,* vol. 1 (2): 81-126.

Smets, F., (2004), Maintaining Price Stability: How Long is the Medium Term, *Journal of Monetary Economics,* 50: 1293-1309.

Smets, F. and R. Wouters (2003) An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area. *Journal of the European Economic Association 1, 1123-75.*

Steinsson, Jon (2003) Optimal Monetary Policy in an Economy With Inflation Persistence. *Journal of Monetary Economics 50, 1425-56*

Svensson, Lars (2003) What Is Wrong with Taylor Rules? Using Judgment in Monetary Policy through Targeting Rules," Version 3.1, January 2003, *Journal of Economic Literature* 41, 426-477.

Woodford, Michael (2003) *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton: Princeton University Press.

Woodford, Michael (2005) Robustly Optimal Monetary Policy with Near Rational Expectations, *NBER* Working Paper 11896, December.

A Optimal policy

The value function problem is given by

$$V(s_t) = \max_{x_t} - \left[\frac{(\pi_t(x_t) - \gamma\pi_{t-1})^2 + \lambda x_t^2}{2} \right] + \beta \mathrm{E}_t V(g(s_t, x_t, u_{t+1}))$$

where the transition functions for the states are given by

$$\begin{aligned}
u_{t+1} &= u_{t+1} \\
\pi_t &= \frac{1}{1+\beta(\gamma - c_{t-1})}(\gamma\pi_{t-1} + \kappa x_t + u_t) \\
R_{t+1} &= R_t + \phi\left(\pi_t^2 - R_t\right) \\
c_t &= c_{t-1} + \phi R_t^{-1}\pi_{t-1}(\pi_t - c_{t-1}\pi_{t-1})
\end{aligned}$$

where the two last equations follow the appropriate form since π_t can be substituted out using the second line.

What we need for the algorithm are the derivatives of the above transition functions with respect to the control, x_t, where we use the notation $s_{t+1}^y = g^y(s_t, x_t, u_{t+1})$.

$$\begin{aligned}
\frac{\partial g^\pi(s_t, x_t, u_{t+1})}{\partial x_t} &= \frac{\kappa}{1+\beta(\gamma - c_{t-1})} \\
\frac{\partial g^R(s_t, x_t, u_{t+1})}{\partial x_t} &= 2\phi\pi_t \frac{\partial g^\pi(s_t, x_t, u_{t+1})}{\partial x_t} \\
\frac{\partial g^c(s_t, x_t, u_{t+1})}{\partial \pi_t} &= \phi R_t^{-1}\pi_{t-1}\frac{\partial g^\pi(s_t, x_t, u_{t+1})}{\partial x_t} \\
\frac{\partial^2 g^\pi(s_t, x_t, u_{t+1})}{\partial x_t^2} &= 0 \\
\frac{\partial^2 g^R(s_t, x_t, u_{t+1})}{\partial x_t^2} &= 2\phi\left(\frac{\partial g^\pi(s_t, x_t, u_{t+1})}{\partial x_t}\right)^2. \\
\frac{\partial^2 g^c(s_t, x_t, u_{t+1})}{\partial x_t^2} &= 0
\end{aligned}$$

We partition the state space using splines of order three, using the following values:

- $\pi_{t-1} : [-0.028, -0.0130, -0.007, -0.002, 0, 0.002, 0.007, 0.0130, 0.028]$
- $R_t : 1.0 * 10^{-3} * [0.0090, 0.0125, 0.015, 0.03, 0.07, 0.15, 0.3]$
- $c_{t-1} : [0.01, 0.100, 0.300, 0.500, 0.700, 0.90, 0.95]$

B Simulations

The model is simulated (in the case of adaptive learning, for the RE case there are analytical solutions) for 500000 periods, for the three cases of policy rules: discretion RE, commitment RE, and fully optimal policy starting from the unconditional averages of the states (recursively found) for each case. The first 10000 observations are dropped when calculating the sample statistics reported in the paper.