Preliminary and incomplete.

Optimal Monetary Policy under Adaptive Learning

Vitor Gaspar, Frank Smets and David Vestin[1]

December 2005

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

JEL Classification System: E52.

[1] 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.

1. Introduction

Rational expectations (Muth, 1961) have become standard in modern macroeconomics. Researchers have systematically explored the implications of rational expectations for economic dynamics and for the conduct of policy. However, rational expectations (paraphrasing Evans and Honkapohja (2001)) assume economic agents who are extremely knowledgeable. An alternative is to limit their knowledge, so that as time goes by and available data changes, so does the agents' forecasting rule. The alternative can be understood as implying bounded rationality, or imperfect information and knowledge. Adaptive learning is a particularly attractive option reckoning with the implications from pervasive structural change and permanent transformation that characterizes the economic environment of modern economies. In other words, adaptive learning may be seen as a minimal departure from rational expectations, a departure which is necessary in order to be able to account for imperfect knowledge about the structure of the economy (model misspecification) and also for structural change. Moreover, some authors, for example Orphanides and Williams (2004) and Milani (2005) have found that adaptive learning models manage to reproduce important features of empirically observed expectations.

Orphanides and Williams (2005) have shown that adaptive learning matters for how monetary policy should be conducted with a view to macroeconomic stability. They show, for the case of linear feedback rules, that 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. Thus, in such a context, a strategy of stricter inflation control helps to reduce both inflation and output gap volatility.

This paper looks at the implications of adaptive learning, on the part of the private sector, for the conduct of optimal monetary policy. In doing so, we build on Svensson's (2003) distinction between "instrument rules" and "targeting rules". An instrument rule expresses the value of the policy instrument as an explicit function of variables in the central bank's information set. A targeting rule, in contrast, expresses it implicitly through the specification of a loss function and the assumption of optimizing, forward-looking behavior on the part of policy-makers. The well-known Taylor (1993) rule is an example of a simple instrument rule. Modern central banking stresses the importance of systematic rule-like behavior. Svensson interprets recent trends in central banking as involving: a) an increasingly precise formulation of the objectives of monetary policy; b) an institutional setting ensuring firm commitment, on the part of the monetary authority, to achieving those objectives. Thus, he argues for a broad concept of monetary policy rules encompassing both instrument rules and targeting rules. In his view, considering target rules is crucial to describe modern central banking. Moreover, when thinking about evaluating monetary policy it is natural to ask: what criteria should be used to determine rational monetary policy rules? A useful benchmark is to consider the minimization of an explicit loss function. Svensson (2003) filled a gap in the literature, providing a systematic discussion of targeting rules. The ambition of our paper is to contribute to bridging the same gap in the adaptive learning literature.

Modeling the optimal behavior of central banks requires specifying its information set. In this paper, we consider the (admittedly) extreme case of sophisticated central banking. Specifically, we assume that the central bank has full information about the structure of the economy (a standard assumption under rational expectations). In our case, the information set includes knowledge about the precise mechanism generating private sector's expectations.

Kydland and Prescott's (1977) seminal contribution opened the way to considering the effects from systematic monetary policy actions and allowed for a theoretical account of important policy concepts such as credibility and reputation. In a world of rational expectations, policy-makers (sufficiently) concerned about their long-run reputation do not yield to short run temptations. The performance of the economy is better as a consequence. Thus, in a rational expectations framework, it is possible to justify primacy of long run goals such as price stability. It is interesting to ask: are similar considerations relevant when we depart from rational expectations? In Gaspar, Smets and Vestin (2005b) we find that optimal policy under adaptive learning responds persistently to cost-push shocks. Through a persistent response to shocks, coupled with the optimal response to other state variables, we have found that optimal central banking under adaptive learning stabilizes inflation expectations, reduces 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. By conducting optimal monetary policy, the central bank establishes a track record of stable inflation and anchors inflation expectations, making it, in turn, easier to maintain price stability. It is a virtuous circle of stability begetting stability. In this paper, we characterize optimal policy under adaptive learning and contrast it with a simple rule, which corresponds to optimal monetary policy under discretion and assuming rational expectations on the part of the private sector. Of course; the dichotomy between the two mechanisms anchoring inflation expectations is not relevant. 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.

The paper is organized as follows. In section 2, we introduce a simple New Keynesian model with adaptive learning. We also present our benchmark calibration assumptions. In section 3, we present the different policy regimes. We characterize the optimal policy to state variables, especially to cost-push shocks, lagged inflation and perceived inflation persistence. In section 4, we discuss how robust the characterization of optimal policy is when we depart from the assumptions made for the benchmark 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

In this section, we introduce a very simple New Keynesian model of inflation dynamics. The New Keynesian model is standard (e.g. Woodford, 2003). We may derive it from the following microeconomic assumptions. First, we assume that preferences are of the Dixit and Stiglitz (1977) type. In such setting, there is limited elasticity of substitution across goods, θ, implying monopolistic competition among a continuum of otherwise identical firms. Imperfect competition is important because it allows producers to be price-setters. Second, we assume that producers set prices in an environment of nominal rigidity, formalized using the Calvo (1983) contracts. Specifically, in any, given, period, all firms will be “allowed” to reset their price optimally, with an exogenous and constant probability. Third, firms produce using a technology that exhibits diminishing returns to labor. Finally, we assume output is demand determined, which means that firms will sell whatever quantity demanded at its current price. These four assumptions create a role for monetary policy, because monetary policy has the ability to affect real output, through its influence over nominal expenditure, and also because, without intervention, markets may produce inefficiently. A simple way, to see the latter, is to note that since all firms are symmetric and marginal cost is increasing in production, the optimal allocation will be such that all firms have the same level of output. However, with some prices fixed, this will typically not be the case. Optimal monetary policy will strive to equalize relative prices of firms of the two groups, to avoid dispersion in the output distribution. The features described so far are, usually, present in new Keynesian models.

In order to examine the issues we are interested in, we introduce two further assumptions. First, motivated by the empirical fact that inflation is relatively persistent, we introduce indexation to lagged inflation among the firms, which do not reset their price optimally along the lines of Christiano et al. (2005), Smets and Wouters (2003) and Woodford (2003). In this case, current inflation will have two components, one coming from the optimally reset prices (the only component in the standard framework) and one due to the fact that all other prices change in proportion to lagged inflation. Specifically, we assume that firms unable to re-set their prices optimally will index their prices (partially) to lagged inflation. Second, we assume a temporary cost-push shock that affects inflation. In terms of microeconomics, this can be motivated by a stochastic intra-temporal elasticity of substitution between goods as in Steinsson (2003), leading to a time-varying mark-up on marginal cost. We introduce this feature as a short cut to get a trade-off for optimal monetary policy, which otherwise will be trivial under the perfect information assumption.

Woodford (2003) shows that under rational expectations these assumptions lead to a Phillips curve of the form

(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 is inflation, x is the output gap, β is the discount rate, u is a cost-push shock (assumed i.i.d.) and κ is a function of the structural parameters including the degree of Calvo price stickiness. Furthermore, up to a second order approximation, the (negative of the) period social welfare function takes the form

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

where $\lambda=\kappa/\theta$ measures the relative weight on output gap stabilization and θ is the elasticity of substitution between the differentiated goods. In the benchmark case, we assume that the central bank uses the social welfare function to guide its policy decisions, both under rational expectations and under private-sector learning.[2] A γ different form one implies that the optimal rate of inflation is zero (otherwise there will be inefficient dispersion of prices in the steady state) and we therefore assume that the inflation target (coinciding with the average level of inflation in the absence of an over-ambitious output-gap target) equals this level. To keep the model simple, we abstract from any explicit representation of the transmission mechanism of monetary policy and simply assume that the central bank controls the output gap directly.

To solve for optimal policy under discretion, we define $z_t = \pi_t - \gamma\pi_{t-1}$. With this notation, we can rewrite (1) and (2) as:

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

(2') $$L_t = z_t^2 + \lambda x_t^2 .$$

[2] It is clear that it matters at which stage of the analysis learning is introduced. 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.

In this formulation, there are no endogenous state variables and since the shocks are iid, the rational expectations solution (which by the way coincides with the standard forward-looking model) must have the property $E_t z_{t+1} = 0$. Thus:

(1'') $\quad z_t = \kappa x_t + u_t$

Hence, the problem reduces to a static optimization problem and we can find optimal monetary policy substituting (1'') into (2') and minimizing the result with respect to the output gap.

Therefore the first order condition is:

$$\kappa(\kappa x_t + u_t) + \lambda x_t = 0$$

implying:

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

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. 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.

Using (3) to substitute for x_t in (1'') allows us to write:

(4) $\quad z_t = \dfrac{\lambda}{\kappa^2 + \lambda} u_t .$

As before the quasi-difference of inflation, z_t, responds only to the cost-push shock. For a given shock, it is allowed to vary as a direct function of the weight of the output gap in the loss function and as an inverse function of the slope of the Phillips curve.

Using again the definition of $z_t = \pi_t - \gamma\pi_{t-1}$

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

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

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

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. The second term in the above equation remains the same when $\gamma>0$. In this case, however, 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 also *should* be high. The reason is that it is price dispersion that drives the welfare criterion. To see the effect of this, suppose that there is no shock in the current period, but that lagged inflation is $\pi_{t-1}>0$ and that all prices were identical at t-1 (this is a probability zero event, but just for the sake of the argument). Then, in order to have no price dispersion, since the fraction of firms who do not reset prices optimally will increase their prices linked to lagged inflation, current prices must rise in proportion to this indexation in order to keep the parity.

Using equations (3) and (4) to substitute for x_t and π_t in the static loss function leads to:

$$L = \frac{\lambda}{\kappa^2 + \lambda} u_t^2 .$$

The loss is an increasing function of the variance of the shocks and of the weight of the output gap in the loss function and a decreasing function of the slope of the Phillips curve.

2.2. Inflation expectations according to adaptive learning.

As already noted above, we consider two assumptions regarding the formation of inflation expectations in equation (1): rational expectations and constant-gain, least-squares learning. For the case of rational expectations, we saw, in equation (3), that optimal monetary policy under discretion only responds to the exogenous shock, and not to lagged inflation. As we saw before, under the optimal discretionary policy, the output gap only responds to the current cost-push shock. In contrast, if the central bank is able to credibly commit to future policy actions, optimal policy will feature a persistent "history dependent" response, as discussed extensively in Woodford (2003). The relevant mechanism relies on the fact that, under optimal policy, perceptions of future policy actions help stabilize current inflation, through expectations. Specifically, by ensuring that, under rational expectations, a decline in inflation expectations is associated with positive cost-push shock, optimal policy manages to spread the impact of the shock over time.

As shown in equation (5), under rational expectations and discretionary monetary policy, the only endogenous state variable is lagged inflation and hence the equilibrium dynamics of inflation will follow a first-order autoregressive process:

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

Moreover, the degree of reduced-form inflation persistence is given by the degree of inflation indexation in (1), i.e. $\rho = \gamma$ and $\tilde{u}_t = \lambda/(\kappa^2+\lambda)u_t$.

We now turn to the assumption of adaptive learning. Specifically, we assume that the private sector believes the inflation process is well approximated by equation (5). 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,[3]

(7) $\pi_t = c_t\pi_{t-1} + \varepsilon_t$.

[3] 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.

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.

$$c_t = c_{t-1} + \phi R_t^{-1}{}_t\pi_{t-1}(\pi_t - \pi_{t-1}c_{t-1}) \tag{8}$$

$$R_t = R_{t-1} + \phi(\pi_{t-1}^2 - R_{t-1}), \tag{9}$$

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), where 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 (discussed further below), but not in updating the parameters. This implies that inflation expectations, in period t, for period t+1 may be written simply as:

$$\mathrm{E}_t\pi_{t+1} = c_{t-1}\pi_t \tag{10}$$

Generally, there is a simultaneity problem in forward-looking models combined with learning. In (1), current inflation is determined, in part, by future expected inflation. However, according to (10), 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 (10) 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, rather 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.[4]

Substituting equation (10) into the New-Keynesian Phillips curve (1) we obtain:

$$\text{(11)} \qquad \pi_t = \frac{1}{1+\beta(\gamma - c_{t-1})}\left(\gamma\pi_{t-1} + \kappa x_t + u_t\right).$$

2.3. Solution method for optimal monetary policy

In the previous sub-sections, we distinguished between two assumptions concerning expectations' formation: rational expectations and adaptive learning. Under adaptive learning we want to distinguish between the case where the central bank follows a simple rule (specifically the rule given in equation (3)) and fully optimal policy. In the first case, the simple rule (3), the Phillips curve (1) and equations (8), (9) and (10), which describe the evolution of private sector expectations, 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. 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 in another paper (see Gaspar, Smets and Vestin (2005a)) labeled such extreme case "sophisticated" central banking "Sophisticated" central banking implies solving the full dynamic optimization problem, where the parameters associated with the estimation process are also state variables.

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

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

(12) $$V(u_{t,}\pi_{t-1,}c_{t-1,}R_t) = \min_{x_t} \frac{(\pi_{t,} - \gamma\pi_{t-1})^2 + \lambda x_t^2}{2} + \beta \mathrm{E}_t V(u_{t+1,}\pi_{t,}R_{t+1,}c_t),$$

subject to the expectations adjusted Phillips curve (1), the recursive parameter updating equations (8) and (9), and, finally the expectations equation (10). Or, alternatively, we may use (11) instead of (1) and (10).

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

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

As shown in the appendix (to be completed), equation (13) can be written as:

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

where $\delta_t = 1 - 2\beta\phi E_t V_R$, $\chi_\mathrm{t} = 1 + \beta(\gamma - \mathrm{c}_{\mathrm{t}-1})$ and V_c, V_π and V_R denote the derivatives of the value function with respect to the variables indicated in the subscript. Equation (14) will be important to interpret the results we will discuss in section 3. Equation (14) does not characterize the optimal solution fully as we have not specified the exact forms of the partial derivatives.

We note that 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 problem is that the value function will not be linear-quadratic in the states and hence employ the collocation-methods described in Judd (1998) and Miranda and Fackler (2002), which amount to approximating the value function with a combination of cubic splines, which translates the problem to a root finding exercise (some details are outlined in the appendix).

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[5] these structural parameters imply that κ=0.019; λ= 0.002. 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). θ=10 corresponds to a mark-up of about 10%. 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 0.03. 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 0.03 corresponds to an average sample length of about 17 quarters. [6] 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.

---

[5] Here we follow the discussion in Woodford (2003). See especially pages 187 and 214-15. For the relevant parameters we rely on ...

[6] See Orphanides and Williams (2004). Similarly, Milani (2005) estimates the gain parameter to be 0.3 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 rule given by equation (3) 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. Finally, we present some sensitivity analysis for different parameters of the economy and a different weight on the output gap in the central bank's loss function.

3.1. Optimal monetary policy, persistence and macroeconomic performance

Table 2 compares, for our benchmark calibration, four cases: optimal policy under commitment and rational expectations (first column); optimal policy under adaptive learning (third column) and the simple rule (equation (3)), both under rational expectations (second column) and adaptive learning (last column). It is instructive to start walking a well-trodden path comparing the outcomes under commitment and discretion, under rational expectations. For such a case it has been shown (see, for example, 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. The reverse is true for inflation. Inflation persistence is considerably higher under the simple rule. Continuing the comparison between optimal policy under commitment and the simple rule, under rational expectations, we see that inflation variance is about 85 % higher under the simple rule, 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 the simple rule, which corresponds to the optimal discretionary monetary policy. Overall, the loss is about is about 30 % higher under the simple rule.

[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 simple monetary policy rule, looking at the second and fourth columns in Table 2. The comparison confirms their findings. Clearly, the autocorrelation of the output gap remains unchanged at zero (given that cost-push shocks are assumed i.i.d.), while, under adaptive learning, the autocorrelation of inflation increases from 0.5 to 0.6. The variance of inflation increases to almost 2.5 times the value under commitment and the variance of the quasi-difference of inflation is more than 1.5 times larger. The volatility of the output gap is the same as under discretion and rational expectations. Thus, the expected welfare loss increases significantly. The intuition is that, under adaptive learning, economic agents perceive inflation as more persistent. Thus, 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.

Optimal central banking under adaptive learning is able to improve outcomes significantly relative to the simple rule (as it is clear from comparing the third and the fourth column in Table 2). Responding persistently to cost-push shocks optimal policy reduces sharply the degree of perceived inflation persistence, to about 0.3, and the persistence of inflation. As before, this is strongly linked with a decline in inflation volatility. Correspondingly, the output gap becomes more persistent and slightly more volatile, than under the simple rule. On balance, the expected welfare loss falls significantly, by about 38 %, under optimal policy against the simple rule (which compares with a gain of about 30 % for commitment over the simple rule, for rational expectations). Interestingly, optimal policy under adaptive learning brings us close to the results under commitment, as we can see from a comparison between the first and the third column in Table 2. Indeed, the output gap exhibits significant persistence and inflation is much less persistent than under the simple rule. Moreover, inflation volatility is sharply reduced at little cost in terms of output gap volatility. Notwithstanding this, optimal monetary policy under adaptive learning is unable, for our benchmark calibration, to reap all the benefits from a commitment regime, under rational expectations.

Finally, in the last line of Table 2, we show the variance of the forecast error relative to the case of commitment. As stated above the average forecast error is zero in all cases. The variance of the forecast errors, under adaptive learning, is inside the range defined under rational expectations, by commitment and the simple rule, for the case of full optimal policy and about 10 % above, for the case of the simple rule. Such results confirm the intuition that the approximation is "reasonable".

Figure 1 provides some additional detail concerning the distribution of the endogenous variables, estimated persistence, inflation, quasi-difference of inflation, output gap and the moment matrix, under optimal policy and the simple policy rule. First, the most interesting comparison relates to the estimated persistence of inflation, c. Panel (a) shows that, not only is the average of the estimated persistence parameter significantly lower under optimal policy, but also that the distribution is more concentrated around the mean. In particular, under optimal policy, the perceived inflation parameter does not go close to one, contrary to what happens under the simple rule. In fact, the combination of the simple policy rule and private sector's perpetual learning at times gives rise to explosive dynamics, when perceived inflation persistence goes to (or above) unity[7]. 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 simple rule. In Gaspar, Smets and Vestin (2005a) we looked at the transition from an economy, regulated by a simple rule, taking off on an explosive path to the anchoring of inflation, through optimal policy. Such explosive dynamics does not occur under the optimal policy rule.

Second, confirming the findings emphasized above, it is clear, from panels (b), (c) and (d) that, under the optimal policy, the distribution of inflation and of the quasi-difference of inflation become more concentrated. Conversely, the distribution of the output gap becomes somewhat wider, although, in panel (d) the two distributions appear almost undistinguishable. Finally, the distribution of the R matrix also shifts to the left and becomes more concentrated reflecting the fact that the variance of inflation falls considerably under the optimal policy.

[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.

3.2. Optimal monetary policy: 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: $(u_{t,}\pi_{t-1,}c_{t-1,}R_t)$. We want to recall equation (14),

$$(14) \qquad x_t = -\frac{\kappa\delta_t}{\kappa^2\delta_t+\lambda\chi_t^2}u_t + \frac{\kappa\gamma(\chi_t-\delta_t)+\beta\kappa\chi_t\phi R_t^{-1}E_tV_c}{\kappa^2\delta_t+\lambda\chi_t^2}\pi_{t-1} + \beta\frac{\kappa\chi_t}{\kappa^2\delta_t+\lambda\chi_t^2}E_tV_\pi$$

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

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 derivatives of the value function with respect to the variables indicated in the subscript.. Equation (14) expresses the policy instrument – here taken to be the output gap x_t – as a function of its determinants. 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 following discussion we assume that $E_t V_R$ is zero, so that the expected marginal benefit from manipulating the moment matrix is zero. In that case, $\delta_t = 1$.[8]

First, if lagged inflation is equal to zero, $\pi_{t-1}=0$, the optimal monetary policy reaction can be simplified to the following expression:

$$(15) \quad x_t = -\frac{\kappa}{\kappa^2 + \lambda\chi_t^2} u_t .$$

Clearly, in such a case the second term in equation (14) is zero. Moreover, it can be shown that for $\pi_{t-1} = 0$, $E_t V_\pi$ is zero. When lagged inflation is zero, the effects of positive and negative cost-push shocks are exactly symmetric. The intuition is that, starting from zero inflation, deflationary shocks and inflationary shocks are exactly as bad. Since the distribution of the shocks is uncorrelated over time and symmetric, it follows that $E_t V_\pi$ is zero.

Thus, when lagged inflation is zero, $\pi_{t-1}=0$, optimal monetary policy amounts to a simple response to the cost-push shock. Moreover, if $\gamma = c_{t-1}$ and as a result $\chi_t^2 = \chi_t = 1$, equation (15) reduces to the simple rule derived under rational expectations given by equation (3). In other words, when lagged inflation is zero and the estimated inflation persistence is equal to the intrinsic persistence, the optimal monetary policy response to a shock under adaptive learning coincides with the optimal response under discretion and rational expectations. The reason for this finding is quite simple. From equation (8), it is clear that, when lagged inflation is zero, the estimated persistence parameter is not going to change. As a result, there is no benefit from trying to affect the perceived persistence

---

[8] In practice δ_t always turns out to be close to one. [to be checked more fully].

parameter. Analogously, it is also possible to show that when ϕ=0 the solution under fully optimal policy coincides with (3), meaning that the simple rule would lead to full optimal policy.

Consider now the case when estimated persistence is lower than intrinsic persistence, $\gamma > c_{t-1}$. This is the case which will on average prevail under optimal policy (see Figure 1 a). For $\gamma > c_{t-1}$, $\chi_t > 1$ and thus $\frac{\kappa}{\kappa^2 + \lambda\chi_t^2} < \frac{\kappa}{\kappa^2 + \lambda}$, showing that the response is more muted then under discretion and rational expectations. The reason is again simple. As shown in equation (11), the smaller the degree of perceived inflation persistence, the smaller the impact of a given cost-push shock on inflation. As a result, it is optimal for the central bank to mute its response to the cost-push shock. This clearly illustrates the first-order benefits of anchoring inflation expectations. Conversely, for $\gamma < c_{t-1}$ and $\chi_t < 1$ and thus $\frac{\kappa}{\kappa^2 + \lambda\chi_t^2} > \frac{\kappa}{\kappa^2 + \lambda}$, the response of optimal policy to cost push shocks becomes stronger than under the simple rule. Again the intuition is clear. When estimated persistence is larger than intrinsic persistence, inflation is threatening to feed on itself, through its impact on expectations. In such circumstances the central bank is well advised to respond stronger to inflationary or deflationary shocks.

We illustrate these results in Figure 2, showing the mean dynamics response of the output gap, inflation and estimated persistence to a one-standard deviation cost push shock, taking lagged inflation to be initially zero. From panel a) it is clear that as estimated persistence increases so does the output gap response. Clearly, the output gap response to a one-standard deviation cost-push shock is more negative when estimated persistence is higher. That helps to mitigate the inflation response, although it is still the case (from panel b) that inflation increases by more when estimated inflation 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]

Let us now return to equation (14) and discuss the second term, which captures part of the optimal response to lagged inflation. Note that the first term in the denominator is zero when $\gamma = c_{t-1}$. 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, in general the second term in the denominator 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 creating unexpected inflation. It happens to be the case that in our simulations the marginal expected value of lowering the degree of inflation persistence 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 both inflation and its semi-difference. 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 (8), 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 denominator cancels out. The overall intuition is simply that a lower estimated persistence parameter makes the overall dynamics of the economy more stable. It is immediately obvious that this effect must be dominant when the estimated persistence parameter has become so large as to threaten explosive dynamics (see discussion in sub-section 3.1.). In our numerical simulations the response to inflation is always negative.

Finally, the third term is also interesting. We have already seen that when $\pi_{t-1}=0$, $E_t(V_\pi)=0$ and this term plays no role. Now, if $\pi_{t-1}>0$, 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 $\pi_{t-1}<0$, 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.

Figures 3a and 3b summarize some of the important features of the shape of the policy function (14) 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. First, 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?

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

The answer, as illustrated, in Figure 4 is yes. The response of the output gap is symmetric. Moreover, from the panel (b) of Figure 4 it is clear that the adjustment of inflation is also symmetric. Finally, panel (c) 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, a different degree of price stickiness and a different degree of intrinsic inflation persistence. Second, we look at the impact of reducing 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, three different degrees of price stickiness and two different degrees of intrinsic inflation persistence. 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 policy 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 favour of lower inflation persistence. The figure also shows that the drop in persistence is larger, the larger the degree of price stickiness. This is interesting because from equation (6) it is clear that a higher κ will not affect the intra-temporal trade-off between inflation and output gap stabilization very much as it affects both the weight on output gap stabilization in the loss function and the slope of the Phillips curve and those effects tend to cancel out. However, a higher degree of price stickiness does increase the cost of cost-push shocks and therefore increases the benefit from reducing the estimated degree of inflation persistence.

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 characterize the conduct of optimal monetary policy in a simple new Keynesian model with adaptive learning. The literature under rational expectations (see, for example, Woodford, 2003) has found that, in case the central bank can commit future policy, it is optimal to follow inertial policy. The reason is that, under rational expectations, the central banker is able, through a gradual but persistent response to cost-push shocks, to reduce the initial impact of the shocks and to spread their impact over time. The stabilizing mechanism operates through the effects that future policy actions have on current inflation expectations and, therefore, on current inflation. It allows for sizeable welfare gains relative to optimal policy under discretion.

In this paper, we show that fully optimal policy, under adaptive learning, lowers perceived inflation persistence relative to its value under a simple rule under either rational expectations or adaptive learning. The result can reasonably be interpreted as meaning that anchoring inflation and inflation expectations is an important feature of optimal policy. Lower perceived persistence contributes to overall stability and delivers significant improvement in economic performance. The stabilizing mechanism is different from the one described above for rational expectations. Future policy actions have no direct effect on expectations. We show that, under adaptive learning, monetary policy responds stronger (and more persistently) to cost-push shocks as the perceived persistence parameter increases.

It is clear that both optimal policy with commitment, under rational expectations, and a sophisticated central banker, under adaptive learning, use the endogeneity of expectations to improve economic outcomes. The mechanisms, however, are very different. In the case of rational expectations it is future policies that matters while, under adaptive learning, the effect comes from past policy actions and outcomes. In the real world, both mechanisms are likely to play out in tandem.

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	0.03	0.019	0.005

Table 2: Summary of macro-economic outcomes

	Rational Expectations		Adaptive Learning	
	Commitment	Discretion	Sophisticated	Simple Rule
Corr(x_t, x_{t-1})	0.65	0	0.54	0
Corr(π_t, π_{t-1})	0.24	0.50	0.34	0.61
Var(x_t)	1	0.98	1.02	0.98
Var($\pi_t - \gamma\pi_{\tau-1}$)	1	1.17	1.10	1.55
E[L_t]	1	1.29	1.08	1.43
Var(error)	1	2.03	1.57	2.23

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, output gap, inflation, and the moment matrix

(a)

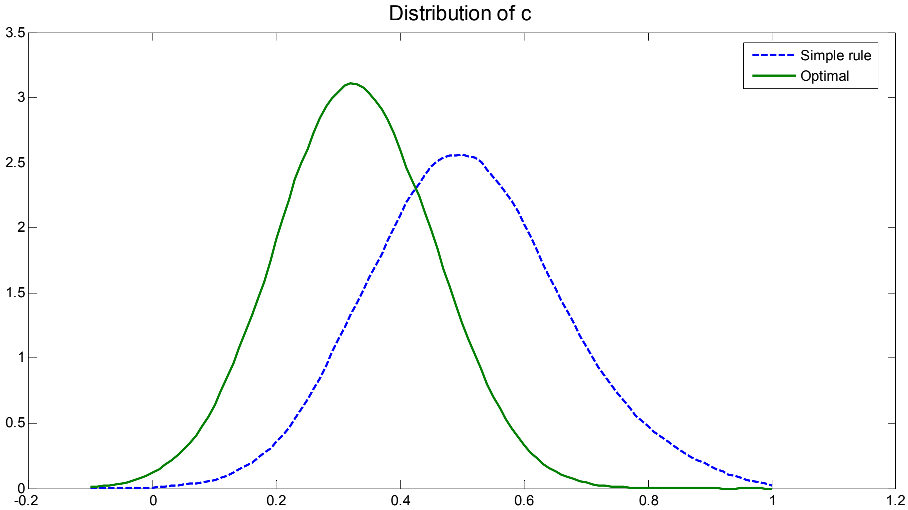

(b)

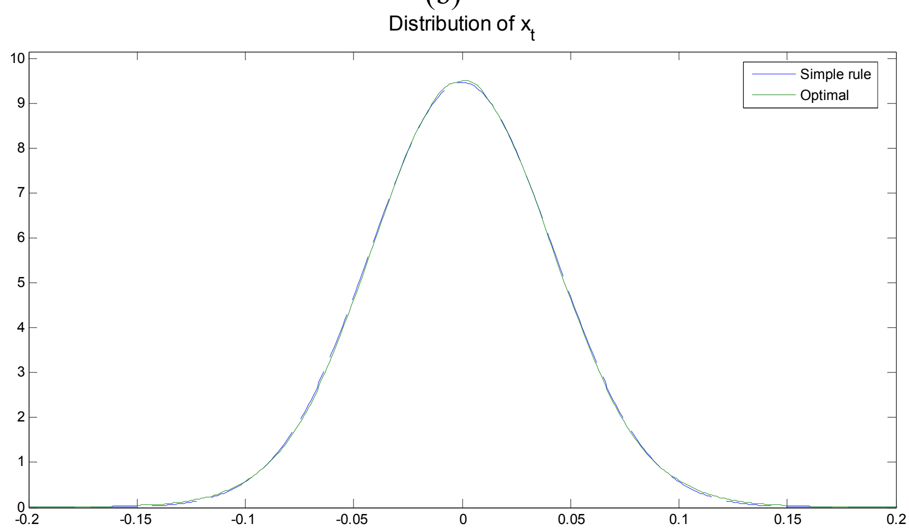

(c)

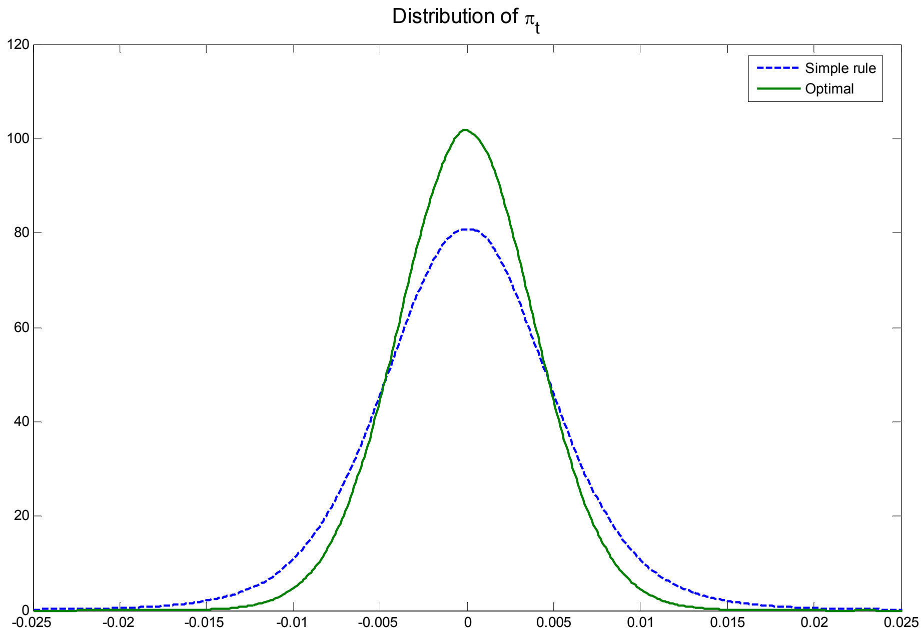

Distribution of π_t
Simple rule
Optimal
120
100
80
60
40
20
0
-0.025
-0.02
-0.015
-0.01
-0.005
0
0.005
0.01
0.015
0.02
0.025

(d)

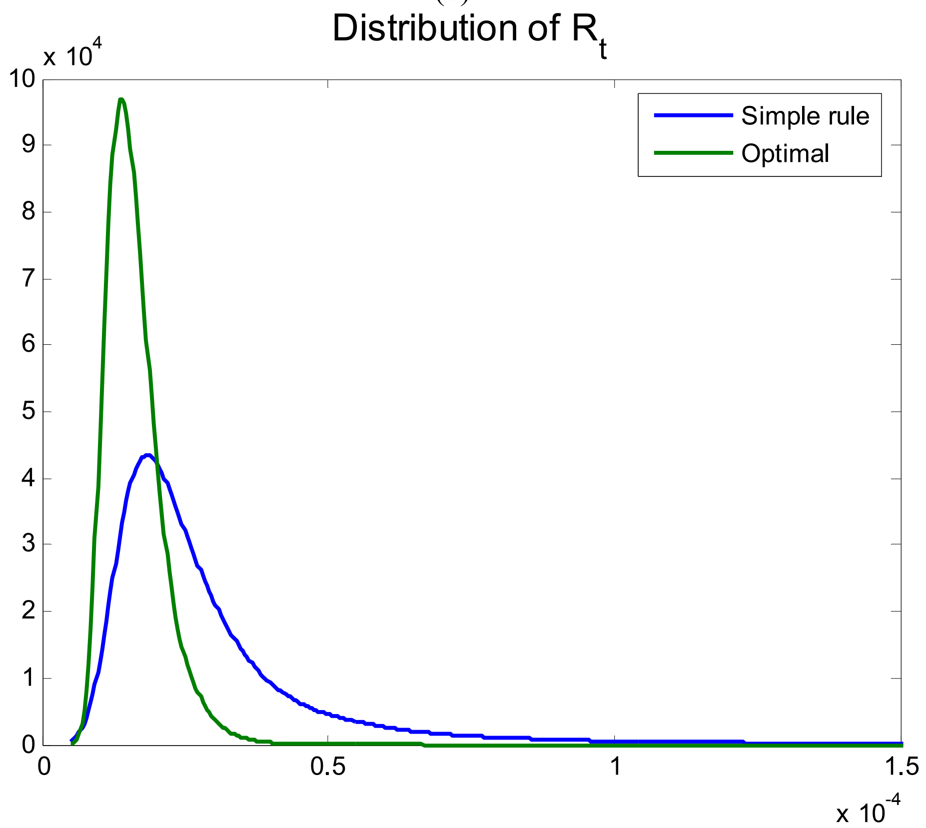

Distribution of R_t
x 10^4
Simple rule
Optimal
10
9
8
7
6
5
4
3
2
1
0
0
0.5
1
1.5
x 10^{-4}

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

Output gap

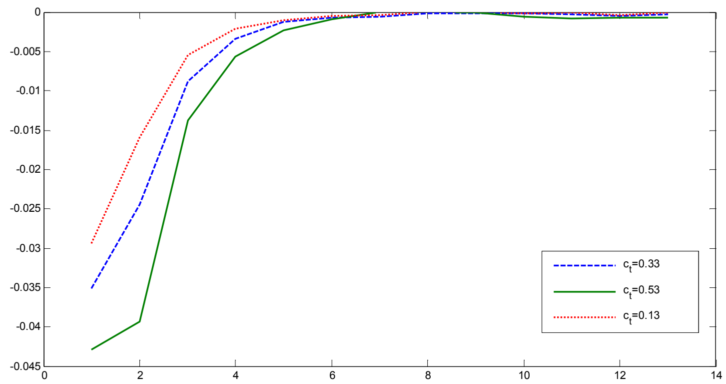

Inflation

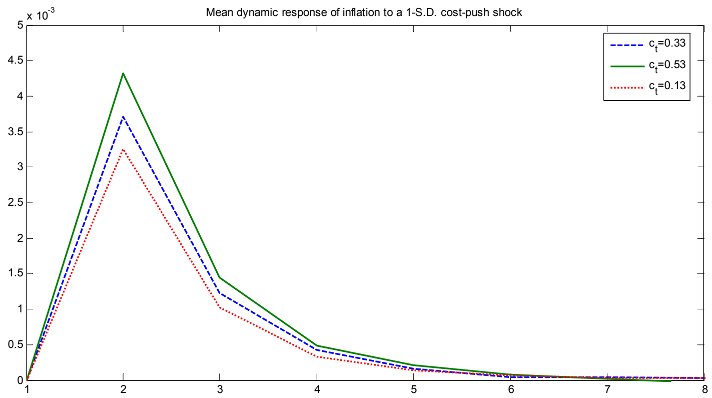

Estimated inflation persistence

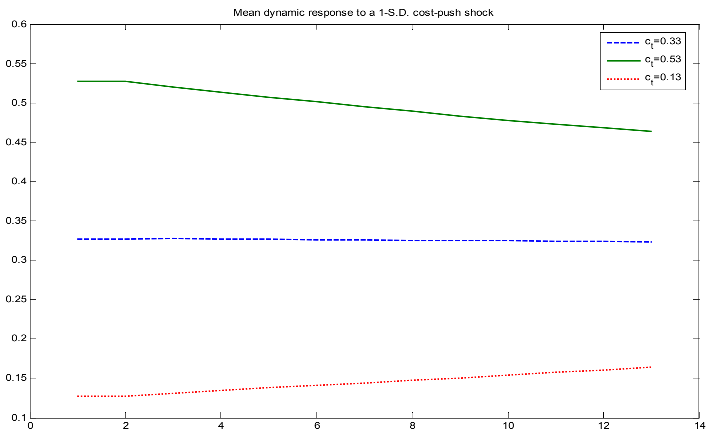

Figure 3: The policy function output gap as a function of lagged inflation and the estimated degree of inflation persistence.

(a)

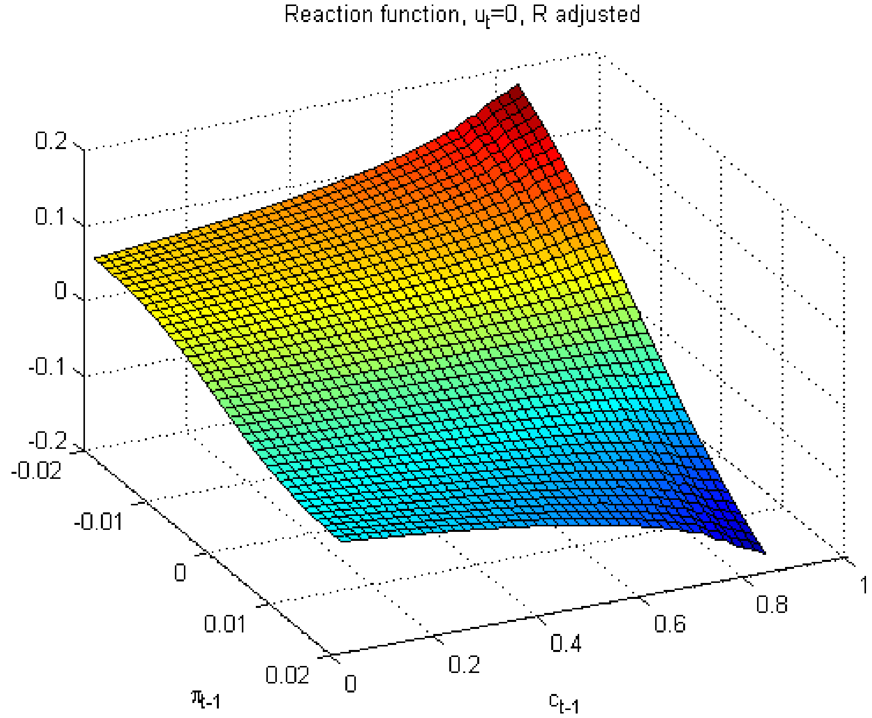

(b)

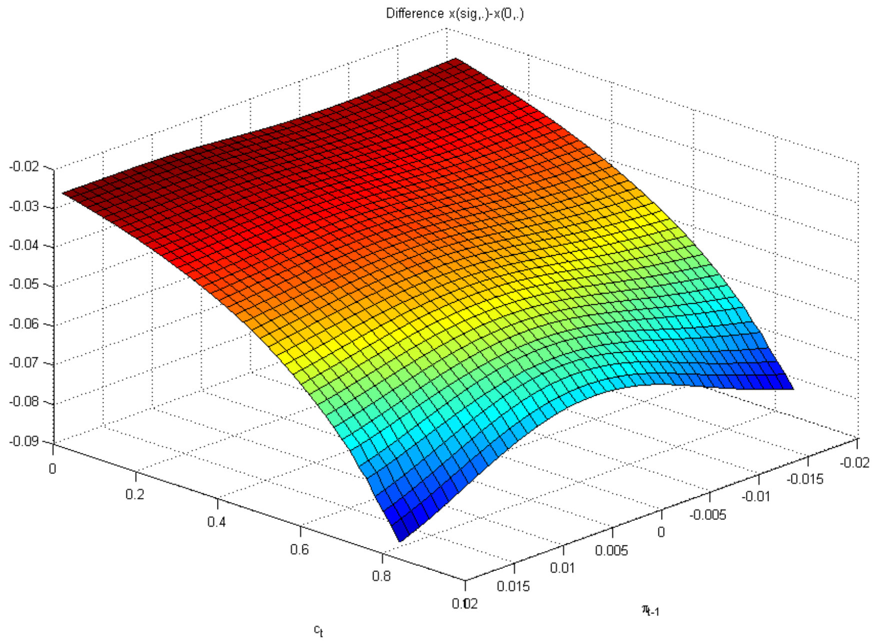

Figure 4: Illustration of symmetry in the response of policy (output gap) to inflation and deflation

(a)

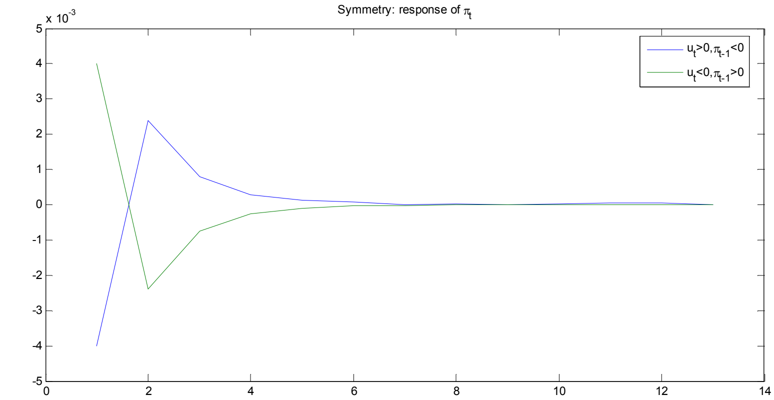

(b)

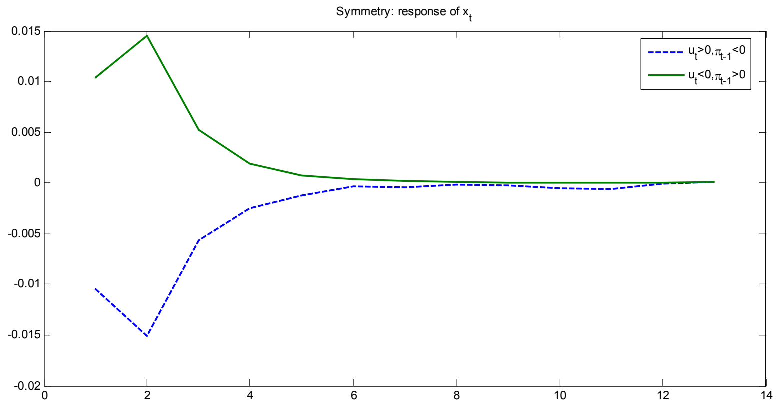

(c)

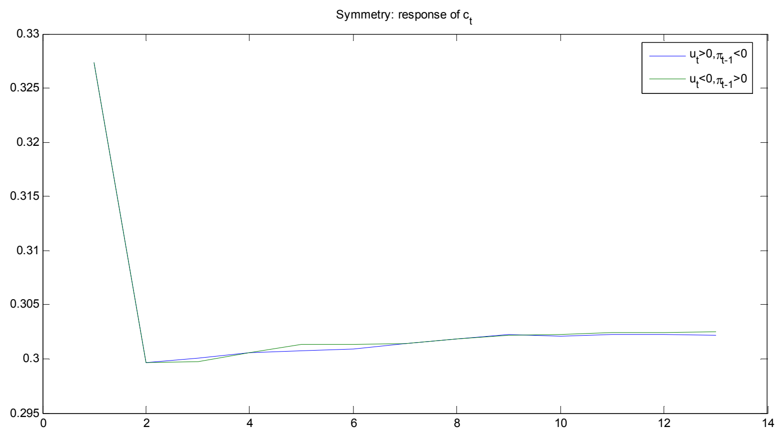

Figure 5: Sensitivity analysis: Average estimated persistence in function of the gain and the degree of price stickiness.

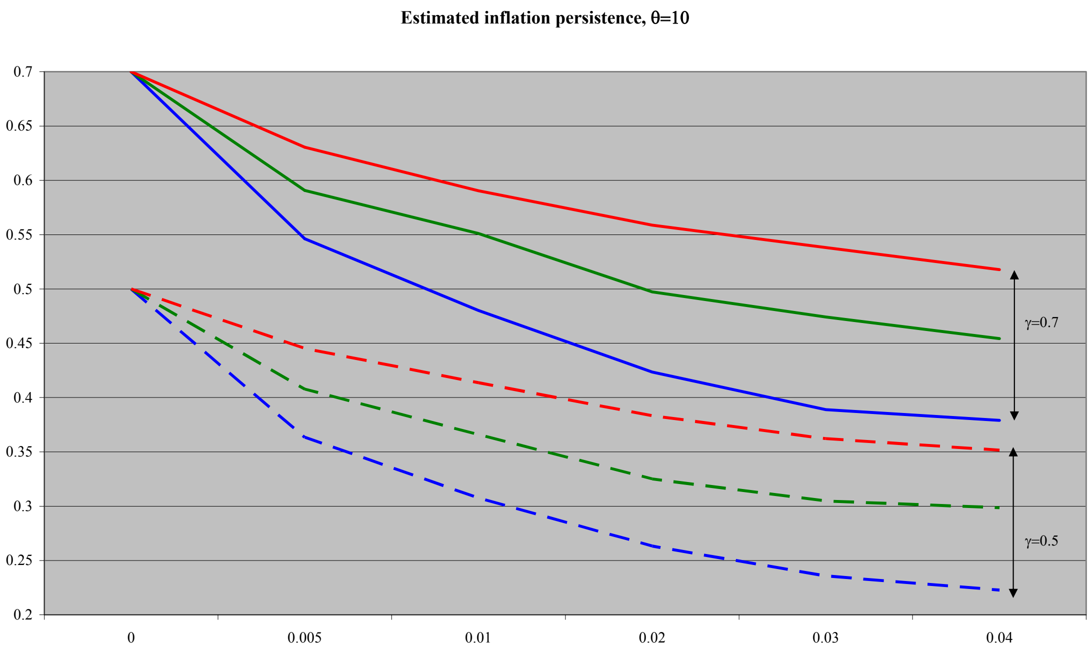

Figure 6: Distribution of estimated inflation persistence as a function of the weight on output gap stabilization.

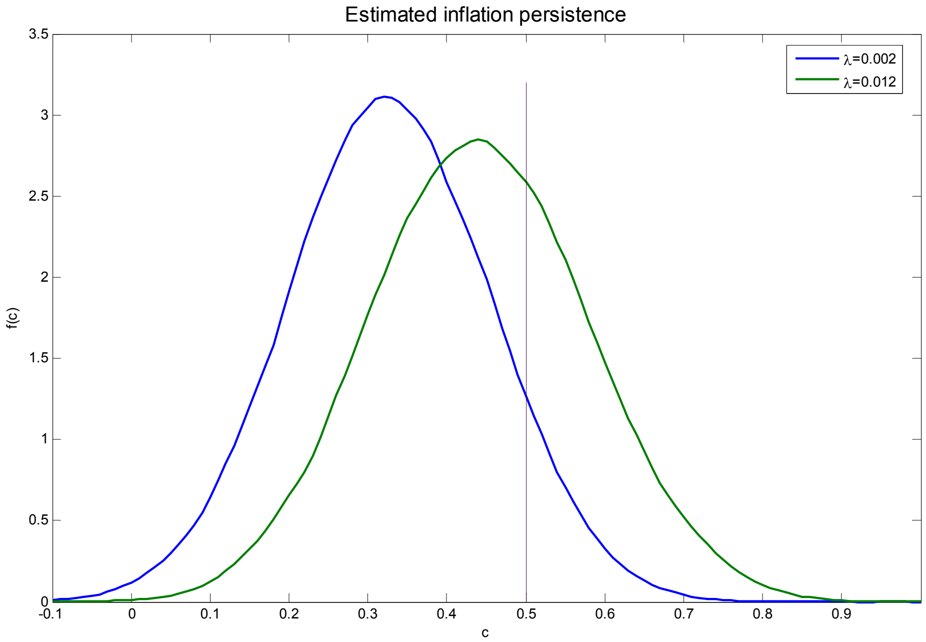

REFERENCES

Bullard, James and Kaushik Mitra (2002) Learning About Monetary Policy Rules. *Journal of Monetary Economics,* 49, 1105-1129.

Calvo, Guillermo (1983) Staggered Prices in a Utility-Maximizing Framework. *Journal of Monetary Economics,* 12, 383-98.

Christiano, Lawrence J., Martin Eichenbaum and Charles Evans (2005) Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy. *Journal of Political Economy, forthcoming.*

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.

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 (2005a), Monetary Policy over Time, *Macroeconomic Dynamics,* forthcoming.

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

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

Kydland, Finn E. and Edward C. Prescott (1977) Rules Rather than Discretion: the Inconsistency of Optimal Plans. *Journal of Political Economy* 85, 3.

Lucas, Robert (1976) Econometric Policy Evaluation: a Critique. In Karl Brunner and Alan Meltzer (eds.), *The Phillips Curve and Labor Markets,* vol.1 of Carnegie-Rochester Conference Series on Public Policy, 19-46, Amsterdam: North-Holland Publishing Company.

McCallum, Bennet T. (2005) A Monetary Policy Rule for Automatic Prevention of a Liquidity Trap. *NBER Working Paper* 11056.

Marcet, A. and T. Sargent (1989) Convergence of Least-Squares Learning Mechanisms in Self-Referential Linear Stochastic Models. *Journal of Economic Theory* 48, 337-368.

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 (2001) Monetary Policy Rules, Macroeconomic Stability and Inflation: A View from the Trenches. *ECB Working Paper Series* 115.

Orphanides, Athanasios (2002) Monetary Policy Rules and the Great Inflation. *American Economic Review,* 92, 115-120.

Orphanides, Athanasios (2003a) Monetary Policy Evaluation with Noisy Information. *Journal of Monetary Economics,* 50, 605-631.

Orphanides, Athanasios (2003b) The Quest for Prosperity without Inflation. *Journal of Monetary Economics,* 50, 633-663.

Orphanides, Athanasios and John C. Williams (2002a) Imperfect knowledge, inflation expectations, and monetary policy. *FEDS Paper* 2002-27.

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

Orphanides, Athanasios and John C. Williams (2003b) Inflation scares and forecast-based monetary policy. *FEDS Paper* 2003-41.

Orphanides, Athanasios and John C. Williams (2003) The Decline of Activist Stabilization Policy: Natural Rate Misperceptions, Learning and Expectations. Presented at the International Research Forum on Monetary Policy, 14&15 November 2003.

Orphanides, Athanasios and John C. Williams (2002) Robust Monetary Policy Rules with Unknown 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.

Orphanides, Athanasios and Simon Van Norden (2002) The Unreliability of Output Gap Estimates in Real Time. *Review of Economics and Statistics,* 84, 569-583.

Phillips, A.W. (1958) The Relation between Unemployment and the Rate of Change in Money Wage Rates in the United Kingdom 1861-1957. *Economica*, 25, 283-94.

Sargent, Thomas (1993) *Bounded Rationality in Macroeconomics.* Oxford: Oxford University Press.

Sargent, Thomas (1999) *The Conquest of American Inflation.* Princeton: Princeton University Press.

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.

Taylor, John, 1993, Discretion vs. Policy Rules in Practice, *Carnegie-Rochester Series on Public Policy*, 39, pp. 195-214.

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

Appendix (to be completed)