Optimal Monetary Policy under Adaptive Learning

Vitor Gaspar, Frank Smets and David Vestin[1]

March 2005

Preliminary and incomplete

---

[1] The views expressed are the authors' own and do not necessarily reflect those of the ECB, Banco de Portugal or the Eurosystem.

1. Introduction

Inflation dynamics crucially depend on how inflation expectations are formed. Modern macroeconomics emphasizes that an engineering approach to policy-making analysis fails, due to its failure to account for human behaviour. In most modern macroeconomics expectations are modelled in accordance with rational expectations (Muth, 1961). Researchers have systematically explored the implications of rational expectations for the conduct of policy. Rational expectations (paraphrasing Evans and Honkapohja (2001)) assume economic agents who are extremely knowledgeable. An alternative approach is to limit their knowledge, so that as time goes by and available data changes, so does the agents' forecasting rule. Thus, the alternative can be understood as implying bounded rationality. In fact it means limiting agents' knowledge about the true structure of the model. Adaptive learning is particularly attractive to reckon with the implications from pervasive structural change and permanent transformation that characterizes the economic environment of modern economies.

The rational expectations revolution had deep implications for our understanding of policy-making. Kydland and Prescott (1977) pointed to strong tensions between optimal and time-consistent policies. Specifically, they have shown that the time-consistent policy-maker will not take the effects of *future* policies on private sector expectations into account. The distinction between optimal and time-consistent policies led to the corresponding distinction between policies under commitment and under discretion. This paper looks at the implications from adaptive learning, on the part of the private sector, for the conduct of optimal monetary policy. Inspired by Svensson (2003), we compare outcomes from "simple rules" with optimal policy that takes expectations formation explicitly into account. Often, learning has been used, in the literature, as an equilibrium selection mechanism. The idea is simply that if a particular rational expectations equilibrium is not learnable then it can be excluded. In most papers, policy is modelled by assuming a simple instrument rule. Svensson (2003) argues that simple instrument rules fail to capture how central banks actually conduct policy. If adaptive learning is (empirically) a good description of expectations formation, then it is important to characterise optimal monetary policy in such a setting.

Modelling the optimal behaviour 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 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 there similar mechanisms at play when we depart from rational expectations?

In this paper, we find that endogenous expectations, under adaptive learning, make anchoring of inflation expectations crucial. We explain the conclusion through a number of systematic comparisons. First, we compare outcomes of "simple rules" under adaptive learning and rational expectations. The "simple rule" we use is simply the optimal discretionary policy under rational expectations. We show that inflation is more stable under rational expectations. Moreover, under adaptive learning, inflation dynamics eventually become explosive under the simple rule.

Second, we compare outcomes, under optimal policy, for the alternatives of rational expectations and adaptive learning. In this case explosive behaviour never occurs. The different expectation formation mechanism implies interesting similarities (and differences) for optimal policy. Anchoring inflation expectations is always key. However, the mechanism through which anchoring occurs are different, as future policies play no role under adaptive learning. Nevertheless, in specific circumstances there are similarities between the monetary policy reaction functions.

Third, we compare the differences relative to optimal policies under the two alternative expectation formation mechanisms. The comparison will allow us to argue that the difference between naïve and optimal central banking (under adaptive learning) does have some similarities with the difference between discretion and commitment (under rational expectations).

The paper is organised as follows. Section 2 will introduce the model and compare outcomes under rational expectations and adaptive learning using a common rule. The rule is the optimal rule under rational expectations (assuming time consistency). Section 3 will make the same comparison for fully optimal policies conditional on the alternative expectation formation mechanisms. Section 4 will conclude.

2. Benchmark: Rational expectations versus adaptive learning

In this section, we use a very simple New Keynesian model of inflation dynamics and analyze the behavior of inflation under a simple monetary policy reaction function under two different assumptions regarding the way the private sector forms its inflation expectations. The New Keynesian model is standard (e.g. Woodford, 2003). It is consistent with the following microeconomic assumptions. Producers set prices in an environment of nominal rigidity, formalized using the Calvo (1983) shortcut. That is, in a given period, a specific firm will be “allowed” to reset its price optimally with an exogenous and constant probability. Firms produce using a technology that exhibits decreasing returns to labour. Furthermore, preferences of Dixit and Stiglitz (1977) type introduce monopolistic competition, with a continuum of otherwise identical firms. Finally, output is assumed to be demand determined, which means that firms will sell whatever quantity demanded at its current price. These four assumptions create a role for monetary policy, as without intervention markets may produce inefficiently. A simple way to see that this is the case 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 almost always present in new Keynesian models.

In order to examine the issues we are interested in, we introduce two further components. 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) and Smets and Wouters (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. Second, we assume that there is 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) and Smets and Wouters (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.

In terms of equations, the discussion above implies a Phillips curve of the form

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

where π is inflation, x is the output gap, γ, β and κ are parameters, u is a cost-push shock (assumed i.i.d.) and $\omega = (1+\beta\gamma)^{-1}$. The period social welfare function is assumed to be of the form

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

where λ is a parameter indicating the relative weight on output gap stabilisation. We will assume here that the central bank uses the social welfare function to guide its policy decisions. Note also that the optimal inflation target is assumed to be zero.

Next, we consider two assumptions regarding the formation of inflation expectations in equation (1). The standard assumption is to assume rational (or model-consistent) expectations. In this case, the private sector knows the structure of the economy as shown in (1) and the monetary policy reaction function implied by the central bank's loss function (2). In this case, it turns out that optimal monetary policy under discretion only responds to the exogenous shock, and not to lagged inflation (in contrast to when the loss function consists of squared inflation and output, see Clarida, Gali and Gertler (1999)). Hence, optimal discretionary policy is described by

(3) $$x_t = -\alpha u_t .$$

We assume, for simplicity, that the output gap is the central bank's policy instrument. Without loss of generality, we could alternatively have used a short-term interest rate as the policy maker's instrument by introducing an IS curve linking the output gap to the real interest rate. 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 and the weight on output gap stabilisation in the loss function. In contrast, if the central bank has access to a commitment technology, optimal policy will be more complicated and will utilize the fact that credible promises of future policy actions can help stabilize current inflation through expectations.

For future reference, Figure 1 shows impulse responses for a one standard deviation cost-push shock (calibration is discussed below) in the case of rational expectations under both discretion and commitment. The two key differences are that in the case of commitment the initial response is lower but then continues to persist even when the transitory shock is out of the economy. The reason is that by creating expectations of deflation, the incentive for current price-setters to increase prices is contained, leading to lower current inflation (see Woodford, 2003, for further discussion of this intuition). At first, it may seem like a small gain that the initial response of both inflation and output is slightly smaller under commitment, because the gain should be balanced against the cost of having to continue moving inflation and output also when the shock has left the economy. However, the welfare differences (at least in terms of relative levels) are large. Table 1 one shows that discretion leads to a 30% higher loss than commitment. The reason is that the loss function is quadratic in inflation and output and hence these relatively small differences lead to important consequences. This will be important when interpreting visually small differences in the tail-behaviour of the distribution of inflation explored below.

It is easy to show that under rational expectations and discretionary monetary policy as in (3), the equilibrium dynamics of inflation will be given by a first-order autoregressive process:

(4) $$\pi_t = \rho\pi_{t-1} + \widetilde{u}_t$$

Moreover, the degree of reduced-form inflation persistence will be given by the degree of inflation indexation in (1), i.e. $\rho = \gamma$.

We now consider the alternative assumption regarding expectations formation and assume that private agents do not know the exact parameters of the economy. However, they do know the reduced-form structure of the economy, i.e. they do know the form of equation (4) and try to estimate this reduced-form recursively. In particular, expectations are formed on the basis of a "constant gain" least squares algorithm implying perpetual learning.

In particular, the agents estimate the following reduced-form equation for inflation,[2]

(5) $\pi_t = c_t \pi_{t-1} + u_t$.

Agents are rational in the bounded sense that they do not take into account the fact that the parameter c varies over time in an optimal way. That is, assuming that the agents know the structure of the equilibrium model, they could instead use a Kalman filter to optimally update their perception of the persistence parameter. One way to rationalize such behaviour of the private sector is in terms of robustness. Many different structural models are consistent with the same reduced form (with possibly different parameter values). Hence, basing the forecast of inflation on a reduced form model is one way to hedge against relying on an incorrect structural model, even though it is less efficient in case one happens to know the correct model.

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

(6) $c_t = c_{t-1} + g_t R_t^{-1} {}_t z_t (y_t - z_t' c_{t-1})$

(7) $R_t = R_{t-1} + g_t (z_t z_t' - R_{t-1}),$

where c is the vector of estimated parameters, R is the moment matrix, y is the dependent variable, z is the vector of explanatory variables and g is the gain. Following equation (5), in our case, inflation and lagged inflation are the dependent and explanatory variable respectively and c is a single parameter capturing the persistence of the AR(1) inflation process. Note that due to the learning dynamics the number of state variables is expanded

[2] 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.

to four: (u_t , π_{t-1}, R_{t-1}, c_{t-1})., where the last two variables are predetermined and known by the central bank at the time they set policy at time t.

Some further observations regarding the updating process are worth noting. First, we assume $g_t = \phi$, a constant gain. There are two reasons for considering the constant gain case. First. only in this case there will be permanent action coming from the learning algorithm. In the case of recursive least squares learning, i.e. when the gain is 1/t, which corresponds to agents running an OLS regression with an increasing sample-length, the estimated parameter will converge to a constant. Second, transition regimes could still be studied, but there would be the additional problem that the value function would not be time invariant. This problem could likely be addressed by treating time as an extra state variable, but we do not want to enter into these complications and therefore focus on the constant gain case.

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:

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

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. But according to (8), 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, if current inflation is used to update the parameter currently used. 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 t+1 inflation. We followed that route in Gaspar et. al. (2005), extending 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 to not allow agents to update the coefficients using current information (i.e. just substitute (8) 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, which is only feasible to solve for smaller problems.[3]

Substituting equation (8) into the New-Keynesian Phillips curve one obtains:

$$(9)\ \pi_t = \frac{1}{1+\beta(\gamma - c_{t-1})}(\gamma\pi_{t-1} + \kappa x_t + u_t).$$

We are now ready to study the dynamics of inflation under rational expectations and adaptive learning when the central bank applies the policy rule (3). In the simulations we use the following set of parameters as a benchmark:

β	γ	λ	ϕ	κ	σ
0.99	0.5	0.05	0.05	0.07	0.005

Gamma is chosen such that there is some inflation persistence in the benchmark calibration. The gain, which corresponds to an average sample length of about 10 years or 40 quarters, is high, but not unreasonably so, to highlight the effects of optimal policy. [4] 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.

In this section, we start by considering the case when the central banker does not take learning explicitly into account, but uses the reaction function from the discretionary rational expectations equilibrium (equation (3)). Figure 2-3 show the distribution of output, inflation and the semi-difference of inflation under both types of expectations formation. Figure 4 shows the distribution of the estimated inflation persistence

[3] It is possible to solve this problem in the current setting. However, we leave this for future research.
[4] See Orphanides and Williams (2004) for a similar calibration of the gains parameter. Milani (2004) estimates the gain parameter to be … using a Bayesian estimation methodology

parameter. A number of observations are worth making. First, somewhat trivially the distribution of the output gap is the same in both cases (Figure 2). This follows directly from the assumption that in both cases the policy maker follows the policy rule given in equation (3). As a result, the output gap is proportional to the cost-push shock in both cases. Second, and more interestingly, as the Kalman gain parameter increases, the distribution of inflation becomes wider and flatter. Concerning inflation, there is considerably more mass in the tails when the Kalman gain is 0.05 compared to the rational expectations case (Figure 3, top panel). This is also the case for the semi-difference of inflation (Figure 3, lower panel). As discussed above, these visually small differences translate into larger welfare differences due to the quadratic loss function. As Table 1 shows, in the benchmark calibration, that learning with an identical policy rule from the rational expectations discretionary equilibrium is associated with a 12% higher loss than the corresponding rational expectations case. For comparison, if the gain instead is substantially smaller, 0.01, the difference shrinks to 3%.[5]

The reason for the higher inflation variance is obvious. Under constant-gain least squares learning, the estimated degree of persistence will vary over time and in a persistent manner. Compared to the rational expectations case, where this parameter is constant, this contributes to a higher variance of inflation. The variability of the estimated parameter is shown in Figure 4. It is also worth noting that there is a slight increase in the mean of the estimated degree of inflation persistence (above 0.5, which is the rational expectations case) and that there is a considerable mass of estimated inflation persistence close to one. Moreover, when inflation and inflation persistence are high, the dynamics of inflation may become explosive. In order to avoid that the simulation breaks down in the event of such explosive behaviour, we follow Orphanides and Williams (2004) and implement a cut-off point at a value of one for the estimated persistence (in other words, the private sector estimates a unit root in inflation). When the estimated parameter is greater than one, we assume that the private sector continues to use a unit root process for inflation. This partly explains the bunching of estimates of inflation persistence at and close to one. In the benchmark calibration, it turns out that the incidence of a binding cut-off point happens about 0.35% of the time.

[5] Under rational expectations, there exist a number of equivalent representations of the reaction function under discretion. In the case of learning, these are no longer equivalent. As the emphasis of this paper is on the optimal policy, we do not exhaust the alternative specifications but settle for the MSV rule.

These findings are also captured by differences in the impulse responses (not shown in the paper) following a cost-push shock. In both cases, the output gap falls for one period. However, the behaviour of inflation is quite different depending on the current estimated inflation persistence. Under rational expectations, the inflation response is always short-lived. In contrast, under adaptive learning the behaviour can vary substantially due to the above discussed effect. In sum, neglecting to take account of learning on the part of the private sector can be costly.

3. Optimal monetary policy under adaptive learning.

In the previous section, we kept the policy rule constant across the two cases. The standard way of using recursive learning is to ask if a certain equilibrium is learnable and analyze which policy rules lead to convergence to rational expectations equilibrium. A different question is, suppose the central bank knows that the agents are learning in this particular way, what is the optimal policy response? In this case the central banker is well aware that policy actions influence expectations formation and thereby inflation dynamics. The central bank is assumed to know the expectation formation mechanism in full. "Sophisticated" central banking implies solving the full dynamic optimisation 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:

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

subject to the expectations adjusted Phillips curve (7) and the recursive parameter updating equations (4) and (5).

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 section 2. 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 we resort to non-linear methods in order to solve the policy problem. We therefore employ the collocation-methods described in Judd (1998) and Miranda and Fackler (2002), which amount to approximating the value function with a combination of non-linear polynomials, which translates the problem to a root finding exercise (some details are outlined in the appendix).

Agents use a forecasting model that is not fully correct as they do not take into account that the estimated persistence parameter has an effect on policy and hence on the evolution of inflation (this is similar to the standard constant gain, but here it is amplified also by the non-linearity). Average forecast errors are, however, zero so it is not the case that the central bank systematically tricks the private sector.

Figures 5 to 7 describe the distribution of inflation, the output gap and the estimated degree of inflation persistence in the case of optimal monetary policy (these figures should be compared with Figures 2 to 4 respectively) and Table 1 shows some statistical features. A number of observations are worth making. First, comparing Figures 3 and 5, it is clear that under sophisticated (optimal) central banking the variance of inflation (and also of the semi-difference of inflation) is lower. In particular, optimal monetary policy reduces the mass of inflation in the tails (i.e. between 0.02 and 0.03 in absolute value). Second, comparing Figure 2 and Figure 6, it is clear that this reduced inflation variability comes at the cost of a higher volatility of the output gap. Overall, the loss under the optimal policy response is, however, reduced by about 15%.

This result is somewhat analogous to what obtains when we compare the results under commitment and discretion in the case of rational expectations. (Table 1) The mechanism behind the improvement is, however, substantially different from the rational expectations case. Under rational expectations, as argued above, commitment allows the central banker to use the future course of policy to spread the impact of economic shocks over time. As we have seen, in section 2 the mechanism may be interpreted as a form of automatic stabilisation. The central bank, by committing to a persistent response to a cost-push shock, induces an undershooting of future inflation. Under rational expectations this brings inflation anticipations down thereby mitigating the impact of the shock.

Under adaptive learning the mechanism must be different as future policy as such plays no role at all. However, a common mechanism is that the central bank realises that it can influence expectations formation. In the case of rational expectations, the central bank realises that credible future policy actions can influence current expectations by making the inflation process mean reverting. In the case of adaptive learning, the central bank realises it can steer the degree of inflation persistence estimated by the private sector.

To illustrate how this affects the optimal policy response, Figures 8 and 9 show the mean dynamic response (the equivalent of impulse responses in a non-linear model) of the output gap and inflation for two different values of the initially estimated inflation persistence in response to a one-standard-deviation cost-push shock. In one case (dashed line), the estimated degree of inflation persistence is 0.3 higher than the mean, whereas in the other case (solid line) the initial estimated inflation persistence is 0.3 lower than the mean. When the initial inflation persistence is estimated to be high, the effect of a cost-push shock on inflation is much higher and more persistent (see Figure 9), then when it is low. The intuition is simple: With a high degree of estimated inflation persistence, agents forecast the effects of the cost-push shock to be persistent and therefore make actual inflation much more persistent. As a result, the central bank needs to respond more aggressively and in a much more prolonged fashion when inflation persistence is high (Figure 8). As the central bank trades off inflation and output variability, high inflation persistence therefore clearly increases both inflation and output gap variability.

This costly trade-off mechanism is, however, only a part of the story of the optimal policy response. As mentioned before, the central bank also realises that by its policy it can steer the future evolution of the estimated inflation persistence and thereby reduce future trade-off problems. In particular, realising that a high degree of estimated inflation persistence is very costly, it will want to bring the estimated degree of inflation persistence down as quickly as possible. It can do this by generating a sequence of negative forecast errors, i.e. deliver inflation rates below the private sector expectations, which through the recursive learning equations leads to a decrease in the estimated persistence. This is exactly what happens in Figure 10, which shows the impulse response of the estimated degree of inflation persistence to the cost-push shock under optimal policy. Of course, because such a policy is costly in terms of creating a persistently negative output gap, the reduction in the estimated inflation persistence is only gradual. However, the mean reversion is higher than when the estimated degree of inflation persistence is unusually low (also Figure 10). Indeed, the reverse argument can be made to explain why the central bank has an interest to slow down the upward convergence process of the estimated inflation persistence in this case. Again, the central bank realises that a low estimated inflation persistence is beneficial in terms of reducing the trade-off problem between inflation and output gap variability in the presence of cost-push shocks. It therefore will respond less aggressively and persistently in this case, thereby prolonging the convergence of the estimated inflation persistence process.

This asymmetric policy behaviour is reflected in the distribution of the estimated inflation persistence shown in Figure 7. First, it results in a downward shift in the mean of the estimated degree of inflation persistence. This is again reminiscent of what happens in the case of commitment under rational expectations. If one generates inflation data from the commitment solution under rational expectations, and estimate an AR(1) process (that is, calculate the first order autocorrelation), the result is a coefficient of 0.3, rather than the 0.5 for the case of rational expectations under discretion. With adaptive learning, the optimal policy leads to a mean estimate of 0.42, which lies in between the two cases. Second, it is also clear from Figure 7 that the distribution of the estimated inflation persistence parameter is now skewed to the left. In particular, realising the cost of high estimated inflation persistence, under the optimal policy the central bank ensures that only under exceptional circumstances the inflation process wanders close to being a unit root process. Comparing Figure 4 and 7, it is clear that the mass of estimated inflation persistence higher than 0.8 is significantly lower under the optimal policy case. Moreover, under the optimal monetary policy there is no longer any instance of explosive inflation behaviour (compared to 0.35% in the case of a naïve central banker). As a result, inflation and inflation expectations are much better anchored under sophisticated (optimal) central banking.

When the gain is decreased, the tendency to pursue persistent policy is reduced (the impulse response of the output-gap "shifts to the left". When the relative weight of output is lower, the degree of stabilization of course is lower and hence the distribution of c is slightly wider and the same is true for inflation. (to be completed)

4. Conclusions

In this paper, we compare the conduct of monetary policy in a simple new Keynesian model, under rational expectations and adaptive learning. We found that, under a simple monetary policy rule, the optimal time-consistent rule under rational expectations, adaptive learning leads to a deterioration in economic performance. While the distribution of the output gap is the same under rational expectations and adaptive learning, the distribution for inflation and for the semi-difference in inflation becomes wider and flatter. Given a quadratic loss function it leads to significant reduction in social welfare. As the Kalman gain is increased, the perceived persistence of inflation becomes more sensitive to contemporary developments, and inflation variance increases. Moreover, the economy, under adaptive learning, eventually displays explosive behaviour.

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, time-consistent policy.

In this paper, we compare full optimal policy, under adaptive learning, with a policy following a simple monetary policy rule (naïve central banking). We find that explosive dynamics never occur under optimal (sophisticated) policy. Moreover, perceived inflation persistence is lower and less dispersed. The results can reasonably be interpreted as meaning that inflation and inflation expectations are better anchored under sophisticated central banking. 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. For low estimated persistence, mean dynamic response to a cost-push shock resembles the impulse response function for optimal time consistent policy. As estimated persistence increases mean dynamics become more inertial, resembling the impulse response function of optimal policy, under commitment.

It is clear that both a central banker with commitment ability, 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.

REFERENCES

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 (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 (2005), Monetary Policy Over Time, forthcoming in Macroeconomic Dynamics.

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.

Milani, F. (2004). Expectations, learning and macro-economic persistence. Mimeo. Princeton University.

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

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

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

Orphanides, Athanasios and John C. Williams (2004). "Robust monetary policy with imperfect knowledge", mimeo, Federal Reserve Board.

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

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

1 Appendix

Towards writing the model in recursive form, note that

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

Note that π_t is used to update c_t, but c_t is not used by agents to forecast until $t+1$. Substituting the second equation into the fourth gives a completely recursive system in the states, $s_{t+1} = g\left(s_t, x_t, \varepsilon_{t+1}\right)$. Derivatives of g are

$$
\begin{aligned}
\frac{d\pi_t}{dx_t} &= \frac{\kappa}{1+\beta\left(\gamma-c_{t-1}\right)} \\
\frac{d^2\pi_t}{dx_t^2} &= 0 \\
\frac{dc_t}{dx_t} &= \frac{\phi\kappa\left(R_{t-1}+\phi\left(\pi_{t-1}^2-R_{t-1}\right)\right)^{-1}\pi_{t-1}}{1+\beta\left(\gamma-c_{t-1}\right)} \\
\frac{d^2c_t}{dx_t^2} &= 0.
\end{aligned}
$$

Finally, the reward function is

$$
U_t = -L_t = -\frac{1}{2}\left[\left(\frac{1}{1+\beta\left(\gamma-c_{t-1}\right)}\left[\gamma\pi_{t-1}+\kappa x_t+u_t\right]-\gamma\pi_{t-1}\right)^2+\lambda x_t^2\right]
$$

such that

$$
\begin{aligned}
\frac{dU_t}{dx_t} &= -\left(\pi_t-\gamma\pi_{t-1}\right)\frac{d\pi_t}{dx_t}-\lambda x_t \\
\frac{d^2U}{dx_t^2} &= -\left(\frac{d\pi_t}{dx_t}\right)^2-\lambda.
\end{aligned}
$$

We then approximate the value function with splines, along the lines of Miranda and Fackler (2002).

Table 1. Variances and expected losses (scaled by 10e-5), $\lambda = 0.05$.

	$\text{Var}(\pi_t - \gamma\pi_{t-1})$	$\text{Var}(x_t)$	$E\left[(\pi_t - \gamma\pi_{t-1})^2 + \lambda x_t^2\right]$
Commitment, RE	1.56	5.77	1.84
Discretion, RE	2.08	4.08	2.28
Learning, D-rule	2.39	4.08	2.57
Learning, Optimal	1.90	5.51	2.18

Figure 1: Impulse responses of inflation and output gap, rational expectations

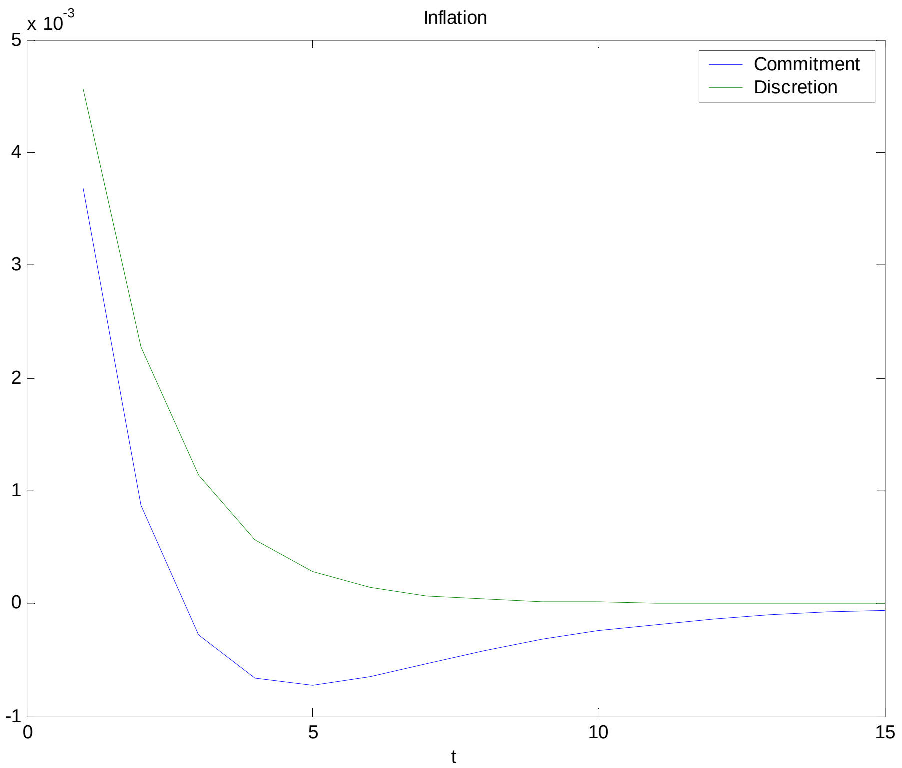

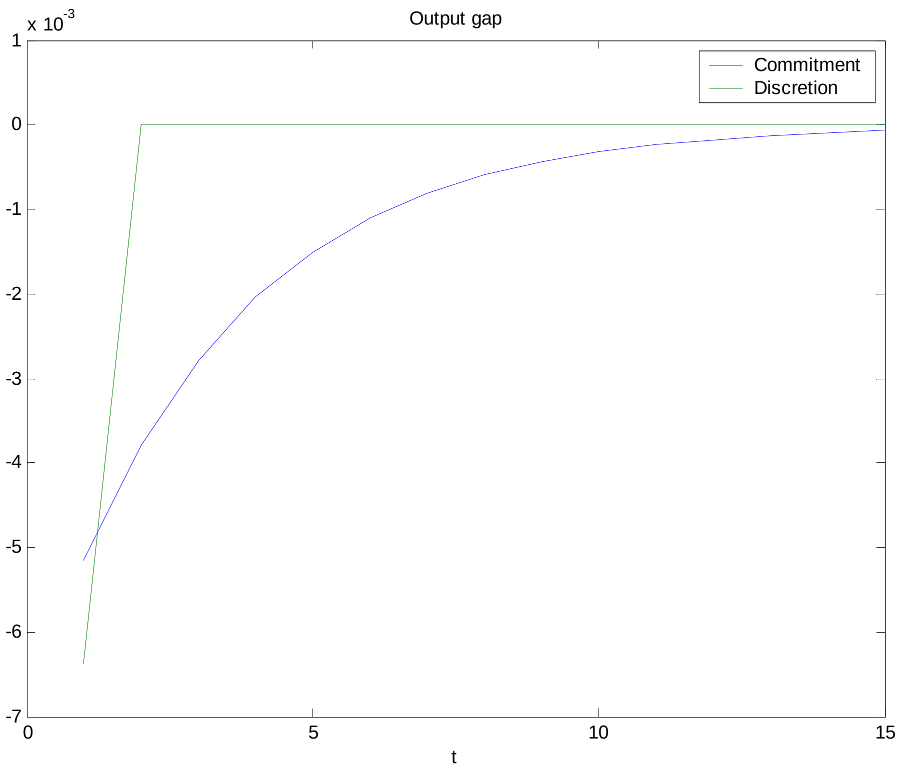

Figure 2: Distribution of output under discretion

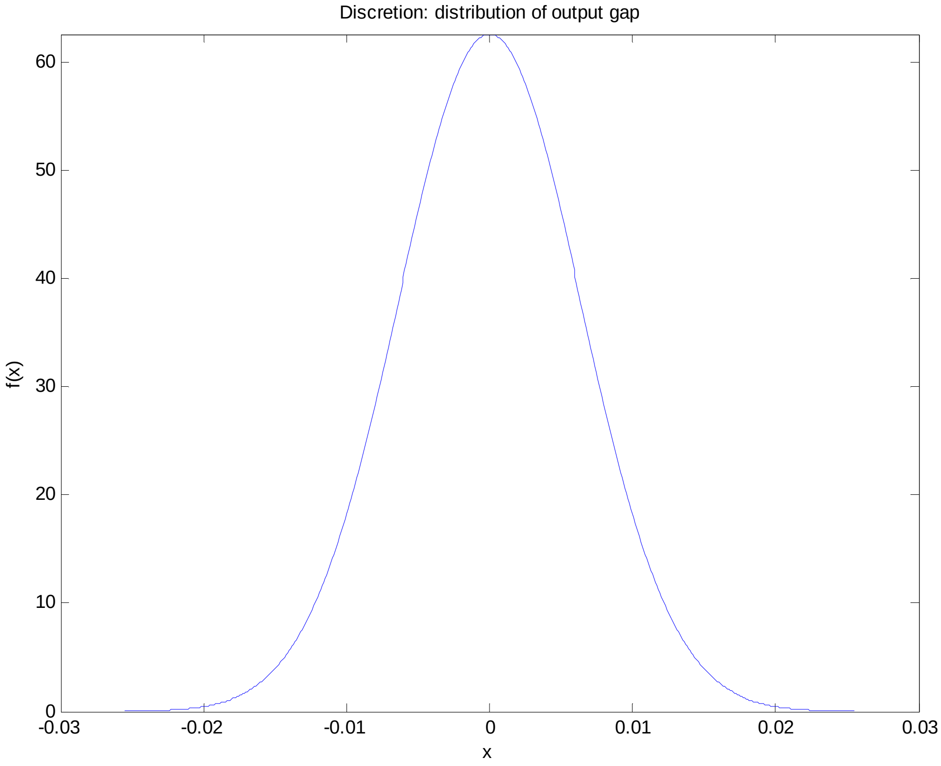

Figure 3: Distribution of inflation and semi-difference of inflation, rational expectations and learning

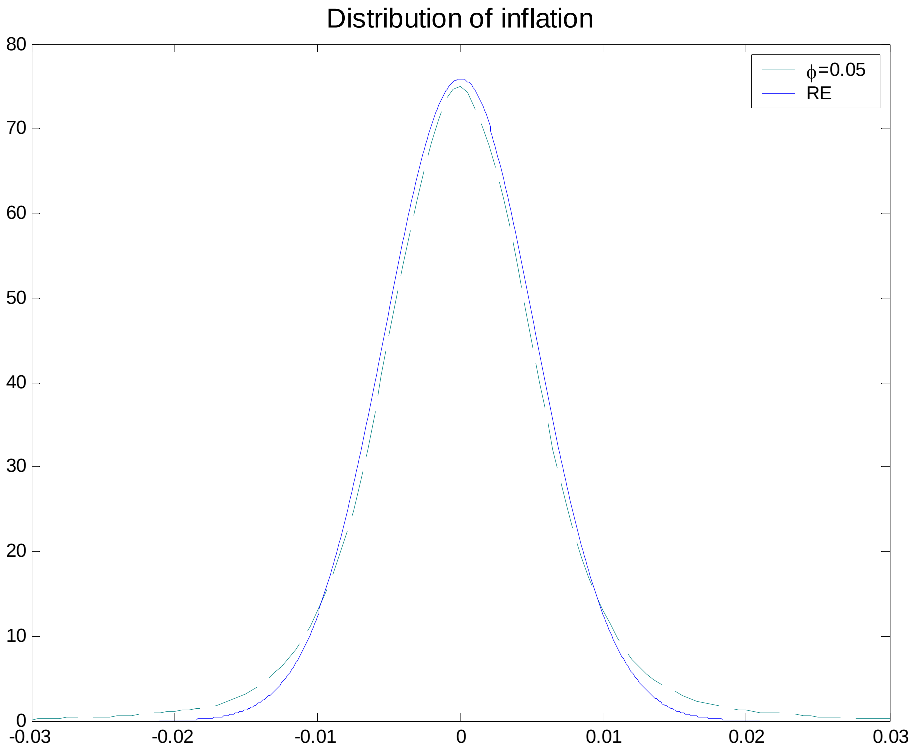

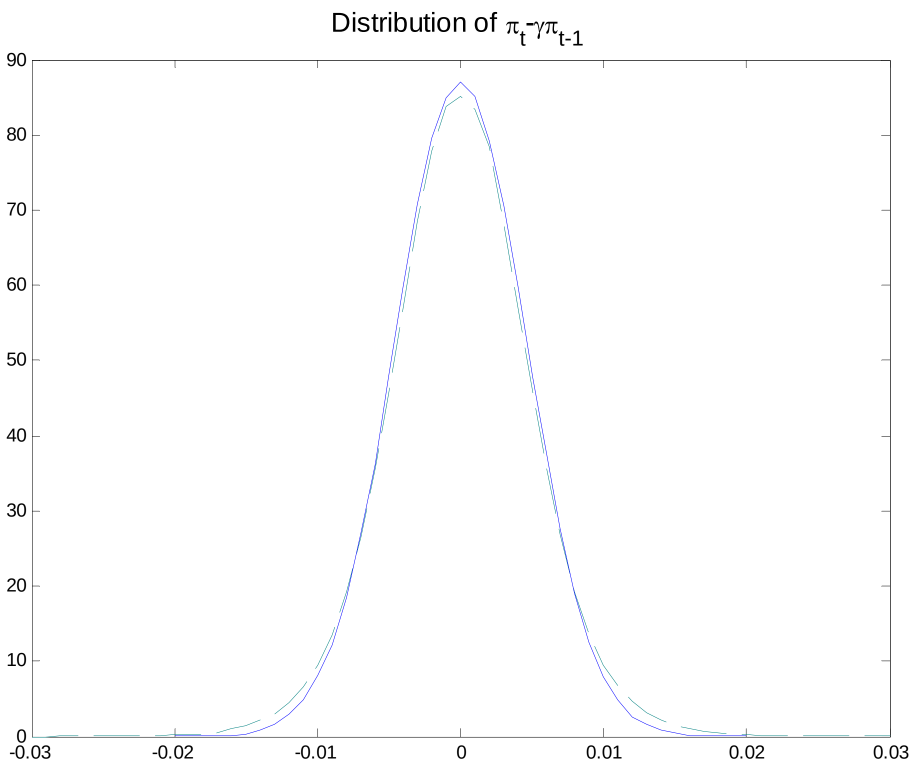

Figure 4: Distribution of estimated parameter, learning (not optimal)

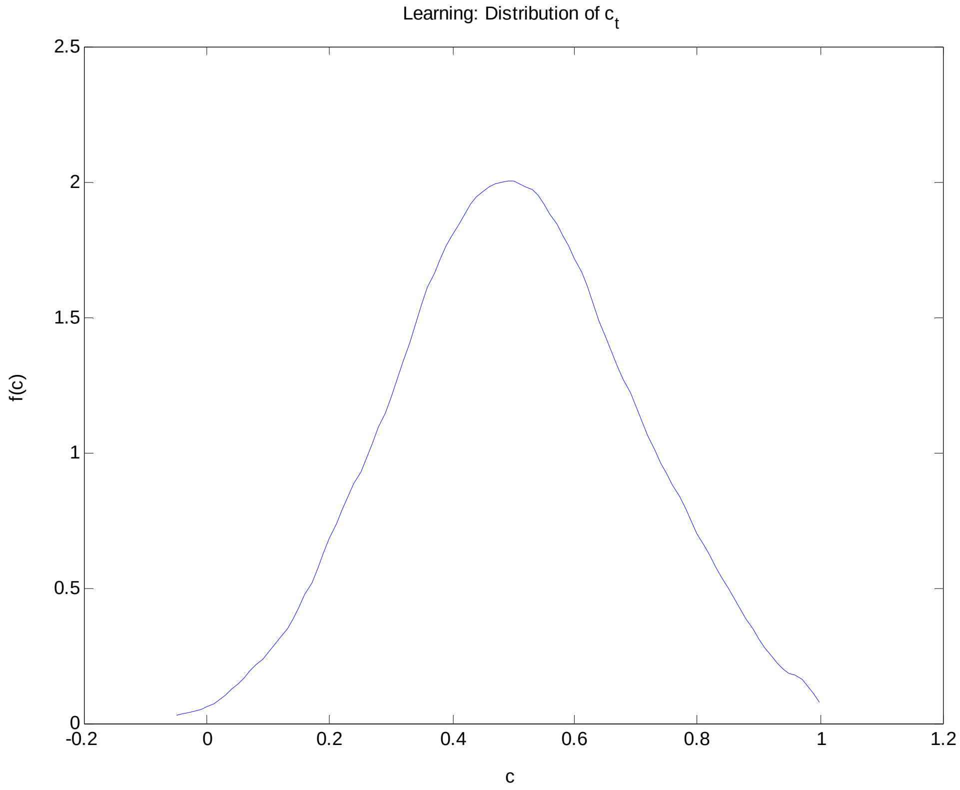

Figure 5: Distribution of inflation, optimal policy

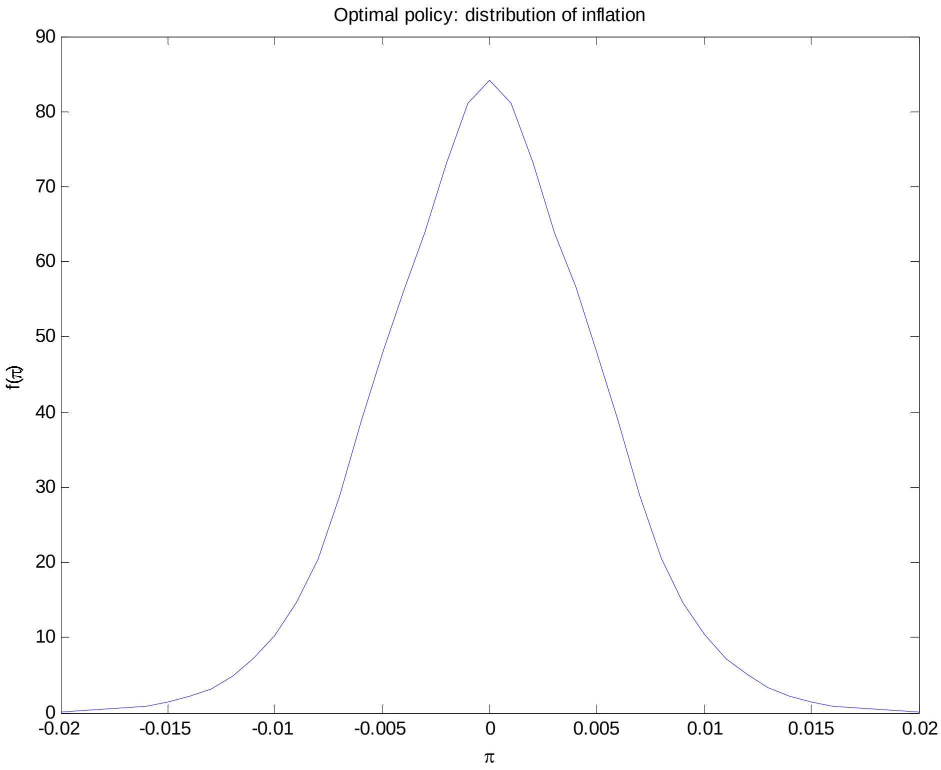

Figure 6: Distribution of output gap, optimal policy

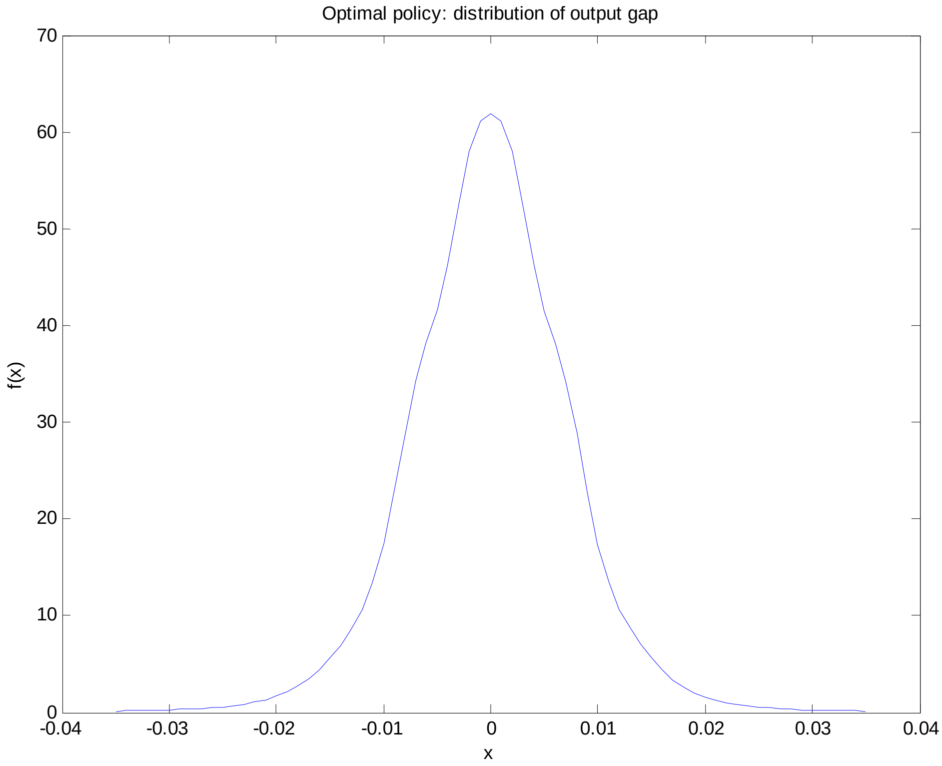

Figure 7: Distribution of the estimated inflation persistence, optimal policy

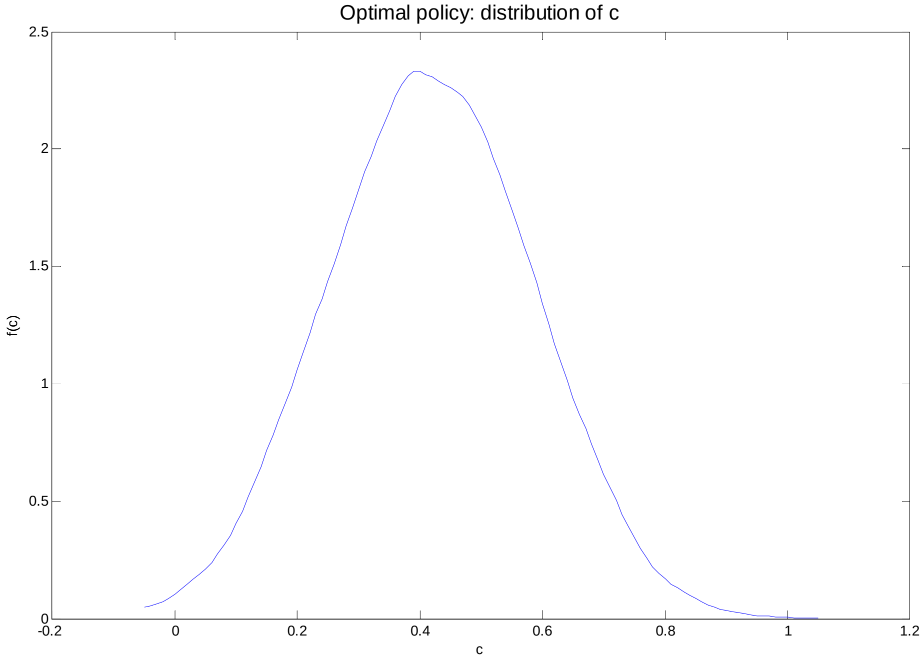

Figure 8: Optimal policy, mean dynamic response of output to a cost push shock

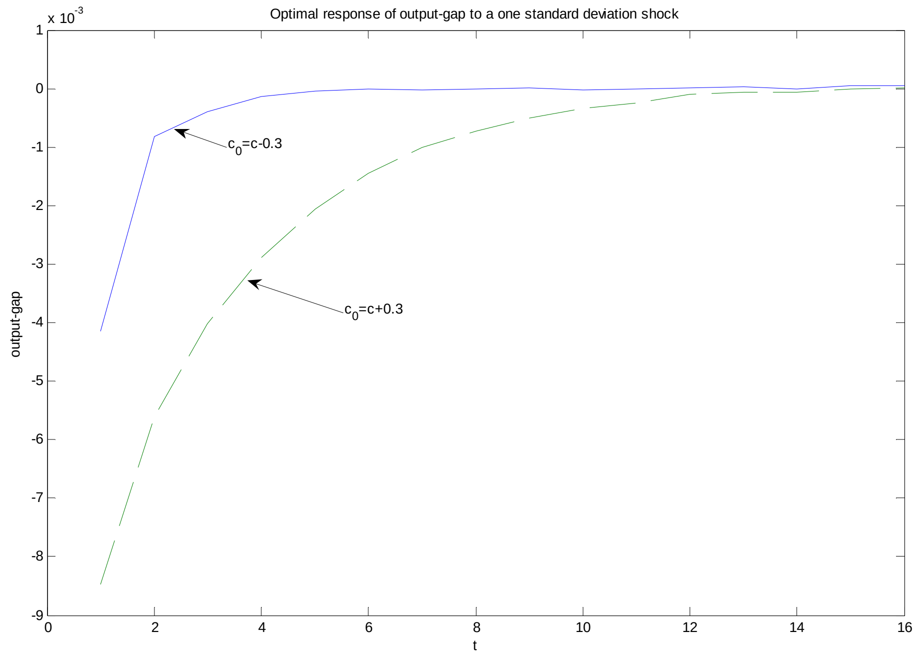

Figure 9: Optimal policy, mean dynamic response of inflation to a cost push shock

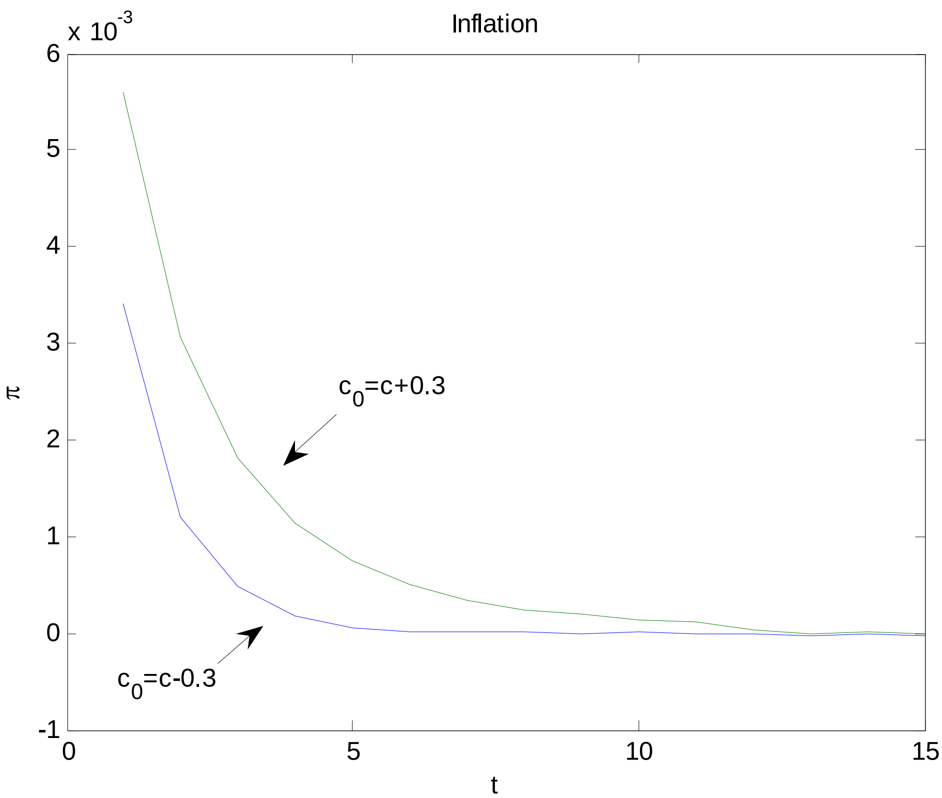

Figure 10: Optimal policy, mean dynamic response of estimated inflation persistence