

Optimal Monetary Policy under Adaptive Learning

Vitor Gaspar, Frank Smets and David Vestin¹

20 February 2006

Abstract: We consider optimal policy when private sector expectations are formed through adaptive learning. Earlier research has found that adaptive learning is consistent with empirical evidence on private sector expectations. In this paper, we consider the (admittedly) extreme case of sophisticated central banking, whereby the central bank has full knowledge about the structure of the economy. Our results confirm that the management of inflation expectations is crucial for the conduct of monetary policy. In particular, when the private sector perceives inflation persistence to be high, optimal policy responds strongly to lagged inflation and inflation shocks thereby stabilizing inflation and anchoring inflation expectations. In our particular example it does so at no cost for output gap stability.

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. A reasonable alternative is the assumption of adaptive learning. In this case, agents have limited knowledge of the precise working of the economy, but as time goes by and available data changes, they update their knowledge and the associated forecasting rule. Adaptive learning may be seen as a minimal departure from rational expectations, which accounts for imperfect knowledge about the structure of the economy and the implications from pervasive structural change that characterizes modern economies. 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 private sector adaptive learning for the conduct of optimal monetary policy. We analyze the monetary policy response to shocks and the associated macro-economic outcomes, when the central bank minimizes an explicit loss

function.² Modeling the optimal behavior of the central bank requires specifying its information set. In this paper, we consider the (admittedly) extreme case of sophisticated central banking, whereby the central bank has full information about the structure of the economy (a standard assumption under rational expectations). The information set therefore includes knowledge about the precise mechanism generating private sector's expectations. The objective is to investigate to what extent the relatively small change in the assumption of how agents form their inflation expectations affects the principles of optimal monetary policy.

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 are (sufficiently) concerned about their long-run reputation so as not to yield to short-run temptations. The performance of the economy is better as a consequence. In a rational expectations framework, it is therefore possible to justify the primacy of long run goals such as price stability. Are similar considerations relevant when we depart from rational expectations? Gaspar, Smets and Vestin (2005b) 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, 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. In this paper, we build on this work by characterizing more fully optimal policy under adaptive learning and

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

contrasting 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; there is no dichotomy between the two mechanisms anchoring inflation expectations. On the contrary, the central bank's ability to influence expectations about the future course of policy rates and its track record in preserving stability are complements.

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 macro-economic outcomes under different policy regimes and characterize the optimal policy to state variables, especially to cost-push shocks, lagged inflation and perceived inflation persistence. Section 4 contains some robustness analysis with respect to different assumptions regarding the calibration. In section 5, we conclude.

2. New Keynesian model with adaptive learning.

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

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

$$(1) \quad \pi_t - \gamma\pi_{t-1} = \beta(E_t\pi_{t+1} - \gamma\pi_t) + \kappa x_t + u_t,$$

where π_t is inflation, x_t is the output gap and u_t is a cost-push shock (assumed i.i.d.).

Furthermore, β is the discount rate, κ is a function of the underlying structural parameters including the degree of Calvo price stickiness, α , and γ captures the degree of intrinsic inflation persistence due to partial indexation in the goods market.

Woodford (2003) has shown that, under the assumed microeconomic assumptions, a quadratic approximation of the (negative of the) period social welfare function takes the following form:

Svensson stresses the importance of looking at optimal policy and targeting rules in order to understand modern central banking.

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

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.³ A γ different from 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.

As discussed in the introduction, we consider two assumptions regarding the formation of inflation expectations in equation (1): rational expectations and adaptive learning. In this subsection, we first solve for optimal policy under discretion and rational expectations. This will serve as a benchmark for the analysis of optimal policy under adaptive learning.

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

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

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

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. Substituting (1'') into (2') and minimizing the result with respect to the output gap, implies the following policy rule:

³ 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.

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

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

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

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

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

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

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 $\gamma=0$), there should be an appropriate balance between inflation and the output gap. The higher the λ , the higher is inflation in proportion to (the negative of) the output gap, because it is more costly to move the output gap. When κ increases, inflation falls relative to the output gap. When $\gamma>0$, it is the balance between the quasi difference of inflation and the output gap that matters. If last periods inflation was high, there is a tendency that current inflation *should* also be high. The reason is that price dispersion drives the welfare criterion. When prices are partially indexed to lagged inflation, other prices must rise in proportion to this indexation in order to avoid price dispersion.

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

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

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

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.

As shown above, under discretion optimal monetary policy only responds to the exogenous shock and not to lagged inflation. 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.

2.2. Inflation expectations according to adaptive learning.

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

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

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

Under adaptive learning, we assume that the private sector believes the inflation process is well approximated by equation (5). However, as they do not know the underlying parameters, they estimate the equation recursively, using a “constant-gain” least squares algorithm, implying perpetual learning.

Thus, the agents estimate the following reduced-form equation for inflation,⁵

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

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.

$$(7) \quad c_t = c_{t-1} + \phi R_t^{-1} \pi_{t-1} (\pi_t - \pi_{t-1} c_{t-1})$$

$$(8) \quad R_t = R_{t-1} + \phi (\pi_{t-1}^2 - R_{t-1}),$$

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

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

$$(9) \quad E_t \pi_{t+1} = c_{t-1} \pi_t$$

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. However, according to (9), 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

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

both the forecast and actual inflation, but not to allow agents to update the coefficients using current information (i.e. just substitute (9) 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.⁶

Substituting equation (9) into the New-Keynesian Phillips curve (1) we obtain:

$$(10) \quad \pi_t = \frac{1}{1 + \beta(\gamma - c_{t-1})} (\gamma\pi_{t-1} + \kappa x_t + u_t).$$

2.3. Solution method for optimal monetary policy

Under adaptive learning 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 under the loss function (2). In the first case, the simple rule (3), the Phillips curve (1) and equations (7), (8) and (9), 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

⁶ It is possible to solve this problem in the current setting. However, we leave this for future research.

“Sophisticated” central banking implies solving the full dynamic optimization problem, where the parameters associated with the estimation process are also state variables. Specifically, in this case the central bank solves the following dynamic programming problem:

$$(11) \quad V(u_t, \pi_{t-1}, c_{t-1}, R_t) = \max_{x_t} \left\{ -\frac{(\pi_t - \gamma\pi_{t-1})^2 + \lambda x_t^2}{2} + \beta E_t V(u_{t+1}, \pi_t, R_{t+1}, c_t) \right\},$$

subject to equation (10) and the recursive parameter updating equations (7) and (8).⁷

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

$$(12) \quad x_t = \psi(u_t, \pi_{t-1}, c_{t-1}, R_t).$$

As shown in the appendix (to be completed), equation (12) can be written as:

$$(13) \quad x_t = -\frac{\kappa\delta_t}{\kappa^2\delta_t + \lambda\chi_t^2} u_t + \frac{\kappa\gamma(\chi_t - \delta_t) + \beta\kappa\chi_t\phi R_t^{-1} E_t V_c}{\kappa^2\delta_t + \lambda\chi_t^2} \pi_{t-1} + \beta \frac{\kappa\chi_t}{\kappa^2\delta_t + \lambda\chi_t^2} E_t V_\pi$$

where $\delta_t = 1 - 2\beta\phi E_t V_R$, $\chi_t = 1 + \beta(\gamma - c_{t-1})$ and V_c , V_π and V_R denote the partial derivatives of the value function with respect to the variables indicated in the subscript. Equation (13) does not characterize the optimal solution fully as we have not specified the exact forms of the partial derivatives, but it will be important to interpret the results in section 3.⁸

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

⁷ The value function is defined as $V(\cdot) = \max_{\{x_j\}} \left\{ -\sum_j \beta^j [(\pi_j - \gamma\pi_{j-1})^2 + \lambda x_j^2] \right\}$ s.t. (1), (8), (9) and (10), that is as

maximizing the negative of the loss. It is important to bear this in mind when interpreting first order conditions.

⁸ Clearly the partial derivatives V_c , V_π and V_R are functions of the state vector $(u_{t+1}, \pi_t, c_t, R_{t+1})$.

a game between the central bank and the private sector (as is the case under discretion and rational expectations). Second, promises of future policy play no role as agents look only at inflation outcomes. Hence, there is no scope for the type of commitment gains discussed in the rational expectations literature. Third, we leave the linear-quadratic world, as the learning algorithm makes the model non-linear.

From a technical perspective, the first two aspects simplify finding the optimal policy whereas the third is a complication. The value function will not be linear-quadratic in the states and hence we employ the collocation-methods described in Judd (1998) and Miranda and Fackler (2002) to solve the model numerically. This amounts to approximating the value function with a combination of cubic splines and translates in a root finding exercise (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 these structural parameters imply that $\kappa=0.019$; $\lambda= 0.002$.⁹ γ is chosen such that there is some inflation persistence in the benchmark calibration. A value of 0.5 for γ is frequently found in empirically estimated new Keynesian Phillips curves (see, for example, Smets, 2002). $\theta=10$ 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 0.02. Orphanides and Williams (2004) found that a value in the range 0.01 to 0.04 is needed to match up the resulting model-based inflation expectations with the Survey of

⁹ Here we follow the discussion in Woodford (2003). See especially pages 187 and 214-15. For the relevant parameters we rely on ...

Professional Forecasters. A value of 0.02 corresponds to an average sample length of about 25 years.¹⁰ In the limiting case, when the gain approaches zero, the influence of policy on the estimated inflation persistence goes to zero and hence plays no role in the policy problem.

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

¹⁰ See Orphanides and Williams (2004). Similarly, Milani (2005) estimates the gain parameter to be 0.03 using a Bayesian estimation methodology.

under commitment, inflation expectations operate as automatic stabilizers in the face of cost-push shocks. Such intuition is clearly present in the results presented in Table 2. Clearly, the output gap is not persistent under the simple rule, (under the assumption that cost-push shocks are i.i.d.). In contrast, under commitment the output gap becomes very persistent with autocorrelation of 0.66. The reverse is true for inflation. Inflation persistence, under the simple rule, is equal to the assumed intrinsic persistence parameter at 0.5. Under commitment it comes down to only 0.24. 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 optimal discretionary monetary policy. Overall, the loss is about 28 % higher under discretion.

[Insert Table 2]

Following Orphanides and Williams (2002), it is also useful to compare the outcomes under rational expectations and adaptive learning for the case of the simple monetary policy rule, looking at the second and fourth columns in Table 2. The comparison confirms their findings. Clearly, the autocorrelation and the volatility of the output gap remain unchanged at zero as under the same simple rule the output gap only responds to the contemporary iid cost-push shock. Nevertheless, under adaptive learning, the autocorrelation of inflation increases from 0.5 to about 0.56. The variance of inflation increases to more than twice the value under commitment, and the variance of the quasi-difference of inflation is about 1.5 times larger. Thus, the expected welfare loss increases significantly to about 37% more, than in the case of commitment, and about 9 percentage points more than under the same rule and rational expectations. 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 0.36, and the persistence of inflation. As before, this is linked with a significant decline in inflation volatility relative to the simple rule. Inflation variance declines 90 percentage points to about only 28% more than in case of commitment. The variance of the quasi-difference of inflation also falls by about 35 percentage points. At the same time, the output gap becomes more persistent, with auto-correlation of 0.51 (still less than under commitment) while the volatility of the output gap is unchanged relative to the simple rule. On balance, the expected welfare loss falls significantly, by about 28 percentage points under optimal policy against the simple rule (which compares with a similar gain of about 28 % 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 (or no) 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 ... % above, for the case of the simple rule. Such results confirm the intuition that our adaptive learning approximation is “reasonable”.

Figure 1 provides some additional detail concerning the distribution of the endogenous variables, i.e. the estimated persistence, output gap, inflation, quasi-difference of inflation, and the moment matrix, under optimal policy and the simple policy rule. First, panel (a) shows not only that the average of the estimated persistence parameter is significantly lower under optimal policy, but also that the distribution is more

concentrated around the mean. It is important to note that, under optimal policy, the perceived inflation parameter does never go close to one, contrary to what happens under the simple 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) unity¹¹. 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. Optimal monetary policy under adaptive learning succeeds in excluding such explosive dynamics.

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

[Insert Figure 1]

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

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

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 (13),

$$(13) \quad x_t = -\frac{\kappa\delta_t}{\kappa^2\delta_t + \lambda\chi_t^2}u_t + \frac{\kappa\gamma(\chi_t - \delta_t) + \beta\kappa\chi_t\phi R_t^{-1}E_tV_c}{\kappa^2\delta_t + \lambda\chi_t^2}\pi_{t-1} + \beta\frac{\kappa\chi_t}{\kappa^2\delta_t + \lambda\chi_t^2}E_tV_\pi$$

where $\delta_t = 1 - 2\beta\phi E_tV_R$, $\chi_t = 1 + \beta(\gamma - c_{t-1})$ and V_c , V_π and V_R denote the partial derivatives of the value function with respect to the variables indicated in the subscript. Equation (13) expresses the policy instrument – here taken to be the output gap x_t – as a function of its determinants. When interpreting equation (14) there are two important points to bear in mind. First, the partial derivatives V_c , V_π and V_R depend on the vector of states $(u_{t+1}, \pi_t, c_t, R_{t+1})$. Second, the value function is defined in terms of a maximization problem. In such a case, a positive partial derivative means that an increase in the state contributes favorably to our criterion. Or, more explicitly, that it contributes to a *reduction* in the loss.

In order to discuss some of the intuition behind the optimal policy reaction function, it is useful to consider a number of special cases. In particular, in the discussion in this subsection, we assume that E_tV_R is zero, so that the expected marginal impact of changes in the moment matrix on the value function is zero. Such assumption provides a reasonable starting point for the discussion for reasons, which we will make clear in the next subsection. If E_tV_R is zero then $\delta_t = 1$, making equation (14) much simpler.

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

$$(14) \quad x_t = -\frac{\kappa}{\kappa^2 + \lambda\chi_t^2} u_t.$$

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 shocks are uncorrelated over time and symmetrically distributed, it follows that $E_t V_\pi$ must be 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 in addition $c_{t-1} = \gamma$ and as a result $\chi_t^2 = \chi_t = 1$, equation (14) 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 (7), it is clear that, when lagged inflation is zero, the estimated persistence parameter is not going to change irrespective of current policy actions. As a result, no benefit can possibly materialize from trying to affect the perceived persistence parameter. The same intuition holds true to explain why when $\phi=0$ 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, $\gamma > c_{t-1}$ (still maintaining the assumption that $\pi_{t-1}=0$). 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 than under discretion and rational expectations. The reason is again simple. As shown in equation (10), the smaller

the degree of perceived inflation persistence, the smaller the impact of a given cost-push shock on inflation, all other things constant. 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 $\gamma < c_{t-1}$, $\chi_t < 1$ 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.

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

[Insert Figure 2]

Returning to equation (13) and departing from the assumption that $\pi_{t-1}=0$, we can discuss the second term, which captures part of the optimal response to lagged inflation. Note that the first term in the numerator is zero when $\gamma = c_{t-1}$ (still using the simplifying assumption that $\delta_t = 1$). In such a case, inflation expectations adjust to past inflation just in line with the partial adjustment of inflation due to its intrinsic persistence (equation 11). Given the loss function this is a desirable outcome. In the absence of any further shock, inflation will move exactly enough so that the quasi-difference of inflation will be zero. Note that when $\gamma > c_{t-1}$ or $\chi_t > 1$ the response of the output gap to past inflation,

according to this effect, is positive. Hence, past inflation justifies expansionary policy. At first sight, this is counter-intuitive. However, the reason is clear, when estimated persistence is below intrinsic persistence, past inflation does not feed enough into inflation expectations, to stabilize the quasi-difference of inflation. In order to approach such a situation an expansionary policy must be followed. This factor is important because it shows that, in the context of our model, there is a cost associated with pushing the estimated persistence parameter too low.

However, in general, the second term in the numerator of the reaction coefficient will be negative and dominate the first term ensuring a negative response of the output gap to inflation. This term reflects the intertemporal trade-off the central bank is facing between stabilizing the output gap and steering the perceived degree of inflation persistence by inducing forecast errors. In our simulations it turns out that the expected marginal cost (the marginal impact on the expected present discounted value of all future losses) of letting estimated inflation persistence increase is always positive, i.e. $V_c < 0$ and large. The intuition is that, as discussed above, a lower degree of perceived persistence will lead to a much smaller impact of future cost-push shocks on inflation, which tends to stabilize inflation, its quasi-difference and the output gap. As a result, under optimal policy the central bank will try to lower the perceived degree of inflation persistence. As is clear from the private sector's updating equation (7), 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.

Finally, the third term in equation (13) 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 (13) in the calibrated model. Figure 3a plots the output gap (on the vertical axis) as a function of lagged inflation and the perceived degree of inflation persistence for a zero cost-push shock and assuming that the moment matrix R equals its average for a particular realization of c . A number of features are worth repeating. First, when lagged inflation and the cost-push shock are zero, the output gap is also zero irrespective of the estimated degree of inflation persistence. Second, when the shock is zero, the response to inflation and deflation is symmetric. Third, as the estimated persistence of inflation increases, the output gap response to inflation (and deflation) rises. It is then interesting to see how the output gap response differs when a positive cost-push shock hits the economy. This is shown in Figure 3b, which plots the differences in output gap response to a positive one-standard deviation cost-push shock and zero cost-push shock as a function of lagged inflation and the perceived persistence parameter. The output gap response is always negative and increases with the estimated degree of inflation persistence. The figure also shows the non-linear interaction with lagged inflation. In particular, the output gap response becomes stronger when inflation is already positive.

[Insert Figure 3]

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

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

The answer is yes, as we illustrate in figure 4, for the case of a positive (negative) cost-push shock when lag inflation is negative (positive). The policy response, apparent in

panel (b), is fully symmetric. Moreover, from the panel (a) of Figure 4 it is clear that the adjustment of inflation is also symmetric. Finally, panel (d) shows that the adjustment of estimated persistence is the same in both cases (the small discrepancy in the figure is due to the numerical accuracy of our numerical procedure). The same would be true of the moment matrix (not shown).

[Insert Figure 4]

3.3. The role of V_R .

When interpreting equation (13) in the previous section, we assumed that $E_t V_R = 0$. In our simulations V_R is, on average, positive but small. Under such conditions, it is clear that although the expected value of V_R may depart from zero δ_t will be close to one most of the time. In fact the average of V_R is about 1/3, which makes $\delta \cong 0.987$. The standard deviation of V_R is about 1.2, implying that $0.939 < \delta < 1.0343$ for values of V_R within a plus or minus 1-SD band. Thus, the arguments in sub-section 3.2 hold approximately in many relevant cases.

To understand the behavior of V_R it is useful to note, from equation (7), that a high value of R reduces the updating of the estimated persistence parameter for a given residual $(\pi_{t+1} - \pi_t c_t)$. Now, the stability of estimated persistence is good when it is low, remember that $V_c < 0$ and large in absolute value. However, when c is high, and optimal policy will imply that it must be reduced, a high value of R is problematic because it implies that extra effort will be necessary to bring it down. Thus, it is intuitive that V_R does register positive and negative values. Furthermore, it is worthwhile to point out that high values of estimated persistence are, in general, associated with large values of R .

The partial derivative of the value function with respect to the moment matrix, V_R , depends on the state variables $(u_{t+1}, \pi_t, c_t, R_{t+1})$. Given the state variables, at period t , the output gap, x_{ty} , is determined by equation (13). The output gap, in turn, determines all relevant state variables in $(u_{t+1}, \pi_t, c_t, R_{t+1})$, except for the shock. Hence, when evaluating

$E_t(V_R)$, the shock at $t+1$ is the only unknown state variable. Let us consider, as an example, the case when circumstances, including the policy response, are such that $(\pi_{t+1} - \pi_t c_t)$ is positive. The higher R_{t+1} the more the impact of $(\pi_{t+1} - \pi_t c_t)$ on estimated persistence will be scaled down (in equation (7) – with a one-period lead - the residual is scaled by the inverse of the moment matrix). This is, in itself, positive. However, there is an additional effect. The moment matrix is very persistent; if R is large today it will remain large in the near future. That, as we have seen before, is good if estimated persistence is low, but unfortunate if it is high.

We illustrate these arguments using two specific examples. In figure 5, we consider the case of a zero cost-push shock and lagged inflation. This implies $\pi_t=0$ as well. In such a case we have already seen that there will be no immediate updating of estimated persistence. Thus, we manage to isolate the intertemporal effects associated with the moment matrix. Figure 5 (c) plots the partial derivative V_R , for different levels of estimated persistence, against the cost-push shock at time $t+1$. As expected the derivative is positive for low values of c , negative for large values of c . There is an intermediate value for which the partial derivative is zero (which happens to be about 0.34 in our example).

[Insert Figure 5]

Alternatively assume that lagged inflation is still zero but the cost-push shock at time t is such that it implies a one standard deviation movement in inflation at t . Clearly, there will be an updating of estimated persistence. Considering the case when $c=0.5$ figure 5 (a) illustrates the association between V_R and the shock at $t+1$. As expected, from our discussion around figure 5 (c), the partial derivative is mostly negative and only changes signs for relatively large values of the shock. As we conjectured before, the opposite is true when estimated persistence is low, as we can see from figure 5 (b). In this case, the partial derivative is mostly positive and only changes sign for negative (and large) cost push shocks.

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 6]

Figure 6 plots the realization of the average perceived inflation persistence in economies with different gains and two different degrees of price stickiness ($\alpha=0.66$, corresponding to our baseline calibration and a higher degree of price stickiness, $\alpha=0.75$). 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 6 shows that an increasing gain leads to a fall in the average perceived degree of inflation persistence. With a higher gain, agents update their estimates more strongly in response to unexpected inflation developments. As a result, the monetary authority can more easily affect the degree of perceived persistence, which affects the trade-off in favor of lower inflation persistence.

Finally, we look at the impact of increasing the weight on output gap stabilization in the central bank's loss function. Figure 7 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 7]

5. Conclusions

In this paper we look at optimal monetary policy when private sector expectations are determined in accordance with adaptive learning. In the literature on the systematic conduct of monetary policy or, in other words, on monetary policy rules there is a tension between two conceptions of the role of model-based policy rules in deliberation, justification and communication. In the first concept, rules (optimal or otherwise) are regarded as simplifications, approximations, or rules of thumb. By being simple, rules, according to this view, fulfill a pedagogical and steering function, while, at the same time, succeeding in summarizing a large number of, historically determined, particular choices. From such a viewpoint, rules work, as descriptions of how policy is actually conducted, as far as they correctly describe *ex post* the decisions taken by successful central bankers. In the same vein, but more modestly, such rules may identify important considerations that a wise policy-maker always takes into account. According to such concept the characterization of monetary policy plays mainly a *descriptive* role.

In the second concept, general rules are used as benchmarks that can be used to assess particular decisions taken. According to such a view the idiosyncratic can never be salient. The correct viewpoint is to look at particulars as instances of the universal approach. According to such a concept optimal monetary policy rules fulfill a *prescriptive* role¹². We have found that, even in the context of an over-simple model, the characterization of optimal policy becomes very involved. It is easy to imagine how much more difficult such a characterization would become if we would try to reckon the complexity of actual policy choices and the prevalence of economic change. Such considerations clearly limit the possibility of using our framework in a *prescriptive* way.

Therefore, we interpret our contribution in the spirit of the first conception. Specifically, we view optimal monetary policy under adaptive learning as illustrating (once more) why medium term price stability and anchoring inflation expectations is key in environments characterized by endogenous inflation expectations. In our set-up, it is the case that

optimal policy will manage private sector expectations, pushing perceived persistence below endogenous persistence in the model. It does so by responding strongly and persistently to past inflation and inflation shocks when perceived persistence is relatively high. Thus, optimal policy succeeds at stabilizing inflation and at anchoring inflation expectations at no cost – in our example – for output gap stability.

In any case, we think that the contrast between the ability of a simple rule and of optimal policy to maintain stability, when the system is buffeted by stochastic disturbances, brings us back to Machiavelli, who wrote: “[Fortune] shows her power where virtue has not been put in order to resist her and therefore turns her impetus where she knows that dams and dikes have not been made to contain her.”¹³

¹³ *Prince*, chapter 25.

Table 1: Relevant parameters for the benchmark case.

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


Table 2: Summary of macro-economic outcomes

	Rational Expectations		Adaptive Learning	
	Commitment	Discretion	Sophisticated	Simple Rule
Corr(x_t, x_{t-1})	0.66	0	0.51	0
Corr(π_t, π_{t-1})	0.24	0.50	0.36	0.56
Var(x_t)	1	0.95	0.95	0.95
Var(π_t)	1	1.85	1.28	2.18
Var($\pi_t - \gamma\pi_{t-1}$)	1	1.37	1.13	1.48
E[L_t]	1	1.28	1.09	1.37
Var(error)	1	tbc	tbc	tbc


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

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


(a) Estimated inflation persistence




(b) Output gap



(c) Inflation



(d) Quasi-difference of inflation



(e) R – moment matrix






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



(c) Estimated inflation persistence





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



(b) Difference $x(\text{sigm},.)-x(0,.)$





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



(c) Estimated c






Figure 5: V_R for different estimated persistence parameters.
 (a) V_R after positive shock, high estimated persistence parameters.



(b) V_R after positive shock, low estimated persistence parameters.



(c) V_R zero shock, zero lag inflation, different estimated persistence parameters.




Figure 6: Sensitivity analysis: Average estimated persistence in function of the gain and the degree of price stickiness.





Figure 7: 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 Unknown 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)