Optimal Monetary Policy When Agents Are Learning*

Krisztina Molnar† and Sergio Santoro‡

January 24, 2006

Abstract

Most studies of optimal monetary policy under learning rely on optimality conditions derived for the case when agents have rational expectations. In this paper, we derive optimal monetary policy in an economy where the Central Bank knows, and makes active use of, the learning algorithm agents follow in forming their expectations. In this setup, monetary policy can influence future expectations through its effect on learning dynamics, introducing an additional tradeoff between inflation and output gap stabilization. Specifically, the optimal interest rate rule reacts more aggressively to out-of-equilibrium inflation expectations and noisy cost-push shocks than would be optimal under rational expectations: the Central Bank exploits its ability to "drive" future expectations closer to equilibrium. This optimal policy closely resembles optimal policy when the Central Bank can commit and agents have rational expectations. Monetary policy should be more aggressive in containing inflationary expectations when private agents pay more attention to recent data. In particular, when beliefs are updated according to recursive least squares, the optimal policy is time-varying: after a structural break the Central Bank should be more aggressive and relax the degree of aggressiveness in subsequent periods. The policy recommendation is robust: under our policy the welfare loss if the private sector actually has rational expectations is much smaller than if the Central Bank mistakenly assumes rational expectations whereas in fact agents are learning.

JEL classification: C62, D83, D84, E0, E5

Keywords: Optimal Monetary Policy, Learning, Rational Expectations

*We are grateful to Seppo Honkapohja, Albert Marcet and Ramon Marimon for very helpful comments and suggestions. All remaining errors are our own. We would also like to thank Fabrice Collard, Jordi Galí, participants at ASSET 2005 Rethymnon, EEA 2005 Amsterdam, KTI summer workshop Budapest, SMYE 2005 Geneva, and in particular our discussants Stephan Sauer and Ákos Valentinyi.

†Universitat Pompeu Fabra, Barcelona; Email: krisztina.molnar@upf.edu.

‡Universitat Pompeu Fabra, Barcelona; Email: sergio.santoro@upf.edu.

1 Introduction

> *Monetary policy makers can affect private-sector expectations through their actions and statements, but the need to think about such things significantly complicates the policymakers' task.* (Bernanke (2004))

How should optimal monetary policy be designed? A particularly influential framework used in studying this question is the dynamic stochastic general equilibrium economy where money has real effects due to nominal rigidities, sometimes referred to as the "New Keynesian" model. Many papers have explored optimal monetary policy in this framework, under the assumption that both agents and policymakers have rational expectations.[1] More recently, the literature has started to explore the robustness of these optimal policies when some of the assumptions of the standard New Keynesian setup are relaxed.[2] An important aspect of this robustness analysis is to model more carefully the process through which the private sector forms expectations. This issue is particularly relevant given that there is a large body of evidence which suggests that agents' forecasts are not consistent with the paradigm of rational expectations.[3] In response, a growing theoretical literature explores the robustness of the optimal policies, which were derived under rational expectations, when instead agents update their expectations according to a learning algorithm.[4] A typical result in this literature is that interest rate rules that are optimal under rational expectations may lead to instability under learning.

Earlier research uses either ad hoc policy rules, as for example Orphanides and Williams (2005a), or optimality conditions derived under rational expectations, like Evans and Honkapohja (2003b), Evans and Honkapohja (2003a) and Evans and Honkapohja (2002). In this paper, we take a normative approach, and address the issue of how in a New Keynesian setup, a rational Central Bank should optimally conduct monetary policy, if the private sector forms expectations following an adaptive learning model.

We are able to analytically derive optimal monetary policy in our theoretical model. One important feature of the optimal policy is that the Central Bank should act more aggressively towards inflation that what a rational expectations model suggests. Earlier work in the literature that uses ad hoc rules has shown similar results computationally (see Ferrero (2003), Orphanides and Williams

[1] See Clarida, Gali, and Gertler (1999) for a survey on this literature, and Woodford (2003) for an extensive treatise on how to conduct monetary policy *via* interest rate rules.

[2] Wieland (2000a) and Wieland (2000b) look at the effects of parameter uncertainty; Aoki (2002) and Orphanides and Williams (2002) explore monetary policy with for data uncertainty, Levin, Wieland, and Williams (2003) and Hansen and Sargent (2001) study model uncertainty.

[3] See Roberts (1997), Forsells and Kenny (2002) and Adam and Padula (2003).

[4] For an early contribution to adaptive learning applied to macroeconomics, see Cagan (1956), Phelps (1967), for early applications to the Muth market model see Fourgeaud, Gourieroux, and Pradel (1986) and Bray and Savin (1986). The modern literature on this topic was initiated by Marcet and Sargent (1989), who were the first to apply stochastic approximation techniques to study the convergence of learning algorithm. Important earlier contributions to the literature on convergence to the rational equilibrium are Bray (1982) and Evans (1985).

(2005a), Orphanides and Williams (2005b)); here we establish that these results extend to the case when the central bank uses the optimal policy, and provide a formal proof. The intuition for the result is that aggressively driving inflation close to equilibrium helps private agents to learn the true equilibrium value of inflation at a faster pace. As is well-known, even with rational expectations the central bank cares about price stability due to nominal rigidities. When, in addition expectations of nominal variables are sluggish because of learning, our results show that monetary policy should be even more aggressive towards inflation. Being aggressive towards inflation generates a welfare cost in terms of an increased volatility of the output gap. We show analytically that the optimal policy involves a more volatile output gap then the rational expectations benchmark; this holds true even if the Central Bank puts a high weight on output gap stabilization.

A second important feature of the optimal policy is that it is time consistent, and qualitatively resembles the commitment solution under rational expectations in the sense that the optimal policy is unwilling to accommodate noisy shocks. As a consequence the impulse response of a cost push shock is also similar to the commitment case. The contemporaneous impact of a cost push shock on inflation is small (compared to the case of discretionary policy rational expectations), and inflation reverts to the equilibrium in a sluggish manner. In both instances this pattern comes from the Central Bank's (CB) ability to directly manipulate private expectations, even if the channels used are quite different. Under commitment the policy maker uses a *credible promise about the future* to obtain an immediate decline in inflation expectations and thus in inflation; the inertia in the optimal solution is due to the commitments carried over from previous periods. In contrast, under learning the pattern results from the sluggishness of expectations: the CB influences private sector's belief through its *past actions*, and the inertia comes from the past realizations of the endogenous variables. We observe a smaller initial response of inflation relative to the RE discretionary case because optimal policy reacts less to the cost push-shock to ease private agents learning. In this sense, we can say that the ability to manipulate future private sector expectations through the learning algorithm plays a role similar to a commitment device under RE, hence easing the short-run trade-off between inflation and output gap.

An analogous investigation, when the model is characterized by a Phillips Curve à la Lucas and private agents follow a constant gain algorithm is performed in Sargent (1999), Chapter 5. A parallel paper of Gaspar, Smets, and Vestin (2005) provides a numerical solution to optimal monetary policy under constant gain learning in the New Keynesian framework with indexation to lagged inflation among firms. They show that an optimally behaving Central Bank aims to decrease the limiting variance of the private sector's inflationary expectations and show that optimal policy qualitatively resembles the commitment solution under rational expectations. In their framework private agents estimate the persistence of inflation. Another important result they find is that, when the degree of estimated persistence is high the central bank should be more aggressive.

The ability to derive analytical solutions allows us to contribute to this literature in several respects. We derive that optimal policy should be more aggressive when private agents heavily discount past data and place more weight on current data. Under constant gain learning this implies that the incentive to decrease volatility of inflationary expectations is more pronounced when the gain parameter is higher. The intuition behind this is: under constant gain learning expectations remain volatile even in the limit, and this limiting variance is higher with a high gain parameter; this volatility in expectations causes welfare losses even in the limit, so it is optimal to conduct monetary policy against it. We also show that optimal policy at the same time allows for higher volatility in output gap expectations. The reason for this is that optimal policy allows for higher variability of the output gap, which translates to higher volatility of output gap expectations. Of course, allowing a higher variance in output gap also causes welfare losses. We analytically determine the extent to which output gap losses should be tolerated.

Our next contribution is to derive optimal policy under decreasing gain learning. We show that our main results are robust to the changing the gain parameter: (1) optimal policy is aggressive on inflation even at the cost of higher output gap volatility, (2) optimal policy under learning qualitatively resembles optimal policy under rational expectations when the Central Bank is able to commit. A new result is that when beliefs are updated according to a decreasing gain algorithm, the optimal policy is time-varying, reflecting the fact that the incentives for the Central Bank to manipulate agents' beliefs evolve over time. After a structural break, for example the appointment of a new central bank governor, the Central Bank should be more aggressive in containing inflationary expectations and decrease the extent of this aggressiveness in subsequent periods. The intuition for this result is that in the first periods after the appointment of a new governor, agents pay more attention to monetary policy actions (place more weight on current data), therefore an optimally behaving central bank should make active use of this by aggressively driving private sector expectations close to the equilibrium inflation.

Finally, we show that when the Central Bank (CB) is uncertain about the nature of expectation formation (within a set relevant for the US economy) the optimal learning rules derived in our paper are more robust than the time consistent optimal rule derived under rational expectations. Optimal learning rules provide smaller expected welfare losses even if the Central Bank assigns only a very small probability to learning and a very high probability to rational expectations in how it believes the private sector forms its expectations.

The rest of the paper is organized as follows: in Section 2 we analyze optimal policy under constant gain learning where there is no exogenous cost-push shock; in Section 3 we study how the introduction of the cost-push shock affects our results; Section 4 relaxes the assumptions that expectations follow constant gain learning, and show that out main results remain valid under decreasing gain learning; Section 5 conducts policy proposal within a set of private sectors expectation formation; Section 6 relaxes the assumption that the policy maker can perfectly observe the fundamental shocks and the beliefs of the agents;

Section 7 concludes.

2 The Model without a cost push shock

We will consider the baseline version of the New Keynesian model, which is by now the workhorse in monetary economics; in this framework, the economy is characterized by two structural equations[5]. The first one is an IS equation:

$$x_t = E_t^* x_{t+1} - \sigma^{-1}(r_t - E_t^* \pi_{t+1} - \overline{rr}_t) + g_t \tag{1}$$

where x_t, r_t and π_t denote time t output gap[6], short-term nominal interest rate and inflation, respectively; σ is a parameter of the household's utility function, representing the intertemporal elasticity of substitution, g_t is an exogenous demand shock and $\overline{rr}_t$ is the natural real rate of interest, i.e. the real interest rate that would hold in the absence of any nominal rigidity. Note that the operator E_t^* represents the (conditional) agents' expectations, which are not necessarily rational. The above equation is derived by loglinearizing the household's Euler equation, and imposing the equilibrium condition that consumption equals output minus government spending .

The second equation is the so-called New Keynesian Phillips Curve (NKPC):

$$\pi_t = \beta E_t^* \pi_{t+1} + \kappa x_t \tag{2}$$

where β denotes the subjective discount rate, and κ is a function of structural parameters; this relation is obtained by assuming that the supply side of the economy is characterized by a continuum of firms that produce differentiated goods in a monopolistically competitive market, and that prices are staggered à la Calvo (Calvo (1983)) [7].

The model with cost push shock will be examined in the next section. We examine the case without cost push separately, because in this case learning introduces an inflation-output gap tradeoff which is not present under rational expectations (see below).

The loss function of the Central Bank (CB) is given by:

$$E_0 \sum_{t=0}^{\infty} \beta^t (\pi_t^2 + \alpha x_t^2) \tag{3}$$

where α is the relative weight put by the CB on the objective of output gap stabilization[8].

[5] For the details of the derivation of the structural equations of the New Keynesian model see, among others, Yun (1996), Clarida, Gali, and Gertler (1999) and Woodford (2003).

[6] Namely, the difference between actual and natural output.

[7] In other words, in each period firm i can reset the price with a constant probability $1-\theta$, and with probability θ it keeps the same price as in the previous period. If firms take this structure into account when deciding the optimal price it can be shown (See Yun (1996)) that the aggregate inflation is given by (2). κ is decreasing in the level of stickiness, the longer are prices fixed in expectation the smaller the effect of the output gap is on inflation.

[8] As is shown in Rotemberg and Woodford (1998), equation (3) can be seen as a quadratic

2.1 Benchmark: discretionary solution under rational expectations and under learning

A key feature of this model is that, if expectations are rational (i.e., if $E_t^* = E_t$), there is *no trade-off between inflation and output gap stabilization*; in fact, following Gali (2003), we can solve forward equation (2) and impose a boundedness condition on π, obtaining:

$$\pi_t = \kappa E_t \sum_{s=0}^{\infty} \beta^s E_t x_{t+s}$$

Therefore, if the CB stabilizes output gap in every period, under RE inflation will also be equal to zero every period; moreover, this plan is time-consistent, in the sense that the optimal plan chosen by the CB if optimizing at period $t+1$ will be equal to the continuation of the optimal plan set when optimizing at t. The absence of inflation bias is due to the fact that, differently from Barro and Gordon (1983) and all the subsequent literature, the target for output chosen by the CB is the natural level of output, and not a higher level; in other words, the target for output gap is zero, as shown in (3). To restore an inflation stabilization/output gap stabilization trade-off it is necessary to modify the NKPC introducing a so-called cost-push shock[9].

The lack of an inflation-output gap tradeoff can be also seen from the discretionary solution. Under discretion private agents take into account how the monetary policy adjusts its policy, given that the monetary authority is free to reoptimize every period. The discretionary rational expectation equilibrium thus has the property that the Central Bank has no incentive to change its policy (it is time consistent).

Since the Central Bank can not credibly manipulate beliefs, in the optimization it takes expectations as given. The policy problem is to choose a time path for the nominal interest rate r_t [10] to engineer a time path of the target variables π_t and x_t such that the social welfare loss (3) is minimized, subject to the structural equations (1) and (2), and given the private sectors expectations.

$$\begin{aligned} &\min_{\{\pi_t, x_t, r_t\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \beta^t (\pi_t^2 + \alpha x_t^2) \qquad (4) \\ &\text{s.t. } (1), (2) \\ &\qquad E_t\pi_{t+1},\ E_t x_{t+1} \text{ given for } \forall\ t \end{aligned}$$

Because there are no endogenous state variables, problem (4) reduces to a sequence of static optimization problems. As shown in Clarida, Gali, and

approximation to the expected household's utility function; in this case, α is a function of structural parameters.

[9]For a discussion of this point, see Gali (2003).

[10]We have chosen the nominal interest rate to be the instrument variable for easier interpretation (as in real life it is usually a primary instrument of central banks). We could have equally chosen π_t or x_t.

Gertler (1999), and the optimality condition to this problem (at time t) is

$$x_t = -\frac{\kappa}{\alpha}\pi_t \ . \tag{5}$$

Under rational expectations (henceforth RE) solving for the fixed point in expectations results that the Central Bank can set $\pi_t = x_t = 0$ in all periods.

Under non-rational expectations (E^*), using (5) the optimal allocations are:

$$\pi_t = \frac{\alpha\beta}{\alpha+\kappa^2} E_t^* \pi_{t+1} \tag{6a}$$

$$x_t = -\frac{\kappa\beta}{\alpha+\kappa^2} E_t^* x_{t+1} \ . \tag{6b}$$

and solving for r_t from the structural equations (1) and (2) yields

$$r_t = \overline{rr}_t + \delta_\pi^{EH} E_t^* \pi_{t+1} + \delta_x^{EH} E_t^* x_{t+1} + \delta_g^{EH} g_t \tag{7}$$

where:

$$\delta_\pi^{EH} = 1 + \sigma \tfrac{\kappa\beta}{\alpha+\kappa^2}$$
$$\delta_x^{EH} = \sigma$$
$$\delta_g^{EH} = \sigma$$

and E_t^* denote non-rational expectations. Throughout the paper we denote the coefficients by EH referring to the paper Evans and Honkapohja (2003b), where the authors derive a rule analogous to (7)[11]. In the terminology introduced in Evans and Honkapohja (2003b), Evans and Honkapohja (2003a) (EH hereafter), this is an *expectations-based reaction function*. EH show that this rule guarantees not only determinacy under RE, but also convergence to the RE equilibrium when expectations E_t^* evolve according to least squares learning.

However, a rational Central Bank, knowing that private agents follow learning, could do even better. In other words the solution (7) under learning is not a full optimum. In the next section we show how optimal monetary policy is modified when the CB optimizes taking into account its effect on private expectations.

2.2 Constant Gain Learning

We will assume that private sector's expectations are formed according to the adaptive learning literature[12]; in particular, we assume that agents' Perceived Law of Motion (PLM) is consistent with the Law of Motion that the CB would implement under RE: in other words, both inflation and output gap are assumed to be constant, and agents use a learning algorithm to find out this constant.

[11]In particular, Evans and Honkapohja (2003a) derive a rule in a framework where a cost push shock is present.

[12]For an extensive monograph on this paradigm, see Evans and Honkapohja (2001).

Throughout this subsection we will assume that expectations evolve following a constant gain algorithm:

$$E_t^* \pi_{t+1} \equiv a_t = a_{t-1} + \gamma(\pi_{t-1} - a_{t-1}) \tag{8}$$

$$E_t^* x_{t+1} \equiv b_t = b_{t-1} + \gamma(x_{t-1} - b_{t-1}) \tag{9}$$

where $\gamma \in (0, 1)$.

The use of constant gain algorithms to track structural changes is well known from the statistics and engineering literature[13]. Analogously, private agents would be likely to use constant gain algorithms if they confidently believe structural changes to occur. This algorithm implies that past data are geometrically downweighted, in other words agents 'believe more' current data. This approach is closely related to using a fixed sample length, or rolling window regressions.

In Section 4 we will relax this assumption, and examine how optimal policy changes when agents follow decreasing gain learning.

To analyze the optimal control problem faced by the CB, we use the standard Ramsey approach, namely we suppose that the policymakers take the structure of the economy (equations (1) and (2)) as given; moreover, we assume that the CB knows how private agents' expectations are formed, and takes into account its ability to influence the evolution of the beliefs. Hence, the CB problem can be stated as follows:

$$\min_{\{\pi_t, x_t, r_t, a_{t+1}, b_{t+1}\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \beta^t (\pi_t^2 + \alpha x_t^2) \tag{10}$$

$$\text{s.t. } (1), (2), (8), (9)$$
$$a_0, b_0 \text{ given}$$

This optimization problem is linear quadratic, the Bellman equation holds, thus the resulting policy is time consistent [14].

The first order conditions at every $t \geq 0$ are:

$$\begin{aligned} \lambda_{1t} &= 0 && (11)\\ 2\pi_t - \lambda_{2t} + \gamma\lambda_{3t} &= 0 && (12)\\ 2\alpha x_t + \kappa\lambda_{2t} - \lambda_{1t} + \gamma\lambda_{4t} &= 0 && (13)\\ E_t\left[\frac{\beta}{\sigma}\lambda_{1t+1} + \beta^2\lambda_{2t+1} + \beta(1-\gamma)\lambda_{3t+1}\right] &= \lambda_{3t} && (14)\\ E_t\left[\beta\lambda_{1t+1} + \beta(1-\gamma)\lambda_{4t+1}\right] &= \lambda_{4t} && (15) \end{aligned}$$

[13] See for example Benveniste and P. (1990), Part I. Chapters 1. and 4.

[14] A problem solved at t is said to be time consistent for $t+1$ if the continuation from $t+1$ on of the optimal allocation chosen at t solves in $t+1$; moreover, in period zero it is time consistent if the problem in period t is time consistent for $t+1$ for all $t \geq 0$.

where λ_{it}, $i = 1, ..., 4$ denote the Lagrange multipliers associated to (1), (2), (8) and (9), respectively. The necessary conditions for an optimum are the first order conditions, the structural equations (1)-(2) and the laws of motion of private agents' beliefs, (8)-(9). Combining equation (11) and (15), we get:

$$\lambda_{4t} = \beta (1-\gamma) E_t [\lambda_{4t+1}]$$

which can be solved forward, implying that the only bounded solution is:

$$\lambda_{4t} = 0 \tag{16}$$

If we put together equations (11)-(14) and (16), we derive the following optimality condition:

$$\frac{\kappa}{\alpha}\pi_t + x_t = \beta E_t \left[\beta\gamma x_{t+1} + (1-\gamma)\left(\frac{\kappa}{\alpha}\pi_{t+1} + x_{t+1}\right)\right] \tag{17}$$

Inflation-Output Gap Tradeoff

A crucial difference from the rational expectations case is that under learning there is an inflation-output gap tradeoff even without a cost push shock. From equation (2) we can see that, if a_t is different from zero, inflation and output gap cannot be set contemporaneously equal to zero, as in the RE case. Hence, the fact that the expectations are not rational, introduces a trade-off between inflation and output gap stabilization that is not present under RE. In particular, we have the contemporaneous presence of two trade-offs. There is an *intratemporal trade-off* between stabilization of inflation at t and output gap at t, determined by the presence of the nonzero term βa_t in the Phillips Curve (2). There is an additional *intertemporal trade-off* between optimal behavior at t and stabilization of output gap at $t+1$, which is generated by the ability of the CB to manipulate future values of a. This can be seen from iterating forward the optimality condition (17):

$$\frac{\kappa}{\alpha}\pi_t + x_t = \beta^2\gamma E_t \left[\sum_{s=1}^{\infty} [\beta(1-\gamma)]^{s-1} x_{t+s}\right].$$

Hence, for a given positive value of x_t, the optimal disinflation is less harsh with respect to the one implied by (5), provided that future output gaps are also expected to be positive. A smaller deflation in turn guarantees that future inflationary expectations will be closer to the rational expectations equilibrium of inflation, zero.

Let us summarize our first result for later reference:

Result 1. *Learning introduces an intertemporal trade-off not present under rational expectations.*

As a result of the intertemporal trade-off, when the CB can manipulate expectations, it renounces to optimally stabilize the economy in period t, in exchange for a reduction in future inflation expectations that allows an ease in the future inflation output gap trade-off embedded in the Phillips Curve.

Optimal allocations

To derive the optimal allocations, we can use (2) to substitute out x_t in (17), then using the evolution of inflationary expectations (8)we get:

$$\pi_t + \frac{\alpha}{\kappa^2}[\pi_t - \beta a_t] - \beta E_t \left[\frac{\alpha}{\kappa^2}(1-\gamma(1-\beta))[\pi_{t+1} - \beta a_{t+1}] + (1-\gamma)\pi_{t+1} \right] = 0 \tag{18}$$

Hence, at an optimum, the dynamics of the economy can be summarized by stacking equations (8), (9) and (18), obtaining the trivariate system[15]:

$$E_t y_{t+1} = A y_t \tag{19}$$

where $y_t \equiv [\pi_t, a_t, b_t]'$, and:

$$A \equiv \begin{pmatrix} \frac{\kappa^2+\alpha+\alpha\beta^2\gamma(1-\gamma(1-\beta))}{\alpha\beta(1-\gamma(1-\beta))+\kappa^2\beta(1-\gamma)} & -\frac{\alpha\beta(1-\beta(1-\gamma)(1-\gamma(1-\beta)))}{\alpha\beta(1-\gamma(1-\beta))+\kappa^2\beta(1-\gamma)} & 0 \\ \gamma & 1-\gamma & 0 \\ \frac{\gamma}{\kappa} & -\frac{\beta\gamma}{\kappa} & 1-\gamma \end{pmatrix}$$

The three boundary conditions of the above system are:

$$\begin{array}{c} a_0,\, b_0 \text{ given} \\ \lim_{t\to\infty} |E_t \pi_{t+1}| < \infty \end{array} . \tag{20}$$

The last one is due to the fact that, if there exists a solution to the problem (10) when the possible sequences $\{\pi_t, x_t, r_t\}$ are restricted being bounded, then this would also be the minimizer in the unrestricted case[16].

Since A is block triangular, its eigenvalues are given by $1-\gamma$ and by the eigenvalues of:

$$A_1 \equiv \begin{pmatrix} \frac{\kappa^2+\alpha+\alpha\beta^2\gamma(1-\gamma(1-\beta))}{\alpha\beta(1-\gamma(1-\beta))+\kappa^2\beta(1-\gamma)} & -\frac{\alpha\beta(1-\beta(1-\gamma)(1-\gamma(1-\beta)))}{\alpha\beta(1-\gamma(1-\beta))+\kappa^2\beta(1-\gamma)} \\ \gamma & 1-\gamma \end{pmatrix} \tag{21}$$

In the Appendix we show that A_1 has one eigenvalue inside and one outside the unit circle, which implies (together with $(1-\gamma) \in (0,1)$) that we can invoke Proposition 1 of Blanchard and Kahn (1980) to conclude that the system (19)-(20) has one and only one solution. In other words, there exists one and only one stochastic process[17] for each of the three variables of y such that (20) is satisfied. Moreover, note that $y_{1t} \equiv [\pi_t, a_t]'$ does not depend on b_t; therefore, the processes for inflation and a that solve (together with the process for b) the system (19)-(20) are also a solution of the subsystem:

$$E_t y_{1t+1} = A_1 y_{1t}$$

[15] Once we have the equilibrium laws of motion for $[\pi_t, a_t, b_t]$, we can use (1) and (2) to derive the equilibrium r_t and x_t.

[16] For a proof, see the Appendix.

[17] Since the system (19) does not depend on the only source of randomness in this economy (i.e., g), in equilibrium the process followed by the endogenous variables turns out to be deterministic, see below.

together with the boundary conditions:

$$a_0 \text{ given}, \quad \lim_{t\to\infty} |E_t \pi_{t+1}| < \infty$$

Since A_1 has the saddle path property, we can express the equilibrium law of motion for inflation as:

$$\pi_t = c_\pi^{cg} a_t \tag{22}$$

Following the adaptive learning terminology, we call (22) the Actual Law of Motion (ALM) of inflation.

We provide a characterization of c_π^{cg} in the following Proposition:

Proposition 1. *Let c_π^{cg} be the feedback coefficient defined in (22); then, the following holds:*

-if $\gamma \in (0,1)$, we have that $0 < c_\pi^{cg} < \frac{\alpha\beta}{\alpha+\kappa^2}$;

-if $\gamma = 0$, i.e. if expectations are constant, we have that $c_\pi^{cg} = \frac{\alpha\beta}{\alpha+\kappa^2}$.

Proof. See the Appendix. □

Under the optimal policy (OP) a positive a_t increases current inflation, but less than proportionally, since $\frac{\alpha\beta}{\alpha+\kappa^2} < 1$.

As is shown in the Appendix, c_π^{cg} depends on all the structural parameters; in particular, its dependence on the constant gain γ is not necessarily monotonic. In fact, a higher value of γ has two effects on c_π^{cg}: on one hand, it increases the effect of current inflation on future expectations, increasing the incentive for the CB to use this influence (i.e., it would determine a lower c_π^{cg}); on the other hand, it reduces the impact of current expectations on future expectations, thus reducing the benefits from a reduction of the expectations, so that there is an incentive to set a higher c_π^{cg}. In Figure 4 we show a numerical example with the calibration found in Woodford (1996), i.e. with $\beta = 0.99$, $\sigma = 0.157$, $\kappa = 0.024$ and $\alpha = 0.04$; in this case, the first effect dominates, so that c_π^{cg} is a monotonically decreasing function of γ.

Using the structural equation (2) we can derive the optimal allocation of the output gap:

$$x_t = c_x^{cg} a_t \tag{23}$$

where:

$$c_x^{cg} = \frac{c_\pi^{cg} - \beta}{\kappa}$$

$c_\pi^{cg} < \frac{\alpha\beta}{\alpha+\kappa^2}$ (see Proposition 1) implies $c_x^{cg} < -\frac{\kappa\beta}{\alpha+\kappa^2}$; if the private sector expects inflation to be positive, the optimal CB response will imply a negative output gap, i.e. the policymaker will contract economic activity (using the interest rate instrument) in order to attain an actual inflation sufficiently smaller than the expected one. Using (22) and (23) in (1) we can derive the nominal interest rate:

$$r_t = \overline{rr}_t + \delta_\pi^{cg} a_t + \delta_x^{cg} b_t + \delta_g^{cg} g_t \tag{24}$$

where:

$$\delta_\pi^{cg} = 1 - \sigma \frac{c_\pi^{cg} - \beta}{\kappa}$$
$$\delta_x^{cg} = \sigma$$
$$\delta_g^{cg} = \sigma$$

The interest rate rule (24) is an *expectations-based reaction function*, which is characterized by a coefficient on inflation expectations that is decreasing in c_π^{cg}: an optimal ALM for inflation that requires a more aggressive undercutting of inflation expectations (a lower c_π^{cg}) calls for a more aggressive behavior of the CB when it sets the interest rate (a higher coefficient on inflation expectations in the rule (24)). Moreover, the coefficients on b_t and g_t are such that their effects on the output gap in the IS curve are fully neutralized.

Since $c_{\pi,t}^{cg} < \beta$ (see Proposition 1) $\delta_{\pi,t}^{cg}$ is always bigger than 1. In response to a rise in expected inflation optimal policy should raise the nominal interest rate sufficiently to increase the real interest rate. An increase in the real rate has a negative effect on current output; this reflects the intertemporal substitution of consumption. Then a contraction in output will decrease current inflation through the Phillips Curve (2), and consequently through Equation (8) inflationary expectations in the next period will decrease. This criteria is also emphasized in Clarida, Gali, and Gertler (1999) under the discretionary rational expectations solution; since this holds both under RE and learning it provides a very simple criteria for evaluating monetary policy[18].

Asymptotically, the system will converge to the RE equilibrium, with inflation and output gap equal to zero, and so do the corresponding expectations; this can be seen from the autonomous, linear, homogeneous system of first-order difference equations (19). The asymptotic properties of this kind of systems are well-known[19], and with two eigenvalues inside and one outside the unit circle, and the set of boundary conditions (20), we have only one non-explosive solution, which is such that in the long run the system converges to the trivial solution $y_t = 0$.

3 Introduction of a cost-push shock

In this section we will change the model, introducing an additional term in the Phillips Curve, called a cost-push shock[20], so that equation (2) becomes:

$$\pi_t = \beta E_t^* \pi_{t+1} + \kappa x_t + u_t \tag{25}$$

[18]Clarida, Gali, and Gertler (2000) estimate that the pre-Volcker area violated this simple criteria.

[19]See for example Agarwal (2002).

[20]For interpretations of this shock, see among others Clarida, Gali, and Gertler (1999), Erceg, Levin, and Henderson (2000), Woodford (2003).

where $u_t \sim N(0, \sigma_u^2)$ is a white noise[21]. In the New Keynesian literature, this shock is introduced to generate a trade-off between inflation and output gap stabilization; because of this, π_t and x_t cannot be set contemporaneously equal to zero in every period. Moreover, the full commitment solution of the optimal monetary policy under RE turns out to be time inconsistent, even if the CB does not have a target for output gap larger than zero. Hence, the time-consistent discretionary solution will be suboptimal, giving rise to what is sometimes called as stabilization bias. There is, however, a crucial difference with the traditional inflation bias problem: the discretion and the commitment solution are not only different in the coefficients of the equilibrium laws of motion of aggregate variables, but even the functional form of these laws of motion differs between the two cases; in particular, under discretion inflation and output gap are linear functions of the cost-push shock only, under commitment an additional dependence on lagged values of output gap is introduced[22].

3.1 Benchmark: discretionary solution under rational expectations and under learning

As shown in Clarida, Gali, and Gertler (1999), when the cost-push shock is iid the discretionary optimal policy implies that the RE solutions for π_t and x_t are:

$$\pi_t^{RE} = \frac{\alpha}{\kappa^2+\alpha} u_t \tag{26a}$$

$$x_t^{RE} = -\frac{\kappa}{\kappa^2+\alpha} u_t \; . \tag{26b}$$

Using optimality conditions of the discretionary rational expectations problem with non-rational expectations (E_t^*) one can derive the following ALM for inflation and output gap:

$$\pi_t = \frac{\alpha\beta}{\alpha+\kappa^2} a_t + \frac{\alpha}{\alpha+\kappa^2} u_t \tag{27a}$$

$$x_t = -\frac{\kappa\beta}{\alpha+\kappa^2} a_t - \frac{\kappa}{\alpha+\kappa^2} u_t \; . \tag{27b}$$

Evans and Honkapohja (2003b) derives the expectations based interest rate rule that implements this allocation:

$$r_t = \overline{rr}_t + \delta_\pi^{EH} E_t^* \pi_{t+1} + \delta_x^{EH} E_t^* x_{t+1} + \delta_g^{EH} g_t + \delta_u^{EH} u_t \tag{28}$$

where:

$$\begin{aligned} \delta_\pi^{EH} &= 1 + \sigma \tfrac{\kappa\beta}{\alpha+\kappa^2} \\ \delta_x^{EH} &= \sigma \\ \delta_g^{EH} &= \sigma \\ \delta_u^{EH} &= \sigma \tfrac{\kappa}{\alpha+\kappa^2} \; . \end{aligned}$$

[21] Note that the cost-push shock is usually assumed to be an AR(1); we instead assume it to be iid to make the problem more easily tractable, see below.

[22] See Woodford (2003), Clarida, Gali, and Gertler (1999) and McCallum and Nelson (1999).

This rule guarantees determinacy under RE, and also convergence to the discretionary RE equilibrium when expectations E_t^* evolve according to least squares learning[23].

3.2 Constant Gain Learning

At the presence of cost push shocks an additional problem arises in designing the optimal monetary policy when agents are learning: namely which PLM the agents are learning. As we explained above the actual law of motion of the discretion and the commitment solution have different functional forms. For analytical simplicity, in this paper we will restrict our attention to the discretionary case. In particular, we assume that agents believe that inflation and output gap are continuous invariant functions of the cost-push shock only, $\pi_t = \pi(u_t)$ and $x_t = x(u_t)$[24]; this hypothesis, together with the iid nature of the shock, implies that the conditional and unconditional expectations of inflation and output gap coincide, and are perceived by the agents as constants. Hence, it is natural to assume that agents will estimate them using their sample means: the stochastic recursive algorithms (8), (9)[25].

We can now follow a procedure analogous to the one used in the model without cost-push shock. The optimality condition we get is the same as before, Equation (17).

With cost push shocks, there is a well known intratemporal tradeoff between inflation and the output gap present under RE; the presence of learning introduces an additional intertemporal tradeoff (Result 1 holds). We can isolate the two tradeoffs from the optimality condition. When $\gamma = 0$ (which implies constant expectations) (17) implies:

$$\frac{\kappa}{\alpha}\pi_t + x_t = \beta E_t \left[\frac{\kappa}{\alpha}\pi_{t+1} + x_{t+1} \right].$$

This can be solved forward, yielding the unique bounded solution:

$$\frac{\kappa}{\alpha}\pi_t + x_t = 0 \ , \tag{29}$$

which is identical to the optimality condition derived in the RE optimal monetary policy literature when the CB sets the optimal plan taking private sector's expectations as given (i.e., in the discretionary case). Clarida, Gali, and Gertler (1999) describe this relation as implying a 'lean against the wind' policy: in other words, if output gap (inflation) is above target, it is optimal to deflate the economy (contract demand below capacity). Because of the presence of the cost push shock in the Phillips Curve, the Central Bank cannot set π_t and x_t equal to zero every period; so an intratemporal tradeoff between inflation and output gap is present (even in the limit). When $\gamma > 0$ iterating forward (17) show

[23] See Evans and Honkapohja (2003b).

[24] In the terminology of Evans and Honkapohja (2001) Chapter 11, the PLM is a noisy steady state.

[25] To be precise, in the algorithms (8), (9) the observations are weighted geometrically, while in the normal sample average they all receive equal weight.

the presence of an intertemporal tradeoff, just like in Section 2.2. When the current output gap is positive, the Central Bank will decrease inflation less then under RE; the Central Bank renounces to fully stabilize the current economy, in exchange of easing future inflation-output gap tradeoffs.

Stacking (17) with the Phillips Curve (25) and the algorithms (8)-(9), we can show that at the optimum the economy evolves according to:

$$\begin{pmatrix} E_t\pi_{t+1} \\ a_{t+1} \\ b_{t+1} \end{pmatrix} = A \begin{pmatrix} \pi_t \\ a_t \\ b_t \end{pmatrix} + \begin{pmatrix} -\frac{\alpha}{\alpha\beta(1-\gamma(1-\beta))+\kappa^2\beta(1-\gamma)} \\ 0 \\ -\frac{\gamma}{\kappa} \end{pmatrix} u_t \tag{30}$$

(where A is defined as in the previous section), plus the boundary conditions (20). The system (30)-(20) is in the form studied in Blanchard and Kahn (1980), so that we can use their results. In particular, since there are two predetermined variables and one non-predetermined, and A has one eigenvalue outside the unit circle and two inside, there exists one and only one solution. Moreover, also the system:

$$\begin{pmatrix} E_t\pi_{t+1} \\ a_{t+1} \end{pmatrix} = A_1 \begin{pmatrix} \pi_t \\ a_t \end{pmatrix} + \begin{pmatrix} -\frac{\alpha}{\alpha\beta(1-\gamma(1-\beta))+\kappa^2\beta(1-\gamma)} \\ 0 \end{pmatrix} u_t \tag{31}$$

(where A_1 is defined as in the previous section) respects the Blanchard-Kahn conditions for existence and uniqueness of a (bounded) solution, and this unique solution can be written as[26]:

$$\pi_t = c_\pi^{cg} a_t + d_\pi^{cg} u_t \tag{32}$$

Combining $E_t\pi_{t+1} = c_\pi^{cg} a_{t+1}$ with the optimality condition (17) and the Phillips Curve (25), and making use of the law of motion of inflation expectations (8), we derive the values of the coefficients c_π^{cg} and d_π^{cg}, which are summarized in the next Proposition.

Proposition 2. *Let the economy evolve according to the system (30), (20); then the ALM for inflation is:*

$$\pi_t = c_\pi^{cg} a_t + d_\pi^{cg} u_t$$

where c_π^{cg} is the same given in Proposition 1, and:

$$d_\pi^{cg} = \frac{\alpha}{\kappa^2 + \alpha + \alpha\beta^2\gamma^2(\beta - c_\pi^{cg}) + \beta\gamma(1-\gamma)\left(\alpha\beta - (\kappa^2+\alpha)\, c_\pi^{cg}\right)}$$

The ALM for output gap and the interest rate rule are given by:

$$x_t = c_x^{cg} a_t + d_x^{cg} u_t \tag{33}$$

$$r_t = \overline{rr}_t + \delta_\pi^{cg} a_t + \delta_x^{cg} b_t + \delta_g^{cg} g_t + \delta_u^{cg} u_t \tag{34}$$

[26]See Blanchard and Kahn (1980), Proposition 1.

where $c_x^{cg}, \delta_\pi, \delta_x^{cg}, \delta_g^{cg}$ are the same as in (24), and:

$$d_x^{cg} = \frac{d_\pi^{cg}-1}{\kappa}$$
$$\delta_u^{cg} = -\sigma\frac{d_\pi^{cg}-1}{\kappa}$$

Plugging (32) into (8), we get:

$$\begin{aligned} a_{t+1} &= a_t + \gamma(c_\pi^{cg} - 1)a_t + \gamma d_\pi^{cg} u_t \\ &= \left(1-\gamma(1-c_\pi^{cg})\right) a_t + \gamma d_\pi^{cg} u_t \end{aligned}$$

which is a stationary[27] AR(1); thus, as is well-known in the literature on adaptive learning, the contemporaneous presence of random shocks in the ALM and of constant gain specification of the updating algorithm, prevents the expectations from converging asymptotically to a precise value: instead, we have that $a_t \sim N\left(0, \frac{\gamma^2 (d_\pi^{cg})^2}{1-(1-\gamma(1-c_\pi^{cg}))^2}\sigma_u^2\right)$.

3.3 Comparison with the myopic rule

In this section we state results regarding how optimal monetary policy under constant gain learning differs from myopic rules used earlier in the literature, where myopic means a rule that considers expectations as given in the optimization problem: in particular we refer to rule (28), derived in Evans and Honkapohja (2003b) (henceforth EH).

It is clear that the coefficients on the output gap expectations and on the demand shock are the same in rule (28) as in rule (34), while the other two coefficients are typically different. Proposition 1 implies $\delta_{\pi,t}^{cg} > \delta_\pi^{EH}$: the interest rate response of OP to out-of-equilibrium inflation expectations is more aggressive than the interest rate response of EH. This is due to the fact that when the CB takes into account its ability to influence agents' beliefs, it optimally chooses to undercut future inflation expectations more than what a myopic CB would do.

From Proposition 1 and 2 it also follows that $\delta_{u,t}^{cg} > \delta_u^{EH}$: optimal policy reacts more aggressively also to cost push shocks. After a positive cost push shock the optimally behaving Central Bank raises the interest rate more aggressively than the myopic one, this in turn decreases output, which has a negative effect on inflation. Thus conducting an aggressive interest rate rule in response to the cost push shock, decreases the influence of the cost push shock on inflation, and this in turn will ease agents learning about the true equilibrium level of inflation.

The inflation and output gap allocations implemented by the two different interest rate rules are also different. Under constant gain learning optimal allocations are characterized by (32) and (33). Under EH allocations are given by (27) with $E_t^* \pi_{t+1} = a_t$.

[27] In fact, since $0 < c_\pi^{cg} < 1$, it immediately follows that $0 < \left(1-\gamma(1-c_\pi^{cg})\right) < 1$.

From Proposition 1 we know that the feedback coefficient under optimal policy c_π^{cg} is smaller than under the EH rule, in order to undercut inflation expectations more. Also the response to the cost push shock is of lesser magnitude when (34) is used instead of (28) (in fact, $c_\pi^{cg} < \frac{\alpha\beta}{\kappa^2+\alpha}$ implies that $d_\pi^{cg} < \frac{\alpha}{\kappa^2+\alpha}$), because the CB is less willing to accommodate noisy shocks, in order to make easier for the private sector to learn what is the long-term value of the conditional expectations of inflation.

Under OP both coefficients in the ALM of x_t are higher in absolute value than under EH. This implies that the CB allows a higher feedback from out of equilibrium expectations and noisy cost push shocks to the output gap then a myopic policymaker.

The difference between OP and the myopic policy can be summarized as follows:

Result 2. *When the CB takes into account its influence on private agents learning it is optimal to decrease the effect of out of equilibrium expectations on inflation (engineering an aggressive interest rate reaction to inflationary expectations) and increase the effect of out of equilibrium expectations on the output gap compared to the myopic policy.*

This way optimal policy undercuts future private sector expectations more aggressively than the myopic policy.

Result 3. *When agents are learning an optimally behaving policymaker accommodates less the effect of noisy shocks to inflation compared to a myopic policymaker, even if it translates into a more volatile output gap.*

This way optimal policy makes it easier for the private sector to learn what is the "true" value of the conditional expectations of inflation.

Similarity to the commitment solution

From Result 2 and 3 it follows that the impact of a given nonzero cost push shock drives inflation (output gap) closer to (further from) target when agents are learning, relative to the discretionary RE case. Interestingly, this behavior *qualitatively resembles the optimal RE equilibrium under commitment* within a simple class of policy rules derived in Clarida, Gali, and Gertler (1999): if the CB can commit to a policy rule that is a linear function of u_t, the solution can be characterized, when compared to the discretionary equilibrium, by inequalities analogous to the ones summarized in the results stated above. However, the (constrained) commitment solution differs from the discretionary one only when the cost-push shock is an AR(1); if u -and consequently, the equilibrium processes for inflation and output gap- is iid, the two solutions coincide, since future (rational) expectations of the agents cannot be manipulated by the CB. Instead, if expectations are backward-looking, the future beliefs can be manipulated also when the shock is iid: the current actions of the CB influence future beliefs through (8) and (9) even if the shock is iid.

In both instances this behavior results from the CB's ability to directly manipulate private expectations, even if the channels used are quite different. In fact, under commitment the policy maker uses a *credible promise on the future* to obtain an immediate decline in inflation expectations and thus in inflation. Under learning we observe a smaller initial response of inflation relative to the RE discretionary case because optimal policy reacts less to the cost push-shock to ease private agents learning. In this sense, we can say that the ability to manipulate future private sector expectations through the learning algorithm plays a role similar to a commitment device under RE, hence easing the short-run trade-off between inflation and output gap.

Another similarity to the commitment solution is the sluggish behavior of inflation after an initial cost push shock. The source of inertia under RE commitment and learning is quite different. Under commitment the policy maker carries commitments made in the past (in other words commits to behave in a past dependent way). Under learning the pattern results from the sluggishness of expectations.

As a result of these two similarities, the impulse response function of inflation to a cost push shock will be also similar under OP and RE commitment. Figure 5 displays the impulse response function of inflation to a unit shock under OP and discretionary RE policy. In the optimal RE discretionary policy, inflation rises on impact and immediately reverts to the steady state once the shock dies out. Instead, under learning the policy maker engineers a smaller initial response of inflation; in subsequent periods inflation gradually converges back to the steady state value. Clarida, Gali, and Gertler (1999) and Gali (2003) show a *similar disinflation path for the Ramsey policy*: a smaller initial inflation compared to the discretionary case, in exchange for a more persistent deviation from the steady state later[28] This behavior of Ramsey policy leads to welfare gains over discretion due to the convexity of the loss function; this preference for slower but milder adjustment to shocks is at the heart of the stabilization bias.

The similarity to the RE commitment solution resembles the analysis carried out in Sargent (1999), Chapter 5, which shows that in the Phelps problem under adaptive expectations[29], the optimal monetary policy drives the economy close to the Ramsey optimum. Moreover, when the discount factor β equals 1, optimal policy under learning replicates the Ramsey equilibrium. In our case, optimal policy under learning cannot replicate the commitment solution even for β going to 1. This result follows from the particular nature of the gains from commitment; commitment calls for an ALM with a different functional form to the discretionary case[30]. In the Phelps problem, on the other hand, the Phillips Curve is such that the discretion and commitment outcome of inflation has the same functional form, but different coefficients. However, also in our

[28] A difference is that commitment policy under RE engineers a sequence of negative inflation after the first period, while a positive sequence under learning.

[29] Phelps (1967) formulated a control problem for a natural rate model with rational Central Bank and private agents endowed with a mechanical forecasting rule, known to the Central Bank.

[30] See Clarida, Gali, and Gertler (1999).

case an increase in the discount factor makes the optimal disinflationary path under learning getting closer to the commitment solution. This can be seen in Table 1, where we summarize the behavior of inflation in response to a unit cost push shock when the model's parameters are calibrated as in Woodford (1996), apart from β which takes several values. As β goes to 1 the initial response of inflation is milder and the path back to the steady state longer.

Table 1: Paths of inflation for different βs after an initial cost push shock

beta	0.5	0.6	0.7	0.8	0.9	1.0
1	0.99	0.99	0.98	0.98	0.96	0.91
2	0.44	0.52	0.61	0.69	0.75	0.73
3	0.24	0.33	0.44	0.55	0.66	0.66
10	0.00	0.01	0.04	0.12	0.27	0.33
50	0.00	0.00	0.00	0.00	0.00	0.01

Woodford (1996) calibration. Cost push shock $u_0 = 1$ in the first period, starting from $a_0 = 0$, $\pi_0 = 0$, $x_0 = 0$, with $\gamma = 0.2$

Welfare Loss Analysis

To have a quantitative feeling of the welfare gains that the use of the optimal rule (34) instead of the EH rule (28) implies, we present a numerical welfare loss analysis.

Let us define the cumulative ext-post losses up to time T under the two interest rate rules as:

$$L_T^{OP} \equiv \sum_{t=0}^{T} \beta^t \left(\left(\pi_t^{OP}\right)^2 + \alpha \left(x_t^{OP}\right)^2 \right)$$

and:

$$L_T^{EH} \equiv \sum_{t=0}^{T} \beta^t \left(\left(\pi_t^{EH}\right)^2 + \alpha \left(x_t^{EH}\right)^2 \right)$$

where the superscripts OP and EH indicates whether the variables are calculated using rule (34) or (28), respectively; then, for a cross sectional sample size of 1000 we compute numerically the value of $\widehat{L}_T \equiv \sum_{i=0}^{1000} L_T^{OP} / \sum_{i=0}^{1000} L_T^{EH}$ for T=10000. We use the calibrations of McCallum and Nelson (1999) (McCN), Woodford (1996) (W) and Clarida, Gali, and Gertler (2000) (CCG). The calibrated coefficients are as follows: in McCN $\kappa = 0.3$, $\alpha = 1.83$, in the Woodford calibration $\kappa = 0.024$, $\alpha = 0.048$ and in the CCG calibration $\kappa = 0.075$, $\alpha = 0.3$ [31] [32]. The initial values for expectations are $a_0 = 0$, and the shocks are drawn

[31] We adjust the CCG calibration for quarterly data, i.e. both the σ and κ values reported by Clarida, Gali, and Gertler (2000) are divided by 4. We would like to thank Seppo Honkapohja for drawing our attention on this difference in calibrations.

[32] The risk aversion parameter σ does not appear in the reduced form for inflation and

from a standardized normal. In all three calibrations $\beta = 0.99$.

Results in Table 2 show that the gain in welfare losses is especially high for constant gain learning with high tracking parameters: for $\gamma = 0.9$ the welfare loss of not using the optimal rule is twice as large as under OP. The intuition behind follows from the fact that, in the presence of a cost push shock, constant gain learning does not settle down to RE, but converges to a limiting distribution. Optimal policy should takes this into account and aims to decrease the limiting variance, while a myopic policy does not[33].

Table 2: Ratio of welfare losses using OP and EH under constant gain $\widehat{L} = L^{OP}/L^{EH}$

Tracking parameter	McCN	W	CCG
0.1	0.82	0.85	0.81
0.2	0.72	0.72	0.68
0.3	0.64	0.62	0.58
0.4	0.61	0.58	0.54
0.5	0.57	0.53	0.50
0.6	0.54	0.50	0.47
0.7	0.52	0.48	0.45
0.8	0.50	0.46	0.43
0.9	0.49	0.45	0.42

McCN: McCallum and Nelson (1999), W: Woodford (1996)
CCG:Clarida et al. (2000)

An increase in the tracking parameter (keeping everything else constant) results in a larger variance of inflation expectations and, consequently, in a larger opportunity cost of adopting a suboptimal rule of the form (28). This is illustrated in Figure 7, which shows that the higher is γ, the higher is the decrease in variance of a under OP compared to EH.

Optimal policy engineers a decrease in the limiting variance of inflationary expectations (Figure 7) at the cost of allowing higher variance in output gap expectations (Figure 8). This result is parallel to results 2 and 3, which both state that optimal policy should focus on decreasing inflation variation even at the cost of higher output gap variation. Table 3 shows that under constant gain learning this incentive to focus on inflation is even more pronounced then under decreasing gain learning. The higher is the tracking parameter, the higher is the limiting variance of expectations and the more incentive the Central Bank has to focus on low variance in inflation allowing for an increase in output gap deviation from the flexible price equilibrium. For $\gamma = 0.9$ the Central Bank

output gap, hence it is not calibrated whatsoever. σ used for the plots of interest rate rule coefficients are: CCG $\sigma = 1/4$, W $\sigma = 0.157$, McCN $\sigma = 1/0.164$.

[33]It is worth noting that the EH rule is designed to ensure learnability of the optimal RE in a decreasing gain environment, and its performance under constant gain is never considered in the EH paper; however, it can be useful to employ a constant gain version of their rule to illustrate potential advantages of fully optimal monetary policy over a myopic rule.

engineers a 75 percent lower welfare loss in inflation when it properly conditions on expectation formation, permitting at the same time 5-10 times more variation in output gap.

Table 3: Ratio of welfare losses using OP and EH under constant gain learning due to inflation and output gap variations

$\widehat{L} = L^{OP}/L^{EH}$

	Inflation			Output gap		
Tracking parameter	McCN	W	CCG	McCN	W	CCG
0.1	0.67	0.72	0.67	3.76	8.37	8.22
0.2	0.54	0.54	0.51	4.34	10.79	9.89
0.3	0.44	0.43	0.40	4.68	11.48	10.36
0.4	0.39	0.38	0.35	5.02	12.08	10.88
0.5	0.34	0.33	0.30	5.09	11.98	10.77
0.6	0.31	0.30	0.28	5.17	12.00	10.77
0.7	0.29	0.28	0.25	5.21	11.94	10.72
0.8	0.27	0.26	0.24	5.21	11.77	10.57
0.9	0.25	0.25	0.23	5.28	11.98	10.75

McCN: McCallum and Nelson (1999), W: Woodford (1996), CCG: Clarida et al. (2000)

Moreover, it is worth noting that the use of a myopic rule under constant gain learning allows for the autocorrelation of inflation to rise, thus increasing the persistence of a shock's effect on inflation expectations. This problem arises from the relatively weak response to inflation expectations which feeds back to current inflation and, in turn, into subsequent expectations and inflations. The optimal rule's strong feedback to inflation expectations dampens this interaction between inflation and expectations[34].

This section has shown that optimal policy under learning is characterized by a more aggressive interest rate reaction to out-of-equilibrium expectations and to the cost push shock than would be optimal when the Central Bank does not make active use of its influence on expectations. This aggressive behavior guarantees that inflation will deviate less from its equilibrium value, thus private agents can learn the true equilibrium level of inflation faster than under myopic policy. Helping inflationary expectations is beneficial, even at the cost of allowing higher deviations in output gap expectations and a higher output gap volatility. Our welfare loss analysis has shown that properly conditioning on private agents expectation formation is especially important in a nonconvergent environment, i.e. when agents follow constant gain learning. Welfare gains from

[34]It can be easily derived that the autocorrelation of inflation under constant gain with EH is $E\pi_t^{EH}\pi_{t-1}^{EH} = \left(\frac{\alpha\beta}{\alpha+\kappa^2}\right)^2 \left(1-\gamma+\gamma\frac{\alpha\beta}{\alpha+\kappa^2}\right)\sigma^2_{a_{EH}} + \frac{\alpha\beta}{\alpha+\kappa^2}\left(\frac{\alpha}{\alpha+\kappa^2}\right)^2 \gamma\sigma_u^2$ while under the optimal rule $E\pi_t^{OP}\pi_{t-1}^{OP} = \left(c_\pi^{cg}\right)^2\left(1-\gamma+\gamma c_\pi^{cg}\right)\sigma^2_{a_{OP}} + c_\pi^{cg}\left(d_\pi^{cg}\right)^2\gamma\sigma_u^2$. We have already seen that $\sigma^2_{a_{OP}} < \sigma^2_{a_{EH}}$, $c_\pi^{cg} < \frac{\alpha\beta}{\alpha+\kappa^2}$ and $d_\pi^{cg} < \frac{\alpha}{\alpha+\kappa^2}$, thus $E\pi_t^{OP}\pi_{t-1}^{OP} < E\pi_t^{EH}\pi_{t-1}^{EH}$.

using the optimal policy are particularly pronounced when private agents use a high tracking parameter (i.e. discount more past data) for forecasting.

4 Decreasing Gain Learning

In this section we relax the assumption of constant gain learning and show that our main results remain valid also with decreasing gain learning (henceforth DG) and show that the time varying nature of expectations imply that during the transition the optimal policy should be time varying even in a stationary environment.

Using a constant gain parameter γ is appropriate when agents believe structural changes to occur. If instead the private sector confidently believes that the environment is stationary it is more reasonable to model their learning behavior with a decreasing gain rule, namely an algorithm of the form:

$$E_t^* \pi_{t+1} \equiv a_t = a_{t-1} + t^{-1}(\pi_{t-1} - a_{t-1}) \tag{35}$$

$$E_t^* x_{t+1} \equiv b_t = b_{t-1} + t^{-1}(x_{t-1} - b_{t-1}) \tag{36}$$

where the only difference with (8)-(9) is the substitution of γ with t^{-1}.

An updating scheme of this form is equivalent[35] to estimating inflation and output gap every period with OLS[36].

4.1 Without cost push shock

Let us first consider the economy without cost push shock. Then the problem of the CB becomes:

$$\min_{\{\pi_t, x_t, r_t, a_{t+1}, b_{t+1}\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \beta^t (\pi_t^2 + \alpha x_t^2) \tag{37}$$

$$\text{s.t. } (1), (2), (35), (36)$$
$$a_0, b_0 \text{ given}$$

The optimization can be solved in a way analogous to the constant gain case; hence, the dynamics of the system can be summarized by the optimality condition:

$$\frac{\kappa}{\alpha}\pi_t + x_t = \beta E_t \left[\beta \frac{1}{t+1} x_{t+1} + \frac{\kappa}{\alpha}\pi_{t+1} + x_{t+1} \right] \tag{38}$$

Iterating it forward we get:

[35] Under certain conditions on the values used to initialize the algorithm, see Evans and Honkapohja (2001).

[36] Note that, since inflation and output gap are assumed by the learners to be constant, the OLS is just the sample averages of the two.

$$\frac{\kappa}{\alpha}\pi_t + x_t = E_t\left[\sum_{s=1}^{\infty}\beta^{s+1}\frac{1}{t+s}x_{t+s}\right] .$$

Similarly to Section 2 our result is that learning introduces an an intratemporal tradeoff between inflation and output that is not present under RE in an economy without a cost push shock and an additional intertemporal tradeoff that is not present in general under rational expectations (Result 1). From the latter it follows that during the transition for a given positive value of x_t, the optimal disinflation is less harsh with respect to the one implied by (5) (optimizing taking expectations as given) provided that the series on the right hand side is expected to be positive. The intuition behind is that when the CB makes active use of the expectation formation, it renounces its ability to optimally stabilize the economy in period t, in exchange for a reduction in future inflation expectations (in absolute value) and this allows an ease in the future inflation-output gap trade-off embedded in the Phillips Curve.

To derive the optimal allocations, we can use (2) to substitute out x_t in (38), then using the evolution of inflationary expectations (35)we get:

$$E_t\left[\pi_{t+1}\right] = A_{11,t}\pi_t + A_{12,t}a_t, \tag{39}$$

where:

$$\begin{aligned} A_{11,t} &\equiv \frac{\kappa^2 + \alpha + \alpha\beta^2\frac{1}{t+1}\left(1+\beta\frac{1}{t+1}\right)}{\alpha\beta(1+\beta\frac{1}{t+1}) + \kappa^2\beta} \\ A_{12,t} &\equiv -\frac{\alpha\beta\left[1-\beta\left(1-\frac{1}{t+1}\right)\left(1+\beta\frac{1}{t+1}\right)\right]}{\alpha\beta(1+\beta\frac{1}{t+1}) + \kappa^2\beta} . \end{aligned}$$

Hence, at an optimum, the dynamics of the economy can be summarized by stacking equations (35), (36) and (39), and obtaining the trivariate system:

$$E_t y_{t+1} = A_t y_t \tag{40}$$

where $y_t \equiv [\pi_t, a_t, b_t]'$, and:

$$A_t \equiv \begin{pmatrix} A_{11,t} & A_{12,t} & 0 \\ \frac{1}{t+1} & 1-\frac{1}{t+1} & 0 \\ \frac{\frac{1}{t+1}}{\kappa} & -\frac{\beta\frac{1}{t+1}}{\kappa} & 1-\frac{1}{t+1} \end{pmatrix} .$$

The three boundary conditions of the above system are (20), the same as in section 2.2.

We can find the solution with the method of undetermined coefficients, with the guess[37]:

$$\pi_t = c^{dg}_{\pi,t} a_t. \tag{41}$$

The sequence $\left\{ c^{dg}_{\pi,t} \right\}$ must satisfy the non-linear, non-autonomous first order difference equation:

$$c^{dg}_{\pi,t} = \frac{c^{dg}_{\pi,t+1}\left(1 - \frac{1}{t+1}\right) - A_{12,t}}{A_{11,t} - c^{dg}_{\pi,t+1}\frac{1}{t+1}} \tag{42}$$

Of course, there exist infinite sequences that satisfy equation (42), one for each initial value $c^{dg}_{\pi,0}$. However, since the boundary conditions require π_t to stay bounded, we will concentrate on the solutions that do not explode.

Proposition 3. *Let* $\left\{ c^{dg}_{\pi,t} \right\}$ *be defined by* (42)*, and assume it is bounded; then,* $\lim_{t\to\infty} c^{dg}_{\pi,t}$ *exists, and is given by:*

$$\lim_{t\to\infty} c^{dg}_{\pi,t} = \frac{\alpha\beta}{\alpha + \kappa^2}$$

Moreover, for any $t < \infty$*, we have:*

$$c^{dg}_{\pi,t} < \frac{\alpha\beta}{\alpha + \kappa^2}$$

Proof. See the Appendix. □

Thus Result 2 holds during the transition: when the CB takes into account its influence on expectations it is optimal to decrease the effect of out-of-equilibrium expectations on inflation compared to the myopic policy (see (6)), in order to undercut future inflation expectations by a larger amount. This relaxes the future inflation-output gap trade-off embedded in the Phillips Curve. The ALM for output gap is:

$$x_t = c^{dg}_{x,t} a_t \qquad c^{dg}_{x,t} = \frac{c^{dg}_{\pi,t} - \beta}{\kappa} \tag{43}$$

If the private sector expects inflation to be positive, the optimal CB will contract economic activity more than EH[38] (using the interest rate instrument); the CB is ready to pay a short-term cost represented by a wider current output gap in order to contain future inflationary expectations.

[37] This guess corresponds to the unique solution under constant gain learning. A proof of uniqueness of a bounded solution for decreasing gain learning is not worked out completely yet.

[38] From $c^{dg}_{\pi,t} < \frac{\alpha\beta}{\alpha+\kappa^2}$ it follows that $c^{dg}_{x,t} < -\frac{\kappa\beta}{\alpha+\kappa^2}$. Compare with the ALM under EH (6).

The nominal interest rate rule is as follows:

$$r_t = \overline{rr}_t + \delta^{dg}_{\pi,t} a_t + \delta^{dg}_x b_t + \delta^{dg}_g g_t \tag{44}$$

where:

$$\delta^{dg}_{\pi,t} = 1 - \sigma \frac{c^{dg}_{\pi,t} - \beta}{\kappa}$$
$$\delta^{dg}_x = \sigma$$
$$\delta^{dg}_g = \sigma$$

Since $c^{dg}_{\pi,t} < \beta$ (see Proposition 3) $\delta^{dg}_{\pi,t}$ is always bigger than 1. In response to a rise in expected inflation optimal policy should raise the nominal interest rate sufficiently to increase the real interest rate. The following proposition pertains to the characteristics of the optimal rule compared to the myopic EH rule (7):

Proposition 4. *Assume that $t < \infty$; then, $\delta^{dg}_{\pi,t} > \delta^{EH}_{\pi}$. Moreover, we have:*
- $\lim_{t\to\infty} \delta^{dg}_{\pi,t} = \delta^{EH}_{\pi}$.

Result 4 under CG is parallelled by our results under DG: the optimal interest rate rule should react more aggressively to out of equilibrium expectations than the EH rule. A CB that knows how its behavior affects private sector expectations should contain more inflationary expectations than a CB that takes expectations as given.

An interesting result is that the coefficient on inflation expectations in the interest rate rule (44) is *time-varying*, reflecting the fact that the Central Bank's incentives to manipulate agents' beliefs evolve over time. This implies that during the transition optimal policy should be time varying even in a stationary environment.

In Figure 1, we show how this coefficient depends on time when the parameters are calibrated according to Clarida, Gali, and Gertler (2000): $\kappa = 0.075$, $\alpha = 0.3$, $\sigma = 1/4$. $\delta^{dg}_{\pi,t}$ is always above its limiting level (see analytical proof in Proposition 4), moreover, it decreases over time. Numerical analysis on the grid $\beta = 0.99$ and $\alpha \in [0.01, 2]$, $\kappa \in [0.01, 0.5]$ shows that this decreasing behavior of $\delta^{dg}_{\pi,t}$ is a robust feature of the model [39]. We find that after the 4th period (from the 4th to the 5th period and so on) $\delta^{dg}_{\pi,t}$ is always decreasing, while in the first 4 periods $\delta^{dg}_{\pi,t}$ might be increasing (hump-shaped) for a combination of low values of α and high values of κ (see Figure 3) [40]. We summarize our new results as:

Result 4. *Optimal policy is time varying even in a stationary environment. It is more aggressive initially, and as time evolves reacts less and less aggressively to out of equilibrium expectations.*

[39] We have chosen the grid to include typical calibrated values for the US and the EURO area.

[40] In fact, $\delta^{dg}_{\pi,t}$ is always decreasing also for other calibrations widely adopted in the New Keynesian Literature, like those taken from Clarida, Gali, and Gertler (2000) and McCallum and Nelson (1999).

To get an intuition, suppose that a structural break occurs. For example there is a policy change because a new central bank governor is appointed, agents know that monetary policy has changed and try to learn how this affects the equilibrium. In this situation is convenient for the CB to react more aggressively to out-of-equilibrium inflation beliefs in the first periods, when agents pay more attention to new information and the CB's possibilities of influencing private expectations are greater. This behavior is beneficial even at the cost of larger short-term losses in terms of output gap variability. As time passes, the expectations will be influenced to a lesser extent by the last realization of inflation, hence determining a CB reaction that closely resembles the optimizing behavior when policymakers cannot manipulate expectations.

The asymptotic behavior of inflation beliefs is given by the following Proposition:

Proposition 5. *Let* $\pi_t = c^{dg}_{\pi,t} a_t$*, where* $c^{dg}_{\pi t}$ *is given by (42); then,* $a_t \to 0$.

Proof. See the Appendix. □

Combining this result with the boundedness of $c^{dg}_{\pi,t}$, the ALM for inflation (41) and output gap (43) tell us that both these variables go to zero asymptotically, restoring the RE allocations. Optimal policy naturally chooses a non-explosive solution (it is E-stable), and drives expectations to the rational expectations equilibrium.

Note that the policy function does not depend on the period when the cb optimizes, even if it is not time invariant. Thus, the optimal policy characterized above is time consistent, in the sense of Lucas and Stokey (1983) and Alvarez, Kehoe, and Neumeyer (2004).

4.2 With Cost Push Shock

In this section we introduce a cost push shock in the New Keynesian Philips Curve.

Proceeding with the same analysis as before, we get the same optimality condition (38). Substituting out x_t using the Phillips Curve (25), and using the evolution of inflation expectations (35) we can show that at the optimum the economy evolves according to:

$$E_t\left[\pi_{t+1}\right] = A_{11,t}\pi_t + A_{12,t}a_t + P_{1,t}u_t \tag{45}$$

where A_{11} and A_{12} are the same as in (39), and

$$P_{1,t} \quad \equiv \quad -\frac{\alpha}{\alpha\beta(1+\beta\frac{1}{t+1})+\kappa^2\beta} \ .$$

Stacking together this condition, and the two learning algorithms (35) and (36), we again have a trivariate system. We can guess and verify that the ALM for inflation is of the form:

$$\pi_t = c^{dg}_{\pi,t}a_t + d^{dg}_{\pi,t}u_t \tag{46}$$

which implies that $E_t\pi_{t+1} = c^{dg}_{\pi,t+1}a_{t+1}$; substituting this expression in (45), and making use of the law of motion of inflation expectations (35), we obtain that the sequence $\left\{c^{dg}_{\pi,t}\right\}$ is identical to (42) and $\left\{d^{dg}_{\pi,t}\right\}$ must satisfy:

$$d^{dg}_{\pi,t} = \frac{P_{1,t}}{c^{dg}_{\pi,t+1}\frac{1}{t+1} - A_{11,t}}. \tag{47}$$

The solution of $\left\{c^{dg}_{\pi,t}\right\}$ is again characterized by Proposition 3. From Proposition 3 and (47) it follows that $0 < d^{dg}_{\pi,t} < \frac{\alpha}{\alpha+\kappa^2}$, a positive cost push shock increases inflation, but less than under myopic policy (compare with (27)). The ALM for output gap and the nominal interest rate rule are given by:

$$x_t = c^{dg}_{x,t}a_t + d^{dg}_{x,t}u_t \tag{48}$$

$$r_t = \overline{rr}_t + \delta^{dg}_{\pi,t}a_t + \delta^{dg}_x b_t + \delta^{dg}_g g_t + \delta^{dg}_{ut}u_t \tag{49}$$

where $c^{dg}_{x,t}, \delta^{dg}_{\pi,t}, \delta^{dg}_x, \delta^{dg}_g$ are the same as in (44), and

$$\begin{aligned} d^{dg}_{x,t} &= \tfrac{d^{dg}_{\pi,t}-1}{\kappa} \\ \delta^{dg}_{ut} &= -\sigma\tfrac{d^{dg}_{\pi,t}-1}{\kappa} \end{aligned}.$$

Since the cost push shock is a new state variable, it enters the interest rate rule. δ^{dg}_{ut} is characterized by the following proposition:

Proposition 6. *Assume that $t < \infty$; then, $\delta^{dg}_{ut} > \delta^{EH}_u$. Moreover, we have:*
- $\lim_{t\to\infty} \delta^{dg}_{ut} = \delta^{EH}_u$.

The inequality $\delta^{dg}_{ut} > \delta^{EH}_u$ is parallel to Result 5: during the transition the optimal policy engineers more aggressive interest rate movements in response to cost push shock variations than EH, and this way it accommodates less the effect of noisy shocks on inflation compared to EH.

δ^{dg}_{ut} is positive and decreasing over time (see Figure 2)[41]. Thus monetary policy should react to the cost push shock in a similar fashion as to out of equilibrium expectations (see Result 4):

Result 5. *Optimal policy reacts aggressively to cost push shocks initially, and dampens its aggressiveness later.*

In response to a positive cost push shock, the Central Bank raises interest rate to contract the output and thus reduce inflation, and future inflationary expectations.

[41] Since $\delta^{dg}_{u,t} < 1$ from (49) it follows that the change of $\delta^{dg}_{u,t}$ through time is identical to that of $\delta^{dg}_{\pi,t}$ and the numerical analysis of Section 4.1 also applies here.

Table 4: Path of cumulative welfare loss ratios under decreasing gain, using OP and EH

	$\widehat{L} = L^{OP}/L^{EH}$		
T	CCG	W	Mc
1	1.62	1.59	1.67
5	1.11	1.12	1.13
10	1.04	1.06	1.06
15	0.99	1.02	1.02
20	0.97	1.00	0.99
30	0.93	0.97	0.96
40	0.91	0.95	0.94
50	0.89	0.94	0.91

McCN: McCallum and Nelson (1999), W: Woodford (1996)
CCG: Clarida et al. (2000), $a_0 = 0$

The asymptotic properties of the ALM (46),(48) depend on the limiting behavior of a_t, which is given by the stochastic recursive algorithm:

$$a_{t+1} = a_t + (t+1)^{-1} \left((c^{dg}_{\pi t} - 1)a_t + d^{dg}_{\pi,t} u_t \right) \tag{50}$$

We study its properties in the Appendix, where we use the stochastic approximation techniques[42] to prove the following Proposition:

Proposition 7. *Let a_t evolve according to (50); then, $a_t \to 0$ a.s.*

This result, together with the boundedness of $c^{dg}_{\pi,t}$, implies that $c^{dg}_{\pi,t} a_t$ goes to zero almost surely; moreover, it is easy to see that $d^{dg}_{\pi,t} \to \frac{\alpha}{\kappa^2+\alpha}$, so that we can conclude that $\pi_t \to \frac{\alpha}{\kappa^2+\alpha} v$ almost surely, where v is a random variable with the same probability distribution as u_t. The equilibrium corresponds to the discretionary rational expectations equilibrium. Optimal policy 'helps' private agents to learn the rational discretionary REE[43].

From Propositions 4 and 6 it follows that the optimal policy converges to the myopic policy; since expectations converge to a constant it is intuitive that in the limit OP behaves as if expectations were fixed.

The difference between OP and EH during the transition amounts to huge welfare gains. To get a quantitative feeling of the welfare gains that the use of rule (49) instead of the EH rule implies along the transition to RE, we perform Monte Carlo simulations up to time T under the two interest rules an calculate the ratio of the cumulative welfare losses similarly to Section 3.3).Results are reported in Table 4.

The Table shows that in the first periods rule (49) yields *ex-post* cumulative welfare losses higher than the EH rule; later, however, our rule starts generating

[42]For an extensive monograph on stochastic approximation, see Benveniste and P. (1990); the first paper to apply these techniques to learning models is Marcet and Sargent (1989).

[43]Note that the PLM of private agents does not nest the commitment REE, only the discretionary REE, so agents have a 'chance' to learn only the latter.

smaller welfare losses[44]. These findings are consistent with our finding that a CB that follows the optimal rule (49) reacts to out-of-equilibrium inflation expectations more aggressively than in the EH case, in order to undercut more future expectations, even if this means allowing a wider output gap in the short run. This implies that in the first periods, when this more aggressive behavior has not yet generated a pay-off in terms of a smaller $|a|$ sufficient to offset the costly output gap variability, our rule performs worse than the EH one; as soon as inflation expectations become small enough, this initial disadvantage is more than compensated. This pattern is magnified by the time-varying behavior of $\delta^{dg}_{\pi,t}$ that we characterized above: the coefficient on inflation expectations in (49) is particularly large in the first periods, hence determining large output gap variations and large welfare losses in the short run, and large gains from the contraction of $|a|$ in the medium and long run. Table 5 shows, that in the long run using the optimal rule dramatically reduces welfare losses: between 29 and 34 percent, depending on the calibration we use.

Both OP and EH are E-stable under learning, so guarantee that the economy converges to the discretionary REE; what changes is the speed of convergence to RE. In particular, Figure 6[45] shows a typical realization of the evolution of expectations under OP and EH. We can observe that inflation expectations converge faster with our rule than with the EH one. This is a consequence of the result derived above: when the CB does take into account its influence on the learning algorithm, it has an incentive to undercut future inflation beliefs more than in the case when it does not. On the other hand, in the bottom panel of Figure 6 we can see that output gap expectations converge more slowly with our rule than with the EH one. It is due to the presence of the intertemporal tradeoff described in Section 2 that the CB is ready to pay a short-term cost in order to undercut future inflation expectations. This short term cost is represented by a wider current output gap and, consequently, by a slower convergence of b to its RE value.

Table 5: Ratio of welfare losses using OP and EH discretion under decreasing gain learning

	$\widehat{L} = L^{OP}/L^{EH}$		
	McCN	W	CCG
Decreasing gain	0.660	0.766	0.719

McCN: McCallum and Nelson (1999), W: Woodford (1996)
CCG: Clarida et al. (2000)

Table 6 confirms that it is optimal to lower inflation's deviation from the target even at the cost of higher output gap variation. For decreasing gain learning, when the Central Bank takes into account its influence on private expectations it engineers an inflation variation 40-50 percent lower, even at the

[44] We report $\widehat{L}_T$ only until period 50; over a longer horizon, the ratio gets smaller.

[45] To obtain Figure 3, we adopted the Woodford (1996) calibration, with the same initial beliefs and the same realization of the cost-push shock process used to produce Table 1.

cost of of allowing a 3-9 times higher welfare loss due to output gap variations. Notice, that the higher is α, the weight on output gap in the welfare loss function, naturally the lower is the additional variation in output gap, compared to the EH rule.

Table 6: Ratio of welfare losses using OP and EH under decreasing gain learning due to inflation and output gap variations

$$\widehat{L} = L^{OP}/L^{EH}$$

	McCN	W	CCG
π	0.55	0.67	0.61
x	2.80	8.97	6.46

McCN: McCallum and Nelson (1999), W: Woodford (1996)
CCG: Clarida et al. (2000)

In this section we have proved that our main results do not depend on what type of learning algorithm private agents follow. Our new results are that under decreasing gain learning optimal policy should be time varying: more aggressive on inflation initially and less in subsequent periods. In the limit, expectations converge to the discretionary rational expectations equilibrium, and optimal policy will be equivalent to the myopic policy. Numerical simulations confirmed that optimal policy under learning engineers dramatically lower welfare losses compared to myopic policy. In the next section we direct out attention to differences between optimal policy under the two learning algorithms, and argue why it is of interest to examine both learning algorithms.

5 Consistent Policy Advice

In the previous sections we have seen that departing from rational expectations and using adaptive expectations instead, has important implications for optimal monetary policy: in general optimal policy should try to contain inflation more than under rational expectations. In this section we argue how expectation formation and accordingly the optimal monetary policy should differ in different environments. We also aim to define consistent policy advice on a relevant set of private agents' expectation formation. In concrete, we ask the question how monetary policy in the US should behave when it is uncertain about expectation formation, but it's uncertainty is restricted to a set which would be reasonable given the economic environment.

The main idea is, that in different different economic environments private agents would be likely to use different forecasting algorithms, and the Central Bank should take this into account in its policy formation [46]. Using a decreasing gain algorithm is appropriate if agents confidently believe that the variable,

[46] In a full-fledged analysis the CB would also understand how its policymaking would affect the private sectors expectation formation. Here, we do not address this issue.

they are trying to forecast has a constant mean over time. On the other hand, the use of constant-gain estimators to deal with structural changes is well known from the statistics and engineering literature [47] and using a constant gain is also optimal when the economy does not converge to REE[48]. In these environments agents would be more likely to use a constant gain algorithm. Accordingly, optimal monetary policy should be different when it operates in a stable environment to when it operates in an environment prone to structural changes[49].

Results in our paper suggest, that in economies with structural breaks optimal monetary policy should try to contain inflationary expectations much more compared to stable economies; in the sense that optimal policy should remain aggressive towards inflation even in the steady state, while in stable economies the Central Bank should be aggressive only during the transition.

During the transition relative aggressiveness of optimal policy under the two learning algorithms changes through time. Under decreasing gain learning optimal policy aims to drive inflationary expectations close to the equilibrium in the first periods, because later its influence on expectations decreases; as a result during the first periods optimal policy under decreasing gain learning is more aggressive in breaking down inflationary expectations than under constant gain learning [50]. Figure 9 shows this difference: during the first periods optimal policy under DG compared to CG has a smaller feedback coefficient of expectations in the actual law of motion of inflation (equations (46) and (32)). As we have seen in the paper this smaller feedback coefficient on inflation expectations translates to a higher coefficient of expectations in the interest rate rule (49)and (34).

Intuitively, if a structural break occurs in a stable environment the Central Bank should raise the interest rate aggressively in response to out of equilibrium expectations to speed up convergence of inflation expectations, even at the cost of slower convergence of output gap expectations (See Graph 6).

On the other hand, initially a less aggressive interest rate increase is needed when the economy is inherently characterized by structural breaks. The reason is that in the latter economy an aggressive interest rate policy has two effects: it implies a smaller volatility of inflation but a higher volatility in output gap (even in the limit). Since expectations remain volatile even in the limit, the Central Bank has to balance between these two effects, and spread the effect of

[47] See for example Benveniste and P. (1990), Part I. Chapters 1. and 4.

[48] The optimal choice of the gain parameter is nonlinear, it depends on the relative importance of tracking versus filtering the observed data. A high value of γ increases tracking, or in other words the responsiveness of the forecasts to the structural change. On the other hand, a high tracking parameter reduces filtering, i.e. the ability of the forecast rule to eliminate noise in the data. Although the choice of the optimal gain parameter is subject to this trade-off, the use of a constant gain is clearly indicated when the structure of the economy is subject to changes.

[49] Recent research addressing the issue of choosing an optimal tracking parameter are Benveniste and P. (1990), Evans and Ramey (2005) and Evans and Honkapohja (1993). Marcet and Nicolini (2003) show that agents might also switch predictor use dynamically: use decreasing gain learning when forecast errors are low, and switch to constant gain learning when forecast errors are high.

[50] This result holds true for all three calibrations used in our paper.

its interest rate policy through time.

Now, the question arises, what is a relevant set of expectation formation? Empirical evidence on this is relatively scarce, and mainly focuses on the US. For the US the estimated γ is typically small. With Bayesian estimation of the New Keynesian model Milani (2005) finds γ to be around 0.02, he also finds γ to be stable through time. Orphanides and Williams (2005b) calibrate γ between 0.01 and 0.4 on the Survey of Professional Forecasters.

Let us conduct an experiment, and suppose that the FED is uncertain about how private sector forms its expectations, but relying on the empirical literature listed above it can define a relevant (but not exhaustive) set of expectations to be: decreasing gain, constant gain with a small gain, and rational expectations. This set of inflation expectations is relevant if private agents confidently believe themselves to be in a stable environment, or if they believe it is a stable environment but, there is a small probability of regime change. We raise the question, whether there is a consistent policy rule, that performs well on this set of expectation formations.

Table (7) shows expected welfare losses [51]. when private agents follow constant gain learning with a small gain, and the Central Bank uses the optimal rule, CG, or mistakenly uses another rule. CG denotes the optimal interest rate rule under constant gain learning (34), CG with bad γ is the same interest rate rule, but we assume the Central Bank misperceives γ. EH denotes the myopic interest rate rule (28), derived in Evans and Honkapohja (2003b); this rule does not take into account the feedback of monetary policy to private expectations (but guarantees stability of learning). DG denotes the optimal rule derived for decreasing gain learning (49). Table (8) shows similar welfare loss calculations for decreasing gain learning and Table (9) for rational expectations[52]. For example the first row in Table (7) shows that if private agents follow constant gain learning with $\gamma = 0.1$ and the Central Bank uses the optimal rule, welfare loss is 177.21. When the Central Bank mistakenly uses the optimal rule derived for $\gamma = 0.2$ the welfare loss is 177.94, under the EH rule the expected welfare loss is 218.16 and finally with DG rule it is 177.21.

Under learning optimal rules derived for learning perform robustly much better that the myopic rule (EH), which does not take into account the central banks effect on private expectations. In other words, when private agents are learning the Central Bank should make active use of this knowledge, even if it is unsure about the exact nature of learning (within the learning rules examined here).

Under rational expectations all of these rules lead to a determinate equilibrium. The myopic rule provides smaller welfare losses than optimal learning rules, and the reason for this is that learning rules allow for too high volatility

[51] Monte Carlo simulations, with length 10,000, crossectional sample size 1000. All results are with the CCG calibration, similar results hold with the W and McN calibrations.

[52] Rational expectations means, substituting the interest rate rule in the IS curve (1), and then using the Phillips Curve (25) solving for the fixed point in expectations. Under all interest rate rules listed above, this results $E_t\pi_{t+1} = E_t x_{t+1} = 0$.

Table 7: Expected welfare losses when private agents are constant gain learners

Interest rate rule / Tracking parameter	CG	CG with bad γ	EH	DG
0.1	177.21	177.94	218.16	195.37
0.2	212.93	214.66	313.98	255.26

CCG: Clarida et al. (2000)

Table 8: Expected welfare losses when private agents are decreasing gain learners

Interest rate rule / Tracking parameter	CG	EH	DG
0.1	200.13	223.01	160.24
0.2	201.50		

CCG: Clarida et al. (2000)

in the output gap [53].

Table 9: Expected welfare losses when private agents have rational expectations

Interest rate rule / Tracking parameter	CG	EH	DG
0.1	96.77	96.44	97.09
0.2	97.98		

CCG: Clarida et al. (2000)

However, welfare losses under RE caused by mistakenly using an optimal learning rule are much less severe than welfare losses due to using a myopic rule when agents are learning. Let us consider the following simple exercise. Assume, that the prior of the Central Bank is that with probability p private agents follow decreasing gain learning, and with probability $1-p$ agents having rational expectations. Then we can calculate the expected welfare loss of using EH p times welfare loss under decreasing gain learning with EH rule, and $1-p$ times welfare loss of using EH under RE. And similarly the expected welfare loss of using DG. Then we can find a cut-off value of p: expected welfare loss of using the DG rule is less than the welfare loss of the myopic rule (EH) even if the Central Bank attributes only 2% probability (or higher) to agents following learning [54].

[53]We would like to note, that since learning rules decrease volatility of inflation and allow for higher volatility in the output gap, for small values of alpha (a small weight on output gap in the welfare loss function) learning rules even outperform the discretionary rule under rational expectations (EH)

[54]The same result holds for constant gain learning with $\gamma = 0.1$. For $\gamma = 0.2$ using the optimal CG rule gives lower expected welfare loss then the EH rule when the Central Bank hypothesizes there is a 10% or higher probability of learning.

In sum, the worst case scenario is using a myopic interest rate rule (discretionary rational expectations rule) when agents are actually learning. Thus when the Central Bank is insecure as to whether agents have rational expectations or are learning, and tries to avoid huge losses it should use an optimal learning rule. This holds true, even if the Central Bank attaches only very small probability to agents following learning.

Our results show, that the "safest" rule to is the DG rule, in the sense that the Central Bank does not make a big mistake if it erroneously uses this rule even if in reality private agents have rational expectations or when agents follow constant gain learning with small gain.

6 Extensions

Up to now, we have supposed that the CB perfectly observes all the relevant state variables of the system, namely the exogenous shocks and the agents' beliefs. In this section we show that our main results extend to a more general framework, where either the shocks or the expectations are not observable. In particular, to make the problem non-trivial, throughout this section we modify the structural equations (1) and (25) with the introduction of unobservable shocks, so that the model is now given by:

$$x_t = E_t^* x_{t+1} - \sigma^{-1}(r_t - E_t^* \pi_{t+1} - \overline{rr}_t) + g_t + e_{x,t} \tag{51}$$

and:

$$\pi_t = \beta E_t^* \pi_{t+1} + \kappa x_t + u_t + e_{\pi,t} \tag{52}$$

where we assume that the CB can observe π_t and x_t only with a lag, and that $e_{x,t}$ and $e_{\pi,t}$ are independent white noise that are not observable, not even with a lag. The rest of the setup is identical to subsection 3.1.

6.1 Measurement Error in the Shocks

We start with the case in which the monetary authority can observe g_t and u_t only with an error; in particular, we assume that it receives the noisy signals g_t^* and u_t^*, where:

$$\begin{aligned} g_t^* &= g_t + \epsilon_t, \qquad \epsilon_t \sim N(0, \sigma_\epsilon^2) \\ u_t^* &= u_t + \eta_t, \qquad \eta_t \sim N(0, \sigma_\eta^2) \end{aligned}$$

To make the problem non-trivial, we also assume that the CB can observe π_t and x_t only with a lag. Note that the shocks do not depend on the policy followed by the CB; hence, the *separation principle* applies, namely, the optimization of the welfare criterion and the estimation of the realizations of the shocks can be solved as separate problems. As is well known, the above signal-extraction problem implies that the expected values of the shocks given the signals are[55]:

$$E\left[g_t/g_t^*\right] \equiv E_t^{CB} g_t = \frac{\sigma_g^2}{\sigma_\epsilon^2+\sigma_g^2} g_t^* \equiv \zeta_g g_t^*$$
$$E\left[u_t/u_t^*\right] \equiv E_t^{CB} u_t = \frac{\sigma_u^2}{\sigma_\eta^2+\sigma_u^2} u_t^* \equiv \zeta_u u_t^*$$

Moreover, the separation principle implies that certainty equivalence holds in designing the optimal interest rate rule, which turns out to be identical to (49), with g_t and u_t replaced by $E_t^{CB} g_t$ and $E_t^{CB} u_t$, respectively:

$$\begin{aligned} r_t &= \overline{rr}_t + \delta_{\pi,t}^{dg} a_t + \delta_x^{dg} b_t + \delta_g^{dg} \zeta_g g_t^* + \delta_{ut}^{dg} \zeta_u u_t^* \\ &= \overline{rr}_t + \delta_{\pi,t}^{dg} a_t + \delta_x^{dg} b_t + \delta_g^{dg} \zeta_g g_t + \delta_g^{dg} \zeta_g \epsilon_t + \delta_{ut}^{dg} \zeta_u u_t + \delta_{ut}^{dg} \zeta_u \eta_t \end{aligned}$$

We can combine the above equation with (51) and (52) to obtain the ALM for inflation and output gap:

$$\pi_t = \mu_{at}^1 a_t + \mu_g^1 g_t + \mu_\epsilon^1 \epsilon_t + \mu_{ut}^1 u_t + \mu_{\eta t}^1 \eta_t + \kappa e_{x,t} + e_{\pi,t}$$
$$x_t = \mu_{at}^2 a_t + \mu_g^2 g_t + \mu_\epsilon^2 \epsilon_t + \mu_{ut}^2 u_t + \mu_{\eta t}^2 \eta_t + e_{x,t}$$

where:

$$\begin{aligned} &\mu_{at}^1 = c_{\pi,t}^{dg}, \quad \mu_{at}^2 = c_{x,t}^{dg} \\ &\mu_g^1 = \kappa\left(1-\zeta_g\right), \quad \mu_g^2 = 1-\zeta_g \\ &\mu_\epsilon^1 = -\kappa\zeta_g, \quad \mu_\epsilon^2 = -\zeta_g \\ &\mu_{ut}^1 = \left(d_{\pi,t}^{dg}-1\right)\zeta_u + 1, \quad \mu_{ut}^2 = \left(\frac{d_{\pi,t}^{dg}-1}{\kappa}\right)\zeta_u \\ &\mu_{\eta t}^1 = \left(d_{\pi,t}^{dg}-1\right)\zeta_u, \quad \mu_{\eta t}^2 = \left(\frac{d_{\pi,t}^{dg}-1}{\kappa}\right)\zeta_u \end{aligned}$$

As a consequence of the measurement error, inflation and output gap now depend on a wider set of state variables; however, it is easy to see that the main findings of the preceding section go through in this modified environment. First of all, the separation principle trivially implies that when the CB takes into account the effect of its decisions on future beliefs, the optimal policy is more aggressive against out-of-equilibrium inflation expectations, compared to the case in which the private sector's expectations are considered as exogenously given[56]; moreover, the analysis of convergence of learning algorithms to the optimal discretionary RE equilibrium[57] does not change in this modified environment.

6.2 Heterogenous Forecasts

As argued in Honkapohja and Mitra (2005) (HM hereafter), the hypothesis that the CB can perfectly observe private sector's expectations is subject to several

[55] E.g., see Hamilton (1994).

[56] For a description of the optimal policy when the CB does not consider its effect on future beliefs, and there is measurement error in the shocks, see Evans and Honkapohja (2003b) section 4.2.

[57] Note that the optimal RE equilibrium is now different from the baseline case, since inflation and output gap depend also on g_t, ϵ_t, η_t, and the unobservable shocks $e_{x,t}$ and $e_{\pi,t}$.

criticisms[58]; it is therefore natural to verify the robustness of our results when this assumption is relaxed. In what follows, we assume that the optimal interest rate rule takes the same form as (49), but the agents' forecasts for inflation and output gap, a_t and b_t, are replaced by the CB internal forecasts, a_t^{CB} and b_t^{CB}[59]; in particular, we suppose that the CB and the private sector forecasts have the same form, and are updated according to the same algorithm, which is given by (35)-(36). The only difference is given by the initial beliefs. Note that this setup corresponds to a situation where the CB, in solving its optimization problem, knows the adaptive algorithm used by the agents to form their expectations, but cannot observe the actual values of these expectations; instead, the CB has a tight prior on a_0 and b_0[60], and forms its internal forecasts accordingly. Plugging the interest rate rule into the structural equations (51) and (52), we get the ALM:

$$\begin{array}{l} \pi_t = \nu_a^1 a_t + \nu_{a^{CB}t}^1 a_t^{CB} + \nu_b^1 b_t + \nu_{b^{CB}}^1 b_t^{CB} + \nu_{ut}^1 u_t + \kappa e_{x,t} + e_{\pi,t} \\ x_t = \nu_a^2 a_t + \nu_{a^{CB}t}^2 a_t^{CB} + \nu_b^2 b_t + \nu_{b^{CB}}^2 b_t^{CB} + \nu_{ut}^2 u_t + e_{x,t} \end{array} \tag{53}$$

where:

$$\begin{array}{l} \nu_a^1 = \beta + \kappa\sigma^{-1}, \quad \nu_a^2 = \sigma^{-1} \\ \nu_{a^{CB}t}^1 = -\kappa\sigma^{-1}\left(1 - \sigma\frac{c_{\pi,t}^{dg} - \beta}{\kappa}\right), \quad \nu_{a^{CB}t}^2 = -\sigma^{-1}\left(1 - \sigma\frac{c_{\pi,t}^{dg} - \beta}{\kappa}\right) \\ \nu_b^1 = \kappa, \quad \nu_b^2 = 1 \\ \nu_{b^{CB}}^1 = -\kappa, \quad \nu_{b^{CB}}^2 = -1 \\ \nu_{ut}^1 = d_{\pi,t}^{dg}, \quad \nu_{ut}^2 = d_{x,t}^{dg} \end{array}$$

Again, our main results are unaffected by this change in the CB information set, both for $t < \infty$ and for $t \to \infty$. In fact, since the parameters in the optimal rule are the same as in rule (49), the results summarized in Propositions 4 and 6 are still valid. On the other hand, we can study E-stability of the system extending Proposition 2 in HM to a time-varying environment. In particular, it is easy to show[61]:

Corollary 1. *Consider the model (53); it is E-stable if and only if the corresponding model with homogenous expectations is E-stable.*

Since E-stability of the homogenous expectations model is ensured by Proposition 7, we conclude that also system (53) is E-stable, and it converges to the optimal discretionary RE equilibrium[62].

---

[58]For example, private expectations and their forecasts produced by different institutions do not necessarily coincide.

[59]This approach is developed in HM, where it is applied to the EH rule and to a simple Taylor rule. Evans and Honkapohja (2003a) use this method in a setup where the CB follows the expectations based interest rule derived in Evans and Honkapohja (2002).

[60]In other words, it believes that $a_0 = a_0^{CB}$ and $b_0 = b_0^{CB}$ with probability one, where a_0^{CB} and b_0^{CB} are given.

[61]The proof is available from the authors upon request.

[62]In fact, the system we are analyzing falls into the class for which E-stability and convergence of real time learning are equivalent, see Evans and Honkapohja (2001).

7 Conclusions

In this paper we analyzed the optimal monetary policy problem faced by a CB that tries to exploit its ability to influence future beliefs of the agents, when they follow adaptive learning to form their expectations. We have shown that monetary policy should be aggressive on inflation, and the reason for this is that in this way private agents learn the true value of steady state inflation faster. We have shown that optimal policy can be implemented by an aggressive interest rate policy, and also that this behavior is optimal even at the cost of higher welfare losses from output gap volatility. We conclude by describing several areas where future research would be useful.

We have shown that learning introduces an additional tradeoff between inflation and output gap stabilization that is not present under rational expectations, namely an intertemporal tradeoff which is generated by the central banks ability to influence future expectations. We analytically show that because of this intertemporal tradeoff, during the transition optimal policy qualitatively resembles the commitment solution under rational expectations. In this sense the Central Bank's desire to influence future expectations by its current action acts as a commitment device.

Optimal policy naturally chooses an E-stable policy, but even though during the transition optimal policy resembles the commitment solution under rational expectations, in our setup it drives expectations to the discretionary rational expectations solution. The reason for this is that agents expectation formation does not nest the commitment solution under rational expectations. Under rational expectations commitment calls for an ALM with a different functional form than the discretionary case (see Clarida, Gali, and Gertler (1999)).

It would very interesting to explore the possibility of reaching the commitment solution with adaptive learning algorithms. This question is particularly interesting as from the backward looking nature of these learning algorithms it follows that such policies are time consistent.

Our analysis was restricted to examining optimal policy given a certain learning algorithm. We have shown that optimal policy under decreasing gain learning performs robustly well both under rational expectations and under constant gain learning with small gain parameters (i.e. using a rolling window regression with a relatively large window). Which suggests that our rule would perform well even if private agents were to switch expectations formation within this set. It would be interesting to examine how monetary should be conducted with endogenous expectation formation, in other words when private agents would change their expectation formation depending on their perception about the underlying economy. Endogenous expectation formation could be formulated for example along the lines of Marcet and Nicolini (2003) where agents dynamically switch between predictor use depending on the last forecast error. An alternative way would be to to model expectation formation as in Molnar (2005) where agents do not switch predictor use, but always a weighted average of predictor forecasts and adjust the weight on predictors dynamically depending on the relative forecasting performance.

A Constant Gain Learning

Lemma 1. *Let the set of all the real bounded sequences be defined as follows:*

$$M^{\infty} \equiv \{\{z_t\} \in R^{\infty} : \{z_t\} \text{ is bounded}\}$$

and let:

$$G \equiv \left\{\{\pi_t, x_t, r_t, a_{t+1}, b_{t+1}\} \in M^{\infty} \times M^{\infty} \times M_{+}^{\infty}\right\}$$

If there exists a sequence $\left\{\pi_t^*, x_t^*, r_t^*, a_{t+1}^*, b_{t+1}^*\right\} \in G$ *that solves the problem:*

$$\min_{\{\pi_t, x_t, r_t, a_{t+1}, b_{t+1}\} \in G} E_0 \sum_{t=0}^{\infty} \beta^t (\pi_t^2 + \alpha x_t^2) \tag{54}$$
$$\text{s.t. (1), (2), (8), (9)}$$
$$a_0,\ b_0 \text{ given}$$

then $\left\{\pi_t^*, x_t^*, r_t^*, a_{t+1}^*, b_{t+1}^*\right\}$ *solves also (10).*

Proof. Let $\left\{\widehat{\pi}_t, \widehat{x}_t, \widehat{r}_t, \widehat{a}_{t+1}, \widehat{b}_{t+1}\right\}$ be an arbitrary unbounded sequence that satisfies the constraints of (10), and such that:

$$\widehat{V} \equiv \sum_{t=0}^{\infty} \beta^t (\widehat{\pi}_t^2 + \alpha \widehat{x}_t^2) < \infty \tag{55}$$

Let $\{\widehat{\pi}_t^n\}$ be defined as:

$$\{\widehat{\pi}_t^n\} \equiv \{\widehat{\pi}_0, \widehat{\pi}_1, ..., \widehat{\pi}_n, \widehat{\pi}_n, \widehat{\pi}_n, ...\}$$

and $\left\{\widehat{x}_t^n, \widehat{r}_t^n, \widehat{a}_{t+1}^n, \widehat{b}_{t+1}^n\right\}$ are defined accordingly to respect the constraints of (10); clearly, $\left\{\widehat{\pi}_t^n, \widehat{x}_t^n, \widehat{r}_t^n, \widehat{a}_{t+1}^n, \widehat{b}_{t+1}^n\right\}$ is bounded, so that:

$$\widehat{V}^n \geq V^*, \qquad \forall n$$

Since this is true for any n, it must be true also in the limit, i.e.:

$$\lim_{n \to \infty} \widehat{V}^n \geq V^*$$

if $\lim_{n \to \infty} \widehat{V}^n$ exists. However, it is easy to see that $\lim_{n \to \infty} \widehat{V}^n = \widehat{V}$; since $\left\{\widehat{\pi}_t, \widehat{x}_t, \widehat{r}_t, \widehat{a}_{t+1}, \widehat{b}_{t+1}\right\}$ was arbitrary, it proves the statement [63]. □

Lemma 2. *Let A_{11} be given by equation (21) in the text; then it has an eigenvalue inside and one outside the unit circle.*

[63] Note that the condition (55) can be imposed without any loss of generality, since any $\left\{\widehat{\pi}_t, \widehat{x}_t, \widehat{r}_t, \widehat{a}_{t+1}, \widehat{b}_{t+1}\right\}$ that does not respect it, for sure cannot do better than $\{\pi_t^*, x_t^*, r_t^*, a_{t+1}^*, b_{t+1}^*\}$.

Proof. First of all, we recall a result of linear algebra that we will use in the proof, i.e. that a necessary and sufficient condition for a 2 by 2 matrix to have an eigenvalue inside and one outside the unit circle, is that[64]:

$$|\mu_1 + \mu_2| > |1 + \mu_1 \mu_2|$$

where μ_1, μ_2 are the eigenvalues of the matrix; in the case of A_{11}, the above condition can be written equivalently:

$$\frac{\kappa^2 + \alpha + \alpha\beta^2\gamma(1-\gamma(1-\beta))}{\kappa^2\beta(1-\gamma) + \alpha\beta(1-\gamma(1-\beta))} + 1 - \gamma >$$
$$1 + \frac{\kappa^2 + \alpha + \alpha\beta^2\gamma(1-\gamma(1-\beta))}{\kappa^2\beta(1-\gamma) + \alpha\beta(1-\gamma(1-\beta))}(1-\gamma) + \frac{\alpha\beta(1-\beta(1-\gamma)(1-\gamma(1-\beta)))}{\kappa^2\beta(1-\gamma) + \alpha\beta(1-\gamma(1-\beta))}\gamma$$

where we have used the fact that the trace is equal to the sum of the eigenvalues, and that the determinant is equal to the product. After simplifying the above inequality, we get:

$$-\gamma > -\gamma\left(\frac{\kappa^2 + \alpha + \alpha\beta^2\gamma(1-\gamma(1-\beta)) - \alpha\beta(1-\beta(1-\gamma)(1-\gamma(1-\beta)))}{\kappa^2\beta(1-\gamma) + \alpha\beta(1-\gamma(1-\beta))}\right)$$

so that all we have to prove is that:

$$\frac{\kappa^2 + \alpha + \alpha\beta^2\gamma(1-\gamma(1-\beta)) - \alpha\beta(1-\beta(1-\gamma)(1-\gamma(1-\beta)))}{\kappa^2\beta(1-\gamma) + \alpha\beta(1-\gamma(1-\beta))} > 1$$

Some tedious algebra shows that this is equivalent to the following expression:

$$\kappa^2(1-\beta(1-\gamma)) + \alpha(1-\beta)(1-\beta(1-\gamma(1-\beta))) > 0$$

which is always true, since β and γ are supposed smaller than one. □

We now prove Proposition 1. First of all, we can guess that inflation follows the ALM (22)[65] and use the optimality condition (18) and the method of undetermined coefficients to verify that c_π^{cg} must satisfy the following quadratic expression:

$$p_2\left(c_\pi^{cg}\right)^2 + p_1 c_\pi^{cg} + p_0 = 0$$

where:

$$\begin{aligned}
p_2 &\equiv \gamma\left[\kappa^2\beta(1-\gamma) + \alpha\beta(1-\gamma(1-\beta))\right] \\
p_1 &\equiv (1-\gamma)\left[\kappa^2\beta(1-\gamma) + \alpha\beta(1-\gamma(1-\beta))\right] - \left[\kappa^2 + \alpha + \alpha\beta^2\gamma(1-\gamma(1-\beta))\right] \\
p_0 &\equiv \alpha\beta(1-\beta(1-\gamma)(1-\gamma(1-\beta)))
\end{aligned}$$

The above polynomial can be equivalently rewritten as follows:

$$c_\pi^{cg} = -\frac{p_0 + p_2\left(c_\pi^{cg}\right)^2}{p_1} \equiv f(c_\pi^{cg})$$

[64]LaSalle (1986).

[65]Which we showed in the text that is the functional form that inflation will have at the optimum.

We will prove that the function $f(\cdot)$, defined on the interval $[0,1]$, is a contraction, so that it admits one and only one fixed point; moreover, since the two roots of the quadratic expression have the same sign (it is due to the fact that both p_2 and p_0 are positive), it follows that the other candidate value for c_π^{cg} is greater than one, which is not compatible with the boundary conditions[66].

First of all, we show that $f(\cdot)$, when defined on the interval $[0,1]$, takes values on the same interval.

Lemma 3. *$f(c_\pi^{cg})$ is strictly monotone increasing on the interval* $[0,1]$.

Proof. Note that:

$$f'(c_\pi^{cg}) = \frac{2\gamma[\alpha\beta(1-\gamma(1-\beta)) + \kappa^2\beta(1-\gamma)]}{\kappa^2 + \alpha + \alpha\beta^2\gamma\,(1-\gamma\,(1-\beta)) - (1-\gamma)[\kappa^2\beta\,(1-\gamma) + \alpha\beta\,(1-\gamma\,(1-\beta))]} c_\pi^{cg}$$

which is positive if and only if the denominator is positive:

$$\kappa^2 + \alpha + \alpha\beta^2\gamma\,(1-\gamma\,(1-\beta)) - (1-\gamma)[\kappa^2\beta\,(1-\gamma) + \alpha\beta\,(1-\gamma\,(1-\beta))] \lesseqgtr 0$$

After rearranging:

$$\kappa^2\left(1-\beta(1-\gamma)^2\right) + \alpha[1-\beta(1-\gamma)(1-\gamma(1-\beta))] + \alpha\beta^2\gamma\,(1-\gamma\,(1-\beta)) \lesseqgtr 0$$

which is always positive. Thus we have proved that $f(c_\pi^{cg})$ is strictly monotone increasing on the interval [0,1]. $\square$

Lemma 4. $f(c_\pi^{cg}) : [0,1] \to [0,1]$

Proof. Since $f(c_\pi^{cg})$ is strictly monotone increasing it suffices to show that $f(0) > 0$ and $f(1) < 1$.

$$f(0) = \frac{\alpha\beta\,(1-\beta\,(1-\gamma)\,(1-\gamma\,(1-\beta)))}{\kappa^2 + \alpha + \alpha\beta^2\gamma\,(1-\gamma\,(1-\beta)) - (1-\gamma)[\kappa^2\beta\,(1-\gamma) + \alpha\beta\,(1-\gamma\,(1-\beta))]}$$

where the denominator is positive (see the preceding proof), and also the numerator is trivially positive. Thus $f(0) > 0$.

$$f(1) = \frac{\gamma\left[\kappa^2\beta\,(1-\gamma) + \alpha\beta\,(1-\gamma\,(1-\beta))\right] + \alpha\beta\,(1-\beta\,(1-\gamma)\,(1-\gamma\,(1-\beta)))}{\kappa^2 + \alpha + \alpha\beta^2\gamma\,(1-\gamma\,(1-\beta)) - (1-\gamma)[\kappa^2\beta\,(1-\gamma) + \alpha\beta\,(1-\gamma\,(1-\beta))]}$$

After rearranging, we get:

$$f(1) \lesseqgtr 1 \Longleftrightarrow 0 \lesseqgtr \kappa^2\,(1-\beta\,(1-\gamma)) + \alpha\,(1-\beta)\,(1-\beta\,(1-\gamma\,(1-\beta)))$$

but, as we argued above, the RHS of the last inequality is always positive; hence, $f(1) < 1$. $\square$

[66]Since it would imply an exploding inflation.

To show that $f(\cdot)$ is a contraction, it suffices to show that its derivative is bounded above by a number smaller than one: in fact, by the Mean Value Theorem, we now that for any a, b, there exists a $c \in (a, b)$ such that:

$$|f(a) - f(b)| \leq |f'(c)|\,|a - b|$$

and if $|f'(c)| \leq M < 1$ for any $c \in [0, 1]$, we have the definition of a contraction.

Lemma 5. *For any $x \in [0, 1]$, $0 < f'(x) \leq f'(1) < 1$.*

Proof. First of all, note that:

$$f'(x) = \frac{2\gamma[\alpha\beta(1-\gamma(1-\beta)) + \kappa^2\beta(1-\gamma)]}{\kappa^2 + \alpha + \alpha\beta^2\gamma\,(1-\gamma\,(1-\beta)) - (1-\gamma)[\kappa^2\beta\,(1-\gamma) + \alpha\beta\,(1-\gamma\,(1-\beta))]}x$$

is positive and increasing in x, so that $\max\limits_{x\in[0,1]} f'(x) = f'(1)$; after some algebraic manipulation, we get:

$$f'(1) \lesseqgtr 1 \Longleftrightarrow (1-\beta\gamma)\,\beta\,(1-\gamma\,(1-\beta)) + \beta\gamma\,(1-\gamma\,(1-\beta)) - 1 \lesseqgtr \frac{\kappa^2}{\alpha}\left(1-\beta\left(1-\gamma^2\right)\right)$$

Since $\beta, \gamma \in (0, 1)$, we have:

$$(1-\beta\gamma)\,\beta\,(1-\gamma\,(1-\beta)) + \beta\gamma\,(1-\gamma\,(1-\beta)) - 1 < 1 - \beta\gamma + \beta\gamma\,(1-\gamma\,(1-\beta)) - 1 < 0$$

so that $f'(1)$ will be smaller than one ($\frac{\kappa^2}{\alpha}\left(1-\beta\left(1-\gamma^2\right)\right)$ is always positive). □

Moreover, we prove the following result.

Lemma 6. *Let $f(\cdot)$ be defined as above; then, $f\left(\frac{\alpha\beta}{\kappa^2+\alpha}\right) \leq \frac{\alpha\beta}{\kappa^2+\alpha}$.*

Proof. Note that:

$$f\left(\frac{\alpha\beta}{\kappa^2+\alpha}\right) = \frac{\alpha\beta\,(1-\beta\,(1-\gamma)\,(1-\gamma\,(1-\beta)))}{\kappa^2 + \alpha + \alpha\beta^2\gamma\,(1-\gamma\,(1-\beta)) - (1-\gamma)[\kappa^2\beta\,(1-\gamma) + \alpha\beta\,(1-\gamma\,(1-\beta))]} +$$
$$+ \frac{\gamma\left[\kappa^2\beta\,(1-\gamma) + \alpha\beta\,(1-\gamma\,(1-\beta))\right]}{\kappa^2 + \alpha + \alpha\beta^2\gamma\,(1-\gamma\,(1-\beta)) - (1-\gamma)[\kappa^2\beta\,(1-\gamma) + \alpha\beta\,(1-\gamma\,(1-\beta))]}\left(\frac{\alpha\beta}{\kappa^2+\alpha}\right)^2$$
$$\gtreqless \frac{\alpha\beta}{\kappa^2+\alpha}$$

if and only if:

$$\frac{\left(\kappa^2+\alpha\right)\alpha\beta\,(1-\beta\,(1-\gamma)\,(1-\gamma\,(1-\beta))) + \gamma\left[\kappa^2\beta\,(1-\gamma) + \alpha\beta\,(1-\gamma\,(1-\beta))\right]\frac{\alpha\beta}{\kappa^2+\alpha}}{\kappa^2 + \alpha + \alpha\beta^2\gamma\,(1-\gamma\,(1-\beta)) - (1-\gamma)[\kappa^2\beta\,(1-\gamma) + \alpha\beta\,(1-\gamma\,(1-\beta))]} \gtreqless 1$$

For $\gamma = 0$ it is easy to verify that $f\left(\frac{\alpha\beta}{\kappa^2+\alpha}\right) = \frac{\alpha\beta}{\kappa^2+\alpha}$. If $\gamma > 0$, since the $\frac{\alpha\beta}{\alpha+\kappa^2} < \beta$, the LHS of the above inequality is smaller than:

$$\frac{\left(\kappa^2+\alpha\right)\alpha\beta\,(1-\beta\,(1-\gamma)\,(1-\gamma\,(1-\beta))) + \beta\gamma\left[\kappa^2\beta\,(1-\gamma) + \alpha\beta\,(1-\gamma\,(1-\beta))\right]}{\kappa^2 + \alpha + \alpha\beta^2\gamma\,(1-\gamma\,(1-\beta)) - (1-\gamma)[\kappa^2\beta\,(1-\gamma) + \alpha\beta\,(1-\gamma\,(1-\beta))]}$$

which is equal to one; in fact:

$$\frac{\left(\kappa^2+\alpha\right)\left(1-\beta\left(1-\gamma\right)\left(1-\gamma\left(1-\beta\right)\right)\right)+\beta\gamma[\kappa^2\beta\left(1-\gamma\right)+\alpha\beta\left(1-\gamma\left(1-\beta\right)\right)]}{\kappa^2+\alpha+\alpha\beta^2\gamma\left(1-\gamma\left(1-\beta\right)\right)-\left(1-\gamma\right)[\kappa^2\beta\left(1-\gamma\right)+\alpha\beta\left(1-\gamma\left(1-\beta\right)\right)]}\gtreqless 1$$

is equivalent to:

$$-\left(\kappa^2+\alpha\right)\beta\left(1-\gamma\right)\left(1-\gamma\left(1-\beta\right)\right)+\left(1-\gamma\left(1-\beta\right)\right)\left[\alpha\beta\left(1-\gamma\left(1-\beta\right)\right)+\kappa^2\beta\left(1-\gamma\right)\right]\gtreqless\alpha\beta^2\gamma\left(1-\gamma\left(1-\beta\right)\right)$$

But the LHS can simplified as:

$$\kappa^2\left(\beta\left(1-\gamma\right)\left(1-\gamma\left(1-\beta\right)\right)-\beta\left(1-\gamma\right)\left(1-\gamma\left(1-\beta\right)\right)\right)+\alpha\beta\left(1-\gamma\left(1-\beta\right)\right)\left(1-\gamma\left(1-\beta\right)-\left(1-\gamma\right)\right)$$

which is equal to:

$$\alpha\beta^2\gamma\left(1-\gamma\left(1-\beta\right)\right)$$

Summing up, we showed that (if $\gamma>0$) the following holds:

$$\frac{\left(\kappa^2+\alpha\right)\left(1-\beta\left(1-\gamma\right)\left(1-\gamma\left(1-\beta\right)\right)\right)+\beta\gamma[\kappa^2\beta\left(1-\gamma\right)+\alpha\beta\left(1-\gamma\left(1-\beta\right)\right)]}{\kappa^2+\alpha+\alpha\beta^2\gamma\left(1-\gamma\left(1-\beta\right)\right)-\left(1-\gamma\right)[\kappa^2\beta\left(1-\gamma\right)+\alpha\beta\left(1-\gamma\left(1-\beta\right)\right)]}=1$$

which implies that:

$$f\left(\frac{\alpha\beta}{\kappa^2+\alpha}\right)<\frac{\alpha\beta}{\kappa^2+\alpha}$$

□

We are now ready to prove the Proposition.

Proof of Proposition 1. Combining the Lemmas 4 and 5 we obtain that $f(\cdot)$ is a contraction when defined on the interval $[0,1]$; moreover, by Lemma 6 we get that f, when defined on $[0,\frac{\alpha\beta}{\kappa^2+\alpha}]$, takes values on the same interval. This result, together with Lemma 5 and with the inequality $\frac{\alpha\beta}{\kappa^2+\alpha}<1$, implies that $f(\cdot)$ is a contraction also when defined on the interval $[0,\frac{\alpha\beta}{\kappa^2+\alpha}]$ and, therefore, that the optimal c_π^{cg} must be between zero and $\frac{\alpha\beta}{\kappa^2+\alpha}$.

Finally, note that when $\gamma=0$, $f(c_\pi^{cg})$ collapses to $\frac{\alpha\beta}{\kappa^2+\alpha}$, which proves also the last statement of the Proposition. □

B Decreasing Gain Learning

Proof of Proposition 3. To prove the first part of the statement; first of all, note that if we solve forward the following difference equation:

$$c_{\pi,t}^{dg}=\beta c_{\pi,t+1}^{dg}+\frac{\alpha\beta}{\kappa^2+\alpha}\left(1-\beta\right)$$

we obtain one and only one bounded solution, i.e.:

$$c_{\pi,t}^{dg}=\frac{\alpha\beta}{\kappa^2+\alpha}\quad\forall t$$

Moreover, we can rewrite the difference equation defining $c_{\pi t}^{dg}$ as:

$$A_{11,t}c_{\pi,t}^{dg} - c_{\pi,t+1}^{dg} \equiv G_t = -\frac{1}{t+1}c_{\pi,t+1}^{dg} - A_{12,t} + \frac{1}{t+1}c_{\pi,t}^{dg}c_{\pi,t+1}^{dg}$$

If c_{π}^{dg} is bounded, it is easy to show that G has a limit:

$$\lim_{t\to\infty} G_t = -\lim_{t\to\infty} A_{12,t} = \frac{\alpha}{\kappa^2+\alpha}(1-\beta)$$

We can also show that the difference equation defined by G converges to:

$$\beta^{-1}c_{\pi,\tau}^{dg} - c_{\pi,\tau+1}^{dg}$$

Summing up, in the limit we have that c_{π}^{dg} evolves according to:

$$c_{\pi\tau}^{dg} = \beta c_{\pi\tau+1}^{dg} + \frac{\alpha\beta}{\kappa^2+\alpha}(1-\beta)$$

which, as we argued before, has one and only one bounded solution:

$$c_{\pi\tau}^{dg} = \frac{\alpha\beta}{\kappa^2+\alpha}$$

We prove the second part of the statement by contradiction. Assume that there exists a $T<\infty$ such that $c_{\pi,t}^{dg} \geq \frac{\alpha\beta}{\alpha+\kappa^2}$; we show that this implies $c_{\pi,t}^{dg} > \frac{\alpha\beta}{\alpha+\kappa^2}$ for any $t>T$. First of all, we can write:

$$\frac{c_{\pi,T+1}^{dg}\left(1-\frac{1}{T+1}\right) - A_{12,T}}{A_{11,T} - c_{\pi,T+1}^{dg}\frac{1}{T+1}} = c_{\pi,t}^{dg} \geq \frac{\alpha\beta}{\alpha+\kappa^2}$$

Rearranging and simplifying, this turns out to be equivalent to:

$$\left(1-\frac{1}{T+1}\left(1-\frac{\alpha\beta}{\alpha+\kappa^2}\right)\right)c_{\pi,t+1}^{dg} \geq \frac{\alpha\beta}{\alpha+\kappa^2}A_{11,T} + A_{12,T} \tag{56}$$

Note that the RHS is equal to:

$$\begin{aligned}
\frac{\alpha\beta}{\alpha+\kappa^2}A_{11,T} + A_{12,T} &= \frac{\alpha\beta}{\alpha\beta(1+\beta\frac{1}{t+1}) + \kappa^2\beta}\left[\beta\left(1+\beta\frac{1}{t+1}\right)\left(1-\frac{1}{T+1}\left(1-\frac{\alpha\beta}{\alpha+\kappa^2}\right)\right)\right] \\
&= \frac{\alpha\beta}{\alpha+\kappa^2\left(1+\beta\frac{1}{t+1}\right)^{-1}}\left(1-\frac{1}{T+1}\left(1-\frac{\alpha\beta}{\alpha+\kappa^2}\right)\right) \\
&> \frac{\alpha\beta}{\alpha+\kappa^2}\left(1-\frac{1}{T+1}\left(1-\frac{\alpha\beta}{\alpha+\kappa^2}\right)\right)
\end{aligned}$$

where the last inequality is due to the fact that $\left(1+\beta\frac{1}{t+1}\right)^{-1} < 1$; putting together the last inequality and (56), we get:

$$c_{\pi,t+1}^{dg} > \frac{\alpha\beta}{\alpha+\kappa^2}$$

Then, we can apply the above argument to $c^{dg}_{\pi,t+2}$ as well and, proceeding by induction, conclude that $c^{dg}_{\pi,t} > \frac{\alpha\beta}{\alpha+\kappa^2}$ for any $t > T$. An immediate consequence is that $\lim_{t\to\infty} c^{dg}_{\pi,t} > \frac{\alpha\beta}{\alpha+\kappa^2}$, which is a contradiction with the result stated in first part of the Proposition, namely $\lim_{t\to\infty} c^{dg}_{\pi,t} = \frac{\alpha\beta}{\alpha+\kappa^2}$. Hence, we have shown that there is no $t < \infty$ such that $c^{dg}_{\pi,t} \geq \frac{\alpha\beta}{\alpha+\kappa^2}$.

□

Proof of Proposition 5. Recall that, as shown in Proposition 3, we have $\lim_{t\to\infty} c^{dg}_{\pi,t} = \frac{\alpha\beta}{\alpha+\kappa^2}$; since $0 < \frac{\alpha\beta}{\alpha+\kappa^2} < 1$, for any $\overline{C}$ with $\frac{\alpha\beta}{\alpha+\kappa^2} < \overline{C} < 1$, there exists a T such that, for any $t \geq T$ we will have $0 < c^{dg}_{\pi,t} < \overline{C}$; moreover, using the ALM for π_t, the law of motion of inflation expectations after T can be rewritten as[67]:

$$a_{t+1} = a_t + (t+1)^{-1}(c^{dg}_{\pi,t} - 1)a_t < a_t + (t+1)^{-1}(\overline{C} - 1)a_t$$

where the RHS of the inequality converges to zero, as shown in Evans and Honkapohja (2000). It is also easy to show that, $\forall t \geq T$ we have $a_{t+1} \geq 0$; thus, invoking the Policemen Theorem, we conclude that $\lim_{t\to\infty} a_t = 0$, i.e. inflation expectations converge to their RE value. □

Finally, we prove Proposition 7. First of all, we will briefly describe some results of stochastic approximation[68] that we will exploit in the proof.

Let's consider a stochastic recursive algorithm of the form:

$$\theta_t = \theta_{t-1} + \gamma_t Q\left(t, \theta_{t-1}, X_t\right) \tag{57}$$

where X_t is a state vector with an invariant limiting distribution, and γ_t is a sequence of gains; the stochastic approximation literature shows how, provided certain technical conditions are met, the asymptotic behavior of the stochastic difference equation (57) can be analyzed using the associated deterministic ODE:

$$\frac{d\theta}{d\tau} = h\left(\theta(\tau)\right) \tag{58}$$

where:

$$h\left(\theta\right) \equiv \lim_{t\to\infty} EQ\left(t, \theta, X_t\right)$$

E represents the expectations taken over the invariant limiting distribution of X_t, for any fixed θ. In particular, it can be shown that the set of limiting points of (57) is given by the stable resting points of the ODE (58).

[67]Without loss of generality, we are assuming that $a_T > 0$; if the opposite were true, a similar argument applies.

[68]Ljung (1977), Benveniste and P. (1990) provide a recent survey.

Proof of Proposition 7. Note that our equation (50) is a special case of (57), where the technical conditions are easily shown to be satisfied; moreover, it is also easy to see that:

$$h\left(a\right)=\lim_{t\to\infty}(c^{dg}_{\pi,t}-1)a=\left(\frac{\alpha\beta}{\alpha+\kappa^{2}}-1\right)a$$

which has a unique possible resting point at $a^{*}=0$. Since $\frac{\alpha\beta}{\alpha+\kappa^{2}}<1$, we have that a^{*} is globally stable, which proves the statement. □

C Comparison with EH Rule

Proof of Propositions 4 and 6. First of all, note that:

$$\delta^{dg}_{\pi,t}\gtreqless\delta^{EH}_{\pi}\Longleftrightarrow\sigma\frac{\beta-c^{dg}_{\pi,t}}{\kappa}\gtreqless\sigma\frac{\kappa\beta}{\alpha+\kappa^{2}}$$

where the second inequality can be rewritten as:

$$\frac{\beta}{\kappa}-\frac{\kappa\beta}{\alpha+\kappa^{2}}\gtreqless\frac{c^{dg}_{\pi,t}}{\kappa}$$

Rearranging the terms, we get:

$$\delta^{dg}_{\pi,t}\gtreqless\delta^{EH}_{\pi}\Longleftrightarrow\frac{\alpha\beta}{\alpha+\kappa^{2}}\gtreqless c^{dg}_{\pi,t}$$

Since we have shown in Proposition 3 that $t<\infty$ implies $c^{dg}_{\pi,t}<\frac{\alpha\beta}{\alpha+\kappa^{2}}$, we conclude that $\delta^{dg}_{\pi t}>\delta^{EH}_{\pi}$. Using a similar argument, it is easy to show that:

$$\delta^{dg}_{ut}\gtreqless\delta^{EH}_{u}\Longleftrightarrow\frac{\alpha}{\alpha+\kappa^{2}}\gtreqless d^{dg}_{\pi,t}$$

which implies, since

$$d^{cg}_{\pi}=\frac{\alpha}{\kappa^{2}+\alpha+\alpha\beta^{2}\gamma^{2}(\beta-c^{cg}_{\pi})+\beta\gamma\left(1-\gamma\right)\left(\alpha\beta-\left(\kappa^{2}+\alpha\right)c^{cg}_{\pi}\right)}<\frac{\alpha}{\alpha+\kappa^{2}},$$

that $\delta^{dg}_{ut}>\delta^{EH}_{u}$ whenever $t<\infty$. Finally, note that Proposition 3 also showed that $\lim_{t\to\infty}c^{dg}_{\pi,t}=\frac{\alpha\beta}{\alpha+\kappa^{2}}$, which trivially yields $\lim_{t\to\infty}\delta^{dg}_{\pi,t}=\delta^{EH}_{\pi}$ and $\lim_{t\to\infty}\delta^{dg}_{ut}=\delta^{EH}_{u}$. □

References

Adam, K., and M. Padula (2003): "Inflation dynamics and subjective expectations in the united states," Discussion paper.

Agarwal, R. P. (2002): *Difference equations and inequalities theory, methods, and applications,*. New York, Marcel Dekker cop.

Alvarez, F., P. J. Kehoe, and P. A. Neumeyer (2004): "The Time Consistency of Optimal Monetary and Fiscal Policies," *Econometrica*, 72(2), 541–567, available at http://ideas.repec.org/a/ecm/emetrp/v72y2004i2p541-567.html.

Aoki, K. (2002): "Optimal Commitment Policy Under Noisy Information," CEPR Discussion Papers 3370, C.E.P.R. Discussion Papers, available at http://ideas.repec.org/p/cpr/ceprdp/3370.html.

Barro, R. J., and D. B. Gordon (1983): "A Positive Theory of Monetary Policy in a Natural Rate Model," *Journal of Political Economy*, 91(4), 589–610, available at http://ideas.repec.org/a/ucp/jpolec/v91y1983i4p589-610.html.

Benveniste, A., M. M., and P. P. (1990): *Adaptive Algorithms and Stochastic Approximations*. Berlin: Springer.

Bernanke, B. S. (2004): "The Logic of Monetary Policy," Speech delivered by federal reserve governor ben bernanke to the national economists club in washington d.c., The Federal Reserve Board.

Blanchard, O. J., and C. M. Kahn (1980): "The Solution of Linear Difference Models under Rational Expectations," *Econometrica*, 48(5), 1305–11, available at http://ideas.repec.org/a/ecm/emetrp/v48y1980i5p1305-11.html.

Bray, M. (1982): "Learning, Estimation, and the Stability of Rational Expectations," *Journal of Economic Theory*, 26(2), 318–339.

Bray, M. M., and N. E. Savin (1986): "Rational Expectations Equilibria, Learning, and Model Specification," *Econometrica*, 54(5), 1129–60, available at http://ideas.repec.org/a/ecm/emetrp/v54y1986i5p1129-60.html.

Cagan, P. (1956): *The Monetary Dynamics of Hyperinflation in 'Studies in the Quantity Theory of Money'*. University of Chicago Press.

Calvo, G. (1983): "Staggered Prices in a Utility Maximizing Framework," *Journal of Monetary Economics*, 12, 383–398.

Clarida, R., J. Gali, and M. Gertler (1999): "The Science of Monetary Policy: A New Keynesian Perspective," *Journal of Economic Literature*, 37, 1661–1707.

——— (2000): "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory," *The Quarterly Journal of Economics*.

Erceg, C. J., A. Levin, and D. W. Henderson (2000): "Optimal Monetary Policy with Staggered Wage and Price Contracts," *Journal of Monetary Economics*, 46, 281–313.

Evans, G. (1985): "Expectational Stability and the Multiple Equilibria Problem in Linear Rational Expectations Models," *The Quarterly Journal of Economics*, 100(4), 1217–33, available at http://ideas.repec.org/a/tpr/qjecon/v100y1985i4p1217-33.html.

Evans, G. W., and S. Honkapohja (1993): "Adaptive forecasts, hysteresis, and endogenous fluctuations," *Economic Review*, pp. 3–13, available at http://ideas.repec.org/a/fip/fedfer/y1993p3-13n1.html.

——— (2000): "Convergence for difference equations with vanishing time-dependence, with applications to adaptive learning," *Economic Theory*, 15(3), 717–725.

——— (2001): *Learning and Expectations in Macroeconomics*. Princeton University Press.

——— (2002): "Monetary Policy, Expectations and Commitment," University of Oregon Economics Department Working Papers 2002-11, University of Oregon Economics Department, available at http://ideas.repec.org/p/ore/uoecwp/2002-11.html.

——— (2003a): "Adaptive Learning and Monetary Policy Design," *The Journal of Money, Credit and Banking*, 35, 1045–1072.

——— (2003b): "Expectations and the Stability Problem for Optimal Monetary Policies," *Review of Economic Studies*, 70(4), 807–824, available at http://ideas.repec.org/a/bla/restud/v70y2003i4p807-824.html.

Evans, G. W., and G. Ramey (2005): "Adaptive expectations, Underparametrization and the Lucas Critique," .

Ferrero, G. (2003): "Monetary Policy and the Transition to Rational Expectations," .

Forsells, M., and G. Kenny (2002): "The rationality of consumers' inflation expectations : survey-based evidence for the euro area," Discussion paper.

Fourgeaud, C., C. Gourieroux, and J. Pradel (1986): "Learning Procedures and Convergence to Rationality," *Econometrica*, 54(4), 845–68, available at http://ideas.repec.org/a/ecm/emetrp/v54y1986i4p845-68.html.

Gali, J. (2003): *New Perspectives on Monetary Policy, Inflation, and the Business Cycle published in Advances in Economic Theory*. Cambridge University Press.

GASPAR, V., F. SMETS, AND D. VESTIN (2005): "Optimal Monetary Policy under Adaptive Learning," .

HAMILTON, J. D. (1994): *Time Series Analysis*. Princeton University Press.

HANSEN, L. P., AND T. J. SARGENT (2001): "Robust Control and Model Uncertainty," *American Economic Review*, 91(2), 60–66, available at http://ideas.repec.org/a/aea/aecrev/v91y2001i2p60-66.html.

HONKAPOHJA, S., AND K. MITRA (2005): "Performance of monetary policy with internal central bank forecasting," *Journal of Economic Dynamics and Control*, 29(4), 627–658, available at http://ideas.repec.org/a/eee/dyncon/v29y2005i4p627-658.html.

LASALLE, J. P. (1986): *The Stability and Control of Discrete Processes*. Berlin: Springer-Verlag.

LEVIN, A., V. WIELAND, AND J. C. WILLIAMS (2003): "The Performance of Forecast-Based Monetary Policy Rules under Model Uncertainty," *American Economic Review*, 93(3), 622–645, available at http://ideas.repec.org/a/aea/aecrev/v93y2003i3p622-645.html.

LJUNG, L. (1977): "Analysis of recursive stochastic algorithms," *IEEE Transactions on Automatic Control*, (AC-22), 551–575.

LUCAS, R. E., AND N. L. STOKEY (1983): "Optimal Fiscal and Monetary Policy in an Economy Without Capital," *Journal of Monetary Economics*, 12, 55–93.

MARCET, A., AND J. P. NICOLINI (2003): "Recurrent Hyperinflations and Learning," *The American Economic Review*, 93(244), 1476–1498.

MARCET, A., AND T. J. SARGENT (1989): "Convergence of Least Squares Learning Mechanisms in Self Referential Linear Stochastic Models," *Journal of Economic Theory*, 48, 337–368.

MCCALLUM, B., AND E. NELSON (1999): *Performance of Operational Optimal Policy Rules in an Estimated Semi-Classical Model, published in Monetary Policy Rules*. Chicago: University of Chicago Press.

MILANI, F. (2005): "Expectations Learning and Macroeconomic Persistence," Job market paper.

MOLNAR, K. (2005): "Learning with Expert Advice," mimeo.

ORPHANIDES, A., AND J. WILLIAMS (2002): "Imperfect Knowledge, Inflation Expectations and Monetary Policy," *FED Working Paper*.

ORPHANIDES, A., AND J. C. WILLIAMS (2005a): "Inflation scares and forecast-based monetary policy," *Review of Economic Dynamics*, 8(2), 498–527, available at http://ideas.repec.org/a/red/issued/v8y2005i2p498-527.html.

——— (2005b): "Robust Monetary Policy with Imperfect Knowledge," Computing in Economics and Finance 400, Society for Computational Economics.

PHELPS, E. S. (1967): "Phillips Curves, Expectations of Inflation and Optimal Unemployment over Time," *Economica*, 34, 254–281.

ROBERTS, J. M. (1997): "Is inflation sticky?," *Journal of Monetary Economics*, 39(2), 173–196.

ROTEMBERG, J. J., AND M. WOODFORD (1998): "Interest-Rate Rules in an Estimated Sticky Price Model," NBER Working Papers 6618, National Bureau of Economic Research, Inc, available at http://ideas.repec.org/p/nbr/nberwo/6618.html.

SARGENT, T. J. (1999): *The Conquest of American Inflation*. Princeton University Press.

WIELAND, V. (2000a): "Learning by doing and the value of optimal experimentation," *Journal of Economic Dynamics and Control*, 24(4), 501–534, available at http://ideas.repec.org/a/eee/dyncon/v24y2000i4p501-534.html.

——— (2000b): "Monetary policy, parameter uncertainty and optimal learning," *Journal of Monetary Economics*, 46(1), 199–228, available at http://ideas.repec.org/a/eee/moneco/v46y2000i1p199-228.html.

WOODFORD, M. (1996): "Optimal Monetary Policy Inertia," *Journal of Monetary Economics*, 37(3), 573–605.

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

YUN, T. (1996): "Nominal price rigidity, money supply endogeneity, and business cycles," *Journal of Monetary Economics*, 37(2-3), 345–370, available at http://ideas.repec.org/a/eee/moneco/v37y1996i2-3p345-370.html.

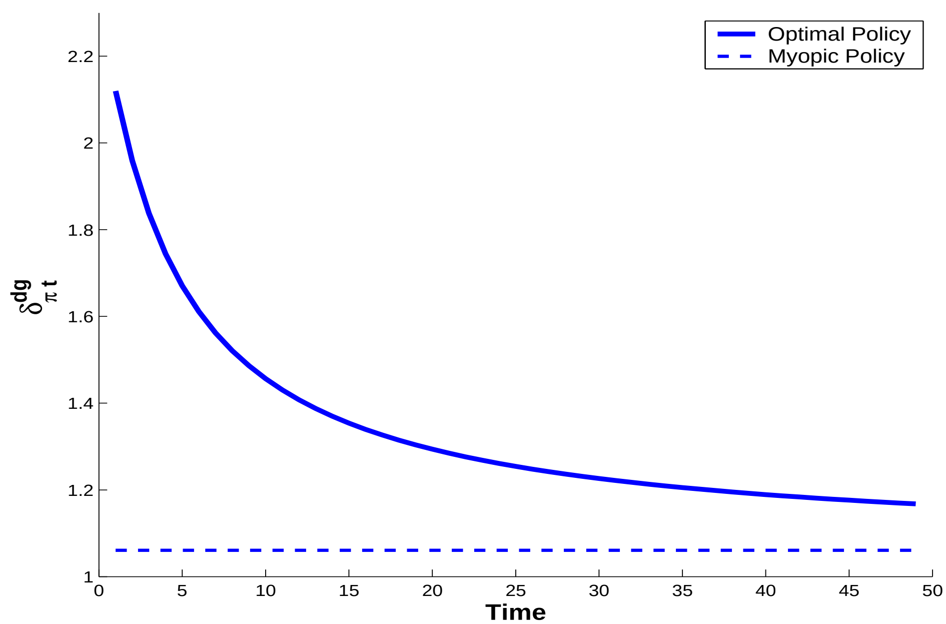

Figure 1: Interest rate rule coefficient on inflation expectations under decreasing gain learning.

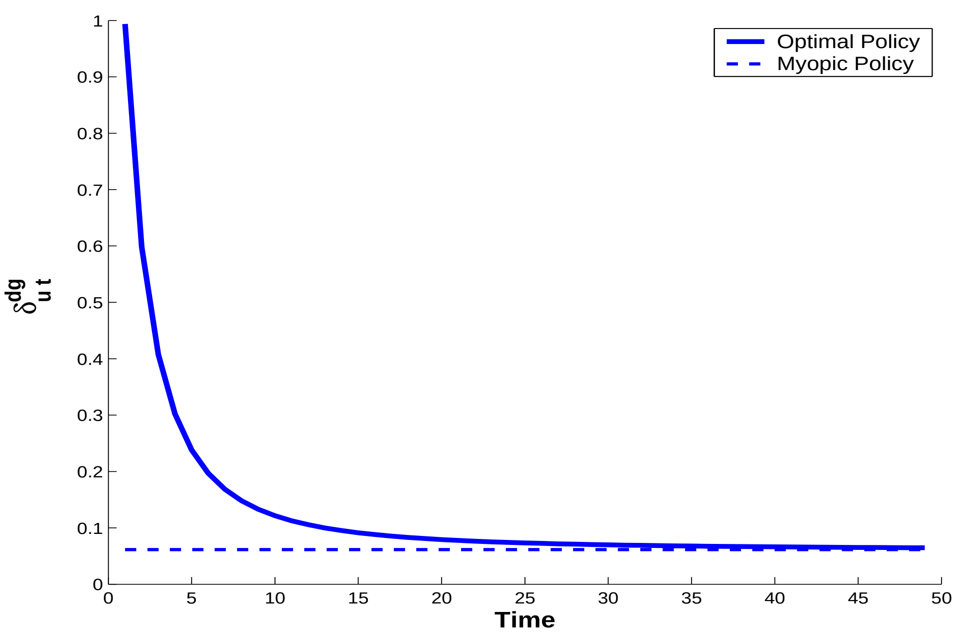

Figure 2: Optimal versus myopic interest rate rule: coefficient or the cost push shock

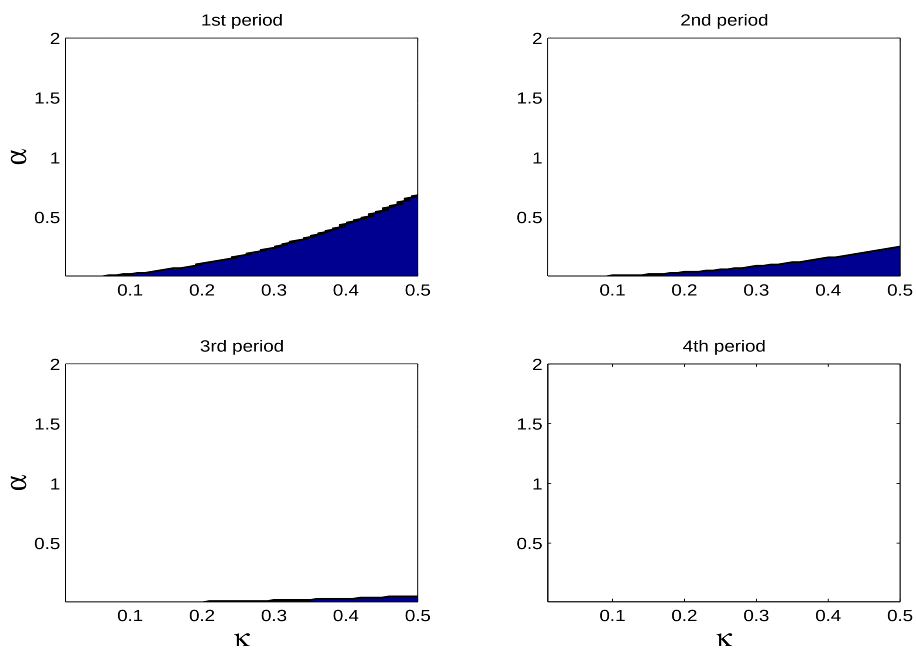

Figure 3: Values of α and κ for which δ_π^{dg} is increasing in the first 4 periods. From the 4th period on δ_π^{dg} is always decreasing. ($\beta = 0.99$)

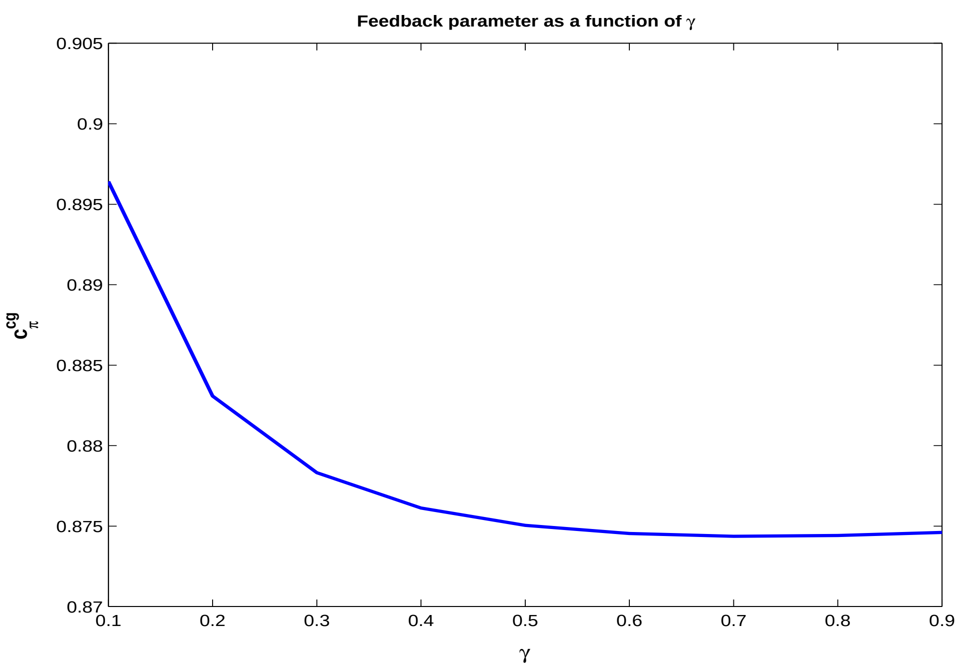

Figure 4: Feedback parameter in the ALM for inflation as a function of γ.

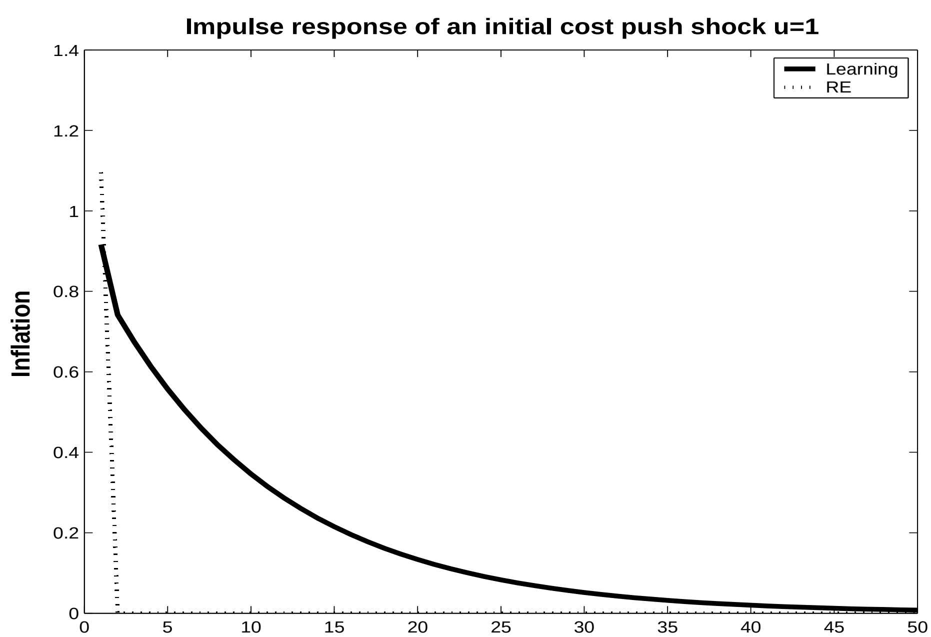

Figure 5: Impulse response of an initial cost-push shock $u = 1$ with optimal policy under learning and optimal discretionary policy under RE, starting from $a_0 = 0$, $\pi_0 = 0$, $x_0 = 0$.

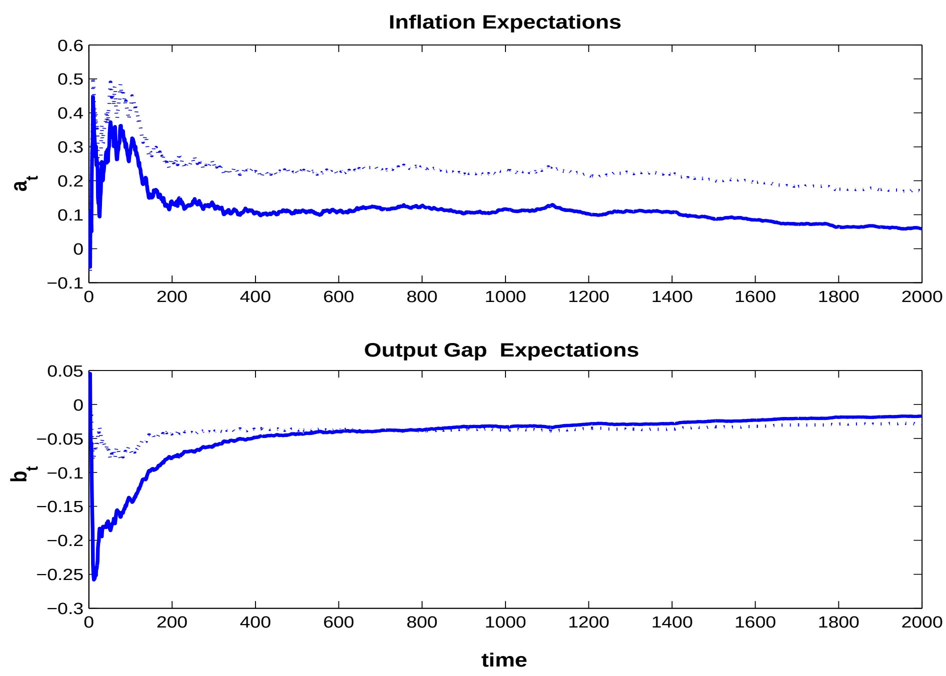

Figure 6: Evolution of inflation and output gap expectations for the optimal (solid line)and the myopic rule (dashed line), when agents follow decreasing gain learning.

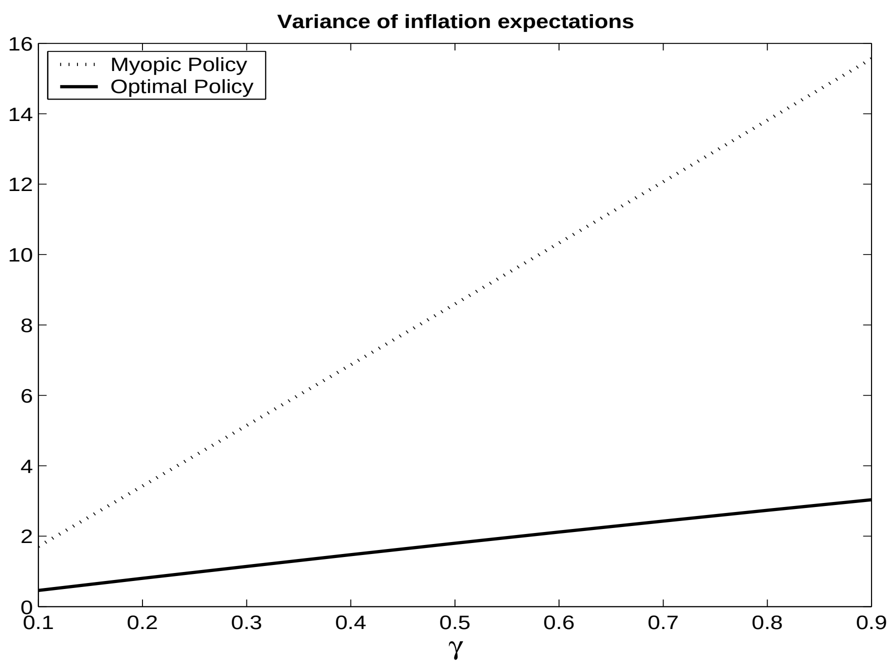

Figure 7: Variance of a_t.

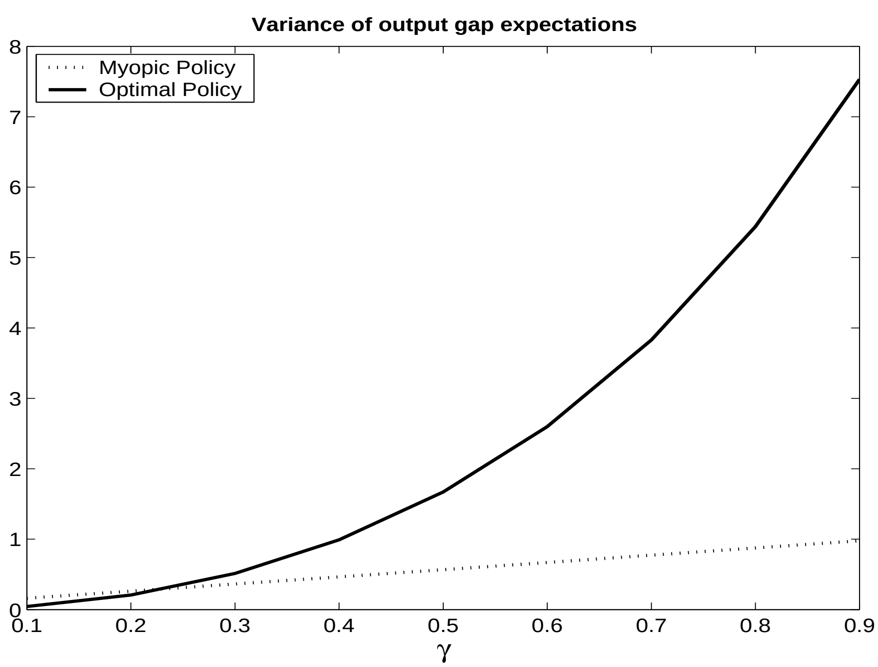

Figure 8: Variance of b_t.

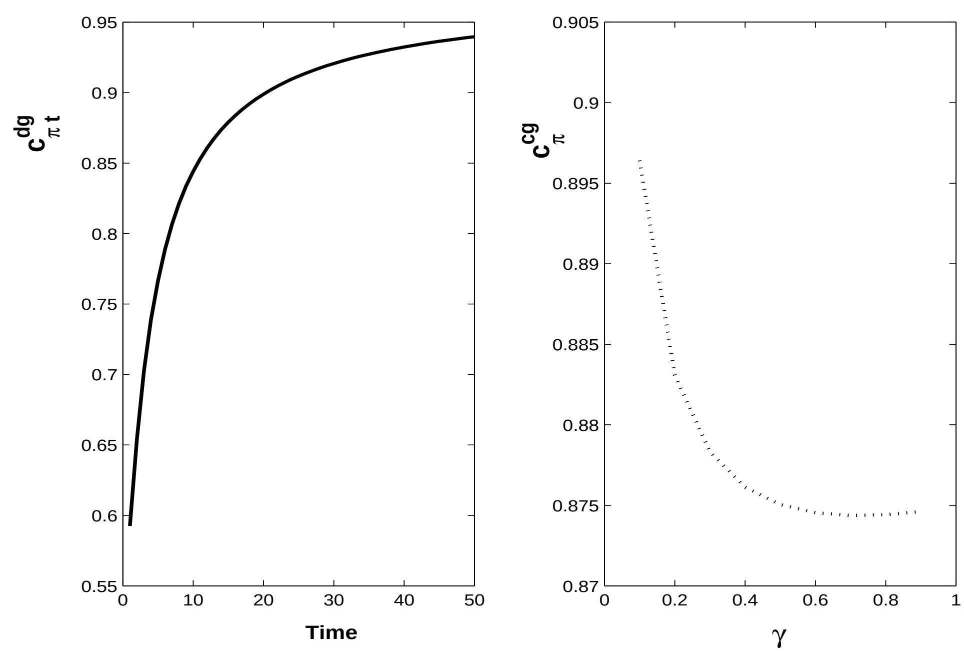

Figure 9: Feedback parameter under decreasing gain learning (DG) and constant gain learning (CG)