



NHH

INSTITUTT FOR SAMFUNNSØKONOMI

DEPARTMENT OF ECONOMICS

SAM 7 2007

ISSN: 0804-6824

FEBRUARY 2007

Discussion paper

Optimal Monetary Policy When Agents Are Learning

BY

KRISZTINA MOLNÁR AND SERGIO SANTORO

This series consists of papers with limited circulation, intended to stimulate discussion.

Norges
Handelshøyskole

NORWEGIAN SCHOOL OF ECONOMICS AND BUSINESS ADMINISTRATION

Optimal Monetary Policy When Agents Are Learning*

Krisztina Molnár[†] and Sergio Santoro[‡]

November 8, 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 trade-off 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 inflation expectations closer to equilibrium. This optimal policy qualitatively resembles optimal policy when the Central Bank can commit and agents have rational expectations. Moreover, 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

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 policy-makers’ task. (Bernanke (2004))

*We are especially grateful to our supervisors Albert Marcet and Ramon Marimon for very helpful comments and encouragements throughout this project. We are also indebted to Kosuke Aoki, Refet Gurkaynak, Seppo Honkapohja, Alessandro Rebucci, José-Víctor Ríos-Rull, Stephan Sauer, and Ákos Valentinyi for useful conversations and suggestions. Previous versions of this paper have been presented at the 2005 Spring Meeting of Young Economist in Geneva, the 2005 ASSET Conference, the 2005 EEA Conference, the 2006 CEF Conference, and several other workshops and seminars; we would like to thank the audiences for their comments. All remaining errors are our own. The views expressed herein are those of the authors, and do not necessarily reflect those of the Banca d’Italia.

[†]Norwegian School of Economics and Business Administration; Email: krisztina.molnar@nhh.no

[‡]Research Department, Banca d’Italia; Email: sergio.santoro@bancaditalia.it

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 However, there is a large body of evidence which suggests that agents’ forecasts are not consistent with the paradigm of rational expectations.² In response, a growing theoretical literature explores the robustness of the optimal policies, which were derived under rational expectations, when instead agents behave as econometricians and estimate a perceived law of motion, updating their expectations according to a learning algorithm.³ In a series of papers, Evans and Honkapohja show that interest rate rules that are optimal under rational expectations may lead to instability under learning, in the sense that agents’ beliefs may fail to converge to rational expectations.⁴ As much of the literature on learning, they are mainly interested in the asymptotic properties of the adaptive algorithms used to update forecasts, leaving aside the implications along the transition of the departure from rational expectations.

Instead, in this paper we consider both the short and the long run consequences that abandoning the rational expectations hypothesis has for monetary policy design. In particular, we take a normative approach, and address the issue of how a rational Central Bank should optimally conduct monetary policy, if the private sector forms expectations following an adaptive learning model. We consider a simple New Keynesian setup, augmented with the law of motion of agents’ beliefs. In the baseline case we assume that the law of motion estimated by the private sector is consistent with the optimal discretionary solution under rational expectations, and that the learning process follows a *constant gain* algorithm, namely a scheme that downweights geometrically past data. The Central Bank knows that agents are learning, and takes their expectations formation scheme into account when solving its control problem.

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 than what a rational expectations model suggests. Earlier work in the literature has shown computationally that, when the Central Bank uses simple monetary policy rules, similar results hold (see Ferrero (2003), Orphanides and Williams (2004b) and Orphanides and Williams (2005)); here we establish that these results extend to the case when the Central Bank uses the optimal policy, and provide a formal proof. In fact, using our analytical solution, we are able to show the existence of an *intertemporal trade-off* between optimal behavior in the current period and in the future; this trade-off, that is not present under rational expectations, induces the monetary authority to be more aggressive towards inflation. The intuition for this result is that aggressively driving inflation close to equilibrium helps private agents to learn the true equilibrium value of inflation at a faster pace. Hence, when the Central Bank takes into account its ability to manipulate future values of private sector beliefs, it is optimal to be more aggressive towards inflation, even if it generates a welfare cost in terms of an increased volatility of the output gap. We show analytically

¹See 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.

²See Roberts (1997), Forsells and Kenny (2002) and Adam and Padula (2003).

³For 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).

⁴See Evans and Honkapohja (2003a), Evans and Honkapohja (2003b) and Evans and Honkapohja (2006).

that 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 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 Central Bank 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 rational expectations 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 rational expectations, hence easing the short run trade-off between inflation and output gap. However, in our setup the optimal policy in the long run drives expectations to the discretionary rational expectations solution, and not to the Ramsey solution. The reason for this is that agents' expectations formation scheme does not nest the commitment solution under rational expectations. Under rational expectations and commitment the law of motion of inflation and output gap has a different functional form than in the discretionary case (see Clarida, Gali, and Gertler (1999)). This finding stresses the point that when we abandon the rational expectation paradigm, several issues arise in monetary policy design that are not present when agents are fully rational; moreover, if we take seriously the assumption that agents' beliefs evolve according to adaptive learning, the implications for policymaking go beyond the asymptotic learnability criterion.

Our next contribution is to derive optimal policy under *decreasing gain* learning, namely when agents' beliefs are formed giving equal weight to all past data. We show that our main results are robust to the change in 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 governor, the Central Bank should be more aggressive in containing inflation 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, when only few data points are available, 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 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.

An early example of optimal monetary policy when agents are learning is Sargent (1999), Chapter 5, where the model is characterized by a Phillips Curve à la Lucas and private agents follow a constant gain algorithm. In this case, the discretionary and commitment solutions under rational expectations have the same functional form, so that Sargent is able to show that a sufficiently patient Central Bank can make the private sector learn the Ramsey equilibrium⁵. A parallel paper of Gaspar, Smets, and Vestin (2005) provides a numerical solution to optimal monetary policy in the New Keynesian framework with indexation to lagged inflation among firms, when private agents estimate the persistence of inflation through a constant gain algorithm. They show that an optimally behaving Central Bank aims to decrease the limiting variance of the private sector's inflation expectations, in order to anchor inflation expectations better. Moreover, they show that optimal policy qualitatively resembles the commitment solution under rational expectations. Another important result they find is that, when the degree of estimated persistence is high the Central Bank should be more aggressive. This framework is used in Gaspar, Smets, and Vestin (2006) to show that the behavior of US economy around the Volcker disinflation is consistent with a regime change from a simple policy that does not take into account expectations formation, to the optimal policy.

The rest of the paper is organized as follows: in Section 2 we briefly recall the discretionary optimal policy when expectations are rational, and analyze optimal policy under constant gain learning ; Section 3 relaxes the assumption that expectations follow constant gain learning, and show that our main results remain valid under decreasing gain learning. Section 4 relaxes the assumption that the policy maker can perfectly observe the fundamental shocks and the beliefs of the agents, and argues that the optimal policy rule derived in the previous Sections is robust to uncertainty about the agents' expectations formation mechanism. Section 5 concludes.

2 The Baseline Model

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 equations⁶. The first one is an IS equation:

$$x_t = E_t^* x_{t+1} - \sigma^{-1}(r_t - E_t^* \pi_{t+1} - \bar{r}r_t) \quad (1)$$

where x_t , r_t and π_t denote time t output gap⁷, short-term nominal interest rate and inflation, respectively; σ is a parameter of the household's utility function, representing risk aversion, and $\bar{r}r_t$ 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 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):

$$\pi_t = \beta E_t^* \pi_{t+1} + \kappa x_t + u_t \quad (2)$$

⁵See subsection 2.4.

⁶For 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).

⁷Namely, the difference between actual and natural output.

where β denotes the subjective discount rate, κ is a function of structural parameters, and $u_t \sim N(0, \sigma_u^2)$ is a white noise cost-push shock⁸; this relation is obtained under the assumption 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))⁹. The coefficient κ is decreasing in the level of stickiness: the longer are prices fixed in expectation the smaller is the effect of the output gap on inflation.

The standard New Keynesian literature imposes the existence of rational expectations (RE), namely that $E_t^* = E_t$. Under this assumption, the full commitment solution of the optimal monetary policy turns out to be time inconsistent, even if the Central Bank (CB) does not have a target for output gap larger than zero. In other words, even if we rule out the possibility of the inflation bias discussed in Barro and Gordon (1983) and all the subsequent literature, there are potential welfare gains associated with the presence of a credible commitment device for the CB. Hence, the time-consistent discretionary solution is 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 introduced¹⁰.

The loss function of the CB is given by:

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

where α is the relative weight put by the CB on the objective of output gap stabilization¹¹.

2.1 Benchmark: discretionary solution under rational expectations and under learning

Let's assume that the CB takes the private sector beliefs as given. In Kreps (1998) terminology, this is equivalent to suppose that the monetary authority is an anticipated utility maximizer.

The policy problem is to choose a time path for the nominal interest rate r_t ¹² to engineer a law of motion 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.

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

⁸Note 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.

⁹In other words, the probability that firm i in period t can reset the price is constant over time and across firms.

¹⁰See Woodford (2003), Clarida, Gali, and Gertler (1999) and McCallum and Nelson (1999).

¹¹As is shown in Rotemberg and Woodford (1998), equation (3) can be seen as a quadratic approximation to the expected household's utility function; in this case, α is a function of structural parameters.

¹²We 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 .

Because there are no endogenous state variables, problem (4) reduces to a sequence of static optimization problems. As shown in Clarida, Gali, and Gertler (1999), the optimality condition to this problem (at time t) is

$$\frac{\kappa}{\alpha} \pi_t + x_t = 0. \quad (5)$$

Combining (5) with the structural equations, one can derive the following law of motion for inflation and output gap:

$$\pi_t^{EH} = \frac{\alpha\beta}{\alpha + \kappa^2} E_t^* \pi_{t+1} + \frac{\alpha}{\alpha + \kappa^2} u_t \quad (6a)$$

$$x_t^{EH} = -\frac{\kappa\beta}{\alpha + \kappa^2} E_t^* \pi_{t+1} - \frac{\kappa}{\alpha + \kappa^2} u_t. \quad (6b)$$

and the interest rate rule that implements this allocations:

$$r_t = \bar{r} \bar{r}_t + \delta_\pi^{EH} E_t^* \pi_{t+1} + \delta_x^{EH} E_t^* x_{t+1} + \delta_u^{EH} u_t \quad (7)$$

where:

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

Throughout the paper we denote the coefficients by EH referring to the paper Evans and Honkapohja (2003b) (EH hereafter), where the authors derive a rule analogous to (7). In the terminology introduced in Evans and Honkapohja (2003b), Evans and Honkapohja (2003a), this is an *expectations-based reaction function*; they 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.

If the agents have RE (i.e., if $E_t^* = E_t$), Clarida, Gali, and Gertler (1999) show that the solution of (4) yields:

$$\begin{aligned} \pi_t^{RE} &= \frac{\alpha}{\kappa^2 + \alpha} u_t \\ x_t^{RE} &= -\frac{\kappa}{\kappa^2 + \alpha} u_t. \end{aligned}$$

Under RE, the assumption that the monetary authority takes private sector beliefs as given has a precise motivation in terms of lack of credibility¹³: if the CB is free to reoptimize every period, agents take it into account ignoring any promise it makes on the future. As a result, the discretionary RE equilibrium has the property that the CB has no incentive to change its policy (it is time consistent).

If private agents follow learning, a fully rational CB could do better than (7). In the next section we show how optimal monetary policy is modified when the CB optimizes taking into account its effect on private expectations.

¹³In the literature this case is known as optimal policy under discretion.

2.2 Constant Gain Learning

We now assume that private sector's expectations are formed according to the adaptive learning literature¹⁴; we assume that agents do not know the exact process followed by the endogenous variables, but recursively estimate a Perceived Law of Motion (PLM) consistent with the law of motion that the CB would implement under RE. As explained above, the optimal allocations of the discretion and the commitment solution under RE have different functional forms, and are therefore associated with different PLMs. 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 $x_t = x(u_t)$ ¹⁵; 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 estimate them using their sample means. Throughout this section we will assume that expectations evolve following the algorithm¹⁶:

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

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

where $\gamma \in (0, 1)$ is the gain parameter, constant through time.

The use of constant gain algorithms to track structural changes is well known from the statistics and engineering literature¹⁷. 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 'trust more' recent data. This approach is closely related to using a fixed sample length, or rolling window regressions.

In Section 3 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:

$$\begin{aligned} \min_{\{\pi_t, x_t, r_t, a_{t+1}, b_{t+1}\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \beta^t (\pi_t^2 + \alpha x_t^2) \\ \text{s.t. (1), (2), (8), (9)} \\ a_0, b_0 \text{ given} \end{aligned} \quad (10)$$

This optimization problem is linear quadratic, the Bellman equation holds, thus the resulting policy is time consistent¹⁸.

¹⁴For an extensive monograph on this paradigm, see Evans and Honkapohja (2001).

¹⁵In the terminology of Evans and Honkapohja (2001) Chapter 11, the PLM is a noisy steady state.

¹⁶To be precise, in the algorithms (8), (9) the observations are weighted geometrically, while in the normal sample average they all receive equal weight.

¹⁷See for example Benveniste, Métivier, and Priouret (1990), Part I. Chapters 1. and 4.

¹⁸A 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 $t \geq 0$.

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

$$\lambda_{1t} = 0 \quad (11)$$

$$2\pi_t - \lambda_{2t} + \gamma\lambda_{3t} = 0 \quad (12)$$

$$2\alpha x_t + \kappa\lambda_{2t} - \lambda_{1t} + \gamma\lambda_{4t} = 0 \quad (13)$$

$$E_t \left[\frac{\beta}{\sigma} \lambda_{1t+1} + \beta^2 \lambda_{2t+1} + \beta(1-\gamma)\lambda_{3t+1} \right] = \lambda_{3t} \quad (14)$$

$$E_t [\beta\lambda_{1t+1} + \beta(1-\gamma)\lambda_{4t+1}] = \lambda_{4t} \quad (15)$$

where λ_{it} , $i = 1, \dots, 4$ 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:

$$\lambda_{4t} = 0 \quad (16)$$

Inflation-Output Gap Trade-off

If we put together equations (11)-(13) and (16), we derive the following optimality condition:

$$2\pi_t + 2\frac{\alpha}{\kappa}x_t + \gamma\lambda_{3,t} = 0, \quad (17)$$

where $\lambda_{3,t}$ is the lagrange multiplier on the evolution of inflation expectations.

From (17) we can isolate two trade-offs faced by the CB in designing the optimal policy. When $\gamma = 0$, namely when expectations are constant and, consequently, cannot be manipulated by the monetary authority, (17) simplifies to:

$$\frac{\kappa}{\alpha}\pi_t + x_t = 0, \quad (18)$$

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). When a cost-push shock is present, (18) represents a well known *intratemporal trade-off* between stabilization of inflation at t and output gap at t : because of the nonzero term u_t in the Phillips Curve (2), π_t and x_t cannot be set contemporaneously equal to zero in every period. Clarida, Gali, and Gertler (1999) describe (18) 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).

Under learning (i.e., when $\gamma > 0$), it turns out that the CB faces an additional *intertemporal trade-off* between optimal behavior at t and later periods, generated by its ability to manipulate future values of a . The CB has to take into account how its choice about inflation/output at time t influences inflation expectations, and thus future intratemporal trade-offs between inflation/output.

The term $\gamma\lambda_{3,t}$ shows an important difference compared to earlier results: the optimal decision should condition on the current stance of inflation expectations. The interpretation of this term is very simple, if π_t changes equation (8) implies that this will influence next periods inflation expectations, a_{t+1} , with a factor γ , and a change in inflation expectations affects welfare losses with a factor $\lambda_{3,t}$. The sign of $\lambda_{3,t}$ depends on current inflation expectations: since the steady state inflation is zero (see the section on this page) an increase in inflation expectations drives them further from the steady state when expectations are positive, this in turn increases welfare loss so the lagrange multiplier on inflation expectations is positive. When inflation expectations are negative, the opposite happens: increasing inflation expectations drives them closer to the steady state, thus $\lambda_{3,t}$ is negative.

When inflation expectations are positive (so $\lambda_{3,t} > 0$) and inflation is positive, the optimal contraction of x_t is harsher than under discretionary policy. It is well documented in the literature that disinflations have real costs¹⁹ and slowly adjusting expectations might give an explanation for this²⁰. Our results show that under learning costs of disinflation are not only a consequence of slowly adjusting expectations, but it is indeed optimal to incur high output losses (compared to discretionary policy) in order to contain inflation expectations. Moreover, the higher are inflation expectations, the higher is $\lambda_{3,t}$ and the bigger is the output loss the CB should engineer in order to bring down inflation.

In fact (17) also implies that the lean against the wind policy is not always optimal. If for example inflation is positive but inflation expectations are sufficiently negative the optimal value of x_t can be zero or even positive.

Let us summarize our first result for later reference:

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

Optimal allocations

We can combine the conditions for an optimum to characterize analytically the optimal allocations implemented by the CB; the results are summarized in the following Proposition.

Proposition 1. *There exists a unique solution of the control problem (10), and the policy function for inflation associated to it has the form:*

$$\pi_t = c_\pi^{cg} a_t + d_\pi^{cg} u_t \quad (19)$$

The coefficient c_π^{cg} can be characterized as follows:

$$\begin{aligned} \text{-if } \gamma &\in (0, 1), \text{ we have that } 0 < c_\pi^{cg} < \frac{\alpha\beta}{\alpha + \kappa^2}, \\ \text{-if } \gamma &= 0, \text{ i.e. if expectations are constant, we have that } c_\pi^{cg} = \frac{\alpha\beta}{\alpha + \kappa^2}, \end{aligned}$$

and:

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

¹⁹For evidence on the costs of ending moderate inflations see for example Ball (1994). Note that our model is valid only around the steady state, so can not be used to model hyperinflationary episodes.

²⁰ See for example Brayton and Tinsley (1996), Erceg and Levin (2003).

Proof. See the Appendix. □

Following the adaptive learning terminology, we call (19) the Actual Law of Motion (ALM) of inflation.

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 1 we show a numerical example with the calibration found in Woodford (1999), 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 . With different parametrization, characterized by a higher κ and a lower α , the relationship would indeed be non monotonic, with c_π^{cg} being a decreasing function of γ for small values of the tracking parameter, and increasing when γ is big; however, the papers that try to estimate the tracking parameter typically find that γ is smaller than 0.1²¹, so that the decreasing brunch of c_π^{cg} as a function of the gain parameter seems the most relevant from an empirical point of view.

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

$$x_t = c_x^{cg} a_t + d_x^{cg} u_t \quad (20)$$

where:

$$\begin{aligned} c_x^{cg} &= \frac{c_\pi^{cg} - \beta}{\kappa} \\ d_x^{cg} &= \frac{d_\pi^{cg} - 1}{\kappa} \end{aligned}$$

$c_\pi^{cg} < \frac{\alpha\beta}{\alpha+\kappa^2}$ (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 (19) and (20) in (1) we can derive the nominal interest rate:

$$r_t = \bar{r}r_t + \delta_\pi^{cg} a_t + \delta_x^{cg} b_t + \delta_u^{cg} u_t \quad (21)$$

where:

$$\begin{aligned} \delta_\pi^{cg} &= 1 - \sigma \frac{c_\pi^{cg} - \beta}{\kappa} \\ \delta_x^{cg} &= \sigma \\ \delta_u^{cg} &= -\sigma \frac{d_\pi^{cg} - 1}{\kappa} \end{aligned}$$

The interest rate rule (21) 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

²¹For examples of estimates of γ , see Milani (2005) and Orphanides and Williams (2004a).

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 (21)). Moreover, the coefficient on b_t is such that its effect on the output gap in the IS curve is 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) inflation expectations in the next period will decrease. This criterion -also known as the ‘‘Taylor principle’’- is 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 criterion for evaluating monetary policy²².

Plugging (19) into (8), we get:

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

which is a stationary²³ 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 $a_t \sim N\left(0, \frac{\gamma^2 (d_{\pi}^{cg})^2}{1 - (1 - \gamma(1 - c_{\pi}^{cg}))^2} \sigma_u^2\right)$.

2.3 Comparison with the EH rule

In this section we state results regarding how optimal monetary policy under constant gain learning differs from rules used earlier in the literature, where the CB is treated as an anticipated utility maximizer (i.e., it considers expectations as given in the optimization problem); in particular we refer to rule (7), derived in EH.

It is clear that the coefficient on the output gap expectations is the same in rule (7) as in rule (21), while the other two coefficients are typically different. Proposition 1 implies $\delta_{\pi}^{cg} > \delta_{\pi}^{EH}$: 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 it would do otherwise.

From Proposition 1 it also follows that $\delta_u^{cg} > \delta_u^{EH}$: optimal policy reacts more aggressively also to cost push shocks. After a positive cost push shock the optimally behaving CB raises the interest rate more aggressively than in the case of an anticipated utility maximizer CB; 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.

An analogous difference emerges when we compare the allocations implemented by the two different interest rate rules; under constant gain learning optimal allocations are characterized by (19)-(20), while EH allocations are given by (6) with $E_t^* \pi_{t+1} = a_t$.

²²Clarida, Gali, and Gertler (2000) estimate that the pre-Volcker era violated this simple criterion.

²³In fact, since $0 < c_{\pi}^{cg} < 1$, it immediately follows that $0 < (1 - \gamma(1 - c_{\pi}^{cg})) < 1$.

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 (21) is used instead of (7) (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.

On the other hand, under OP both coefficients in the ALM of x_t are higher in absolute value than under EH, hence allowing a higher feedback from out of equilibrium expectations and noisy cost push shocks to the output gap.

The difference between (7) and (21) 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 inflation expectations) and increase the effect of out of equilibrium expectations on the output gap compared to the EH policy; moreover, it accommodates less the effect of noisy shocks to inflation compared to the EH policy, even if it translates into a more volatile output gap.*

Welfare Loss Analysis

To have a quantitative feeling of the welfare gains that the use of the optimal rule (21) instead of the EH rule (7) implies, we present a numerical welfare loss analysis.

Since welfare losses in utility terms are hard to interpret we report consumption equivalents: for a given monetary policy rule we calculate the cumulative utility losses resulting from deviations from the steady state allocation and then express what is the equivalent percentage decrease of the steady state consumption that results in the same cumulative utility loss (For details of the calculation see the Appendix.). We use the calibration of Woodford (1999): $\beta = 0.99$, $\kappa = 0.024$, $\alpha = 0.048$ and $\sigma = 0.157$. We perform Monte Carlo with a simulation length 10,000 and a cross sectional sample size of 1000, with the initial condition $a_0 = b_0 = 0$. Cost push shocks are drawn from a normal distribution with 0 mean and variance 0.1.

Table 1 reports consumption equivalents for a range of tracking parameters. For small tracking parameters the results are in the range of Lucas' original estimates²⁴: consumption losses resulting from cyclical fluctuations are small. For higher tracking parameters the consumption equivalents are also higher, which results from the fact in the presence of a cost push shock, constant gain learning does not settle down to RE, but converges to a limiting distribution and the limiting variance of inflation expectation increases in γ (keeping other coefficients constant). This is illustrated in Figure 3. A higher variance of inflation expectations implies higher variance of inflation and output gap, thus a higher welfare loss both under OP (see equation (19) and (20)) and under EH (see equation (6)).

We would like to note that inflation and output gap variance can be expressed as a linear function of the variance of the cost push shock, therefore the absolute value of consumption equivalents is bigger for a bigger σ_u^2 , but the ratio of consumption equivalents under OP and EH are not sensitive to the choice of σ_u^2 .

Optimal policy decreases consumption equivalents relative to the EH rule (7) (see the third column in Table 1). Even for tracking parameters below 0.05²⁵ the gain from using an optimal

²⁴See Lucas (1987).

²⁵Estimates for the US are typically in this range. 0.0187 is the estimation of Milani (2005) with Bayesian estimation, for a calibration of the tracking parameter see Orphanides and Williams (2004a).

interest rate rule can be around 1 – 3%. The gain in consumption equivalents is higher the higher is the gain parameter. For a very high tracking parameter $\gamma = 0.9$ the welfare loss in consumption terms of not using the optimal rule is twice as large as under OP. This follows from the fact that, optimal policy takes into account that expectations have a limiting variance while the EH policy considers expectations to be fixed²⁶. As a result, optimal policy aims to decrease the limiting variance of inflation expectations while EH does not, and the higher is the tracking parameter the bigger is the decrease in the limiting variance OP engineers compared to EH (see Figure 3).

Table 1: Consumption equivalents using OP and EH under constant gain learning

Tracking parameter	p^{OP}	p^{EH}	p^{OP}/p^{EH}
0.0187	0.0129	0.0129	0.9990
0.05	0.0148	0.0151	0.9759
0.08	0.0171	0.0185	0.9243
0.1	0.0188	0.0213	0.8830
0.3	0.0371	0.0619	0.5996
0.5	0.0554	0.1122	0.4935
0.9	0.0910	0.2217	0.4106

Woodford (1999) calibration

It is interesting to examine the composition of welfare losses coming from inflation variation and output gap variation. For this we calculate the equivalent permanent consumption decrease for welfare losses caused by only inflation variation or output gap variation respectively, and report the ratios of OP and EH in Table 2. The table demonstrates Result 2: compared to EH, the optimal policy lowers inflation variation even at the cost of higher output gap variation. The higher is the tracking parameter, the higher is the incentive of the Central Bank to focus on lowering inflation variance and allowing for an increase in output gap deviation from the flexible price equilibrium. For $\gamma = 0.9$ compared to EH an optimally behaving Central Bank engineers a 78% lower welfare loss in inflation when it properly conditions on expectation formation, permitting at the same time 15 times more variation in output gap.

Table 2: Ratio of consumption equivalents of losses due to inflation and output gap variations using OP and EH under constant gain learning

Tracking parameter	Inflation	Output gap
0.0187	0.9962	1.2296
0.05	0.9441	3.6263
0.08	0.8511	7.0185
0.1	0.7853	9.0290
0.3	0.4187	15.6711
0.5	0.3073	16.0060
0.9	0.2286	15.5719

Woodford (1999) calibration

²⁶It 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.

Moreover, it is worth noting that the use of the rule (7) 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 expectations²⁷.

This section 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 using the optimal policy are particularly pronounced when private agents use a high tracking parameter (i.e. discount more past data). Optimal policy under learning is characterized by a more aggressive interest rate reaction to out of equilibrium expectations and to cost push shocks than would be optimal when the CB does not make active use of its influence on expectations. This aggressive interest rate policy guarantees that inflation will deviate less from its equilibrium expected value, thus private agents can learn the true expected value of inflation faster than under EH policy. Containing inflation expectations is beneficial, even at the cost of allowing higher deviations in output gap expectations and a higher output gap volatility.

2.4 Comparison with the commitment solution

In this section we show that the optimal policy response to a supply shock under learning is qualitatively similar to that of the commitment solution under RE. However, despite of the similarities in short run behavior, in the limit the two equilibria are different. The learning equilibrium intrinsically depends on how private agents learn.

There is a qualitative similarity of optimal policy under learning to the commitment RE solution in the following sense, after a one time cost push shock (1) inflation rises on impact less then under discretionary policy and (2) after the cost push shock dies out inflation and the output gap remains more persistent compared to discretionary policy.

Both under learning and RE the small response of contemporaneous inflation to a temporary cost push shock results from the CB's ability to directly manipulate private expectations, even if the channels used are quite different. 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* (Result 2). 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 and output after an initial cost push shock. Under commitment the the source of inertia is that the policy maker carries commitments made in the past (in other words commits to behave in a past dependent way). Under learning it is the expectations of the private sector which is past dependent and changes sluggishly. So the source of inertia in both instances is something linking the periods: under

²⁷It can be easily derived that the autocorrelation of inflation under constant gain with EH is $E\pi_t^{EH}\pi_{t-1}^{EH} = \left(\frac{\alpha\beta}{\alpha+\kappa^2}\right)^2 \left(1 - \gamma + \gamma\frac{\alpha\beta}{\alpha+\kappa^2}\right)\sigma_{a_{EH}}^2 + \frac{\alpha\beta}{\alpha+\kappa^2} \left(\frac{\alpha}{\alpha+\kappa^2}\right)^2 \gamma\sigma_u^2$ while under the optimal rule $E\pi_t^{OP}\pi_{t-1}^{OP} = (c_\pi^{cg})^2 (1 - \gamma + \gamma c_\pi^{cg})\sigma_{a_{OP}}^2 + c_\pi^{cg} (d_\pi^{cg})^2 \gamma\sigma_u^2$. We have already seen that $\sigma_{a_{OP}}^2 < \sigma_{a_{EH}}^2$, $c_\pi^{cg} < \frac{\alpha\beta}{\alpha+\kappa^2}$ and $d_\pi^{cg} < \frac{\alpha}{\alpha+\kappa^2}$, thus $E\pi_t^{OP}\pi_{t-1}^{OP} < E\pi_t^{EH}\pi_{t-1}^{EH}$.

commitment and RE it is the commitment of the CB while under learning it is the expectation formation of the private sector.

As a result of these two similarities, the impulse response function of inflation to a temporary cost push shock will be also similar under OP and RE commitment. Figure 2 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 out²⁸. 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. 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. 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.

A difference compared to the impulse response of inflation under full commitment RE is that there is no overshooting of inflation under learning. Commitment policy under RE engineers a sequence of negative inflation after the first period, while a positive sequence under learning. In this sense the learning impulse response is more consistent with central bank practices: after a positive supply shock central banks gradually bring inflation back to the target²⁹. A second difference is that the full commitment is characterised by a smaller output decrease compared to RE discretionary policy (see Clarida, Gali, and Gertler (1999)), while under learning the initial decrease of output is bigger than under discretion and RE. Under RE the commitment of the central bank can improve the current terms of inflation output trade-off, while under learning monetary policy can only influence the future private expectations and improve future inflation output trade-offs.

Interestingly both the inflation and output gap impulse response show a *closer qualitative similarity with the optimal RE equilibrium under commitment within a simple class of policy rules*³⁰. Clarida, Gali, and Gertler (1999) derives that if the CB can commit to a policy rule that is a linear function of the cost push shock $-\pi_t = cu_t$, $c \in \mathbb{R}$ - and u is AR(1) the solution can be characterized by inequalities analogous to the Result 2. 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.

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 expectations³¹, 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. The intuition is simple: since in the Phelps problem the discretion and commitment outcome of inflation have the same functional form, but different coefficients, a CB patient enough is willing to trade off higher short term losses for the opportunity that in the long run the private sector learns the welfare-improving Ramsey equilibrium. As we stressed above, in our more general case discretionary

²⁸Recall that u is *i.i.d.*

²⁹This argument was made by Steinsson (2003).

³⁰We would like to note that for smaller γ s the impulse responses resemble more the discretionary RE ones, while for higher γ s they resemble more the limited commitment ones.

³¹Phelps (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.

and commitment solutions under RE show different functional form; hence, Sargent’s result does not hold anymore: even if the PLM of the agents is reasonable, in the sense that it coincides with the optimal RE equilibrium under discretion, it is impossible for the CB to drive the economy towards the Ramsey equilibrium, irrespective of its patience. This finding strengthens the point that when we abandon the RE paradigm, several issues arise in monetary policy design that are not present when agents are fully rational; moreover, if we take seriously the assumption that agents’ beliefs evolve according to adaptive learning, the implications for policymaking go beyond the asymptotic learnability criterion. However, also in our case an increase in the discount factor makes the optimal disinflationary path under learning get closer to the commitment solution. This can be seen in Table 3, 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 (1999), 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 3: Paths of inflation for different β s 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 (1999) calibration. Cost push shock $u_0 = 1$ in the first period, starting from $a_0 = 0$, $\pi_0 = 0$, $x_0 = 0$, with $\gamma = 0.2$

3 Decreasing Gain Learning

In the previous section we assumed that private sector beliefs are updated according to a constant gain algorithm; this assumption is particularly appropriate when agents believe structural changes to occur. If instead the private sector confidently believes that the environment is stationary³², it is more reasonable to model their learning behavior with a decreasing gain rule, namely an algorithm of the form:

$$E_t^* \pi_{t+1} \equiv a_t = a_{t-1} + t^{-1}(\pi_{t-1} - a_{t-1}) \quad (22)$$

$$E_t^* x_{t+1} \equiv b_t = b_{t-1} + t^{-1}(x_{t-1} - b_{t-1}) \quad (23)$$

where the only difference with (8)-(9) is the substitution of γ with t^{-1} . An updating scheme of this form is equivalent³³ to estimating inflation and output gap every period with OLS³⁴.

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

³²Note that in our model there are no structural breaks, or other non-stationary elements.

³³Under certain conditions on the values used to initialize the algorithm, see Evans and Honkapohja (2001).

³⁴Note that, since inflation and output gap are assumed by the learners to be constant, the OLS is just the sample averages of the two.

nature of expectations imply that during the transition the optimal policy should be time varying even in a stationary environment.

An additional reason to investigate if our previous analysis is robust to this change in the learning algorithm is rooted in the literature on predictor switching³⁵, which emphasizes that agents are likely to change their expectations formation scheme according to past performance, measured in terms of forecast errors. In particular, Marcet and Nicolini (2003) argue that agents are likely to use constant gain when the prediction errors are large, but they are likely to use decreasing gain when the prediction errors are small; if the private sector beliefs start far from equilibrium, constant gain may generate good forecast errors at the beginning but decreasing gain would generate good forecast errors at the end. Even if a rigorous study of a model with endogenously changing learning scheme is beyond the scope of this paper, the fact that our main results are common to both constant and decreasing gain algorithms strengthens our conclusions.

With DG, the problem of the CB becomes:

$$\begin{aligned} & \min_{\{\pi_t, x_t, r_t, a_{t+1}, b_{t+1}\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \beta^t (\pi_t^2 + \alpha x_t^2) \\ & \text{s.t. (1), (2), (22), (23)} \\ & \quad a_0, b_0 \text{ given} \end{aligned} \tag{24}$$

The optimization can be solved in a way analogous to the constant gain case; hence, the first order conditions imply the following optimality condition:

$$2\pi_t + 2\frac{\alpha}{\kappa}x_t + \frac{1}{t+1}\lambda_{3,t} = 0, \tag{25}$$

Similarly to Section 2 our result is that learning introduces an intratemporal trade-off between inflation and output that is not present under RE in an economy without a cost push shock and an additional intertemporal trade-off that is not present in general under rational expectations (Result 1). From the latter it follows that, during the transition, when inflation expectations are positive and inflation is positive, the optimal contraction of x_t is harsher than under discretionary policy. 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, note that we can use first order conditions to rewrite (25) in the following way:

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

Using (2) to substitute out x_t in the above equation, and then using the evolution of inflation expectations (22) we get:

$$E_t [\pi_{t+1}] = A_{11,t}\pi_t + A_{12,t}a_t + P_{1,t}u_t \tag{26}$$

where:

³⁵For theoretical papers on predictor switching see for example Adam (2005), Branch and Evans (2006), Brock and Hommes (1997), Sethi and Franke (1995). For an empirical analysis see Branch (2004).

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

Hence, at an optimum, the dynamics of the economy can be summarized by stacking equations (22), (23) and (26), and obtaining the trivariate system:

$$E_t y_{t+1} = A_t y_t + P_t u_t \quad (27)$$

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

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

We can find the solution with the method of undetermined coefficients, with the guess³⁶:

$$\pi_t = c_{\pi,t}^{dg} a_t + d_{\pi,t}^{dg} u_t \quad (28)$$

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

$$c_{\pi,t}^{dg} = \frac{c_{\pi,t+1}^{dg} \left(1 - \frac{1}{t+1}\right) - A_{12,t}}{A_{11,t} - c_{\pi,t+1}^{dg} \frac{1}{t+1}} \quad (29)$$

and the sequence $\{d_{\pi,t}^{dg}\}$ is defined as:

$$d_{\pi,t}^{dg} = \frac{P_{1,t}}{c_{\pi,t+1}^{dg} \frac{1}{t+1} - A_{11,t}}.$$

Of course, there exist infinite sequences that satisfy equation (29), one for each initial value $c_{\pi,0}^{dg}$. However, since the boundary conditions require π_t to stay bounded, we will concentrate on the solutions that do not explode. The properties of the coefficients in (28) are characterized in the following Proposition.

Proposition 2. *Let $\{c_{\pi,t}^{dg}\}$ be defined by (29), and assume it is bounded; then, $\lim_{t \rightarrow \infty} c_{\pi,t}^{dg}$ exists, and is given by:*

$$\lim_{t \rightarrow \infty} c_{\pi,t}^{dg} = \frac{\alpha\beta}{\alpha + \kappa^2}$$

³⁶This guess corresponds to the unique solution under constant gain learning.

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

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

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 equation (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:

$$x_t = c_{x,t}^{dg} a_t + d_{x,t}^{dg} u_t \quad (30)$$

where:

$$\begin{aligned} c_{x,t}^{dg} &= \frac{c_{\pi,t}^{dg} - \beta}{\kappa} \\ d_{x,t}^{dg} &= \frac{d_{\pi,t}^{dg} - 1}{\kappa} \end{aligned}$$

If the private sector expects inflation to be positive, the optimal CB will contract economic activity more than EH³⁷ (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 inflation expectations.

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).

The nominal interest rate rule is as follows:

$$r_t = \bar{r}r_t + \delta_{\pi,t}^{dg} a_t + \delta_x^{dg} b_t + \delta_{u_t}^{dg} u_t \quad (31)$$

where:

$$\begin{aligned} \delta_{\pi,t}^{dg} &= 1 - \sigma \frac{c_{\pi,t}^{dg} - \beta}{\kappa} \\ \delta_x^{dg} &= \sigma \\ \delta_{u_t}^{dg} &= -\sigma \frac{d_{\pi,t}^{dg} - 1}{\kappa} \end{aligned}$$

Since $c_{\pi,t}^{dg} < \beta$ (see Proposition 2) $\delta_{\pi,t}^{dg}$ 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 EH rule (7):

Proposition 3. *Assume that $t < \infty$; then:*

$$-\delta_{\pi,t}^{dg} > \delta_{\pi}^{EH}, \quad \delta_{u_t}^{dg} > \delta_u^{EH}.$$

Moreover, we have:

$$-\lim_{t \rightarrow \infty} \delta_{\pi,t}^{dg} = \delta_{\pi}^{EH}, \quad \lim_{t \rightarrow \infty} \delta_{u_t}^{dg} = \delta_u^{EH}.$$

Proof. See the Appendix. □

³⁷From $c_{\pi,t}^{dg} < \frac{\alpha\beta}{\alpha + \kappa^2}$ it follows that $c_{x,t}^{dg} < -\frac{\kappa\beta}{\alpha + \kappa^2}$. Compare with the ALM under EH (6).

Result 2 under CG is paralleled 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 inflation expectations more than a CB that takes expectations as given.

An interesting result is that the coefficient on inflation expectations in the interest rate rule (31) 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 5, we show how this coefficient depends on time when the parameters are calibrated according to Woodford (1999). $\delta_{\pi,t}^{dg}$ is always above its limiting level (see analytical proof in Proposition 3), 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 $\delta_{\pi,t}^{dg}$ is a robust feature of the model³⁸. We find that after the 4th period (from the 4th to the 5th period and so on) $\delta_{\pi,t}^{dg}$ is always decreasing, while in the first 4 periods $\delta_{\pi,t}^{dg}$ might be increasing (hump-shaped) for a combination of low values of α and high values of κ (see Figure 7)³⁹. We summarize our new results as:

Result 3. *Optimal policy is time varying even in a stationary environment. Initially it reacts more aggressively to both out of equilibrium expectations and cost push shocks and dampens its aggressiveness later.*

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 inequality $\delta_{ut}^{dg} > \delta_u^{EH}$ is parallel to the second part of Result 2: 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_{ut}^{dg} is positive and decreasing over time (see Figure 6)⁴⁰. In response to a positive cost push shock, the Central Bank raises interest rate to contract the output and thus reduce inflation, and future inflation expectations.

The asymptotic properties of the ALM (28),(30) depend on the limiting behavior of a_t , which is given by the stochastic recursive algorithm:

$$a_{t+1} = a_t + (t+1)^{-1} \left((c_{\pi t}^{dg} - 1)a_t + d_{\pi,t}^{dg} u_t \right) \quad (32)$$

We study its properties in the Appendix, where we use the stochastic approximation techniques⁴¹ to prove the following Proposition:

³⁸We have chosen the grid to include typical calibrated values for the US and the EURO area.

³⁹In fact, $\delta_{\pi,t}^{dg}$ 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).

⁴⁰Since $\delta_{u,t}^{dg} < 1$ from (31) it follows that the change of $\delta_{u,t}^{dg}$ through time is identical to that of $\delta_{\pi,t}^{dg}$.

⁴¹For an extensive monograph on stochastic approximation, see Benveniste, Métivier, and Priouret (1990); the first paper to apply these techniques to learning models is Marcet and Sargent (1989).

Proposition 4. *Let a_t evolve according to (32); then, $a_t \rightarrow 0$ a.s.*

This result, together with the boundedness of $c_{\pi,t}^{dg}$, implies that $c_{\pi,t}^{dg}a_t$ goes to zero almost surely; moreover, it is easy to see that $d_{\pi,t}^{dg} \rightarrow \frac{\alpha}{\kappa^2 + \alpha}$, so that we can conclude that $\pi_t \rightarrow \frac{\alpha}{\kappa^2 + \alpha}v$ 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 REE⁴².

From Proposition 3 it follows that the optimal policy converges to the EH policy; since expectations converge to a constant it is intuitive that in the limit OP behaves as if expectations were fixed. Below we provide a numerical analysis on how the difference during the transition translates into welfare losses. Similarly to Section 2 we report consumption equivalents⁴³.

Tables 4 and 5 show that similarly to the constant gain case in the long run OP engineers a lower consumption equivalent than the EH policy, and OP engineers lower variation of inflation at the cost of allowing higher variation in output. The last row of the first column in Table 4 shows that if we start the economy from the RE equilibrium, $a_0 = 0$, in the long run the consumption equivalent of OP is about 10% lower than that of EH. Table 5 reports the composition of these losses: the optimal policy engineers an inflation variation 20 percent lower than EH and allows a 3-9 times higher welfare loss due to output gap variations.

These long run gains of OP result from the different transition path towards the RE equilibrium this policy engineers compared to EH.

Let us first examine the path of expectations. Both OP and EH are E-stable under learning, so guarantee that expectations converge to the discretionary REE, the difference is the speed of convergence. Figure 8 shows a typical realization of the evolution of expectations under OP and EH. We can observe that inflation expectations converge faster and output gap expectations converge more slowly with our rule than with the EH one. This is a consequence of the intertemporal trade-off (Result 1): when the CB does take into account its influence on the learning algorithm, it has an incentive to undercut future inflation beliefs. The way the central bank can achieve this, is to keep inflation close to its RE value; since inflation expectations are formed as averages of past inflation data, this policy undercuts future inflation expectations. Because of the intratemporal trade-off between inflation and output, the cost of keeping inflation closer to its RE value is a wider output gap and consequently a slower convergence of b to its RE value.

We report how the ratios of OP and EH consumption equivalents evolve during the transition in Table 4. In the first periods the optimal interest rate rule (31) yields *ex-post* higher cumulative welfare losses expressed in consumption terms than the EH rule; later, however, our rule starts generating smaller welfare losses. These findings are consistent with our finding that a CB that follows the optimal rule (31) 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 generated yet 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 $\delta_{\pi,t}^{dg}$ that we characterized above: the coefficient on inflation expectations

⁴²Note 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.

⁴³We report the permanent percentage decrease in the steady state consumption that is equivalent to the cumulative welfare losses up to time T under OP as a ratio of the same measure under EH (See Appendix.). Results are obtained by Monte Carlo simulations.

Table 4: Path of cumulative consumption equivalent ratios under decreasing gain, using OP and EH

T	p^{OP}/p^{EH}			
	$a_0 = 0$	$a_0 = 1$	$a_0 = 2$	$a_0 = 3$
1	2.086	4.145	4.327	4.362
5	1.511	2.241	2.325	2.344
10	1.279	1.574	1.609	1.617
20	1.104	1.116	1.117	1.118
26	1.057	0.993	0.986	0.984
27	1.050	0.978	0.969	0.967
40	0.997	0.841	0.821	0.817
43	0.989	0.820	0.799	0.795
49	0.975	0.786	0.763	0.758
10,000	0.899	0.583	0.542	0.533

Woodford (1999) calibration

in (31) 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.

Since the main advantage of OP is that it helps private agents' inflation expectations to converge faster, the advantage of OP over EH increases the further away initial expectations are from the RE values. The different columns of Table 4 report ratios of consumption equivalents for different initial inflation expectations. The higher is a_0 the bigger is the consumption cost of OP compared to EH in the first periods: OP allows for higher welfare losses in order to keep inflation closer to the SS in order to help inflation expectations converge faster. As time goes on, inflation expectations converge closer to 0 under OP than under EH; the further away a_0 is from the equilibrium, also the further away future inflation expectations remain from the equilibrium under EH. Consequently the inflation output gap trade-off remains worse under EH and consumption equivalents remain also higher than under OP. The bigger is a_0 the bigger is the gain in decreasing inflation variation of OP over EH, and the higher is the output gap variation OP allows compared to EH (See Table 5)⁴⁴.

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 RE equilibrium, and optimal policy is equivalent to the one derived under the assumption of constant expectations. Numerical simulations confirmed the relevance of the welfare gains induced by the implementation of the optimal policy.

4 Robustness Analysis

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

⁴⁴Similarly to Section 2 ratios of consumption equivalents do not depend on the choice of σ_u^2 .

Table 5: Ratio of consumption equivalents of losses due to inflation and output gap variations using OP and EH under decreasing gain learning

	$a_0 = 0$	$a_0 = 1$	$a_0 = 2$	$a_0 = 3$
	Inflation			
p^{OP}	0.016	0.054	0.166	0.353
p^{EH}	0.019	0.125	0.439	0.961
p^{OP}/p^{EH}	0.838	0.432	0.379	0.368
	Output gap			
p^{OP}	0.029	0.411	1.547	3.434
p^{EH}	0.005	0.031	0.110	0.241
p^{OP}/p^{EH}	6.044	13.188	14.100	14.279

Woodford (1999) calibration

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 (2) with the introduction of unobservable shocks, so that the model is now given by:

$$x_t = E_t^* x_{t+1} - \sigma^{-1}(r_t - E_t^* \pi_{t+1} - \bar{r}r_t) + e_{x,t} \quad (33)$$

and:

$$\pi_t = \beta E_t^* \pi_{t+1} + \kappa x_t + u_t + e_{\pi,t} \quad (34)$$

where we assume 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 Section 3. Hence, the control problem of the CB is given by 24 where the first two constraints are now given by 33-34, instead of 1-2; it is easy to show that the interest rate rule that implements the optimal allocations should now include also terms in $E_t^{CB} e_{\pi,t}$ and $E_t^{CB} e_{x,t}$, where E_t^{CB} denotes time- t expectations conditional on the CB information set. To simplify the problem, and since these additional terms do not affect our main results, we assume that $E_t^{CB} e_{\pi,t} = E_t^{CB} e_{x,t} = 0$, even if it does not correspond to the optimal filtering.

4.1 Measurement Error in the Shocks

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

$$u_t^* = u_t + \eta_t, \quad \eta_t \sim N(0, \sigma_\eta^2)$$

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 cost push shock given the signal is⁴⁵:

$$E[u_t | u_t^*] = \frac{\sigma_u^2}{\sigma_\eta^2 + \sigma_u^2} u_t^* \equiv \zeta_u u_t^*$$

⁴⁵See Hamilton (1994).

Note that in the above equation we are not using all the available information; to keep the problem simpler, we assume that the posterior beliefs of the CB, $E_t^{CB}u_t$, are not fully rational, but are instead conditional only on the signals. In other words:

$$E_t^{CB}u_t \equiv E[u_t | u_t^*]$$

Moreover, the separation principle implies that certainty equivalence holds in designing the optimal interest rate rule, which turns out to be identical to (31), with u_t replaced by $E_t^{CB}u_t$:

$$\begin{aligned} r_t &= \bar{r}_t + \delta_{\pi,t}^{dg} a_t + \delta_x^{dg} b_t + \delta_{ut}^{dg} \zeta_u u_t^* \\ &= \bar{r}_t + \delta_{\pi,t}^{dg} a_t + \delta_x^{dg} b_t + \delta_{ut}^{dg} \zeta_u u_t + \delta_{ut}^{dg} \zeta_u \eta_t \end{aligned}$$

We can combine the above equation with (33) and (34) to obtain the ALM for inflation and output gap:

$$\begin{aligned} \pi_t &= \mu_{at}^1 a_t + \mu_{ut}^1 u_t + \mu_{\eta t}^1 \eta_t + \kappa e_{x,t} + e_{\pi,t} \\ x_t &= \mu_{at}^2 a_t + \mu_{ut}^2 u_t + \mu_{\eta t}^2 \eta_t + e_{x,t} \end{aligned}$$

where:

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

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 given⁴⁶; moreover, the analysis of convergence of learning algorithms to the optimal discretionary RE equilibrium⁴⁷ does not change in this modified environment.

To make the exposition clearer, in this subsection we assumed that the CB has partial information on the cost push shock, but can perfectly observe the natural real rate of interest; disposing of this assumption would not alter our results.

4.2 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 criticisms. For example, private sector expectations and their forecasts produced by different institutions do not necessarily coincide. In this case, the CB could use a proxy for the agents' beliefs. In what follows, we assume that the optimal interest rate rule takes the same form as (31), but the agents' forecasts for inflation and

⁴⁶For 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.

⁴⁷Note that the optimal RE equilibrium is now different from the baseline case, since inflation and output gap depend also on η_t , and the unobservable shocks $e_{x,t}$ and $e_{\pi,t}$.

output gap, a_t and b_t , are replaced by the CB internal forecasts, a_t^{CB} and b_t^{CB} ⁴⁸; 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 (22)-(23). 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 ⁴⁹, and forms its internal forecasts accordingly. Plugging the interest rate rule into the structural equations (33) and (34), we get the ALM:

$$\begin{aligned}\pi_t &= \nu_a^1 a_t + \nu_{a^{CB}t}^1 a_t^{CB} + \nu_b^1 b_t + \nu_{b^{CB}t}^1 b_t^{CB} + \nu_{ut}^1 u_t + \kappa e_{x,t} + e_{\pi,t} \\ x_t &= \nu_a^2 a_t + \nu_{a^{CB}t}^2 a_t^{CB} + \nu_b^2 b_t + \nu_{b^{CB}t}^2 b_t^{CB} + \nu_{ut}^2 u_t + e_{x,t}\end{aligned}\quad (35)$$

where:

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

Again, our main results are unaffected by this change in the CB information set, both for $t < \infty$ and for $t \rightarrow \infty$. In fact, since the parameters in the optimal rule are the same as in rule (31), the results summarized in Proposition 3 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 show⁵⁰:

Corollary 1. *Consider the model (35); 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 4, we conclude that also system (35) is E-stable, and it converges to the optimal discretionary RE equilibrium⁵¹.

4.3 Policy Advice

We have seen that assuming adaptive learning instead of RE fundamentally changes the Ramsey solution for optimal policy. However assuming that the CB knows the exact learning algorithm followed by private agents, is clearly a strong assumption to make. Therefore in this section we examine what is the policy recommendation when the monetary authority is uncertain about the nature of private sector expectations.

In particular we aim to define consistent policy advice on a relevant set of private agents' expectation formation. First we investigate the issue of how monetary policy in the US should behave when the FED is uncertain about expectation formation, but its uncertainty is restricted to a set

⁴⁸This 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 (2006).

⁴⁹In 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.

⁵⁰The proof is available from the authors upon request.

⁵¹In 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).

which would be reasonable given the empirical evidence. Then we examine the policy recommendation on a set of more volatile learning algorithms.

Empirical evidence on learning is relatively scarce, and mainly focuses on the US. For the US constant gain algorithms with a small tracking parameter is a good approximation for the data. With Bayesian estimation of the New Keynesian model with adaptive learning Milani (2005) finds γ to be 0.0187, he also finds γ to be stable through time. Orphanides and Williams (2004a) take a different approach, they calibrate γ on the Survey of Professional Forecasters and find that tracking parameters between 0.01 and 0.04 fit well survey expectations.

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 set of expectations to be: constant gain with a small gain, and RE⁵². Moreover, it has no probability distribution over these two possible realizations of the agents' expectations formation mechanism; instead, we use robust control theory and look for the policy that minimizes the maximum loss⁵³. We perform numerical Monte Carlo analysis to examine welfare losses when private expectations are taken from this set and the CB interest rate rule is either an optimal rule for a given small gain parameter or the EH rule (7), derived in Evans and Honkapohja (2003b). In order to present the results in a compact way, for a given private expectation we compare consumption equivalents of using the optimal rule to the consumption equivalents of using a wrong rule and report increases in consumption equivalents. Note that when the structural equations are given by (33) and (34), the ALM for inflation and output gap depend also on the non observable shocks $e_{\pi,t}$ and $e_{x,t}$; hence, the values of the welfare losses are functions of the variances of these shocks too. However, it is easy to see that the ranking between different rules is independent of these variances⁵⁴. Thus, since any parametrization of the non observable terms would be equally arbitrary and would deliver the same results in terms of relative performance of alternative rules, in tables 6-7 we consider the extreme case with the variances of $e_{\pi,t}$ and $e_{x,t}$ that go to zero.

Table 6 reports results when initial expectations coincides with RE: $a_0 = 0$. In this case, constant gain expectations with a small gain will be really close to the rational forecasts. We can think of this economy as populated by agents who are making only very small mistakes compared to the rational forecasts.

The main result is that the worst case scenario is using the EH rule when private agents are learning. A min-max rule (following Hansen and Sargent (2006)), which minimizes the maximum loss is always a learning rule.

Under RE all of these rules lead to a determinate equilibrium. The EH rule provides smaller losses than optimal learning rules (see last line of table 6), and the reason for this is that learning rules allow for too high volatility in the output gap⁵⁵.

However, losses under RE caused by mistakenly using an optimal learning rule are smaller than losses due to using the EH rule when agents are learning.

⁵²Rational expectations means, substituting the interest rate rule in the IS curve (33), and then using it and the Phillips Curve (34) to solve for the fixed point in expectations. Under all interest rate rules listed above, this results $E_t \pi_{t+1} = E_t x_{t+1} = 0$.

⁵³For an extensive treatise on the use of robust control techniques in economics, see Hansen and Sargent (2006).

⁵⁴In other words, the variances of $e_{\pi,t}$ and $e_{x,t}$ do not alter the ordinal properties of the loss function; to see this, an inspection of equation (44) in the Appendix suffices.

⁵⁵We 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).

Table 6: Consumption equivalents under the optimal or a wrong rule, initial inflation expectations at RE

Expectations	0.0187	0.02	0.03	0.04	RE
Interest rate rule					
$\gamma = 0.0187$	0.01302	0.01307	0.01353	0.01412	0.01265
$\gamma = 0.02$	0.01302	0.01307	0.01353	0.01412	0.01265
$\gamma = 0.03$	0.01302	0.01307	0.01352	0.01410	0.01265
$\gamma = 0.04$	0.01303	0.01308	0.01353	0.01409	0.01265
EH	0.01303	0.01308	0.01359	0.01426	0.01265
Maximum percentage increase compared to optimal rule	EH	EH	EH	EH	$\gamma = 0.04$
	0.09	0.12	0.47	1.18	0.02

Woodford (1999) calibration. Starting from RE: $a_0 = 0$.

Consumption equivalents for a given underlying private sector expectation formation and a given interest rate rule.

When private agents are learning and the FED uses a bad learning rule, consumption equivalents increase but the loss is always smaller than losses of using the EH rule. The bigger is the misperception of the monetary policy about γ the bigger is the increase in consumption equivalents. When for example agents follow constant gain with $\gamma = 0.04$ and the central bank uses an optimal rule with $\gamma = 0.03$ consumption equivalent is 0.03% higher than it is when the optimal interest rate rule is used. While if the FED uses $\gamma = 0.02$ which is further from the true tracking parameter, the loss increases to 0.17%. The percentage increase in loss achieved using the EH rule is 1.18%, which is bigger than with any of the learning rules.

When we initialize the economy at the RE equilibrium, beliefs stay close to the RE. This way our analysis does not take into account an advantage of the optimal learning rules, which is that it helps private agents to learn faster the rational expectations forecasts. Therefore in table 7 we report numerical results for $a_0 = 1$. Our results show that the gain of using a learning rule over the EH rule is much bigger in this case, since the EH rule increases consumption equivalents compared to the optimal policy by 3 – 11%. Learning rules on the other hand result in smaller losses under learning, even if they are misspecified.

We now assume that the monetary authority is able to formulate a probability distribution over the mechanism used by the private sector to form its forecasts. In particular, let's assume that the prior of the FED is that with probability p private agents follow constant gain learning with a given tracking parameter, and with probability $1 - p$ agents have RE. Then we can calculate the expected welfare loss of using EH p times the consumption equivalent under constant gain learning with EH rule, and $1 - p$ times consumption equivalent of using EH under RE. Then we can find a cut-off value of p for which the expected loss in consumption terms of using the OP rule is less than the welfare loss of the EH rule.

A surprising result is that the cutoff value of p is between 1 – 1.5%⁵⁶. This means that it is optimal to use the learning rule even if the CB attributes only a very small 1 – 2% probability (or higher) to agents following learning and a very high probability to RE.

In sum, our “policy advice” for the FED is to choose an optimal learning rule even if it attributes only very small probability to learning.

⁵⁶Cutoff values for p are slightly lower when we initialize the economy out of RE.

Table 7: Consumption equivalents under the optimal or a wrong rule, initial inflation expectations out of RE

Expectations	0.0187	0.02	0.03	0.04	RE
Interest rate rule					
$\gamma = 0.0187$	0.1249	0.1240	0.1175	0.1120	0.0127
$\gamma = 0.02$	0.1250	0.1240	0.1173	0.1116	0.0127
$\gamma = 0.03$	0.1257	0.1245	0.1167	0.1102	0.0127
$\gamma = 0.04$	0.1271	0.1258	0.1170	0.1098	0.0127
EH	0.1287	0.1282	0.1246	0.1215	0.0127
Maximum percentage increase compared to optimal rule	EH	EH	EH	EH	$\gamma = 0.04$
	2.99	3.37	6.75	10.64	0.02

Woodford (1999) calibration. Starting from RE: $a_0 = 1$.

Consumption equivalents for a given underlying private sector expectation formation and a given interest rate rule.

Evans and Ramey (2006) shows that in an economy with a high probability of structural changes it is optimal to use constant gain algorithms with high tracking parameters. Therefore we could perform the same analysis on a set of constant gain algorithms with much higher tracking parameters. In this case the results⁵⁷ are even more against the use of the EH rule. For example, when agents follow CG with $\gamma = 0.2$ and $a_0 = 0$ using the EH rule results in 40% higher losses compared to the optimal learning rule. This is intuitive, because the higher is the gain parameter, the bigger is the influence that monetary policy can have on expectations if the CB makes active use of the learning algorithm; thus, the EH rule makes a bigger mistake when it does not take this into account⁵⁸.

Interestingly, different assumptions on the initial beliefs yields different solutions to the min-max problem: the robust rule is always a learning rule, as we stressed above, but when $a_0 = 0$ the min-max corresponds to the optimal rule when γ is the largest in the set, and when $a_0 = 1$ to the optimal rule when γ is the smallest. This reflects the tension between *tracking* and *accuracy* present in any adaptive algorithm, and analyzed in Benveniste, Métévier, and Priouret (1990): a higher γ implies a better tracking of the mean dynamics of the underlying process, and a larger variance around this mean dynamics. If the process starts close to the mean dynamics (in our model, the RE equilibrium), the second effect would prevail over the first one, inducing a positive relation between γ and the welfare losses; consequently, for any rule used by the CB, the maximum losses are attained when the gain parameter is the largest possible (0.04), and the robust rule coincides with the optimal one for that value of γ . If the process starts far from the mean dynamics, the opposite line of reasoning applies.

In this subsection we presented numerical evidence suggesting that, when the CB is insecure as to whether agents have RE or are learning, it should use an optimal learning rule, unless it attaches a tiny probability to agents following learning; we are aware that a full-fledged robustness analysis would require a larger set of possible expectations' formation mechanisms, but this is beyond the scope of this paper.

⁵⁷ Available from the authors upon request.

⁵⁸ The cutoff value of p is between 1-2 percent. Thus the CB should use a learning rule even if it attributes only a very small probability to learning.

5 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 in this framework the implications for policymaking go beyond the asymptotic learnability criterion. In the short run the optimal policy under learning resembles more the commitment solution under rational expectations than the discretionary solution under rational expectations. Both the commitment solution under rational expectations and the Ramsey solution under learning aims to anchor inflation expectations, thus it accommodates less the effect of noisy supply shocks on inflation. The intuition behind is simple and stems from the presence of a new intertemporal inflation output trade-off, that is not present under rational expectations. Under learning the central bank has to take into account how its policy affects future inflation expectations, since out of equilibrium expectations worsen the future inflation output trade-offs. As a result optimal policy is aggressive towards inflation, in order to induce private agents to learn faster the equilibrium expected value of inflation.

In the long run the equilibrium depends on how private agents learn. 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)).

For future research it would be very interesting to explore the possibility whether optimal policy under adaptive learning can drive the economy to the commitment solution under rational expectations. This question is particularly interesting as from the backward looking nature of these learning algorithms it follows that such policies are time consistent, so the commitment solution could be reached in a time consistent fashion.

A large body of research in learning focused on how to design rules that are stable under learning; a typical result is that a strong reaction to out of equilibrium inflation expectations is necessary. We would like to note that under optimal policy examining E-stability is not necessary, since optimal policy naturally chooses an E-stable solution; moreover optimal policy is similar to the consensus reached in earlier papers on the desirability of the monetary policy being aggressive towards inflation, in order to anchor private agents' inflation expectations.

An additional message of our paper for policy conduct is that optimal policy should closely monitor private sectors expectations. Actually this is what is happening in real life: central banks do pay close attention to private expectations. Under rational expectations this is not justified, since expectations are pinned down by the model and the monetary policy rule, however once we depart from rationality expectations become a natural state variable.

Since optimal policy depends on the way private agents learn, we think a particularly important area of future research would be to estimate how agents learn. Even if expectations are not rational, expectations should be endogenous, one should allow agents to abandon their ad hoc learning rule if they can do better. Empirical research on examining differences in learning behavior in different environments is still missing. It would be also interesting to examine how monetary policy 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 Marcat and

Nicolini (2003) where agents dynamically switch between different predictors depending on the last forecast error. An alternative way would be to model expectation formation as in Molnar (2005) where agents do not change the predictor used, but always use a weighted average of the forecasts generated by different predictors, and adjust the weight on predictors dynamically depending on the relative forecasting performance.

A Constant Gain Learning

In this section we give the outline of the derivation of the inflation law of motion (19), and prove Proposition 1.

We start from the optimality condition (17), that we recall below:

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

Using the Phillips curve (2) and the evolution of inflation expectations (8), we get:

$$E_t [\pi_{t+1}] = A_{11}\pi_t + A_{12}a_t + P_1u_t \quad (36)$$

where:

$$\begin{aligned} A_{11} &\equiv \frac{\kappa^2 + \alpha + \alpha\beta^2\gamma(1-\gamma(1-\beta))}{\alpha\beta(1-\gamma(1-\beta)) + \kappa^2\beta(1-\gamma)} \\ A_{12} &\equiv -\frac{\alpha\beta(1-\beta(1-\gamma)(1-\gamma(1-\beta)))}{\alpha\beta(1-\gamma(1-\beta)) + \kappa^2\beta(1-\gamma)} \\ P_1 &\equiv -\frac{\alpha}{\alpha\beta(1-\gamma(1-\beta)) + \kappa^2\beta(1-\gamma)} \end{aligned}$$

Hence, at an optimum, the dynamics of the economy can be summarized by stacking equations (8), (9) and (36), and obtaining the trivariate system:

$$E_t y_{t+1} = A y_t + P u_t \quad (37)$$

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

$$A \equiv \begin{pmatrix} A_{11} & A_{12} & 0 \\ \gamma & 1-\gamma & 0 \\ \frac{\gamma}{\kappa} & -\frac{\beta\gamma}{\kappa} & 1-\gamma \end{pmatrix}, \quad P = \begin{pmatrix} P_1 \\ 0 \\ -\frac{\gamma}{\kappa} \end{pmatrix}.$$

The three boundary conditions of the above system are:

$$\begin{aligned} &a_0, b_0 \text{ given} \\ &\lim_{s \rightarrow \infty} |E_t \pi_{t+s}| < \infty \end{aligned} \quad (38)$$

The last one is due to the fact that, if there exists a solution to the problem (10) when the possible stochastic processes $\{\pi_t, x_t, r_t, a_{t+1}, b_{t+1}\}$ are restricted to be bounded, then this would also be the minimizer in the unrestricted case⁵⁹.

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

$$A_1 \equiv \begin{pmatrix} A_{11} & A_{12} \\ \gamma & 1-\gamma \end{pmatrix} \quad (39)$$

In the following Lemma 1 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 (37)-(38) has one and only one solution. In other words,

⁵⁹The proof is available from the author upon request.

there exists one and only one stochastic process for each of the three variables of y such that (38) 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 (37)-(38) are also a solution of the subsystem:

$$E_t y_{1t+1} = A_1 y_{1t} + (P_1, 0)' u_t$$

together with the boundary conditions:

$$a_0 \text{ given, } \lim_{s \rightarrow \infty} |E_t \pi_{t+s}| < \infty$$

By Lemma 1, we can invoke Proposition 1 of Blanchard and Kahn (1980) to conclude that the law of motion for inflation can be written in the form:

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

as stated in Proposition 1.

Lemma 1. *Let A_1 be given by equation (39) in the text; then it has an eigenvalue inside and one outside the unit circle.*

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 that⁶⁰:

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

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

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

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. □

⁶⁰LaSalle (1986).

We now prove the rest of Proposition 1. First of all, we can guess that inflation follows the ALM (19) and use the optimality condition (36) and the method of undetermined coefficients to verify that c_π^{cg} must satisfy the following quadratic expression:

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

where:

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

and that:

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

The polynomial in c_π^{cg} can be equivalently rewritten as follows:

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

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 conditions⁶¹.

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

Lemma 2. $f(c_\pi^{cg})$ 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:

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

After rearranging:

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

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

Lemma 3. $f(c_\pi^{cg}) : [0, 1] \rightarrow [0, 1]$

⁶¹Since it would imply an exploding inflation.

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 $f(0) > 0$.

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

After rearranging, we get:

$$f(1) \leq 1 \iff 0 \leq \kappa^2(1-\beta(1-\gamma)) + \alpha(1-\beta)(1-\beta(1-\gamma(1-\beta)))$$

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

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 4. 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 $\max_{x \in [0, 1]} f'(x) = f'(1)$; after some algebraic manipulation, we get:

$$f'(1) \leq 1 \iff (1-\beta\gamma)\beta(1-\gamma(1-\beta)) + \beta\gamma(1-\gamma(1-\beta)) - 1 \leq \frac{\kappa^2}{\alpha}(1-\beta(1-\gamma^2))$$

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 ($\frac{\kappa^2}{\alpha}(1-\beta(1-\gamma^2))$ is always positive). \square

Moreover, we prove the following result.

Lemma 5. 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:

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

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

$$\stackrel{\geq}{\leq} \frac{\alpha\beta}{\kappa^2+\alpha}$$

if and only if:

$$\frac{(\kappa^2+\alpha)\alpha\beta(1-\beta(1-\gamma)(1-\gamma(1-\beta))) + \gamma[\kappa^2\beta(1-\gamma)+\alpha\beta(1-\gamma(1-\beta))]}{\kappa^2+\alpha+\alpha\beta^2\gamma(1-\gamma(1-\beta))-(1-\gamma)[\kappa^2\beta(1-\gamma)+\alpha\beta(1-\gamma(1-\beta))]} \stackrel{\geq}{\leq} 1$$

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:

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

which is equal to one; in fact:

$$\frac{(\kappa^2+\alpha)(1-\beta(1-\gamma)(1-\gamma(1-\beta))) + \beta\gamma[\kappa^2\beta(1-\gamma)+\alpha\beta(1-\gamma(1-\beta))]}{\kappa^2+\alpha+\alpha\beta^2\gamma(1-\gamma(1-\beta))-(1-\gamma)[\kappa^2\beta(1-\gamma)+\alpha\beta(1-\gamma(1-\beta))]} \stackrel{\geq}{\leq} 1$$

is equivalent to:

$$-(\kappa^2+\alpha)\beta(1-\gamma)(1-\gamma(1-\beta)) + (1-\gamma(1-\beta))[\alpha\beta(1-\gamma(1-\beta)) + \kappa^2\beta(1-\gamma)] \stackrel{\geq}{\leq} \alpha\beta^2\gamma(1-\gamma(1-\beta))$$

But the LHS can be simplified as:

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

which is equal to:

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

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

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

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 3 and 4 we obtain that $f(\cdot)$ is a contraction when defined on the interval $[0, 1]$; moreover, by Lemma 5 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 4 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 completes the proof. \square

B Decreasing Gain Learning

In this section we prove Propositions 2 and 4.

Proof of Proposition 2. To prove the first part of the statement, note that if we solve forward the following difference equation:

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

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:

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

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

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

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:

$$\left(1 - \frac{1}{T+1} \left(1 - \frac{\alpha\beta}{\alpha + \kappa^2}\right)\right) c_{\pi T+1}^{dg} \geq \frac{\alpha\beta}{\alpha + \kappa^2} A_{11,T} + A_{12,T} \quad (40)$$

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 (40), we get:

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

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

Finally, we prove Proposition 4. First of all, we briefly describe some results of stochastic approximation⁶² that we will exploit in the proof.

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

$$\theta_t = \theta_{t-1} + \gamma_t Q(t, \theta_{t-1}, X_t) \quad (41)$$

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 (41) can be analyzed using the associated deterministic ODE:

$$\frac{d\theta}{d\tau} = h(\theta(\tau)) \quad (42)$$

where:

$$h(\theta) \equiv \lim_{t \rightarrow \infty} EQ(t, \theta, X_t)$$

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 (41) is given by the stable resting points of the ODE (42).

⁶²Ljung (1977), Benveniste, Métivier, and Priouret (1990) provide a recent survey.

Proof of Proposition 4. Note that our equation (32) is a special case of (41), where the technical conditions are easily shown to be satisfied; moreover, it is also easy to see that:

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

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. \square

C Comparison with EH Rule

Proof of Proposition 3. First of all, note that:

$$\delta_{\pi,t}^{dg} \geq \delta_{\pi}^{EH} \iff \sigma \frac{\beta - c_{\pi,t}^{dg}}{\kappa} \geq \sigma \frac{\kappa\beta}{\alpha + \kappa^2}$$

where the second inequality can be rewritten as:

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

Rearranging the terms, we get:

$$\delta_{\pi,t}^{dg} \geq \delta_{\pi}^{EH} \iff \frac{\alpha\beta}{\alpha + \kappa^2} \geq c_{\pi,t}^{dg}$$

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

$$\delta_{ut}^{dg} \geq \delta_u^{EH} \iff \frac{\alpha}{\alpha + \kappa^2} \geq d_{\pi,t}^{dg}$$

which implies, since

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

that $\delta_{ut}^{dg} > \delta_u^{EH}$ whenever $t < \infty$. Finally, note that Proposition 2 also showed that $\lim_{t \rightarrow \infty} c_{\pi,t}^{dg} = \frac{\alpha\beta}{\alpha + \kappa^2}$, which trivially yields $\lim_{t \rightarrow \infty} \delta_{\pi,t}^{dg} = \delta_{\pi}^{EH}$ and $\lim_{t \rightarrow \infty} \delta_{ut}^{dg} = \delta_u^{EH}$. \square

D Derivations of Consumption Equivalents

In this section we follow derivations of Adam and Billi (2006).

Woodford (2003) chapter 6 shows that the second order approximation of the representative agents discounted utility flow is given by

$$U = -\bar{Y}U_c L^P, \quad (43)$$

where \bar{Y} denotes the steady state level of output associated with zero inflation in the absence of disturbances, U_c is the marginal utility of consumption at \bar{Y} and

$$L^P = \frac{1}{2} \frac{\sigma + \omega}{\alpha} \sum_{i=1}^{\infty} \beta^i (\pi_{t+i}^2 + x_{t+i}^2),$$

where (σ) is the households Arrow-Pratt Measure of relative risk aversion and ω is the elasticity of a firm's real marginal cost with respect to its own output, L^P denotes L generated under a policy rule P.

Assuming a permanent reduction in consumption from \bar{Y} by $p \geq 0$ percent, a second order approximation of the utility loss is

$$\begin{aligned} \frac{1}{1-\beta} \left(-U_c \bar{Y} \frac{p}{100} + \frac{U_{cc}}{2} \left(\bar{Y} \frac{p}{100} \right)^2 \right) &= \frac{-U_c \bar{Y}}{1-\beta} \left(\frac{p}{100} - \frac{U_{cc} \bar{Y}}{2U_c} \left(\bar{Y} \frac{p}{100} \right)^2 \right) = \\ \frac{-U_c \bar{Y}}{1-\beta} \left(\frac{p}{100} + \frac{\sigma}{2} \left(\bar{Y} \frac{p}{100} \right)^2 \right), \end{aligned}$$

where U_{cc} is the second derivative of utility of utility with respect to consumption evaluated at \bar{Y} . Equating this utility loss to (43), the welfare loss generated under policy rule P gives

$$\frac{p}{100} + \frac{\sigma}{2} \left(\bar{Y} \frac{p}{100} \right)^2 - (1-\beta)(L^P) = 0.$$

The percentage loss in steady state consumption equivalent to the decrease in utility generated by following rule P is

$$p = 100\sigma \left(-1 + \sqrt{1 + \frac{2(1-\beta)(L^P)}{\sigma}} \right).$$

Since x and π are expressed in percentage points we have to rescale the losses and use

$$p = 100\sigma \left(-1 + \sqrt{1 + \frac{2(1-\beta)(L^P)\sigma}{100^2}} \right). \quad (44)$$

References

- ADAM, K. (2005): “Learning to Forecast and Cyclical Behavior of Output and Inflation,” *Macroeconomic Dynamics*, 9(1), 1–27.
- ADAM, K., AND R. BILLI (2006): “Discretionary Monetary Policy and the Zero Lower Bound on Nominal Interest Rates,” *Journal of Monetary Economics*, forthcoming.
- ADAM, K., AND M. PADULA (2003): “Inflation dynamics and subjective expectations in the United States,” Working paper series 222, European Central Bank, available at <http://ideas.repec.org/p/att/eurcbw/2003222.html>.
- ALVAREZ, F., P. J. KEHOE, AND P. A. NEUMEYER (2004): “The Time Consistency of Optimal Monetary and Fiscal Policies,” *Econometrica*, 72(2), 541–567.
- BALL, L. (1994): “What Determines the Sacrifice Ratio?,” in *Monetary Policy*, ed. by N. G. Mankiw. Chicago: University Chicago Press.
- 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.
- BENVENISTE, A., M. MÉTIVIER, AND P. PRIOURET (1990): *Adaptive Algorithms and Stochastic Approximations*. Berlin: Springer-Verlag.
- 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.
- BRANCH, W. A. (2004): “The Theory of Rationally Heterogeneous Expectations: Evidence from Survey Data on Inflation Expectations,” *The Economic Journal*, 114(497), 592–621.
- BRANCH, W. A., AND G. W. EVANS (2006): “Intrinsic heterogeneity in expectation formation,” *Journal of Economic Theory*, 127(1), 264–295.
- BRAY, M. 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.
- BRAYTON, F., AND P. TINSLEY (1996): “A guide to FRB/US: a macroeconomic model of the United States,” Finance and Economics Discussion Series 96-42, Board of Governors of the Federal Reserve System (U.S.), available at <http://ideas.repec.org/p/fip/fedgfe/96-42.html>.
- BROCK, W. A., AND C. H. HOMMES (1997): “A Rational Route to Randomness,” *Econometrica*, 65(5), 1059–1096.
- CAGAN, P. (1956): “The Monetary Dynamics of Hyperinflation,” in *Studies in the Quantity Theory of Money*, ed. by M. Friedman. Chicago: University of Chicago Press.

- CALVO, G. (1983): “Staggered Prices in a Utility Maximizing Framework,” *Journal of Monetary Economics*, 12(3), 383–398.
- CLARIDA, R., J. GALI, AND M. GERTLER (1999): “The Science of Monetary Policy: A New Keynesian Perspective,” *Journal of Economic Literature*, 37(2), 1661–1707.
- (2000): “Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory,” *The Quarterly Journal of Economics*, 115(1), 147–180.
- ERCEG, C. J., AND A. T. LEVIN (2003): “Imperfect Credibility and Inflation Persistence,” *Journal of Monetary Economics*, 50(4), 915–944.
- EVANS, G. W. (1985): “Expectational Stability and the Multiple Equilibria Problem in Linear Rational Expectations Models,” *The Quarterly Journal of Economics*, 100(4), 1217–33.
- EVANS, G. W., AND S. HONKAPOHJA (2001): *Learning and Expectations in Macroeconomics*. Princeton: Princeton University Press.
- (2003a): “Adaptive Learning and Monetary Policy Design,” *The Journal of Money, Credit and Banking*, 35(6), 1045–1072.
- (2003b): “Expectations and the Stability Problem for Optimal Monetary Policies,” *Review of Economic Studies*, 70(4), 807–824.
- (2006): “Monetary Policy, Expectations and Commitment,” *Scandinavian Journal of Economics*, 108(1), 15–38.
- EVANS, G. W., AND G. RAMEY (2006): “Adaptive expectations, Underparametrization and the Lucas Critique,” *Journal of Monetary Economics*, 53(2), 249–264.
- 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,” Working paper series 163, European Central Bank, available at <http://ideas.repec.org/p/att/eurocbw/2002163.html>.
- FOURGEAUD, C., C. GOURIEROUX, AND J. PRADEL (1986): “Learning Procedures and Convergence to Rationality,” *Econometrica*, 54(4), 845–68.
- GALI, J. (2003): “New Perspectives on Monetary Policy, Inflation, and the Business Cycle,” in *Advances in Economic Theory*, ed. by M. Dewatripont, L. Hansen, and S. Turnovsky. Cambridge: Cambridge University Press.
- GASPAR, V., F. SMETS, AND D. VESTIN (2005): “Optimal Monetary Policy under Adaptive Learning,” .
- (2006): “Monetary Policy over Time,” *Macroeconomic Dynamics*, 10(2), 207–229.
- HAMILTON, J. D. (1994): *Time Series Analysis*. Princeton: Princeton University Press.
- HANSEN, L. P., AND T. J. SARGENT (2006): *Robustness*. Forthcoming Princeton University Press.

- 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.
- KREPS, D. (1998): “Anticipated Utility and Dynamic Choice,” in *Frontiers of Research in Economic Theory*, ed. by D. P. Jacobs, E. Kalai, and M. Kamien. Cambridge: Cambridge University Press.
- LASALLE, J. P. (1986): *The Stability and Control of Discrete Processes*. Berlin: Springer-Verlag.
- LJUNG, L. (1977): “Analysis of recursive stochastic algorithms,” *IEEE Transactions on Automatic Control*, (AC-22), 551–575.
- LUCAS, R. E. (1987): *Models of Business Cycles*. Oxford: Basil Blackwell.
- LUCAS, R. E., AND N. L. STOKEY (1983): “Optimal Fiscal and Monetary Policy in an Economy Without Capital,” *Journal of Monetary Economics*, 12(1), 55–93.
- MARCET, A., AND J. P. NICOLINI (2003): “Recurrent Hyperinflations and Learning,” *American Economic Review*, 93(5), 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(2), 337–368.
- MCCALLUM, B., AND E. NELSON (1999): “Performance of Operational Optimal Policy Rules in an Estimated Semi-Classical Model,” in *Monetary Policy Rules*, ed. by J. Taylor. 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. C. WILLIAMS (2004a): “The Decline of Activist Stabilization Policy: Natural Rate Misperceptions, Learning, and Expectations,” *Board of Governors of the Federal Reserve System (U.S.), International Finance Discussion Papers*, (804).
- (2004b): “Imperfect Knowledge, Inflation Expectations and Monetary Policy,” in *The Inflation-Targeting Debate*, ed. by B. Bernanke, and M. Woodford. Chicago: Chicago University Press.
- (2005): “Inflation scares and forecast-based monetary policy,” *Review of Economic Dynamics*, 8(2), 498–527.
- PHELPS, E. S. (1967): “Phillips Curves, Expectations of Inflation and Optimal Unemployment over Time,” *Economica*, 34(135), 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: Princeton University Press.

- SETHI, R., AND R. FRANKE (1995): “Behavioural Heterogeneity under Evolutionary Pressure: Macroeconomic Implications of Costly Optimisation,” *Economic Journal*, 105(430), 583–600.
- STEINSSON, J. (2003): “Optimal monetary policy in an economy with inflation persistence,” *Journal of Monetary Economics*, 50(7), 1425–1456.
- WOODFORD, M. (1999): “Optimal Monetary Policy Inertia,” *The Manchester School, Supplement*, 67(0), 1–35.
- (2003): *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton: Princeton University Press.
- YUN, T. (1996): “Nominal price rigidity, money supply endogeneity, and business cycles,” *Journal of Monetary Economics*, 37(2-3), 345–370.

E Figures




Figure 1: Feedback parameter in the ALM for inflation as a function of γ .




Figure 2: Impulse response of inflation for an initial cost-push shock $u = 1$. Solid line: optimal policy under learning and private agents following learning with $\gamma = 0.9$. Dashed line: optimal discretionary policy under RE with private agents have rational expectations. Initial conditions: $a_0 = 0$, $\pi_0 = 0$, $x_0 = 0$.




Figure 3: Variance of inflation expectations




Figure 4: Variance of output gap expectations




Figure 5: Interest rate rule coefficient on inflation expectations under decreasing gain learning.




Figure 6: Interest rate rule coefficient on the cost push shock under decreasing gain




Figure 7: 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 8: Evolution inflation and output gap expectations under the optimal (solid line) and the EH rule (dashed line), when private agents follow decreasing gain learning