2010 | 08

Working Paper

Research Department

Optimal monetary policy when agents are learning

By Krisztina Molnár and Sergio Santoro

NB NORGES BANK

Working papers fra Norges Bank, fra 1992/1 til 2009/2 kan bestilles over e-post:
servicesenter@norges-bank.no
eller ved henvendelse til: Norges Bank, Abonnementsservice
Postboks 1179 Sentrum
0107 Oslo
Telefon 22 31 63 83, Telefaks 22 41 31 05

Fra 1999 og senere er publikasjonene tilgjengelige på www.norges-bank.no

Working papers inneholder forskningsarbeider og utredninger som vanligvis ikke har fått sin e4ndelige form.
Hensikten er blant annet at forfatteren kan motta kommentarer fra kolleger og andre interesserte.
Synspunkter og konklusjoner i arbeidene står for forfatternes regning.

Working papers from Norges Bank, from1992/1 to 2009/2 can be ordered by e-mail:
servicesenter@norges-bank.no
or from Norges Bank, Subscription service,
P.O.Box. 1179 Sentrum
N-0107Oslo, Norway.
Tel. +47 22 31 63 83, Fax. +47 22 41 31 05

Working papers from 1999 onwards are available on www.norges-bank.no

Norges Bank's working papers present research projects and reports (not usually in their final form)
and are intended inter alia to enable the author to benefit from the comments of colleagues and other interested
parties. Views and conclusions expressed in working papers are the responsibility of the authors alone.

ISSN 1502-8143 (online)
ISBN 978-82-7553-556-4 (online)

Optimal Monetary Policy When Agents Are Learning*

Krisztina Molnár[†] and Sergio Santoro[‡]

May 27, 2010

Abstract

We derive the optimal monetary policy in a sticky price model when private agents follow adaptive learning. We show that this slight departure from rationality has important implications for policy design. The central bank faces a new intertemporal trade-off, not present under rational expectations: it is optimal to forego stabilizing the economy in the present in order to facilitate private sector learning and thus ease the future intratemporal inflation-output gap trade-offs. The policy recommendation is robust: the welfare loss entailed by the optimal policy under learning if the private sector actually has rational expectations is much smaller than if the central bank mistakenly assumes rational expectations when in fact agents are learning.

JEL classification: C62, D83, D84, E52
Keywords: optimal monetary policy, learning, rational expectations

*We are especially grateful to our supervisors Albert Marcet and Ramon Marimon for very helpful comments and encouragement throughout this project. We are also indebted to Kosuke Aoki, Gernot Doppelhofer, Refet Gurkaynak, Seppo Honkapohja, Thijs van Rens, Alessandro Rebucci, José-Víctor Ríos-Rull, Stephan Sauer, and Ákos Valentinyi for useful conversations and suggestions. Any errors are our own. The views expressed herein are those of the authors, and do not necessarily reflect those of the Bank of Italy or Norges Bank.

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

[‡]Department for International Economic Analysis and Relations, Bank of Italy; Email: sergio.santoro@bancaditalia.it

1 Introduction

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

Optimal monetary policy design is extensively studied under the assumption of rational expectations (RE). Despite the fact that the role of deviations from RE is emphasized in several theoretical and empirical papers[1], the influence of less-than-rational expectations on the optimal policy conduct is not yet well understood. Instead, earlier literature examined the robustness of Taylor rules derived under RE, and have shown that slight deviations from rationality are important for policy design. Taylor rules that are optimal or guarantee determinacy under RE, can lead to instability if private expectations follow adaptive learning (see Bullard and Mitra (2002), Evans and Honkapohja (2003a), Evans and Honkapohja (2003b) and Evans and Honkapohja (2006)).

In this paper, we investigate the interaction between departures from RE and monetary policy from a different angle: instead of examining the asymptotic behavior of Taylor rules, we address the issue of how a rational central bank (CB) should optimally conduct monetary policy if the private sector forms expectations with adaptive learning. We assume the CB is rational within the model, knows how private agents form their expectations, and takes their expectations formation scheme into account when solving its control problem. We conduct our analysis in a standard dynamic stochastic general equilibrium (DSGE) model with nominal rigidities, in order to facilitate comparison with the earlier literature.

The main contribution of this paper is to derive analytically the optimal solution. The advantage of closed-form solutions is to provide a better understanding of policy trade-offs. There is a well known *intratemporal* inflation-output gap trade-off. We show that a slight departure from RE introduces a new *intertemporal* trade-off. In period t the CB renounces to its ability to stabilize the economy in the way that would be optimal under RE and discretion, in order to reduce future inflation expectations, hence ease the future intratemporal inflation-output gap trade-off. Hence a slight departure from rationality is not only relevant for the limiting stability of the equilibrium, but inherently changes policy design. Our quantitative analysis shows that incorporating the intertemporal tradeoff into policymaking increases welfare substantially even if the departure from RE equilibrium is small.

Our policy recommendation is that stabilizing private inflation expectations is more important when these deviate from rationality than under RE. Earlier literature analyzing the welfare effect of different Taylor rules have also shown that

[1] See for example Marcet and Nicolini (2003), Milani (2007), Slobodyan and Wouters (2009).

the CB should act against inflation beliefs more aggressively than what is suggested by an RE model (see for example Ferrero (2007), Orphanides and Williams (2005b) and Orphanides and Williams (2005c)). Our analytical solutions rationalize these earlier numerical results.

Our results also provide a rationale for the general practice by CB to closely monitor private sector expectations. Under RE this is not justified, since expectations are pinned down by the model and the monetary policy rule. Instead, once we depart from rationality, expectations become a state variable, therefore optimal policy should condition on private expectations.

When expectations are rational, a credible CB can manipulate them by committing to a future course of action; instead, under adaptive learning there is no such role for promises, since beliefs are affected only by past occurrences. Nevertheless, Sargent (1999), chapter 5, obtains the remarkable result that the optimal policy in the Phelps problem[2] is such that a CB patient enough can replicate asymptotically the commitment solution under RE. This finding partly generalizes to our setup. Optimal policy does not replicate the commitment solution, but there is a qualitative similarity: the impulse response to a cost-push shock is similar to the commitment case, in the sense that the contemporaneous impact of a cost-push shock on inflation is small (compared to the case of discretionary policy under RE), and inflation reverts to the equilibrium in a sluggish manner. This similarity is stronger when the CB is more patient. Both under RE and learning, this pattern comes from the CB's ability to directly manipulate private expectations, even if the channels used are quite different. Under commitment, the policymaker uses *credible promises about the future*, while under learning, the pattern results from the impact that *past actions* have on beliefs. Thus, the ability to manipulate future private sector expectations through the learning algorithm plays a role similar to a commitment device under RE, hence eases the future short-run trade-off between inflation and the output gap.

Assuming that the CB knows and makes active use of the exact form of private expectations is undoubtedly a very strong hypothesis. In reality, there is still a lively debate about how to model private sector expectations; therefore we also perform two kinds of robustness checks, one under Knightian and the other under probabilistic uncertainty. We compare the optimal learning rule derived in our paper to the time consistent optimal rule derived under RE. When the CB is uncertain about the nature of expectations formation in the sense of Knight (1921),[3] the optimal learning rules derived in our paper are more robust. When, instead, the CB has a probability distribution defined over the set of possible forms

[2]Phelps (1967) formulates a control problem for a natural rate model with a rational CB and private agents endowed with a mechanical forecasting rule, known to the CB.

[3]Knightian uncertainty refers to the impossibility of forming a probability assessment of the possible states of the world.

of private expectations, the expected welfare losses are smaller under the optimal learning rules even if the CB assigns only a very small probability to the possibility that agents use learning instead of RE.

A relevant topic for future research is to examine how robust this policy recommendation is to different deviations from rationality. In a paper closely related to ours, Gaspar, Smets, and Vestin (2006) focus on the case when private agents learn about the persistence of inflation, when firms index to lagged inflation. They show numerically that an optimally behaving CB aims to anchor inflation expectations better. This result is analogous to ours and suggests that if the private sector is not fully rational, an increased concern for stabilizing inflation expectations is an important policy advice independently of the exact rule followed by private agents to form their expectations.

The rest of the paper is organized as follows. In Section 2, after briefly recalling the discretionary optimal policy when expectations are rational, we show the existence of the new intertemporal trade-off under learning. Section 3 characterizes the optimal allocations (and the interest rate rule that supports them) when agents use constant gain learning, underlining how the presence of the intertemporal trade-off increases the CB aggressiveness against inflation beliefs. Section 4 relaxes the assumption that expectations follow constant gain learning, and shows that our main results remain valid under decreasing gain learning. Section 5 argues that the optimal policy rule derived in the previous sections is robust to uncertainty about the agents' expectations formation mechanism, and Section 6 concludes.

2 The model

We consider the baseline version of the New Keynesian model; in this framework, the economy is characterized by two structural equations.[4] The first one is an IS equation:

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

where x_t, r_t and π_t denote the time t output gap (i.e. the difference between actual and natural output), the short-term nominal interest rate and inflation, respectively. σ is a parameter of the household's utility function, representing risk aversion, and $\overline{rr}_t$ is the natural real rate of interest, i.e. the real interest rate that would hold in the absence of any nominal rigidity. We assume that it is distributed as an AR(1) process:

$$\overline{rr}_t = \rho \overline{rr}_{t-1} + \varepsilon_t, \tag{2}$$

[4]For 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).

where $\varepsilon_t \sim N(0, \sigma_\varepsilon^2)$. Note that the operator E_t^* represents the private agents' expectation conditional on the time t information set, which is not necessarily rational. The above equation is derived by loglinearizing the household's Euler equation and imposing the equilibrium condition that consumption equals output.

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

$$\pi_t = \beta E_t^* \pi_{t+1} + \kappa x_t + u_t, \tag{3}$$

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 shock[5]; this relation is obtained from optimal pricing decisions of monopolistically competitive firms whose prices are staggered à la Calvo (1983).[6]

The loss function of the CB is given by:

$$E_0 \sum_{t=0}^{\infty} \beta^t \left(\pi_t^2 + \alpha x_t^2 \right), \tag{4}$$

where α is the relative weight put by the CB on the objective of output gap stabilization.[7]

2.1 Benchmark: discretionary solution under rational expectations and learning

Under adaptive learning credibility of the CB has no role, because promises about the future do not influence expectations. Therefore, our benchmark under RE is discretionary monetary policy, when the CB takes private sector beliefs as given. In Kreps (1998) terminology, this is equivalent to assuming that the monetary authority is an anticipated utility maximizer. It can be argued that in real life beliefs have both a backward looking component, sensitive to past occurrences, and a forward looking one, which can be influenced by commitments of a credible CB. Hence, both of these aspects can be relevant for monetary policymaking. There has been extensive research on the topic of central banks' credibility under RE. We think it is important to understand also the other extreme, when the CB credibility plays no role, because expectations are backward looking.

[5] Note that the cost-push shock is usually assumed to be an AR(1) process, however we instead assume it to be *iid* to make the problem more tractable. This assumption is also supported by Milani (2006), who shows that learning can endogenously generate persistence in inflation data, and assuming a strongly autocorrelated cost-push shock becomes redundant.

[6] In other words, the probability that a firm in period t can reset the price is constant over time and across firms.

[7] As is shown in Rotemberg and Woodford (1997), equation (4) can be obtained as a quadratic approximation to the expected household's utility function; in this case, α is a function of structural parameters.

The policy problem is to minimize the social welfare loss (4), subject to the structural equations (1) and (3), and given the private sector's expectations:

$$\begin{aligned} &\min_{\{\pi_t, x_t, r_t\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \beta^t \left(\pi_t^2 + \alpha x_t^2\right) \\ &\text{s.t. } (1), (3) \\ &\qquad E_t^* \pi_{t+1},\ E_t^* x_{t+1} \text{ given for } \forall\ t \end{aligned} \tag{5}$$

As shown in Clarida, Gali, and Gertler (1999), the optimality condition to this problem (at time t) is:

$$\frac{\kappa}{\alpha}\pi_t + x_t = 0. \tag{6}$$

Using (6), Evans and Honkapohja (2003b) derive the following law of motion for inflation and the output gap, and the interest rate rule that implements these allocations:

$$\pi_t^{EH} = \frac{\alpha\beta}{\alpha + \kappa^2} E_t^* \pi_{t+1} + \frac{\alpha}{\alpha + \kappa^2} u_t \tag{7a}$$

$$x_t^{EH} = -\frac{\kappa\beta}{\alpha + \kappa^2} E_t^* \pi_{t+1} - \frac{\kappa}{\alpha + \kappa^2} u_t\ . \tag{7b}$$

$$r_t = \overline{rr}_t + \delta_\pi^{EH} E_t^* \pi_{t+1} + \delta_x^{EH} E_t^* x_{t+1} + \delta_u^{EH} u_t, \tag{7c}$$

where:

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

In the terminology introduced in Evans and Honkapohja (2003b), 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 agents have RE (i.e., if $E_t^* = E_t$), the system of equations (7) collapses to:

$$\pi_t^{RE} = \frac{\alpha}{\kappa^2 + \alpha} u_t, \quad x_t^{RE} = -\frac{\kappa}{\kappa^2 + \alpha} u_t\ ,$$

which is the optimal policy under discretion derived in Clarida, Gali, and Gertler (1999).

2.2 Optimal policy under learning

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

We assume that the private sector's expectations are formed according to the adaptive learning literature.[8] 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 shown in Clarida, Gali, and Gertler (1999), the optimal allocations of the discretion and the commitment solution under RE have different functional forms and are therefore associated with different PLMs. In this paper, we restrict our attention to the discretionary case. In particular, we assume that agents believe that inflation and the output gap are continuous invariant functions of the cost-push shock only, $\pi_t = \pi(u_t)$ and $x_t = x(u_t)$.[9] 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.[10] Throughout the paper we will assume that expectations evolve following the algorithm:

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

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

where γ_t is a deterministic sequence of gains in the interval $(0, 1)$, which governs how responsive estimate revisions are to new data. In the next two sections, we will be more explicit on the precise form taken by γ_t.

We choose equations (8)-(9) to model the private sector's PLM since they are consistent with the optimal discretionary RE solution in our setup; hence, it is the correct PLM if the CB has no credibility, which is the case under adaptive learning.[11]

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

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

[10] To be precise, in the algorithms (8) and (9), the observations are weighted geometrically if $\gamma_t = \gamma$, while if $\gamma_t = 1/t$ all observations receive equal weight.

[11] If we had assumed a *hybrid* NKPC, motivated by indexation to past inflation among firms, a model consistent PLM of private agents should also have included lagged inflation (as in Gaspar, Smets, and Vestin (2006)). We think the Gaspar, Smets, and Vestin (2006) analysis is important

To analyze the optimal control problem faced by the CB, we suppose that the policymakers take the structure of the economy (equations (1) and (3)) as given; moreover, we assume that the CB knows how private agents' expectations are formed, and takes into account its ability to influence the evolution of the beliefs. Hence, the CB problem can be stated as follows:

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

$$\text{s.t. } (1), (3), (8), (9),$$
$$a_0, b_0 \text{ given.}$$

Note that contrary to our benchmark problem (5), the CB now also takes first order conditions with respect to private expectations. When expectations depart from rationality and follow a law of motion, they become a natural state variable.

Assuming that the CB knows the exact learning algorithm followed by private agents is a strong hypothesis. In real life, there is still no consensus about how we should model private expectations. Nevertheless, we think it is important to examine how the policy recommendation changes if private agents depart slightly from rationality, and the monetary authority takes this departure into account. In Section 5, we relax this assumption and examine the robustness of our results when the CB is uncertain about how the private sector forms its expectations.

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

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

where λ_{it}, $i = 1, ..., 4$ denote the Lagrange multipliers associated with (1), (3), (8) and (9), respectively. The necessary conditions for an optimum are the first-order conditions, the structural equations (1)-(3) and the laws of motion of private

and more research is needed on how the exact nature of expectation formation modifies the optimal policy recommendation. Nevertheless, not assuming indexation not only enables us to derive closed-form solutions, but is also supported by empirical evidence. There is a recent strand of empirical literature that argues that the presence of indexation is not a robust feature of the data; see Benati (2008) and Cogley and Sbordone (2005), among others. Furthermore, Woodford (2007) questions the necessity (and the correctness) of price indexation to replicate inflation dynamics, especially when expectations are not rational.

agents' beliefs, (8)-(9). Note that the optimality conditions are not time invariant if the γ_t depends on time; however, because they are exogenous and deterministic, the policy function that solves the optimality conditions 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).[12] Combining equations (11) and (15), we get:

$$\lambda_{4t} = \beta \left(1 - \gamma_{t+2}\right) E_t \left[\lambda_{4t+1}\right],$$

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

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

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

$$2\pi_t + 2\frac{\alpha}{\kappa}x_t + \gamma_{t+1}\lambda_{3,t} = 0, \tag{17}$$

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 $\gamma_{t+1} = 0$, namely when expectations are constant and, consequently, cannot be manipulated by the monetary authority, (17) simplifies to:

$$\frac{\kappa}{\alpha}\pi_t + x_t = 0\ , \tag{18}$$

which is identical to the optimality condition derived in the RE optimal monetary policy literature when the CB sets the optimal plan taking the 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 the output gap at t: because of the nonzero term u_t in the Phillips Curve (3), 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 the output gap (inflation) is above target, it is optimal to deflate the economy (contract demand below capacity).

Under learning (i.e., when $\gamma_{t+1} > 0$), the CB faces an additional *intertemporal trade-off* between optimal behavior in t and in later periods, generated by its ability to manipulate future values of inflation expectations. The CB has to take

[12] A problem solved at t is said to be time consistent for $t+1$ if the continuation from $t+1$ of the optimal allocations chosen at t solves it 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 $t \geq 0$.

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 $\gamma_{t+1}\lambda_{3,t}$ shows an important difference compared with earlier results: the optimal decision should be conditional on the current stance of inflation expectations. The interpretation of this term is very simple: equation (8) implies that a change in π_t will influence the next period's inflation expectations, a_{t+1}, by a factor γ_{t+1}, and a change in inflation expectations affects welfare losses by a factor $\lambda_{3,t}$. The sign of $\lambda_{3,t}$ depends on current inflation expectations: because target inflation is zero, an increase in inflation expectations drives them further away from the target 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 occurs: 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 costs.[13] Brayton and Tinsley (1996) and Erceg and Levin (2003) argue that disinflation can be costly because of slowly adjusting expectations. Our results show that, under learning, it is indeed optimal to incur high output losses (compared with discretionary policy) in order to contain inflation expectations. Moreover, the higher inflation expectations are, the higher $\lambda_{3,t}$ is and the bigger the output loss the CB should engineer in order to bring down inflation.

When inflation expectations are negative ($\lambda_{3,t} < 0$), (17) 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.*

3 Constant gain learning

In this section, we assume that agents' beliefs are updated according to a *constant gain* algorithm, namely that $\gamma_t = \gamma \in (0, 1)$ for any t.[14] In Section 4 we will

[13] For 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 it cannot be used to model hyperinflationary episodes.

[14] As discussed extensively in the learning literature, private agents are likely to use such a learning scheme if they believe structural changes are going to occur.

relax this assumption and examine how optimal policy changes when agents follow decreasing gain learning.

We can combine the conditions for an optimum derived in Section 2.2, specialized to the constant gain specification, 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) *with* $\gamma_t = \gamma$*, and the policy function for inflation associated to it has the form:*

$$\pi_t = c_\pi^{cg} a_t + d_\pi^{cg} u_t. \tag{19}$$

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

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

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

and

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

Following the adaptive learning terminology, we call (19) the actual law of motion (ALM) of inflation.

Under the optimal policy, increasing a_t increases current inflation, but less than proportionally, because $\frac{\alpha\beta}{\alpha+\kappa^2} < 1$. As is shown in the Appendix, c_π^{cg} depends on all the structural parameters. 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 the one hand, a higher γ increases the effect of current inflation on future expectations, therefore the CB has a higher incentive to engineer a lower feedback from inflation expectations to inflation (i.e. a lower c_π^{cg}). On the other hand, a higher γ reduces the impact of current expectations on future expectations, which reduces the benefits from a reduction of the expectations, so there is an incentive to set a higher c_π^{cg}. In Figure 1 we show a numerical example with the calibration of Woodford (1999), with $\beta = 0.99$, $\sigma = 0.157$, $\kappa = 0.024$ and $\alpha = 0.04$. In this case, the first effect dominates, therefore c_π^{cg} is a monotonically decreasing function of γ. With different parameterizations, characterized by a higher κ and a lower α, the relationship would indeed be nonmonotonic, with c_π^{cg} being a decreasing function of γ for small values of the tracking parameter, and increasing when γ is big. However, empirical estimates of the tracking parameter find that γ is typically smaller than 0.1, therefore the decreasing brunch of c_π^{cg} as a function of the gain parameter seems the most relevant from an empirical point of view.[15]

[15] For examples of estimates of γ, see Milani (2007), Orphanides and Williams (2005a), and Branch and Evans (2006).

Using the structural equation (3) we can derive the optimal allocation of the output gap:

$$x_t = c_x^{cg} a_t + d_x^{cg} u_t, \tag{20}$$

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 rate sufficiently lower than the expected one. Using (19) and (20) in (1) we can derive the nominal interest rate:

$$r_t = \overline{rr}_t + \delta_\pi^{cg} a_t + \delta_x^{cg} b_t + \delta_u^{cg} u_t, \tag{21}$$

where:

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

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 that requires a more aggressive undercutting of inflation expectations (a lower c_π^{cg}) calls for more aggressive behavior from 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.

Because $c_\pi^{cg} < \beta$ (see Proposition 1) δ_π^{cg} is always greater than one. In response to a rise in expected inflation, optimal policy should raise the nominal interest rate sufficiently to increase the real interest rate. In other words, the Taylor principle emphasized in Clarida, Gali, and Gertler (1999) holds.

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

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

which is a stationary AR(1).[16] Thus, as is well known in the literature on adaptive learning, the contemporaneous presence of random shocks in the ALM and of a constant gain specification of the updating algorithm prevent the expectations from converging asymptotically to a precise value: instead, $a_t \sim N\left(0, \frac{\gamma^2 \left(d_\pi^{cg}\right)^2}{1-\left(1-\gamma(1-c_\pi^{cg})\right)^2}\sigma_u^2\right)$.

3.1 Comparison with the EH rule

In this section, we compare optimal monetary policy under constant gain learning to the rules used earlier in the literature, where the CB is treated as an anticipated utility maximizer. In particular, we refer to the rule (7c), derived in EH.

In the optimal interest rate rule (21), the coefficient on the output gap expectations is the same as in the discretionary rule (7c), while the other two coefficients are typically different. Proposition 1 implies $\delta_\pi^{cg} > \delta_\pi^{EH}$: the optimal interest rate response to out of equilibrium inflation expectations is more aggressive than the interest rate response of EH, hence inducing a smaller increase in inflation in response to an increase in a ($c_\pi^{cg} < c_\pi^{EH}$). 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 it would do otherwise.

From Proposition 1, it also follows that $\delta_u^{cg} > \delta_u^{EH}$: optimal policy reacts more aggressively to cost-push shocks then the EH rule. After a positive cost-push shock, the optimally behaving CB raises the interest rate more aggressively than what an anticipated utility maximizer CB would do; this in turn decreases output, which has a negative effect on inflation. Thus, an aggressive interest rate rule in response to the cost-push shock decreases the influence of the cost-push shock on inflation (in fact, $c_\pi^{cg} < \frac{\alpha\beta}{\kappa^2+\alpha}$ implies that $d_\pi^{cg} < \frac{\alpha}{\kappa^2+\alpha}$), and this in turn eases agents' learning about the true equilibrium level of inflation.

On the other hand, under optimal policy 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 (7c) and (21) can be summarized as follows:

Result 2. *When the central bank takes into account not only the intratemporal trade-off but also the intertemporal trade-off, it accommodates less the effect of out of equilibrium inflation expectations and noisy cost-push shocks on inflation. In this way, optimal policy facilitates learning of the private sector.*

It is also worth noting that optimal policy decreases the autocorrelation of inflation compared with EH.[17] The optimal rule's strong feedback to inflation expectations dampens the interaction between inflation and expectations. This

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

[17] It can be easily derived that the autocorrelation of inflation under constant gain with EH is

lowers the persistence of a shock's effect on expectations and on inflation. This result is analogous to the findings of Gaspar, Smets, and Vestin (2006) in a different model. They show that when firms index their prices to past inflation it is optimal to decrease inflation persistence.

Welfare loss analysis

To obtain a quantitative measure of the welfare gains of using the optimal learning rule (21), we present a numerical welfare loss analysis. Because welfare losses in utility terms are hard to interpret, we report consumption equivalents (following Adam and Billi (2007)): for a given monetary policy rule we calculate the cumulative utility losses resulting from deviations from the steady state allocation and then express the equivalent percentage decrease of the steady state consumption that results in the same cumulative utility loss. We use the calibration of Woodford (1999): $\beta = 0.99$, $\kappa = 0.024$, $\alpha = 0.048$ and $\sigma = 0.157$.[18] We perform a Monte Carlo with simulation length 10,000 and a cross-sectional sample size of 1,000. Cost-push shocks are drawn from a normal distribution with 0 mean and variance 0.1. Initial beliefs are the RE equilibrium: $a_0 = b_0 = 0$.

Table 1 reports consumption equivalents when agents use constant gain learning, both under the corresponding optimal learning rule (CG) and under the EH rule (7c).[19] For small tracking parameters, the results are in the range of the original estimates of Lucas (1987): consumption losses resulting from cyclical fluctuations are small.[20] The higher the tracking parameter, the higher the consumption equivalents are, both under optimal policy and under the EH rule, because of higher variance of inflation expectations (see also Figure 2). This in turn implies higher variance of inflation and output, both under CG (see equation (19) and (20)) and under EH (see equation (7)), and higher consumption equivalents.[21]

$E\pi_t^{EH}\pi_{t-1}^{EH} = \left(\frac{\alpha\beta}{\alpha+\kappa^2}\right)^2\left(1-\gamma+\gamma\frac{\alpha\beta}{\alpha+\kappa^2}\right)\sigma_{a_{EH}}^2 + \frac{\alpha\beta}{\alpha+\kappa^2}\left(\frac{\alpha}{\alpha+\kappa^2}\right)^2\gamma\sigma_u^2$ while under the optimal rule $E\pi_t^{CG}\pi_{t-1}^{CG} = (c_\pi^{cg})^2\left(1-\gamma+\gamma c_\pi^{cg}\right)\sigma_{a_{CG}}^2 + c_\pi^{cg}\left(d_\pi^{cg}\right)^2\gamma\sigma_u^2$. We have already seen that $\sigma_{a_{CG}}^2 < \sigma_{a_{EH}}^2$, $c_\pi^{cg} < \frac{\alpha\beta}{\alpha+\kappa^2}$ and $d_\pi^{cg} < \frac{\alpha}{\alpha+\kappa^2}$, thus $E\pi_t^{CG}\pi_{t-1}^{CG} < E\pi_t^{EH}\pi_{t-1}^{EH}$.

[18] Similar consumption equivalents are obtained using other standard calibrations, like Clarida, Gali, and Gertler (2000) and McCallum and Nelson (1999).

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

[20] Consumption equivalents are higher if we start the economy away from the RE equilibrium. Also, the gain of using the optimal rule is higher when initial expectations are further away from the RE equilibrium, because the main advantage of the optimal rule is that it helps private agents to learn the equilibrium faster then the EH rule.

[21] 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

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

γ	p^{CG}	p^{EH}	p^{CG}/p^{EH}	p_π^{CG}/p_π^{EH}	p_x^{CG}/p_x^{EH}
0.0183	0.013	0.0130	0.9991	0.9966	1.2097
0.05	0.0148	0.0151	0.9774	0.9464	3.5609
0.08	0.0171	0.0184	0.9280	0.8561	6.9223
0.1	0.0188	0.0211	0.8881	0.7914	8.941
0.3	0.0369	0.0608	0.6068	0.4246	15.7893
0.5	0.0551	0.1104	0.4994	0.3114	16.1679
0.9	0.0908	0.2187	0.4151	0.2311	15.7401

Woodford (1999) calibration, $a_0 = 0$.

The gain from using optimal policy over the EH rule can be nonnegligible even if initial inflation expectations are at the RE equilibrium, and expectations stay close to the RE equilibrium. For tracking parameters below 0.05, which is a typical range of estimates for the US[22], the gain from using an optimal interest rate rule can be around $1-3\%$. The higher the gain parameter, the more optimal policy decreases inflation expectations variance compared with EH (also see Figure 2), and the bigger is the gain in consumption equivalents. For the extreme case of $\gamma = 0.9$, the steady state consumption loss of using the EH rule is 60% higher than under the the optimal rule.

The long-run gains of containing inflation expectations come at a cost in the short run. Figure 4 plots the transition path of cumulative consumption equivalents. In the first periods, the optimal interest rate rule (21) 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 results 1 and 2: because of the intertemporal trade-off, it is optimal to react to out of equilibrium inflation expectations more aggressively than the EH rule in order to undercut more future expectations, even if it results in short-term output gap losses. As soon as inflation expectations become small enough, this initial loss is more than compensated.[23]

Another way to gauge what the intertemporal trade-off implies for welfare is to calculate separately the equivalent permanent consumption decrease because of losses caused by only inflation or output gap variation (see Table 1). The main

σ_u^2, but the ratio of consumption equivalents under CG and EH is not sensitive to the choice of σ_u^2.

[22]See for example Branch and Evans (2006), Milani (2007) and Orphanides and Williams (2005a).

[23]Results not reported here show that the further away initial expectations are from the RE equilibrium, the larger the long-run gains are and the bigger are the short-run costs of using the optimal rule. For details, see the working paper version of this paper.

result is that optimal policy lowers inflation variation at the cost of higher output gap variation. The higher is the tracking parameter, the more inflation variation is lowered: for $\gamma = 0.9$ an optimally behaving CB engineers a 77% lower welfare loss in inflation when it properly conditions on expectation formation, permitting at the same time 15 times more variation in the output gap compared with the EH rule.[24]

In this section, we derived the fully optimal monetary policy when agents follow constant gain learning and compared it to the optimal discretionary rule, when the CB does not make active use of its influence on expectations. The next section shows similarities to the commitment solution under RE.

3.2 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 the similarities in short-run behavior, in the limit, the two equilibria are different. The learning equilibrium intrinsically depends on how private agents learn.

Figure 5 displays the impulse response function of inflation to a unit shock under CG and discretionary RE policy. In the optimal RE discretionary policy, inflation rises on impact and immediately reverts to the steady state once the *iid* shock dies out. Under learning the policymaker engineers a smaller initial response of inflation; in subsequent periods inflation gradually converges back to the steady state value. Gali (2003) shows a *similar disinflation path for the optimal policy under RE and commitment*: a smaller initial inflation compared with the discretionary case, in exchange for a more persistent deviation from the steady state later.[25]

These similarities arise because under both learning and RE commitment the CB can directly manipulate private expectations, even if the channels used are quite different. Under commitment, the policymaker uses a *credible promise on the future* to obtain an immediate decline in inflation expectations and thus in inflation; moreover, the necessity to fulfill past commitments introduces additional

[24] In the framework of Gaspar, Smets, and Vestin (2006), the CB engineers a lower welfare loss in inflation without a significant cost in output. This result is difficult to compare with ours because the presence of indexation changes their setup along three important dimensions: the CB wants to stabilize a *quasi-difference* of inflation instead of inflation itself, the NKPC is of the hybrid type, and in our model agents learn about the expected value of inflation while in Gaspar, Smets, and Vestin (2006) agents learn about the persistence of inflation.

[25] This behavior of optimal policy under commitment leads to welfare gains over discretion because of the convexity of the loss function; this preference for slower but milder adjustment to shocks is at the heart of the stabilization bias.

inertia in inflation and output. Under learning, we observe a smaller initial response of inflation relative to the RE discretionary case because optimal policy dampens the inflation response to the cost-push shock to *ease private agents' learning* (Result 2), and the past-dependent nature of private sector beliefs imparts sluggishness on the system. 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 the output gap.

One difference compared with 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, yet a positive sequence under learning. A second difference is that the full commitment is characterized by a smaller output decrease compared with RE discretionary policy (see Clarida, Gali, and Gertler (1999)), while under learning the initial decrease of output is bigger than under discretion and RE. The reason for this is that while under RE the commitment of the CB can improve the current terms of the inflation-output trade-off, under learning monetary policy can only influence future expectations and can improve only future inflation-output trade-offs.

Sargent (1999), chapter 5, shows a similarity between the optimal policy under adaptive learning and the RE commitment solution in the Phelps problem: optimal monetary policy drives the economy close to the Ramsey optimum, and when the discount factor β equals 1, optimal policy under learning replicates the Ramsey equilibrium. In the Phelps problem, the discretion and commitment outcome of inflation have the same functional form, therefore when agents learn in this functional form they can converge to both equilibria. A sufficiently patient CB is willing to incur higher short-term losses for the opportunity to drive private expectations to the welfare-improving Ramsey equilibrium.

In our model, discretionary and commitment solutions under RE have a different functional form; hence the equilibrium depends on how agents learn, and Sargent's result does not hold anymore. However, in our case an increase in the discount factor also makes the optimal disinflationary path under learning move closer to the commitment solution. This can be seen in Figure 5: as β gets closer to 1, the initial response of inflation becomes milder and the path back to the steady state longer.

The findings in this subsection strengthen 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, and the implications for policymaking go beyond the asymptotic learnability criterion: as we showed, the equilibrium law of motion of optimal inflation can be significantly affected by the way agents learn, and careful consideration of private sector beliefs can play a role qualitatively similar to a commitment device, even in the absence of CB credibility.

4 Decreasing gain learning

In this section, we relax the assumption of constant gain learning and show that our main results remain valid also with decreasing gain learning .[26]

We assume agents use the following decreasing gain learning rules (henceforth DG):

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

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

where the only difference from (8)-(9) is the substitution of γ with t^{-1}. Under certain conditions on the values used to initialize this algorithm (see Evans and Honkapohja (2001)), it is equivalent to estimating the conditional expectations of inflation and output gap every period with OLS.[27]

In the Appendix, we derive the following optimal allocations.

Proposition 2. *The solution of the control problem* (10) *with* $\gamma_t = 1/t$ *yields the following policy function for inflation:*

$$\pi_t = c_{\pi,t}^{dg} a_t + d_{\pi,t}^{dg} u_t, \quad (24)$$

where $c_{\pi,t}^{dg}$ *and* $d_{\pi,t}^{dg}$ *are deterministic functions of time characterized as follows:*

- $\lim_{t\to\infty} c_{\pi,t}^{dg}$ *exists, and is given by* $\lim_{t\to\infty} c_{\pi,t}^{dg} = \dfrac{\alpha\beta}{\alpha+\kappa^2}$;
- *for any* $t < \infty$, *we have that* $c_{\pi,t}^{dg} < \dfrac{\alpha\beta}{\alpha+\kappa^2}$

and

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

where the matrices $P_{1,t}$ *and* $A_{11,t}$ *are defined in the Appendix.*

[26]Decreasing gain algorithms place equal weight on all observations, which is optimal in stationary environments.

[27]Note that, because the conditional expectations of inflation and output gap are assumed by the learners to be constant, the OLS estimate is just the sample averages of the two.

With decreasing gain, during the transition Result 2 holds: there is a new intertemporal trade-off, therefore it is optimal to decrease the effect of out of equilibrium expectations on inflation compared with the EH rule (equation (7)) in order to drive future inflation expectations closer to the equilibrium. This relaxes the future intratemporal inflation-output gap trade-off embedded in the Phillips curve. The ALM for the output gap is:

$$x_t = c_{x,t}^{dg} a_t + d_{x,t}^{dg} u_t, \tag{25}$$

where

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

As in the constant gain case, if the private sector expects inflation to be positive, the optimal CB will contract economic activity more than the EH rule.[28] The CB is ready to pay a short-term cost represented by a wider current output gap in order to contain future inflation expectations.

The nominal interest rate rule is:

$$r_t = \overline{rr}_t + \delta_{\pi,t}^{dg} a_t + \delta_x^{dg} b_t + \delta_{ut}^{dg} u_t, \tag{26}$$

where

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

Because $c_{\pi,t}^{dg} < \beta$ (see Proposition 2) $\delta_{\pi,t}^{dg}$ is always bigger than 1; hence, the Taylor principle holds. In the Appendix, the following results are derived.

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

- $\delta_{\pi,t}^{dg} > \delta_{\pi}^{EH}$, $\delta_{ut}^{dg} > \delta_{u}^{EH}$,
- $\lim_{t\to\infty} \delta_{\pi,t}^{dg} = \delta_{\pi}^{EH}$, $\lim_{t\to\infty} \delta_{ut}^{dg} = \delta_{u}^{EH}$.

During the transition, the optimal interest rate rule is similar to the constant gain rule: it reacts more aggressively to out of equilibrium expectations (and cost-push shocks) than the EH rule.

[28] From $c_{\pi,t}^{dg} < \frac{\alpha\beta}{\alpha+\kappa^2}$ it follows that $c_{x,t}^{dg} < -\frac{\kappa\beta}{\alpha+\kappa^2}$. Compare with the ALM under EH (7).

An interesting result is that the coefficient on inflation expectations in the interest rate rule (26) is *time-varying*, reflecting the fact that the CB'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. This coefficient can be characterized as follows:

Proposition 4. *Let $\delta_{\pi,t}^{dg}$ be given by $1-\sigma\frac{c_{\pi,t}^{dg}-\beta}{\kappa}$; then, there exists a $T<\infty$ such that $\left\{\delta_{\pi,t}^{dg}\right\}_{t=T}^{\infty}$ is a monotonic decreasing sequence.*

After time T, the bank dampens its aggressiveness in reacting to out of equilibrium inflation expectations (and cost-push shocks).[29] For empirically relevant coefficient estimates, time T is maximum a few quarters. Numerical analysis on the grid $\beta=0.99$ and $\alpha\in[0.01,2]$, $\kappa\in[0.01,0.5]$ shows that T is typically very small.[30] We find that after the fourth period (from the fourth to the fifth period and so on) $\delta_{\pi,t}^{dg}$ always decreases, while in the first four periods $\delta_{\pi,t}^{dg}$ might increase (hump-shaped) for a combination of low values of α and high values of κ (see Figure 6). Figures 7 and 8 show that for the Woodford (1999) calibration, $\delta_{\pi,t}^{dg}$ and $\delta_{u,t}^{dg}$ always decrease over time (i.e., $T=0$).[31]

To further clarify this issue, consider the following example: a new CB governor is appointed, and agents start learning how this affects the equilibrium. In this situation it is optimal for the CB to react more aggressively to out of equilibrium inflation beliefs in the first period, when agents pay more attention to new information and the CB's possibilities of influencing private expectations are therefore greater. This policy is beneficial even at the cost of larger short-term losses in terms of output gap variability. As time passes, expectations will be influenced to a lesser extent by the most recent realizations of the inflation rate, hence the CB's reaction will more closely reflect the case where the policymakers cannot manipulate expectations.

The asymptotic properties of the ALM, (24) and (25), depend on the limiting behavior of a_t, which is given by the stochastic recursive algorithm:

$$a_{t+1}=a_t+(t+1)^{-1}\left((c_{\pi t}^{dg}-1)a_t+d_{\pi,t}^{dg}u_t\right). \tag{27}$$

We study its properties in the Appendix, where we use the stochastic approximation techniques[32] to prove the following Proposition:

[29]From (26) it is easy to see that the change in $\delta_{u,t}^{dg}$ through time has the same sign as $\delta_{\pi,t}^{dg}$.

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

[31]$\delta_{\pi,t}^{dg}$ is always decreasing also for other calibrations widely adopted in the New Keynesian literature, such as those taken from Clarida, Gali, and Gertler (2000) and McCallum and Nelson (1999).

[32]For an extensive monograph on stochastic approximation, see Benveniste, Métivier, and

Proposition 5. *Let a_t evolve according to (27); then, $a_t \to 0$ 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 $d_{\pi,t}^{dg} \to \frac{\alpha}{\kappa^2+\alpha}$, therefore we can conclude that $\pi_t \to \frac{\alpha}{\kappa^2+\alpha} v$ almost surely, where v is a random variable with the same probability distribution as u_t. The equilibrium corresponds to the discretionary RE equilibrium, and private agents learn the unconditional expectation of inflation and output under discretionary RE.[33]

It follows from Proposition 3 that the optimal policy converges to the EH policy. Because in the limit expectations converge to a constant, it is intuitive that in the limit optimal policy behaves as if expectations were fixed. However, during the transition optimal policy results in substantially lower welfare losses. For the Woodford (1999) calibration, even if we start inflation expectations from the RE equilibrium, $a_0 = 0$, in the long run the consumption equivalent under the optimal rule is about 10% lower than that of EH. If initial expectations are slightly different from the long-run equilibrium then gains are even higher. For $a_0 = 1$ the welfare losses under the optimal policy are 42% lower than under EH. In the first period, the optimal interest rate rule (26) yields *ex-post* higher cumulative welfare losses expressed in terms of consumption than the EH rule; later, however, our rule starts generating smaller welfare losses. These findings are similar to the numerical results of the constant gain.

An alternative way to examine the mechanisms at work when the CB employs the optimal rule instead of the EH rule is to look at the path of expectations. Both the optimal and the EH rule are E-stable under learning, so expectations converge to the discretionary REE; the difference is the speed of convergence. Figure 9 shows a typical realization of the evolution of expectations under both rules. 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. However, 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.

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

Priouret (1990); the first paper to apply these techniques to learning models was Marcet and Sargent (1989).

[33]Note that the PLM of private agents does not nest the commitment REE, only the discretionary REE.

converge to the discretionary RE equilibrium, and optimal policy is equivalent to the one derived under the assumption of constant expectations.

5 Robust policy advice

In this section, we relax the assumption that the CB has perfect knowledge about the learning algorithm followed by private agents and ask what the policy recommendation is when the monetary authority is uncertain about the nature of private sector expectations.[34] In particular, we aim to define consistent policy advice on a set of private agents' expectations formation schemes empirically relevant for the US.

Let us conduct an experiment, in which we assume that the US Federal Reserve is uncertain about how the private sector forms its expectations but, relying on the empirical literature, it can define a relevant set of expectations, which includes both constant gain with a small gain and RE. The empirical literature on the US shows that a constant gain algorithm with a small tracking parameter is a good approximation of the data. For example, Milani (2007) estimates the New Keynesian model with adaptive learning using Bayesian methods and finds γ to be 0.0183. Orphanides and Williams (2005a) and Branch and Evans (2006) calibrate γ to fit the Survey of Professional Forecasters and find that tracking parameters between 0.01 and 0.04 fit survey expectations well. We therefore consider tracking parameters in this range.

In addition, let us assume that the FED has no probability distribution over the possible forms of private expectations. Instead, it uses robust control and looks for the policy that minimizes the maximum loss.[35]

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 discretionary rule under RE (7c). We assume initial inflation expectations coincide with RE ($a_0 = 0$), so constant gain expectations with a small gain will stay close to the rational forecasts. We can think of this economy as populated by agents who are making only very small mistakes compared with the rational forecasts.

Table 2 reports consumption equivalents. In order to present the results in a compact way, the last line of Table 2 shows percentage increases in consumption

[34] We have assumed that the CB perfectly observes all the relevant state variables of the system, namely the exogenous shocks and the agents' beliefs. It is possible to show that our main results extend to a more general framework, where the shocks or the expectations are not observable, and the CB has to solve a signal extraction problem to learn about them. For details, see the working paper version of this paper.

[35] For an extensive treatise on the use of robust control techniques in economics, see Hansen and Sargent (2007).

equivalents of the worst case compared with the optimal rule for a given private expectation. The main result is that the worst-case scenario is using the EH rule

Table 2: Consumption equivalents under the optimal or an incorrect rule, initial inflation expectations at RE

Expectations	$\gamma = 0.0183$	$\gamma = 0.03$	$\gamma = 0.04$	RE
Interest rate rule				
$\gamma = 0.0183$	0.013006	0.013548	0.014157	0.012642
$\gamma = 0.03$	0.013009	0.013542	0.014130	0.012642
$\gamma = 0.04$	0.013015	0.013545	0.014126	0.012643
EH	0.013017	0.013606	0.014294	0.012641
The worst rule	EH	EH	EH	$\gamma = 0.04$
% increase of cons.eq.*	0.13	0.47	1.18	0.02

* worst rule compared with the optimal

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. A min-max rule (following Hansen and Sargent (2007)), which minimizes the maximum loss, is always a rule that is optimal under learning.

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 2), and the reason for this is that learning rules allow for volatility in the output gap that is too high.[36]

However, losses under RE caused by mistakenly using an optimal learning rule are smaller than losses associated with the use of the discretionary EH rule when agents are learning.

When private agents are learning and the FED uses a misspecified learning rule, consumption equivalents increase but the loss is always smaller than losses associated with using the EH rule. The bigger the misperception of the monetary policy about γ is, the bigger the increase in consumption equivalents. When, for example, agents follow constant gain with $\gamma = 0.04$ and the CB uses the optimal interest rate rule, the consumption equivalent is 0.014126. When the optimal rule with $\gamma = 0.03$, is used the consumption equivalent increases to 0.01413, which is a 0.03% increase. If the FED uses $\gamma = 0.0183$, which is further from the true tracking parameter, the consumption increases by 0.22% to 0.014157. The percentage increase in loss achieved using the EH rule is 1.18%, which is bigger than with any of the learning rules.

[36]Since learning rules decrease the volatility of inflation and allow for higher volatility in the output gap, for small values of α learning rules do outperform the EH rule even under RE.

The main advantage of the optimal learning rules compared with EH is that they help private agents to learn the RE forecasts faster. When we initialize the economy at the RE equilibrium, it is likely that this advantage would be quantitatively less relevant, since beliefs stay close to RE. Therefore, we repeat the numerical analysis for $a_0 = 0.1$. Table 3 reports that the gain associated with using a learning rule over the EH rule is bigger in this case. The EH rule increases consumption equivalents compared with the optimal policy by $0.4-2\%$. Learning rules, on the other hand, result in smaller losses under learning, even if they are misspecified.

Table 3: Consumption equivalents under the optimal or a wrong rule, initial inflation expectations not RE

Expectations	$\gamma = 0.0183$	$\gamma = 0.03$	$\gamma = 0.04$	RE
Interest rate rule				
$\gamma = 0.0183$	0.014132	0.014593	0.01514	0.012642
$\gamma = 0.03$	0.014143	0.014577	0.015094	0.012642
$\gamma = 0.04$	0.014163	0.014584	0.015086	0.012643
EH	0.01418	0.014722	0.015372	0.012641
The worst rule	EH	EH	EH	$\gamma = 0.04$
% increase of cons.eq.*	0.39	0.98	1.87	0.03

* worst rule compared with the optimal

Woodford (1999) calibration. Starting out of RE equilibrium: $a_0 = 0.1$. Consumption equivalents for a given underlying private sector expectation formation and a given interest rate rule.

Let us now assume that the monetary authority is able to formulate a probability distribution over private expectations, and ask what is the minimum probability of the private agents following learning that makes the FED choose a learning rule over a discretionary rule. In particular, let's assume that the FED's prior 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 as the sum of two components: p times the consumption equivalent under constant gain learning with the EH rule, and $1-p$ times the consumption equivalent of using EH under RE. Then, find a cutoff value of p for which the expected loss in consumption terms of using the CG rule is less than the welfare loss of the EH rule.

A surprising result is that the cutoff value of p is between 1 and 1.5% even when initial expectations are at the RE equilibrium.[37] This means that it is optimal to use the learning rule even if the CB attributes only a very small probability to agents following learning and a very high probability to RE.

[37] Cutoff values for p are lower when we initialize the economy out of RE.

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

6 Conclusions

We have shown that deviations of private sector beliefs from RE can have important policy implications that go beyond the possible failure to converge asymptotically to RE. Besides the well-known intratemporal inflation-output gap trade-off, when expectations are less than rational the CB faces an intertemporal trade-off between stabilizing the economy in the present and in the future. Optimal policy does not stabilize the economy in the present in order to influence future inflation expectations and ease future inflation-output gap trade-offs.

Importantly, if there is uncertainty about the way the private sector forms its beliefs, optimal policy derived under less than RE is more robust than optimal policy under RE. Because there is an ongoing debate regarding the degree of rationality embedded in private expectations, our results suggest that monetary policy should take into account the possibility of an intertemporal trade-off arising from the nature of private expectations formation.

Overall, our results suggest that the way private expectations are formed is a significant issue for policy design, and monetary policy should closely monitor private expectations. We hope this result will motivate more research interest in understanding how private expectations are formed in different environments. In most of the paper, we took the extreme view that the CB has perfect knowledge; because it is a strong assumption, it seems important to learn by how much optimal policy changes when the CB does not have full knowledge of the way private expectations are formed, such as for example in Woodford (2009).

A Constant gain learning

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

Combining the optimality conditions (11-14) and (16) we can write:

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

Using the Phillips curve (3) and the evolution of inflation expectations (8), we get:

$$E_t [\pi_{t+1}] = A_{11}\pi_t + A_{12}a_t + P_1 u_t, \ , \tag{28}$$

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 (28) and obtaining the trivariate system:

$$E_t y_{t+1} = A y_t + P u_t \ , \tag{29}$$

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:

$$\begin{array}{c} a_0, \, b_0 \text{ given} \\ \lim_{s\to\infty} |E_t \pi_{t+s}| < \infty \end{array} \ . \tag{30}$$

The last one is a result of 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 case.[38]

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

$$A_1 \equiv \begin{pmatrix} A_{11} & A_{12} \\ \gamma & 1-\gamma \end{pmatrix} \ . \tag{31}$$

[38] The proof is available from the author upon request.

In 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 (29)-(30) has one and only one solution. In other words, there exists one and only one stochastic process for each of the three variables of y such that (30) 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 (29)-(30) 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

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

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 (31) in the text; then it has an eigenvalue inside and one outside the unit circle.*

Proof. First of all, we recall that a necessary and sufficient condition for a 2 by 2 matrix to have one eigenvalue inside and one outside the unit circle, is that:[39]

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

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

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

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

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:

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

[39] LaSalle (1986).

so that all we have to prove is that:

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

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

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

which is always true, because β and γ are assumed to be smaller than one. □

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 (28) and the method of undetermined coefficients to verify that c_π^{cg} must satisfy the following quadratic expression:

$$p_2\left(c_\pi^{cg}\right)^2+p_1c_\pi^{cg}+p_0=0\ ,$$

where

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

and that:

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

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

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

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, because the two roots of the quadratic expression have the same sign (this is due to the fact that both p_2 and p_0 are positive), it follows that the other candidate value for c_π^{cg} is greater than one, which is not compatible with the boundary conditions.[40]

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* $[0,1]$.

[40]Because it would imply an exploding inflation.

Proof. Note that:

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

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

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

After rearranging:

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

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

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

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

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

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

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

After rearranging, we get:

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

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

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 know that for any a, b, there exists a $c \in (a, b)$ such that:

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

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:

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

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

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

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

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

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

Moreover, we prove the following result.

Lemma 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:

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

if and only if:

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

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

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

which is equal to one; in fact:

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

is equivalent to:

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

But the LHS can simplified as:

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

which is equal to:

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

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

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

which implies that:

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

□

We are now ready to prove the Proposition.

Proof of Proposition 1. Combining the lemmas 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.

□

B Decreasing gain learning

In this section, we prove Propositions 2 and 5.

Proof of Proposition 2. To derive the optimal allocations, note that we can use first-order conditions (11-14) and (16) and $\gamma_t=1/t$ to rewrite:

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

Using (3) to substitute out x_t in the above equation, and then using the evolution of inflation expectations (22) we get:

$$E_t\left[\pi_{t+1}\right]=A_{11,t}\pi_t+A_{12,t}a_t+P_{1,t}u_t\ ,\tag{32}$$

where:

$$A_{11,t} \equiv \frac{\kappa^2 + \alpha + \alpha\beta^2 \frac{1}{t+1}\left(1 + \beta \frac{1}{t+1}\right)}{\alpha\beta(1 + \beta\frac{1}{t+1}) + \kappa^2\beta},$$

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

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

Hence, at an optimum, the dynamics of the economy can be summarized by stacking equations (22), (23) and (32), and obtaining the trivariate system:

$$E_t y_{t+1} = A_t y_t + P_t u_t \ , \tag{33}$$

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

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

We can find the solution with the method of undetermined coefficients with the guess:[41]

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

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

$$c^{dg}_{\pi,t} = \frac{c^{dg}_{\pi,t+1}\left(1 - \frac{1}{t+1}\right) - A_{12,t}}{A_{11,t} - c^{dg}_{\pi,t+1}\frac{1}{t+1}}, \tag{34}$$

and the sequence $\left\{d^{dg}_{\pi,t}\right\}$ is defined as

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

as stated in the Proposition. Clearly, once we solve for $c^{dg}_{\pi,t}$, finding the value of $d^{dg}_{\pi,t}$ is a trivial task. Of course, there exist infinite sequences that satisfy equation (34), one for each initial value $c^{dg}_{\pi,0}$. However, because the boundary conditions require π_t to stay bounded, we will concentrate on the solutions that do not explode. To characterize its properties, first note that if we solve forward the following difference equation:

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

[41] This guess corresponds to the unique solution under constant gain learning.

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

$$c^{dg}_{\pi t} = \frac{\alpha\beta}{\kappa^2+\alpha} \quad \forall t \ .$$

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

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

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

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

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

$$\beta^{-1} c^{dg}_{\pi,\tau} - c^{dg}_{\pi,\tau+1} \ .$$

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

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

which, as we state in the Proposition, has one and only one bounded solution:

$$c^{dg}_{\pi\tau} = \frac{\alpha\beta}{\kappa^2+\alpha} \ .$$

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

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

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

$$\left(1 - \frac{1}{T+1}\left(1-\frac{\alpha\beta}{\alpha+\kappa^2}\right)\right) c^{dg}_{\pi T+1} \geq \frac{\alpha\beta}{\alpha+\kappa^2} A_{11,T} + A_{12,T} \ . \tag{35}$$

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 a result of the fact that $\left(1+\beta\frac{1}{t+1}\right)^{-1}<1$; putting together the last inequality and (35), 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 $\lim_{t\to\infty}c_{\pi t}^{dg}>\frac{\alpha\beta}{\alpha+\kappa^2}$, which is a contradiction to the result stated before, namely $\lim_{t\to\infty}c_{\pi t}^{dg}=\frac{\alpha\beta}{\alpha+\kappa^2}$. Hence, we have showed that there is no $t<\infty$ such that $c_{\pi t}^{dg}\geq\frac{\alpha\beta}{\alpha+\kappa^2}$. □

To prove Proposition 4 we first state and prove the following technical lemma:

Lemma 6. *Let λ_1 be the smallest root of the second-order polynomial:*

$$\rho(p)\equiv\omega_2p^2+\omega_1p+\omega_0\ ,$$

where:

$$\begin{aligned}\omega_2&\equiv-\gamma\left[\left(\kappa^2+\alpha\right)\beta+\alpha\beta^2\gamma\right],\\ \omega_1&\equiv\left[\left(\kappa^2+\alpha\right)(1-\beta(1-\gamma))+\alpha\beta^2\gamma^2(1+\beta)\right],\\ \omega_0&\equiv-\alpha\beta\left[1-\beta\left(1-\gamma+\beta\gamma-\beta\gamma^2\right)\right]\ ,\end{aligned}$$

and where the restrictions on the parameters α, β and κ are the same as those imposed in the rest of the paper. Then, there exists a $\overline{\gamma}\in(0,1)$ such that when $\gamma\in(0,\overline{\gamma}]$, the following holds:

$$\frac{\partial}{\partial\gamma}\lambda_1<0\ .$$

Proof. First of all, note that applying the implicit function theorem, we have:

$$\frac{\partial}{\partial\gamma}\lambda_1=-\left.\frac{\partial\rho/\partial\gamma}{\partial\rho/\partial p}\right|_{p=\lambda_1}\gtrless0\Leftrightarrow\left.\frac{\partial\rho}{\partial\gamma}\right|_{p=\lambda_1}\lesseqgtr0\ ,\tag{36}$$

where we used the fact that, because $\omega_2<0$, $\left.\frac{\partial\rho}{\partial p}\right|_{p=\lambda_1}>0$. Moreover, we have:

$$\frac{\partial\rho}{\partial\gamma}\equiv\psi(p)=\varpi_2p^2+\varpi_1p+\varpi_0\ ,$$

where

$$\begin{aligned}\varpi_2&\equiv-\left[\left(\kappa^2+\alpha\right)\beta+2\alpha\beta^2\gamma\right],\\ \varpi_1&\equiv\left[\left(\kappa^2+\alpha\right)\beta+2\alpha\beta^2\gamma(1+\beta)\right],\\ \varpi_0&\equiv\alpha\beta^2\left[\beta-2\beta\gamma-1\right]\ .\end{aligned}$$

It is easy to show that, (i) there exists a $\overline{\gamma}_1$ such that, for $\gamma<\overline{\gamma}_1$, the largest root of $\rho(\cdot)$ is bigger than the largest root of $\psi(\cdot)$ (actually, the former goes to infinity as γ goes to zero); (ii) there exists a $\overline{\gamma}_2$ such that, for $\gamma<\overline{\gamma}_2$, the quadratic polynomial:

$$\rho(p)-\psi(p)\,,$$

has one positive and one negative root. Combining this result with the fact that both $\rho(\cdot)$ and $\psi(\cdot)$ are concave, we obtain that, for $\gamma < \overline{\gamma} \equiv \min\{\overline{\gamma}_1, \overline{\gamma}_2\}$, the smallest root of $\rho(\cdot)$ lies between the two roots of $\psi(\cdot)$; in other words:

$$\left.\frac{\partial \rho}{\partial \gamma}\right|_{p=\lambda_1} > 0 \ .$$

Using this result in (36) completes the proof. □

An immediate corollary of the above lemma is the following:

Corollary 1. *Let λ_{1t} be the smallest root of the second-order polynomial:*

$$\rho_t(p) \equiv \omega_{2t} p^2 + \omega_{1t} p + \omega_{0t} \ ,$$

where

$$\begin{aligned}
\omega_{2t} &\equiv -\tfrac{1}{t+1}\left[\left(\kappa^2+\alpha\right)\beta + \alpha\beta^2 \tfrac{1}{t+1}\right], \\
\omega_{1t} &\equiv \left[\left(\kappa^2+\alpha\right)\left(1-\beta\left(1-\tfrac{1}{t+1}\right)\right) + \alpha\beta^2\left(\tfrac{1}{t+1}\right)^2(1+\beta)\right], \\
\omega_{0t} &\equiv -\alpha\beta\left[1-\beta\left(1-\tfrac{1}{t+1}+\beta\tfrac{1}{t+1}-\beta\left(\tfrac{1}{t+1}\right)^2\right)\right] \ .
\end{aligned}$$

Then, there exists a $T < \infty$ such that $\{\lambda_{1t}\}_{t=T}^{\infty}$ is a monotonic increasing sequence.

Proof. First of all, note that λ_{1t} and ω_{it}, $i = 1, 2, 3$, are defined as the correspondent coefficient in the statement of Lemma 6, with γ replaced by $(t+1)^{-1}$; hence, $t+1 \geq 2$ is equivalent to $\gamma \leq \overline{\gamma}$ implies $t+1 \geq T+1$, where $T+1$ is the integer part of $\frac{1}{\gamma}$. Invoking the result of Lemma 6, we get that λ_{1t} increases as $(t+1)^{-1}$ decreases. □

We are now ready to prove Proposition 4.

Proof of Proposition 4. First of all, note that $\delta_{\pi t}^{dg}$ is decreasing if and only if $c_{\pi t}^{dg}$ is increasing; hence, we prove this latter statement. Recall that:

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

which means that, for any finite t, we have:

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

Because $1 - \frac{1}{t+1}\left(1 - c_{\pi,t}^{dg}\right)$ is a positive expression, $c_{\pi t+1}^{dg} - c_{\pi t}^{dg} \gtrless 0$ is equivalent to the second-order inequality:

$$\omega_{2t}\left(c_{\pi t}^{dg}\right)^2 + \omega_{1t} c_{\pi t}^{dg} + \omega_{0t} \gtrless 0 \ ,$$

where

$$\omega_{2t} \equiv -\tfrac{1}{t+1}\left[\left(\kappa^2+\alpha\right)\beta + \alpha\beta^2\tfrac{1}{t+1}\right],$$
$$\omega_{1t} \equiv \left[\left(\kappa^2+\alpha\right)\left(1-\beta\left(1-\tfrac{1}{t+1}\right)\right) + \alpha\beta^2\left(\tfrac{1}{t+1}\right)^2(1+\beta)\right],$$
$$\omega_{0t} \equiv -\alpha\beta\left[1-\beta\left(1-\tfrac{1}{t+1}+\beta\tfrac{1}{t+1}-\beta\left(\tfrac{1}{t+1}\right)^2\right)\right] .$$

Let λ_{1t}, λ_{2t} be the two roots of the above quadratic expression, such that $\lambda_{1t} < \lambda_{2t}$; because ω_{2t},$\omega_{0t} < 0$ for any t, and $-(\omega_{1t}/\omega_{2t})$ can be easily shown to be positive, we know that λ_{1t},$\lambda_{2t} > 0$ and that:

$$\lambda_{1t} < c^{dg}_{\pi t} < \lambda_{2t} \Longleftrightarrow c^{dg}_{\pi t+1} - c^{dg}_{\pi t} > 0 .$$

It is easy to see that $\lambda_{2t} > \frac{\alpha\beta}{\alpha+\kappa^2}$ for any t, which implies that:

$$\lambda_{1t} < c^{dg}_{\pi t} < \frac{\alpha\beta}{\alpha+\kappa^2} \Longleftrightarrow c^{dg}_{\pi t+1} - c^{dg}_{\pi t} > 0 ,$$

because we showed in Proposition 2 that $c^{dg}_{\pi t} < \frac{\alpha\beta}{\alpha+\kappa^2}$ for any finite t. Now assume, for the sake of contradiction, that $c^{dg}_{\pi t} \leq \lambda_{1t}$ for some $\tau \geq T$, where T is the one defined in Corollary 1; then, $c^{dg}_{\pi\tau+1} \leq c^{dg}_{\pi\tau}$ and, for Corollary 1, $\lambda_{1\tau+1} > \lambda_{1\tau}$.

Combining these two inequalities yields the conclusion that $c^{dg}_{\pi\tau+1} \leq \lambda_{1\tau+1}$. Repeating the preceding line of reasoning infinitely many times implies that a subsequence of $\left\{c^{dg}_{\pi t}\right\}$ moves monotonically away from $\frac{\alpha\beta}{\alpha+\kappa^2}$, so that $\lim_{t\to\infty} c^{dg}_{\pi t}$, if it exists, is definitely smaller than $\frac{\alpha\beta}{\alpha+\kappa^2}$, contradicting Proposition 2. This completes the proof. □

Finally, we prove Proposition 5. First of all, we briefly describe some results of the stochastic approximation[42] that we will exploit in the proof.

Let us consider a stochastic recursive algorithm of the form:

$$\theta_t = \theta_{t-1} + \gamma_t Q\left(t, \theta_{t-1}, X_t\right) , \tag{37}$$

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

$$\frac{d\theta}{d\tau} = h\left(\theta(\tau)\right) , \tag{38}$$

where

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

E represents the expectations taken over the invariant limiting distribution of X_t, for any fixed θ. In particular, it can be shown that the set of limiting points of (37) is given by the stable resting points of the ODE (38).

[42]Ljung (1977), Benveniste, Métivier, and Priouret (1990) provide a recent survey.

Proof of Proposition 5. Note that our equation (27) is a special case of (37), where the technical conditions are easily shown to be satisfied; moreover, it is also easy to see that:

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

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

C Comparison with EH rule

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

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

where the second inequality can be rewritten as

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

Rearranging the terms, we get:

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

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

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

which implies, because:

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

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

References

Adam, K., and R. Billi (2007): "Discretionary Monetary Policy and the Zero Lower Bound on Nominal Interest Rates," *Journal of Monetary Economics*, 54(3), 728–752.

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: Chicago University Press.

Benati, L. (2008): "Investigating Inflation Persistence Across Monetary Regimes," *The Quarterly Journal of Economics*, 123(3), 1005–1060.

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., and G. W. Evans (2006): "A simple recursive forecasting model," *Economics Letters*, 91(2), 158–166.

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.

Bullard, J., and K. Mitra (2002): "Learning about Monetary Policy Rules," *Journal of Monetary Economics*, 49(6), 1105–1129.

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.

Cogley, T., and A. M. Sbordone (2005): "A Search for a Structural Phillips Curve," Staff Report 203, Federal Reserve Bank of New York.

Erceg, C. J., and A. T. Levin (2003): "Imperfect Credibility and Inflation Persistence," *Journal of Monetary Economics*, 50(4), 915–944.

Evans, G. W., and S. Honkapohja (2001): *Learning and Expectations in Macroeconomics*. Princeton: Princeton University Press.

——— (2003a): "Adaptive Learning and Monetary Policy Design," *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," *The Scandinavian Journal of Economics*, 108(1), 15–38.

Ferrero, G. (2007): "Monetary Policy, Learning and the Speed of Convergence," *Journal of Economic Dynamics and Control*, 31(9), 3006–3041.

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 (2006): "Adaptive Learning, Persistence, and Optimal Monetary Policy," *Journal of the European Economic Association*, 4(2-3), 376–385.

Hansen, L. P., and T. J. Sargent (2007): *Robustness*. Princeton: Princeton University Press.

Knight, F. H. (1921): *Risk, Uncertainty and Profit*. Boston, MA: Houghton Mifflin Company.

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*. New York: 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. (2006): "A Bayesian DSGE Model with Infinite-Horizon Learning: Do "Mechanical" Sources of Persistence Become Superfluous?," *International Journal of Central Banking*, 2(3), 87–106.

——— (2007): "Expectations, Learning and Macroeconomic Persistence," *Journal of Monetary Economics*, 54(7), 2065–2082.

ORPHANIDES, A., AND J. C. WILLIAMS (2005a): "The Decline of Activist Stabilization Policy: Natural Rate Misperceptions, Learning, and Expectations," *Journal of Economic Dynamics and Control*, 29(11), 1927–1950.

——— (2005b): "Imperfect Knowledge, Inflation Expectations and Monetary Policy," in *The Inflation-Targeting Debate*, ed. by B. Bernanke, and M. Woodford. Chicago: Chicago University Press.

——— (2005c): "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.

ROTEMBERG, J. J., AND M. WOODFORD (1997): "An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy," in *NBER Macroeconomics Annual 12*, ed. by B. Bernanke, and J. J. Rotemberg. Cambridge (MA): MIT Press.

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

SLOBODYAN, S., AND R. WOUTERS (2009): "Learning in an Estimated Medium-Scale DSGE Model," CERGE-EI Working Papers wp396, The Center for Economic Research and Graduate Education - Economic Institute, Prague.

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.

——— (2007): "Interpreting Inflation Persistence: Comments on the Conference on "Quantitative Evidence on Price Determination"," *Journal of Money, Credit and Banking*, 39(Supplement 1), 203–210.

——— (2009): "Robustly Optimal Monetary Policy under Near-Rational Expectations," *American Economic Review*, forthcoming.

Yun, T. (1996): "Nominal price rigidity, money supply endogeneity, and business cycles," *Journal of Monetary Economics*, 37(2-3), 345–370.

D Figures

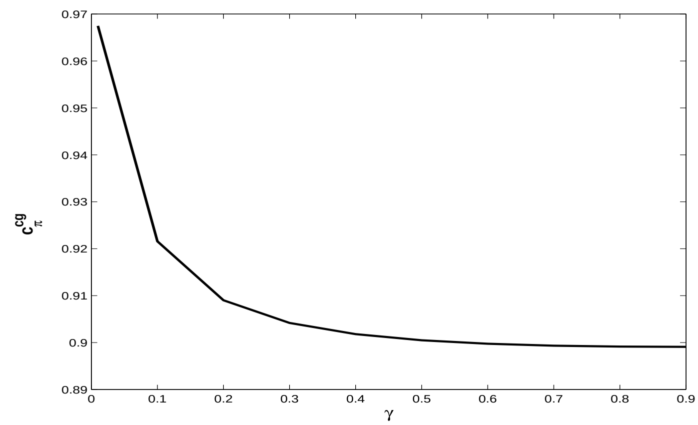

Figure 1: Feedback parameter in the ALM for inflation as a function of γ

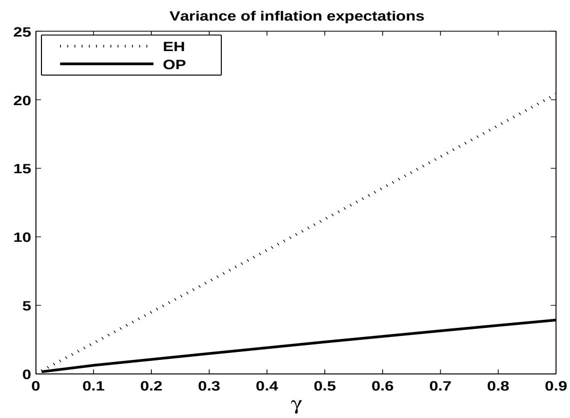

Figure 2: Variance of inflation expectations

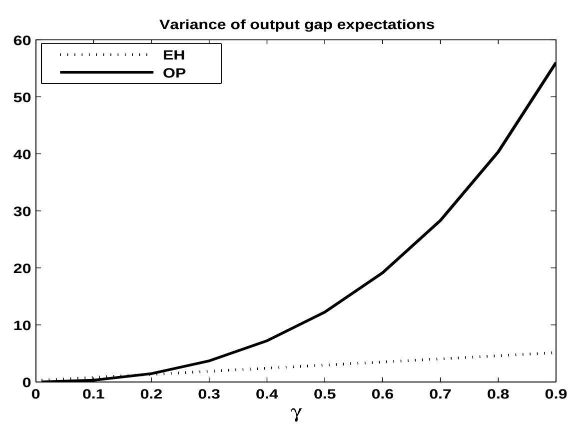

Figure 3: Variance of output gap expectations

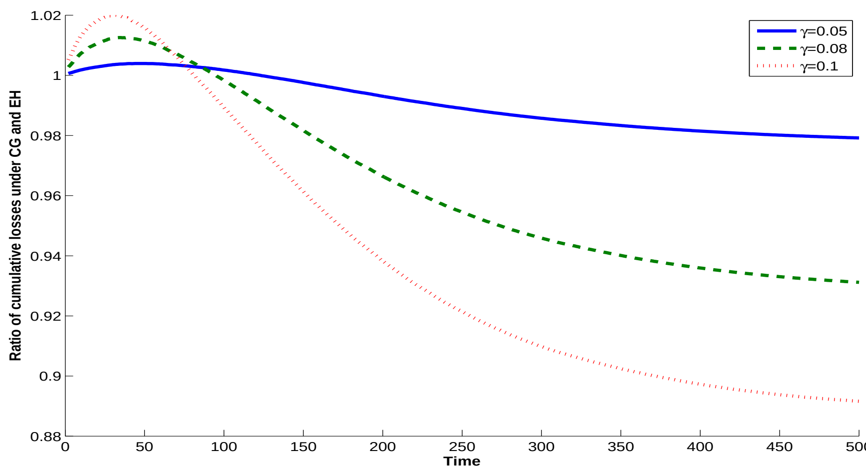

Figure 4: The dynamics of welfare losses (ratio of cumulative consumption equivalents under CG and EH) under constant gain learning, with initial beliefs at the rational equilibrium

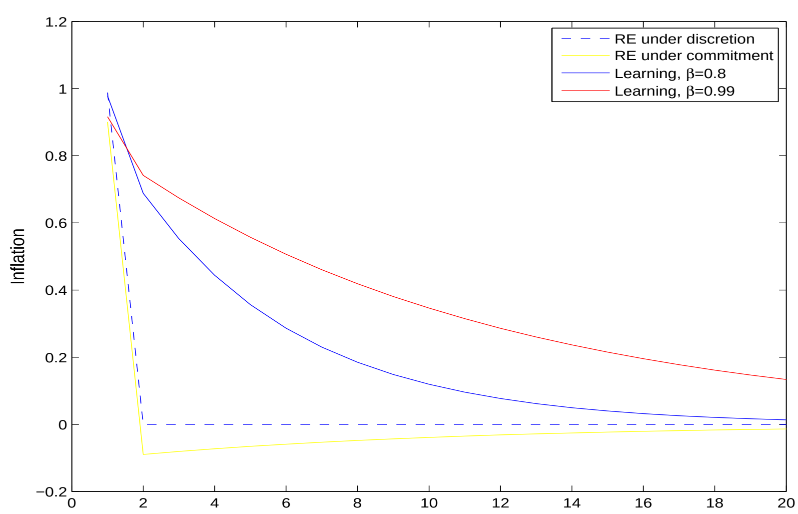

Figure 5: Impulse response of inflation for an initial cost-push shock $u = 1$. Solid line: optimal policy under learning and private agents following learning with $\gamma = 0.9$. Dashed line: optimal discretionary policy under RE with private agents having rational expectations. Initial conditions: $a_0 = 0$, $\pi_0 = 0$, $x_0 = 0$.

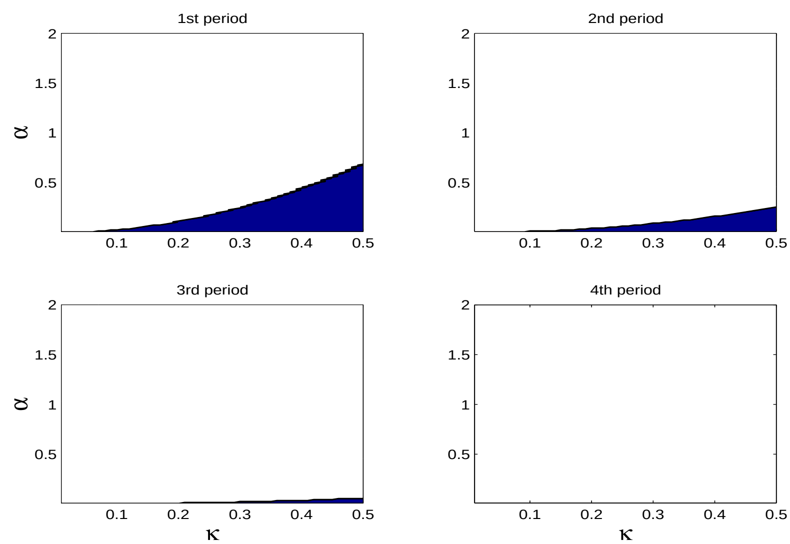

Figure 6: Values of α and κ for which δ_π^{dg} is increasing in the first four periods. From the fourth period on δ_π^{dg} is always decreasing. ($\beta = 0.99$)

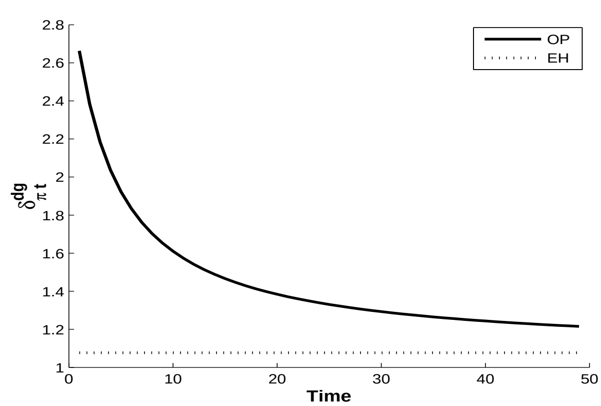

Figure 7: Interest rate rule coefficient on inflation expectations under decreasing gain learning

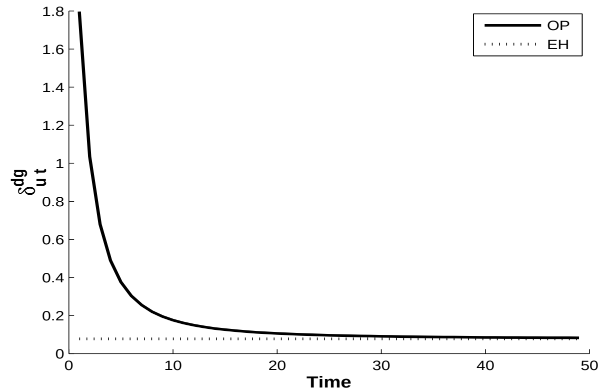

Figure 8: Interest rate rule coefficient on the cost-push shock under decreasing gain

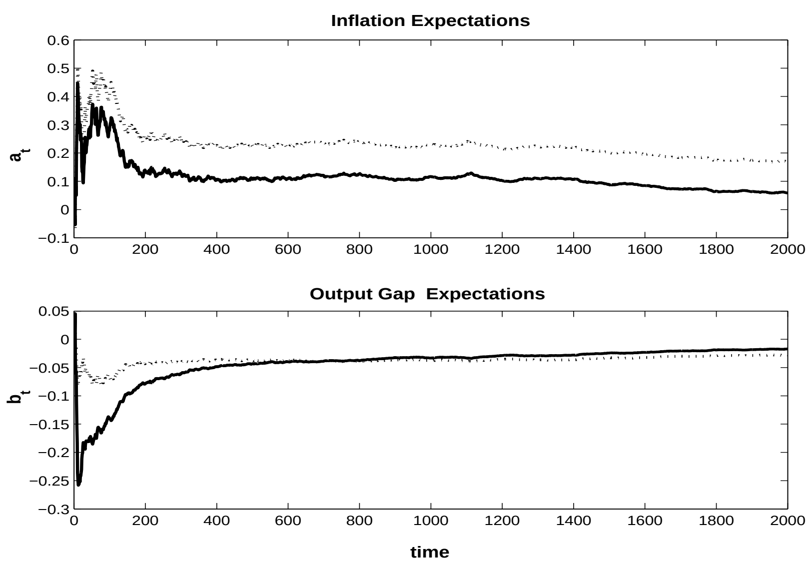

Figure 9: Evolution inflation and output gap expectations under the optimal (solid line) and the EH rule (dashed line), when private agents follow decreasing gain learning