Essays on Monetary Policy and Learning

Krisztina Molnár
DEPARTMENT OF ECONOMICS
UNIVERSITAT POMPEU FABRA
Doctoral Thesis

Supervisor: Albert Marcet

Barcelona, April 2006

Dipòsit legal: B.34201-2006
ISBN: 84-690-0324-0

Contents

List of Tables

List of Figures

Acknowledgements

Doing my Ph.D. in Barcelona was a beautiful period, I would like to thank many people for this.

My first thank goes to my supervisor Albert Marcet. A large part of my thesis can be attributed to what I learned from him. He taught me much of what I know about academic work. Every single time he looked at my work or we talked was very useful - even listening to his questions on seminars. By the way, I always waited for his questions on the seminars.

I am also grateful to Ramon Marimon, who is like my second supervisor. He introduced me to the theory of learning and inspired me and my coauthor Sergio by directing us towards interesting research. Not less importantly, I would like to thank him for the projects he inspired me not to do. I have been saved lots of wasted hours.

I am especially grateful to both Albert and Ramon for keeping me enthusiastic about my work and helping me through the more difficult periods with teaching me to distinguish between what is important and what is not.

I need not say how much I learned from Kosuke Aoki, Jordi Galí who taught me monetary economics and whose door was always open for me. I had very fruitful conversations about research with Gábor Lugosi and Thijs Van Rens. Fabrice Collard, Christian Haefke, Rosemarie Nagel, Francesc Ortega all gave me valuable advice at some stage of my thesis.

I would also like to thank people who had great influence on me before my PhD.; Max Gillman and János Vincze. They were the first to teach me modern macroeconomics, and who made me want to accomplish a Ph.D. .

I would like to thank my coauthor and great friend Sergio Santoro; we think the same about what is really interesting and what is not and I realized this is very rare. I have met many other exceptional people whom I thank for discussions about economics and other things: Árpád Ábrahám, Anton Nakov, Chiara Forlatti, Norbert Maier, Ádám Szeidl, and Lutz Weinke.

I cannot give enough thanks for my parents to keep me together even when I had to face things that hurt me. Finally this last line is for my boyfriend, András, a big kiss for everything.

Foreword

Recently a number of papers argue for modelling expectations as some type of adaptive learning algorithm. This provides foundations for rational expectations models, since under certain conditions adaptive agents can eventually learn a rational expectations equilibrium; this way one can limit focus on the equilibria which are learnable under some type of learning rule. My thesis builds on the results of the least squares learning literature, which models individual agents as econometricians: agents are running least squares regressions using available data in order to form their expectations. Of course there are many ways to depart from rationality, and this makes the model hardly falsifiable. Therefore in the first chapter of my thesis I examine whether least squares learning is a reasonable departure from rationality: I show that the presence of learners in an economy can be rationalized even in coexistence with rational agents. In this chapter I also build on the literature "forecasting with expert advice" to model how agents dynamically select between forecasting algorithms. In the second chapter, I argue that learning is not only learning about coefficients and not only a useful tool for equilibrium selection but have important implications for policy design: optimal monetary policy under learning introduces new features of policy behavior that are not present under rational expectations.

Chapter 1
Learning with Expert Advice

Surveys of inflation forecasts show that expectations combine forward-looking and backward-looking elements. This contradicts "conventional wisdom": in the presence of rational agents adaptive agents would be driven out of the market. In the first chapter of the thesis, I rationalize this finding in an equilibrium framework. My model has two types of agents: one having rational expectations and the other using adaptive learning. The proportion of these agents in the population evolves according to their past forecasting performance. I show that (1) Even an underparametrized learning algorithm survives competition with rational expectations; (2) The presence of rational agents speeds up the convergence of the learning algorithm. This finding is particularly interesting since it is often argued that learning has such slow convergence that doubt is cast as to the validity of the final limit. In sum, these findings strengthen the case for using learning models enriched with rational agents to model expectations.

Chapter 2
Optimal Monetary Policy when Agents are Learning

joint with Sergio Santoro

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

Chapter 1

Learning with Expert Advice

The importance of forward-looking behavior in economic decision making has long been recognized in economics. However the modelling of expectations remains a matter of controversy. Rational expectations is criticized for placing unreasonable computational and informational demands on economic agents. Moreover, a vast empirical literature on testing survey data rejects rational expectations[1] and economic models with rational expectations often perform very badly when tested on data[2]. These criticisms highlight the importance of seeking alternative ways to model expectations.

However, once we depart from fully rational expectations, there are many ways to do so. Thus, agents should be allowed to endogenously choose between them, and abandon their ad-hoc expectation rule, if they could do better. *"payoffs of optimizers ...exceed those of non-optimizers...Normally, one would expect this to exert some pressure, however slight, on the population composition...non-optimizers would be driven to extinction in the long run"* [3]. Indeed, early literature on rational expectations motivated rational expectations saying that if agents did not behave rationally they would disappear from the market.

Least squares learning departs from rationality in a way that still attributes a lot of rationality to agents: they "behave as working economists or econometricians" [4] [5]. Still, the choice of a learning algorithm is necessar-

[1]See for example, Lovell (1986) , Baghestani (1992) Ball and Croushore (1995)

[2]See Roberts (2001)

[3]Sethi and Franke (1995) pp.584

[4]Sargent (1993) p.22.

[5]For a survey of adaptive learning and boundedly rational modelling see Marimon

ily arbitrary, and subject to the caveat raised above: agents would abandon their learning rule when they have access to a better forecasting algorithm.

Another problem with least squares learning is that although it can converge to the rational expectations equilibrium slow convergence can significantly affect finite sample behavior of the economy. In this sense slow convergence might even cast doubt on the validity of the final limit.

In our paper we address these criticisms of least squares learning. We raise the question: "What happens if agents follow least squares learning but they have access to the forecast of an 'expert' who can actually do better?" This expert will have rational expectations, in the sense that he will know the structure of the underlying economy and also the algorithm used by the learners. Both least squares and rational forecasts will be costless. We establish a sort of 'forecasting competition' between the least squares learner and the rational agent in a self referential model, and examine how agents switch between the forecasting rules, whether the learner survives and whether the economy converges to the rational expectations equilibrium or a different equilibrium. We are also interested in whether the presence of a well informed rational agent alters convergence properties of learning. In particular, how conditions of stability change and whether the presence of an expert can "help" to increase the rate of convergence.

Our paper uses dynamic predictor selection, the ratio of learners and rational agents in the population is not imposed exogenously, but depends on past performance. Whoever made the better forecasts in the past will have a higher weight in the population. An important feature of this weighting is that heterogeneity can be an equilibrium outcome.

We propose a simple weighting algorithm that is analytically tractable for least squares learning; thus it gives a clear intuition about convergence results. In the limit it measures how the relative forecasting success of one predictor over the other is appreciated. In the limiting case when agents pay attention to even tiny differences in forecasting performance, weights correspond to the probability of one predictor having better forecasts than the other. In Adam (2004) agents also use the forecast error to choose between predictors. A major difference is that in Adam (2004) agents switch to either one or the other predictor, while in our paper agents use a combination of the forecasting models, adjusting their weights only partially. This is appropriate in a stochastic environment as: if today one predictor was better,

(1996).

there is a positive probability that the other predictor will produce a smaller forecast error next period. A recent survey Branch (2004) provides evidence for a similar expectation formation: agents adjust their weight on predictors dynamically as a function of the relative mean squared forecasting errors.

Our weighting algorithm is linked to a recent approach in learning theory, known as "prediction with expert advice". This literature considers problems of sequential decision making under uncertainty. Prediction with expert advice avoids making assumptions about the data generating process. Instead, the decision maker competes with a wide set of forecasting experts, and the goal is to design decision algorithms that are almost as good as the best expert in a wide benchmark class. In other words, the goal of the predictor is not to minimize his loss function, but to minimize his relative loss (or regret) compared to the best "expert" in a fixed, possibly infinite, set of experts. The main focus is to provide general upper bounds on the cumulative regret [6] [7].

In addition to the weighting algorithm, our paper differs in two main respects from previous papers with endogenous weighting. First, we use different types of predictors and second, we assume all of them are costless.

We consider two types of agents: a rational agent and a least squares learner. In Brock and Hommes (1997) and Sethi and Franke (1995) there are costly rational and costless naive predictors. Results of Brock and Hommes (1997) are generalized by Branch (2002) by expanding the set of predictors with simple adaptive agents. Thus, our least squares learners exhibit more rationality than naive agents and simple adaptive agents in these papers. Adam (2004) and Branch and Evans (2004) are examples for learning with endogenous weighting. In Adam (2004) agents choose between two learning algorithms. Branch and Evans (2004) build a heterogeneous learning model where learners are constrained to underparametrize.

Also, in contrast to Brock and Hommes (1997), Sethi and Franke (1995) and Branch (2002) our agents do not have to pay for the forecasts of the

[6] In the terminology of this literature experts in our paper would be the two forecasting algorithms (forecasting with least squares regression and rational expectations). To avoid confusion we will call only the rational agent to be an expert. (In the terminology of this literature predictors in our paper are "*simulatable experts*": functions that use data accessible to the forecaster himself, thus the forecaster can simulate the experts' future reactions.)

[7] This idea of forecasting with the use of expert advice is applied in several branches of economic theory (for a summary see Cesa-Bianchi and Lugosi (2006)). In game theory, the concept of correlated equilibria takes experts to be pure strategies; in finance, portfolio choice models regard experts as different portfolio strategies.

rational agent. In Sethi and Franke (1995) costly optimization of the rational agent is needed to guarantee survival of naive agents, with costless rational forecasts naive agents are driven out of the economy. In Brock and Hommes (1997) coexistence may prevail under zero optimization cost when agents do not fully optimize [8]. We believe that assuming zero cost for the rational predictor is realistic. For example the central bank inflation forecasts are typically publicly available and the closest to rationality.

One main result of our paper is that least squares learners can survive forecasting competition even with a costless rational predictor, and even if agents pay attention to any infinitesimal difference in the forecasting performance of the learning algorithm and the rational agents. This means that with a positive probability least squares forecasts are closer than the the rational expert's forecasts to the actual outcome, even in the limit. In other words the weights on least squares learning will not collapse to zero, in equilibrium both learners and rational agents are present in the population. Interestingly, this conclusion holds even if the learning algorithm is underparametrized. This result rationalizes empirical work on survey expectations which suggest that in reality expectations combine backward- and forward-looking elements.

A second main result is that the presence of a rational agent increases the rate of convergence of least squares learning. This is important since in practice slow convergence can mean no convergence at all: for example, when there are regime shifts or changes in policy and it takes several years for agents to learn the new equilibrium. We believe our speed of convergence results also support modelling adaptive learning together with rational expectations, since the slow convergence problem can be avoided. [9]

We close the introduction by discussing some related empirical and theoretical research that supports the relevance of our setup.

Recent work on survey inflation expectations detect a behavior similar to the equilibrium of our model. For instance Roberts (1998), Baak (1999) and Chavas (2000) detect an intermediate degree of rationality: inflation expectations can be well approximated by a weighted average of forward-looking and backward-looking expectations. Moreover, modelling expectations as a weighted average of adaptive and rational expectations is supported by an-

[8]In the terminology of this literature: agents choose between predictors with a finite intensity of choice

[9]For research on speed of convergence of least squares learning see Marcet and Sargent (1995), of Bayesian learning see Vives (1993).

other strand of empirical literature. These papers show that the empirical performance of standard models improves when: instead of rational expectations they are modelled as a mixture of backward-looking and forward-looking expectations (see Roberts (2001) or with survey expectations (see Roberts (1997)) [10].

Our paper supports modelling expectations in line with the empirical evidence -as a mixture of backward- and forward looking expectations- with showing that the coexistence of these expectations can be the limiting outcome[11].

Section 1 presents the model. Section 2 establishes properties of the equilibrium and conditions of convergence to the equilibrium. Section 3 examines speed of convergence. We establish speed of convergence results to least squares learning, and derive the conditions under which the presence of a well informed expert can speed up convergence. Then, section 4 provides some numerical results on finite sample speed of convergence with Monte Carlo simulations.

1.1 The Model

The starting point of the analysis is a simple self-referential model. We keeping the model simple and focus on the expectations side.

Let the endogenous variable p be the price level. Assume it is determined by the price level expected for the next period and by the nominal money stock m. The money stock is exogenous and follows an AR(1) process.

$$p_t = \lambda \widetilde{E}_t p_{t+1} + m_t \tag{1.1}$$

$$m_t = \varrho m_{t-1} + \varepsilon_t \qquad \varrho \in [0,1),\ \varepsilon \sim iid\ N(0,\sigma_\varepsilon^2) \tag{1.2}$$

[10] Survey evidence is subject to the caveat that survey respondents may not have incentives to provide accurate information. So survey expectations are at best a noisy measure of inflation expectations and at worst tell us nothing about actual inflation expectations. However it may boost confidence in the usefulness of survey expectations that they are helpful in modelling inflation and in predicting wages and interest rates (Roberts (1995),Roberts (1997), Englander and Stone (1989)).

[11] Modelling expectations as a mixture of adaptive and rational agents was argued for example, by Sethi and Franke (1995), Roberts (1997), Caroll (2003). Modelling expectations as a mixture of learning and rational agents was argued for example, by Nunes (2004).

Equation (1.1) and (1.2) together with the assumptions about the expectation term completely determine the price level. There are several economic models that fit this model: for example, the cobweb model or a simple version of the Cagan model of inflation (Cagan (1956))[12].

We assume that there are two types of agents: one following least squares learning (henceforth LS) and one having rational expectations (henceforth RE). The rational agent is rational not in the sense of the traditional rational expectations equilibrium (this would provide very bad forecasts) but he is rational within the learning equilibrium: he knows the structure of the economy and forms expectations conditioning on the forecasts of learners. The expectation of the least squares learner we denote by E^{LS} and the expectation of the rational agent by E^{RE}. The aggregate expectation about the price level of next period $\widetilde{E}_t p_{t+1}$ is formed as an average of the forecasts of the two types of agents:

$$\widetilde{E}_t p_{t+1} = \omega_t E_t^{LS} p_{t+1} + (1 - \omega_t) E_t^{RE} p_{t+1} \qquad \omega_t \in [0,1] \tag{1.3}$$

Predictor weights - ω - evolve over time, depending on the past forecasting performance of the two types of agents. The representative agent evaluates the forecasting performance of the two competing predictors - LS and RE - and dynamically adjusts the weights that he applies to each predictor. He does not have any sophisticated model at hand. He simply observes the forecasts of two predictors, and assigns weights according to how well they have forecasted previously. For example, agents could simply read the official inflation forecasts of the central bank and of the ministry of finance, and then decide which one to believe more. If the ministry of finance continuously underestimated inflation compared to good forecasts from the central bank, agents will believe the ministry's forecasts to be less reliable, and assign a lower weight to its forecasts.

The model can be also interpreted as a heterogeneous agents model. Then the weights are population weights, and the model has an evolutionary interpretation: the type with more successful forecasts will be more the dominant type. An example for this could be a population of firms choosing between

[12] The basic model of asset pricing under risk neutrality takes the same form, with p_t interpreted as the price of stock, m_t as its dividend. $\lambda = \frac{1}{1+r}$ is the one period discount factor, r the rate of return on the riskless asset.

two pricing algorithms. If one pricing algorithm performs worse than the other, more and more firms would switch to the better one. To avoid confusion, throughout the paper we will resort mainly to the representative agent interpretation.

The representative agent adjusts the predictor weights in every period. The weight on LS evolves in a recursive fashion: weight in period t equals the weight in the previous period, adjusted by a measure of the forecasting performance of LS in $t-1$. Adjusting the weights only partially is suitable in a stochastic environment: when even if today one predictor was closer to the true outcome, with a positive probability it will be the worse one next period. The forecasting performance is measured by the function F, which compares the forecast error of LS and RE; and the time t weight on LS is adjusted towards the value of F taken in time $t-1$.

$$\omega_t = \omega_{t-1} + \frac{1}{t}\left\{F[\,(E^{RE}_{t-1}p_t - p_t)^2 - (E^{LS}_{t-1}p_t - p_t)^2\,] - w_{t-1}\right\}, \qquad (1.4)$$
$$\text{where } F:\mathbb{R}\rightarrow[0,1], for x \leq y F(x) \leq F(y)$$
$$\text{Given} E_0^{LS}p_1, E_0^{RE}p_1, \omega_0 = 1.$$

A key feature of F that the "expert" literature imposes is *symmetry* around 0, formally

$$F(x) = 1 - F(-x)$$

First, let us first consider an example of F as an indicator function which takes the value 1 whenever LS has a smaller or equal forecast error, 0 when RE has a smaller forecast error. Then F at time $t-1$ simply indicates whether LS was better or not, and ω measures how many times LS forecasted better than RE in the past. In the limit ω has a very intuitive interpretation, ω converges to the probability that LS has smaller forecast error than RE.

For some theorems we need continuity of F, which together with the symmetry condition implies $F(0) = 0.5$.

As $F(\cdot)$ is *monotone*, a better LS forecast implies a bigger F. When the LS agent is better $F(\cdot)$ is higher then 0.5, and the weight on LS is adjusted towards a higher value then the weight on RE. Similarly, when RE is better $F(\cdot)$ is lower than 0.5, and the weight on LS is adjusted towards a number lower than 0.5. When RE and LS make equally good forecasts $F(\cdot)$ takes the value 0.5 and their weight is adjusted towards 0.5.

When F is the indicator function, any infinitesimal difference between forecast errors is rewarded. Also, any small difference is rewarded in the same way as bigger differences in forecast errors. By choosing another functional form for F we can also give a measure to how the representative agent evaluates the relative forecasting success of LS. Let us consider an example: F_1 : CDF of $\mathcal{N}(0, \sigma_1^2)$ and F_2 : CDF of $\mathcal{N}(0, \sigma_2^2)$, $\sigma_1 < \sigma_2$. Then, whenever LS makes a better forecast F_1 gives a higher value than F_2, so F_1 adjusts the weight of LS to a greater degree than F_2. A different example could be $F(x) = \frac{1}{\pi}\arctan(\alpha x) + \frac{1}{2}$, $\alpha > 0$. Then the larger α is, the more a good forecasting performance is rewarded in F. In the limiting case $F(\cdot)$ is an indicator function [13] taking the value 1 whenever LS is better, 0.5 when LS and RE made equal forecasts and 0 whenever the rational agent is better; in this case any infinitesimal difference in forecasting performance is considered important.

Similar dynamic expectation formation was found by Branch (2004) in the Michigan Survey of inflationary expectations. He found evidence that agents switch predictor use as the relative mean squared errors change: agents' predictor choices respond negatively to increases in relative mean square error.

The dynamic predictor selection algorithm developed in this paper has the advantage that the joint dynamics of the learning algorithm and the weight can be analytically examined with stochastic approximation, which makes the intuition clearer. The first papers in dynamic predictor selection used either a multinomial logit [14] (see for example, Brock and Hommes (1997)) or an algorithm which closely resembles the replicator dynamic in evolutionary game theory (see Sethi and Franke (1995)); in a learning environment none of these algorithms allow using stochastic approximation [15]. Learning algorithms with dynamic predictor selection were examined before in Branch and Evans (2004) who used multinomial logit law of motion, and the competing forecasting algorithms were misspecified models. They show that intrinsic heterogeneity is possible, where in equilibrium agents use several underparametrized forecasting rules.

[13] The normal density function in the limit of zero variance is the Dirac delta function; the cumulative distribution function of it is the unit step function. The value at 0 is often set to 0.5 by convention. Considering the function $F(x) = \frac{1}{\pi}\arctan(\alpha x) + \frac{1}{2}$, $\alpha > 0$, with $\alpha \to \infty$, it is easy to show that for $x > 0$ the limit is 1, for $x < 0$ the limit is 0 and for $x = 0$ the limit is 0.5 (with the convention $\lim_{\alpha \to \infty} 0\ \alpha = 0$).

[14] In a random utility model, under certain assumptions on shocks, the limiting probability that a given individual chooses an alternative is given exactly by multinomial logit.

[15] In these papers competing predictors were a rational and a naive agent.

Our modelling choice for the relative predictor success is the most recent relative forecast error. In the above mentioned papers the underlying economy (1.1)-(1.2) is the cobweb model and they use last period relative profit [16] as a fitness measure. In Brock and Hommes (1997) this yields similar results to choosing the relative forecast error. Similarly our results remain unchanged with relative profit [17].

For the analytical proofs we choose $F(\cdot)$ to be the CDF of the normal distribution with 0 mean and σ^2 variance, similar results obtain with other examples of $F(\cdot)$ mentioned in this section.

1.1.1 Least Squares Learning

This section describes the learning algorithms used to model the first type of agents.

Note that under *rational expectations the MSV solution* of (1.1) (1.2) is $\mathrm{E}_t p_{t+1} = \frac{\varrho}{1-\lambda\rho} m_t$. So a learning algorithm has a chance to learn the MSV rational expectations equilibrium only if it conditions its expectations on m. We consider two different learning algorithms. The first, which is less "clever", uses only past price levels to formulate the forecast. The second, which is more sophisticated takes into account that the price level should depend on m, and an other, .

Denote the less sophisticated algorithm by LS_1. He observes only past price levels, thus the best he can do is to run regression on a constant, i.e. take averages of past price levels.

$$LS_1 \qquad E_t p_{t+1}^{LS_1} = \alpha_t \quad \alpha_t = \frac{\sum_{i=1}^{t-1} p_i}{t-1} \tag{1.5}$$

Or in a recursive formulation

$$\alpha_t = \alpha_{t-1} + \frac{1}{t}(p_{t-1} - \alpha_{t-1}). \tag{1.6}$$

This learner can not converge to the rational expectations equilibrium, since he is running an underparametrized regression (not conditioning m).

[16]Except Branch and Evans (2004) who use unconditional expected relative profit. This choice allows them to consider the fixed point of a map rather than a solution to a difference equation as in Brock and Hommes (1997).

[17]When (1.1) is developed from profit maximizing firms with quadratic cost function.

Still, he can do quite well, he can learn the true unconditional expectation of the price level.

The second learning algorithm is 'more clever', he also observes the money supply m and runs his regression in the Minimum State Variable (MSV) form (henceforth LS_2). So LS_2 correctly hypothesizes that the last period's money supply is a leading variable of the price level today. He runs a regression of price on lagged money supply, and then makes his forecast of next period's price level with his latest estimated coefficient and the current period's money supply [18].

$$LS_2 \qquad E_t p_{t+1}^{LS_2} = \beta_t m_t \quad \beta_t = \frac{\sum_{i=1}^{t-1} p_i m_{i-1}}{\sum_{i=1}^{t-1} m_{i-1}^2} \tag{1.7}$$

The recursive formulation of the regression coefficient is

$$\beta_t = \beta_{t-1} + \frac{1}{t-1}\frac{1}{R_{t-1}} m_{t-2}\left(p_{t-1} - m_{t-2}\beta_{t-1}\right) \tag{1.8a}$$

$$R_t = R_{t-1} + \frac{1}{t}\left(m_{t-1}^2 - R_{t-1}\right) \tag{1.8b}$$

where R_t is the moment matrix[19].

1.1.2 Rational Expert

Our next modelling choice concerns the second type of agents the rational experts. Again, we consider two specifications, one more sophisticated than the other. Rational expectations in this context does not mean forecasting the rational expectations equilibrium. In the presence of learners the traditional REE forecast would perform very badly. First, because the the economy might converge to a different equilibrium. Second, because even if the economy converges to the REE these forecasts would be wrong during the transition. We assume rational expectations are rational within the learning equilibrium: rationals know the model with its parameters, the stochastic process of m, and condition their expectations on the LS forecasts. The difference between the more and less clever rational agent is that the first

[18] Equilibrium remains unchanged if a constant is also included in the regression.

[19] $R_t = \frac{\sum_{i=1}^{t} m_i m_i'}{t}$

observes the evolution of the weights while the second does not. So, the less clever rational has a misperception about the economy.

We assume the less clever expert, RE_1 hypothesizes that the whole population follows least squares learning, i.e. $\omega_t = 1\ \forall t$. He would be fully rational only if the weight on least squares learners turned out to be 1 in every period. In the simulations we indeed set the initial weight on LS to 1, thus at the beginning RE_1 is rational and remains close to rationality in the first periods when the ratio of LS agents is still high. RE_1 can be interpreted as an agent who investigated a lot in discovering the true parameters of the economy but considers himself too small to influence expectations. One could hypothesize for example that a forecasting agency must have the capacity to investigate thoroughly the underlying economy, but does not believe or does not know exactly how his forecasts are influencing aggregate expectations, i.e. to what extent agents "believe" his forecast.

Expert RE_1 forms his expectations in the following way. Using $\omega_t = 1$ and $E_t^{LS} p_{t+1}$ in (1.1) he can calculate $p_t = \lambda E_t^{LS} p_{t+1} + m_t$. He knows m follows an AR(1) process (equation (1.2)), and forms his forecast as:

$$E_t^{RE_1} p_{t+1} = \lambda E_t^{LS} p_{t+1} + \varrho m_t \tag{1.9}$$

Notice that for simplicity we assume the rational agent is myopic, in the sense that he assumes that the expectation of LS he observes today remains unchanged in the next period: $E_t^{LS} p_{t+1} = E_{t+1}^{LS} p_{t+2}$.

When the least squares learner runs a regression on a constant $E_t^{LS_1} p_{t+1} = \alpha_t$ the forecast of RE_1 is

$$E_t^{RE_1} p_{t+1} = \lambda \alpha_t + \varrho m_t \tag{1.10}$$

When the least squares learners regress p on m $E_t^{LS_2} p_{t+1} = \beta_t m_t$ the forecast of RE_1 is

$$E_t p_{t+1}^{RE_1} = E_t[(\lambda \beta_t + 1) m_{t+1}] = (\lambda \beta_t + 1) \rho m_t \tag{1.11}$$

Next consider the more clever expert. RE_2 knows exactly how his forecasts are followed by agents: he observes the weights denoted to him and the least squares learner. Think for example, of a central bank, who does not have full credibility, but it has to gain credibility by publishing its forecasts, and making good forecasts. In this interpretation the weight on the rational agent will be higher when private agents give a bigger credibility to its forecasts. Furthermore, this central bank can investigate the 'credibility' of his forecasts (observes what is the weight of its forecast in private sectors expectations) and uses this knowledge it to improve his future forecasts. Again, we assume he is myopic and thinks expectation of LS and weights remain unchanged in the next period, mathematically $E_t^{LS} p_{t+1} = E_{t+1}^{LS} p_{t+2}$ and $\omega_t = \omega_{t+1}$. One can find his forecasting function with guess and verify [20]. In general the guess should be $p_t = a(\omega_t) E_t^{LS} p_{t+1} + b(\omega_t) m_t$. When the LS expectations are in the form $E_t^{LS} p_{t+1} = \beta_t m_t$ it is equivalent to guess $p_t = c(\omega_t) m_t$.

With $E_t^{LS_1} p_{t+1} = \alpha_t$ it can be easily verified that $a = \frac{\lambda \omega_t}{1-\lambda(1-\omega_t)}$ and $b = \frac{1}{1-\lambda\varrho(1-\omega_t)}$. Thus forecast of RE_2, using $E_t m_{t+1} = \varrho m_t$ is:

$$E_t^{RE_2} p_{t+1} = \frac{\lambda \omega_t}{1-\lambda(1-\omega_t)} \alpha_t + \frac{1}{1-\lambda\varrho(1-\omega_t)} \varrho m_t \tag{1.12}$$

With $E_t^{LS_1} p_{t+1} = \beta_t m_t$ it can be similarly verified that $c(w_t) = \frac{1+\lambda\omega_t\beta_t}{1-\lambda\varrho\ (1-\omega_t)}$ [21]. The forecast of RE_2, using $E_t m_{t+1} = \varrho m_t$ is:

$$E_t^{RE_2} p_{t+1} \quad = \quad \frac{1+\lambda\omega_t\beta_t}{1-\lambda\varrho\ (1-\omega_t)} \varrho m_t \tag{1.13}$$

Summing up, our model consists of the underlying economy, equations (1.1) and (1.2) and the aggregate inflationary expectation, equation (1.3). This latter in turn is based on the forecasts of two types of predictors: least squares learners and rational agents. Predictor weights evolve according to their past performance, equation (1.4). We examine convergence properties conditional on how sophisticated these forecasts are: whether the rational agent observes the weight or not, and the least squares learner has a correctly specified regression (1.8) or an underparametrized regression (2.8). We examine whether the introduction of a well informed expert implies different model behavior compared to the standard least squares learning case. Does it converge to a different equilibrium? What are the equilibrium weights on

[20] This method is used, with the guess of Nunes (2004), where he models constant gain learning with rational expectations in a New Keynesian framework, with fixed weights.

[21] Or equivalently $p_t = \frac{\lambda\omega_t}{1-\lambda\varrho\ (1-\omega_t)} E_t^{LS} p_{t+1} + \frac{1}{1-\lambda\varrho\ (1-\omega_t)} m_t$

least squares learners and rational agents? Is its rate of convergence different from that of LS learning?

1.2 Equilibrium under Least Squares Learning with Rational Experts

In this section we examine convergence of our economy. As a benchmark we compare convergence of LS_1 and LS_2. After adding rational agents there are several interesting questions to be examined. Establishing convergence results for the weights answers whether least squares learning 'survives' in the presence of rational agents. Further, we examine whether the presence of rational agents modifies the equilibrium and the conditions of convergence. We also examine how different specifications for the rational agents affect our results.

1.2.1 Benchmark: Convergence Under Least Squares Learning

As a benchmark we now study the economy without rational agent ($\omega_t = 1 \forall t$). The two learning algorithms can potentially learn very different equilibria depending on their specification. LS_1 is not conditioning on the state variable m, so the equilibrium he can achieve is a restricted perceptions equilibrium [22]. Since LS_1 is only taking averages of past inflation , the best he can achieve is to find the true unconditional expectation of the price level. LS_2 on the other hand is running a regression in the MSV form, therefore has a 'chance' to learn the MSV rational expectations equilibrium.

The model with only LS_2 learners is examined by several papers, for a proof of convergence we refer to these[23]. It can be shown that β_t converges to the rational expectations solution $\beta_f = \frac{\varrho}{1-\lambda\varrho}$ (henceforth β^{RE}), given $\lambda\varrho < 1$. Further, the equilibrium is the MSV rational expectations equilibrium

$$p_t = \frac{\varrho}{1 - \lambda\varrho} m_t. \tag{1.14}$$

[22] In a restricted perceptions equilibrium expectations are optimal within a restricted class misspecified beliefs. See Evans and Honkapohja (2001a).

[23] See for example, Marcet and Sargent (1989).

When the economy is populated only with LS_1 agents aggregate expectations are given by their Perceived Law of Motion (henceforth PLM) $E_t^{LS_1} p_{t+1} = \alpha_t$. Using this expectation in (1.1) gives the Actual Law of Motion (henceforth ALM) for the economy

$$p_t = \lambda \alpha_t + m_t \tag{1.15}$$

Heuristically an equilibrium α is one where agents perceptions about the future price level become justified by the actual outcome; i.e. the possible resting point of α is the fixed point of a mapping from the PLM to the ALM. In the appendix we show that α_t converges to $\alpha_f = 0$ if $\lambda < 1$. Thus, the equilibrium is $p_t = m_t$, which is different from the equilibrium with LS_2 agents.

The interesting thing in this simple exercise is to see how the equilibrium is intrinsically different depending on the learning algorithm agents follow.

1.2.2 Least Squares Learning in the MSV form -LS_2- in the Presence of a Rational Expert

We now introduce a rational agent and examine how the economy evolves with endogenous weighting: what are the equilibrium weights, how the presence of an expert affects the equilibrium and the conditions for convergence. Let us first examine the case when LS learning is specified in the MSV form.

Let us first consider a population with LS_2 and RE_1 agents. Recall that this expert knows that initially all agents follow LS_2, knows the underlying economy, so at the beginning he is indeed able to calculate rational expectations. However, he does not observe further evolution of the weights and mistakenly thinks in all periods the whole population consists of learners. It follows from the the evolution of the weights that there is an interesting feedback from RE_1's forecast to his forecasting performance. If RE_1 has more accurate forecasts he will be more credible, his weight will increase. On the other hand, a higher weight on him (and a smaller weight on LS) will cause his forecasts to deteriorate as he thinks weight on LS is 1 in all periods. His equilibrium weight will be a fixed point where these two opposite effects cancel each other.

To find the equilibrium we have to determine the Actual Law of Motion of the economy. For this one has to substitute the PLM of the learner (1.7) and the rational agent (1.11), into aggregate expectations (1.3), then aggregate expectations to the underlying economy (1.1). It can be easily shown that for the economy with LS_2 and RE_1 The Actual Law of Motion for the price level is

$$p_t = [\lambda (\omega_t \beta_t + (1 - \omega_t)(\lambda \beta_t + 1)\varrho) + 1]\, m_t \tag{1.16}$$

The ALM for the population with LS_2 and RE_2 is identical to the perception of RE_2:

$$p_t = \frac{1 + \lambda \omega_t \beta_t}{1 - \lambda \varrho\, (1 - \omega_t)} m_t. \tag{1.17}$$

This is not surprising, since RE_2 knows everything to calculate the current price level, in this sense he is fully rational[24].

In equilibrium agents perceptions should become true, i.e. the equilibrium is a fixed point of a mapping from the PLM to the ALM. Since agents are forecasting one period ahead price level we have to compare the ALM of p_t to the PLM $E^{LS}_{t-1} p_t$. [25].

Proposition 1. *Let the economy* (1.1)-(1.4) *be populated with*
i.) *LS_2 and RE_1*
ii.) *or LS_2 and RE_2*

then the vector $[\beta_t, \omega_t]$ converges to $[\beta^{RE}, \frac{1}{2}]$, where $\beta^{RE} = \frac{\varrho}{1-\lambda\varrho}$ is the β corresponding to the rational expectations equilibrium if
i.) $\lambda \in (-\frac{2}{\varrho}, \frac{1}{\varrho})$
ii.) $\lambda < \frac{1}{\varrho}$.

Proof. See the Appendix. □

The above propositions show that even though LS_2 has much less information than a rational agent, he will not die out as in the limit $\omega_t \rightarrow \frac{1}{2}$. In

[24] To forecast p_{t+1}, RE_2 could even do better in the following sense. Since RE_2 observes the weights he can calculate p_t and the current forecasting error of LS. Then using (1.8) in principle he could calculate how β will be updated next period. However, this would lead to a complicated solution for the rational expectation. Since solving this is not in the focus of the paper, we apply the simplifying assumption that rational agents are myopic.

[25] For details see the Appendix.

equilibrium half of the population will be learners and the other half will be rationals. The intuition behind is that in the limit the expectations of the learner and the expert are negligibly close to each other. As the shocks are symmetric around the mean, in the limit the learner and the expert have the same probability of producing the better forecast [26]. This way shocks guarantee that in some periods learners in other periods rationals get closer to the actual outcome.

The above propositions also state a striking difference between an economy with RE_1 and RE_2: compared to the benchmark with only LS_2 agents, RE_2 does not alter the condition for convergence while RE_1 does. There is a range of coefficient values $\lambda < -\frac{2}{\varrho}$ for which LS_2 converges without an expert, but in the presence of RE_1 the LS_2 does not converge. An example for an economy where λ can be negative is the cobweb model[27]. In the cobweb model, with naive expectations, supply lags behind the price, and this introduces periodic fluctuations and might also introduce instability. In our model the source of instability is different: the presence of a rational agent who places too high weight on learners forecast ($\omega_t = 1 \forall t$ in our case) introduces a non-linearity, the rational agent perceives that the price level will depend negatively on the LS forecast, but the actual price will depend negatively also on his forecast. When λ is very low, the aggregate effect of LS forecasts on the price level might turn out to be positive and explosive. In other words the presence of a rational agent who does not have correct perceptions might introduce instability.

1.2.3 Least Squares Learning About a Constant -LS_1- in the Presence of a Rational Expert

Now let us turn to LS_1 when least squares learners are running an underparametrized regression, with PLM $E_t^{LS_1} p_{t+1} = \alpha_t$. Previously we showed that for $\lambda < 1$ learning converges to $\alpha_f = 0$: the true unconditional expectation of the price under rational expectations. This section, as the previous one, considers convergence in the presence of a rational expert.

With expert RE_1, the Actual Law of Motion of the economy is obtained by substituting the forecast of the learner (1.5) and the forecast of RE_1 (1.10)

[26] See proof of proposition 2 footnote 42

[27] The cobweb model was identified by the Hungarian economist, Nicholas (Miklos) Kaldor.

into aggregate expectations (1.3), and (1.3) into the price level equation (1.1). The Actual Law of Motion with expert RE_1 is:

$$p_t = \left(\lambda\omega_t + \lambda^2 (1-\omega_t)\right)\alpha_t + (\lambda\varrho(1-\omega_t)+1)\, m_t \tag{1.18}$$

The ALM with the second expert is obtained similarly by using (1.5), (1.12), (1.3) and (1.1). The ALM in this case is equal to what RE_2 perceives about p:

$$p_t = \frac{\lambda\omega_t}{1-\lambda(1-\omega_t)}\alpha_t + \frac{1}{1-\lambda\varrho(1-\omega_t)} m_t \tag{1.19}$$

Proposition 2. *Let the economy* (1.1)-(1.4) *be populated with*
i.) *LS_1 and RE_1*
ii.) *or LS_1 and RE_2.*
Then $[\alpha_t, \omega_t]$ converges to a fixed point $[0, \omega_f]$

i.) *In the first case ω_f is the solution of*

$$\omega = \iint F(\ m^2\varrho^2 - 2m\varrho(m\varrho+\varepsilon)[\lambda\varrho(1-\omega)+1]\ )\ d\phi_m\ d\phi_\varepsilon$$

and the fixed point is stable if the following condition holds:

$$\lambda \in (\frac{-1}{1-\omega_f}, 1)\ \text{and}\ \frac{2|\lambda|\varrho^2}{\sqrt{2\pi}\sigma}\frac{\sigma_\varepsilon^2}{1-\varrho^2}\left(\varrho + \frac{2\sqrt{1-\varrho^2}}{\pi}\right) < 1$$

ii.) *In the second case ω_f is the solution of the following equation*

$$\omega = \iint F\left(-\frac{\varrho^2 m^2 + 2\varrho m\varepsilon}{(1-\lambda(1-\omega))^2}\right) \mathrm{d}\phi_m \mathrm{d}\phi_\varepsilon\ .$$

and the fixed point is stable if

$$\lambda < 1\ \text{and}\ \frac{2|\lambda|}{(1-\lambda(1-\omega_f))^3}\frac{\sigma_\varepsilon^2}{\sqrt{2\pi}\sigma(1-\varrho^2)}\left(\varrho^2 + \frac{2\sqrt{1-\varrho^2}}{\pi}\right) < 1.$$

Where F is the c.d.f. of $N(0,\sigma)$, ϕ_m, ϕ_ε are distribution functions of m and ε.

When $\varrho = 0$ the equilibrium weight both in case i.) and ii.) is $\omega_f = \iint F(0)\, \mathrm{d}\phi_m \mathrm{d}\phi_\varepsilon = 0.5$.

Proof. See the Appendix. □

Propositions 2 states analytical upper bounds for conditions of stability, since closed form solutions cannot be established[28]. With Monte carlo simulation (Tables 1.4 and 1.5)we show that these upper bounds are indeed always satisfied[29], thus conditions for stability simplify to:

$$\textbf{i.)} \qquad \lambda \in (\frac{-1}{1-\omega_f}, 1) \tag{1.20}$$

$$\textbf{ii.)} \qquad \lambda < 1 \tag{1.21}$$

The most surprising result of Proposition 2 is that the weight on LS_1 is strictly positive even in the limit. In other words even an underparametrized learning algorithm survives forecasting competition with a rational agent [30]. Of course LS_1 is a worse predictor then a rational agent (when $\varrho \neq 0$), but he is still quite clever in the sense that he eventually learns the true unconditional expectation of p under rational expectations. Thus in the limit there is a positive probability that LS_1 makes a forecast that is closer to the actual outcome than the rational agent's forecast.

When $\varrho = 0$, m is a random noise, so conditioning on m does not help forecasting more then just taking averages of p: the weights of LS_1 will be equal to the weights of the rational expert, $\omega_f = 0.5$. Figure 1.1 shows that when $\varrho \neq 0$ the underparametrized learning algorithm naturally performs worse than a rational agent: the equilibrium weight on LS_1 is never bigger than the equilibrium weight on the expert.

An interesting result of Proposition 2 is that the equilibrium weights depend on the persistence of autoregressive process [31]. Figure 1.1 shows that the higher is ϱ the smaller is the weight on LS_1. The intuition behind this

[28]This is necessary since in both cases one eigenvalue of the Jacobian does not have a closed form solution.

[29]We examined $0 \leq \varrho \leq 0.9$, $-2 < \lambda < 1$. Tables 1.4 and 1.5 summarize results $0 < \lambda < 1$.

[30]This result holds true also for $\sigma \to \infty$, so if the representative agent pays attention to any small difference in forecasting performance.

[31]For LS_1 finding the exact value of equilibrium weight cannot be derived analytically. Solving for ω_f is a complicated fixed point problem of a function inside the normal cumulative distribution function. (See Appendix.) Since the latter does not have a closed form solution we cannot give exact analytical results. However it is easy to get numerical solutions.

is simple: the more persistent is m, the longer time it takes for a shock on m to die out, so the LS_1 makes a bigger mistake by not using data on m to forecast. So, a more persistent stochastic process for m makes the forecasts of LS_1 worse compared to RE who conditions on m.

In the limit LS_1 learns the same restricted perceptions equilibrium as without the presence of a rational agent: $\alpha_f = 0$. A result similar to the previous section is that a rational agent with a misspecified model might introduce instability. The presence of RE_1 again decreases the parameter set for which convergence to the equilibrium applies. RE_2 on the other hand does not affect stability conditions: convergence of LS_1 with or without RE_2 has the same condition.

However, the equilibrium changes; the equilibrium depends on the specification of the rational expert, and on the equilibrium weight.

The equilibrium under LS_1 and RE_1 is: $p_t = (\lambda\varrho(1-\omega_f)+1)m_t$, thus the equilibrium will be different from the MSV solution.

Interestingly, the equilibrium can get close to the MSV rational expectation solution with a correctly specified expert. With LS_1 and RE_2 the equilibrium is: $p_t = \frac{1}{1-\lambda\varrho(1-\omega_f)}m_t$.When ϱ is sufficiently high, so that ω_f is close to zero, the equilibrium will be close to the rational equilibrium. In other words, when the persistence of m is very high, LS_1 forecasters will perform so badly that experts will dominate the equilibrium.

In this section we have shown that with dynamic predictor selection, when the weights on predictors are adjusted depending on their last forecasting performance, learning survives even in the presence of a costless rational agent. Surprisingly, this is true even if the learning algorithm is underparametrized. This happens because the underparametrized learning algorithm learns the true unconditional mean of inflation; in the limit the price level varies around its mean, thus the forecasts of the learner will be better than the rational forecast with a positive probability.

We have also shown that the equilibrium depends intrinsically on the type of agents in the economy. When the learning algorithm nests the MSV rational expectations equilibrium, the economy indeed converges to this equilibrium. When the learning algorithm is underparametrized, the equilibrium will be different. In this case, the forecasting performance of the learning algorithm will depend on the stochastic characteristics of the variable that is missing from its regression; its forecasting performance will in turn determine its population weight, and thus the equilibrium.

Stability of learning might change in the presence of a rational agent. When a well informed rational agent is present it does not alter convergence properties. However when the rational agent has a misperception he may induce instability.

1.3 Speed of Convergence

In the previous section we have shown that the presence of a rational agent does not alter to where the beliefs of learners converge (even though the equilibrium itself might change), as long as the stability conditions are satisfied. Still, there can be differences in how quickly learning does converge. This section will examine whether an 'expert advice' can speed up convergence of least squares learning.

Learning is often criticized for its slow convergence, which can cause problems for economic decision making on several grounds. Slow convergence implies that the asymptotic distribution for test statistics can be very different when agents follow LS learning compared to when agents have rational expectations. When speed of convergence to rational expectations is very slow, the confidence intervals will be larger than the confidence intervals from classical econometrics. This means that an econometrician who derives confidence intervals assuming agents are rational will reject the null hypothesis too often if agents actually follow LS learning. Also, when learning converges slowly, expectations will be out of the rational equilibrium for a considerable time. Then making decisions based on rational expectations when expectations actually follow learning would be erroneous, even if we know that in the long run learning converges to the rational expectations equilibrium and learning had been present for a long time.

In this section first we derive analytical speed of convergence results, then we provide finite sample rate of convergence results by Monte Carlo simulations. Throughout this section we assume convergence is not an issue: we set $\lambda \in (0, 1)$ thus conditions for convergence are always met.

1.3.1 Analytical speed of convergence results

We examine speed of convergence applying the theorem of Benveniste and P. (1990) (theorem 3, page 110).

Let θ_t be the vector of parameter estimates, X_t the state vector, and γ_t the deterministic sequence of gains. The function $\mathcal{Q}$ expresses the way in which the estimates of θ are updated from period $t-1$ to t. In our case $\gamma_t = \frac{1}{t}$. With LS_1 $\theta_t = [a_t, \omega_t, R_t]$ and with LS_2 $\theta_t = [\beta_t, \omega_t, R_t]$. The state vector is $X_t = [m_{t-2}, \varepsilon_{t-1}]$. $\mathcal{Q}$ is the updating term in the recursive formulation in equations θ (2.8),(1.8),(1.4) respectively.

From the recursive formulation

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

define

$$h(\theta) = E[\mathcal{Q}(\theta, X_t)]$$

for fixed θ. Let θ_f be such that $h(\theta_f) = 0$. The theorem of Benveniste et al. concludes that if the Jacobian of $h(\theta)$ evaluated at θ_f has all eigenvalues less than $-\frac{1}{2}$ in real part then

$$t^{0.5}(\theta_t - \theta_f) \quad \xrightarrow{\mathcal{D}} \quad N(0, P)$$

where the matrix P satisfies

$$\left[\frac{1}{2}h_\theta(\theta_f)\right] P + P\left[\frac{1}{2}h_\theta(\theta_f)\right]' + E\mathcal{Q}(\theta, X_t)\mathcal{Q}(\theta, X_t)' = 0$$

Thus if the above conditions are met we have root t convergence to θ_f. Moreover, for higher eigenvalues of the Jacobian convergence is slower, in the sense that the asymptotic variance-covariance is higher. So even when there is root-t convergence, higher eigenvalues of $h_\theta(\theta_f)$ imply slower convergence.

Benchmark: if there is only least squares learning

To understand better the theorem of Benveniste et al. and also as a point of comparison, let us first examine speed of convergence under LS_2. Substituting the ALM of inflation (1.14) into (1.8), the recursive formulation for $\theta = [\beta \; R]$, it is easy to show:

$$h\begin{pmatrix} \beta \\ R \end{pmatrix} = \begin{bmatrix} \frac{1}{R}M_m[(\lambda\beta + 1)\varrho - \beta] \\ M_m - R \end{bmatrix}$$

Where $E[m_{t-2}m_{t-2}] = \frac{\sigma_\varepsilon^2}{1-\varrho^2} = M_m$. The Jacobian of $h(\theta)$ is

$$\frac{\partial h(\theta)}{\partial \theta} = \begin{bmatrix} \frac{1}{R} M_m(\lambda\varrho - 1) & \frac{-1}{R^2} M_m[(\lambda\beta + 1)\varrho - \beta] \\ 0 & -1 \end{bmatrix}$$

We have to evaluate this at the fixed point $\theta_f = [\beta^{RE}\ M_m]$. The two eigenvalues are: -1 and $\lambda\varrho - 1$. For $\sqrt{t}$ convergence these have to be smaller than $-\frac{1}{2}$, so under LS_2 we need $\lambda\varrho < \frac{1}{2}$. [32]

Similar calculations for LS_1 yield: for $\sqrt{t}$ convergence under LS_1 we need $\lambda < \frac{1}{2}$.

One immediate difference between LS_1 and LS_2 is that the persistence of the autoregressive process influences speed of convergence in the latter but not in the former. When least squares learning does not condition on m the stochastic properties of m will have no effect on how quickly learning gets to the equilibrium. On the other hand when least squares learning conditions on m the more persistent is m the less information can be gained from variations in m (the PLM will be closer to the ALM) and the slower learning will converge to the equilibrium.

Speed of convergence of learning with an 'expert'

We now derive analytical speed of convergence results of least squares learners in the presence of a rational expert.

Learning in the MSV form:LS_2
In Section 1.3.1 we derived that for root-t convergence of LS_2 $\lambda\varrho < \frac{1}{2}$ has to hold. We now examine how this condition changes when the learner has an 'expert' advice at hand. Analytical results are derived in Proposition 3 and 4 in the Appendix. Here we summarize the conditions for speed of convergence in Table 1.1.

Table 1.1: Conditions for $\sqrt{t}$ convergence - learning in the MSV form

LS_2	LS_2 with RE_1	LS_2 with RE_1
$\lambda\varrho < \frac{1}{2}$	$\frac{1}{2}\lambda\varrho + \frac{1}{2}\lambda^2\varrho^2 < \frac{1}{2}$	$\lambda\varrho < \frac{2}{3}$

Results show that in the presence of an expert root-t convergence applies for a wider range of coefficients. With RE_1, least squares learning in the

[32]For further discussion see Marcet and Sargent (1995).

MSV form has $\sqrt{t}$ convergence if $\frac{1}{2}\lambda\varrho + \frac{1}{2}\lambda^2\varrho^2 < \frac{1}{2}$, while in the presence of RE_2 if $\lambda\varrho < \frac{2}{3}$. For $\varrho = 0.7$ for example, LS_2 has $\sqrt{t}$ convergence for $\lambda < 0.71$, while in the presence of RE_1 upper bound for λ increases to 0.88 and the presence of RE_2 increases it even further to 0.95.

Learning about a constant:LS_1
In the Appendix Propositions 5 and 6 derive sufficient analytical conditions for $\sqrt{t}$ convergence of LS_1 in the presence of a rational agent. Since one eigenvalue does not have closed form solution, we use Monte Carlo simulations - see Table 1.4 and 1.5 - and show that this eigenvalue is indeed always smaller than $-\frac{1}{2}$, thus a necessary and sufficient condition for $\sqrt{t}$ convergence of LS_1 is that the other eigenvalue should be smaller than $-\frac{1}{2}$. Results are summarized in Table 1.2.

Table 1.2: Conditions for $\sqrt{t}$ convergence - learning about a constant

LS_1	LS_1 with RE_1	LS_1 with RE_1
$\lambda < \frac{1}{2}$	$\lambda\omega_f + \lambda^2(1-\omega_f) < \frac{1}{2}$	$\lambda < \frac{1}{1+\omega_f}$

Similarly to the case of LS_2, the presence of an expert increases speed of convergence of LS_1. We can see this using $\omega_f \in [0, 1]$.

In the case of RE_1 $\lambda\omega_f + \lambda^2(1-\omega_f) < \frac{1}{2} hastohold.$ It is easy to see that $\lambda\omega_f + \lambda^2(1-\omega_f)$ is increasing in ω_f and has its maximum at $\frac{1}{2}$

$$\lambda\omega_f + \lambda^2(1-\omega_f) < \frac{1}{2} \; . \tag{1.22}$$

Thus we can conclude that the presence of RE_1 increases speed of convergence of LS_1: in the presence of RE_1 the LS_1 algorithm has square root-t convergence on a wider parameter set.

With RE_2 using $\omega_f \in [0, 1]$ we have

$$\frac{1}{2} < \frac{1}{1+\omega_f} \; ; \tag{1.23}$$

thus LS_1 has root-t convergence on a wider parameter set when RE_2 is also present in the economy. Moreover, using (1.24) and (1.23) we can conclude that

$$\lambda\omega_f + \lambda^2(1-\omega_f) < \frac{1}{1+\omega_f} \; ; \tag{1.24}$$

the cleverer the expert advice is the more it increases the convergence speed of LS_1.

For $\varrho = 0$ there is a closed form solution for the weight $\omega_f = 0.5$ (see Proposition 2), so we also have closed form solutions for speed of convergence. LS_1 in the presence of RE_1 has $\sqrt{t}$ convergence if $\lambda < 0.61$ and in the presence of RE_2 if $\lambda < 0.66$.

It is interesting to examine the role of ϱ in speed of convergence results. A result already known in the literature is that the more persistent the stochastic process is the slower is the convergence of LS learning in the MSV form. Marcet and Sargent (1995) derived analytically this result for the same model with LS_2. Timmermann (1996) found the same result numerically[33]. Our results show that this remains true in the presence of a rational expert.

However, when learning does not condition on the stochastic process m this previous result is reversed: with a rational expert the more persistent is m the quicker learning will converge to the equilibrium. The intuition behind this result is simple. LS_1 does not condition its forecast on m, so when the economy is populated only with LS_1 agents, its speed of convergence does not depend on the autoregressive parameter of m. However, in the presence of a rational agent ϱ affects the speed of convergence through the equilibrium weight. The more persistent the money process is the bigger mistake it is not to condition on it: a higher ϱ leads to a lower equilibrium weight on LS_1 (see Figure 1.1). A lower weight of the learner means that the 'expert' - who in turn conditions on m - has higher influence on the actual law of motion of inflation. Through the forecasts of the 'expert' m has higher and more persistent influence on inflation, and the least squares learner will discover he is 'doing the wrong thing' sooner. In other words his Perceived Law of Motion of the learner will be very different from the Actual Law of Motion and this leads to a higher speed of convergence.

Analytical results of this section show that speed of convergence of least squares learning is faster when learners have access to the forecast of a well

[33]Timmermann (1996) examined a stock pricing equation with learning where the exogenous process was an autoregressive dividend process. By simulations he found that a more persistent dividend process results slower convergence of learning.

informed expert. Moreover convergence is faster when the rational expert has a better information set. Learning in the MSV form has slower speed of convergence the more persistent the autoregressive process m is. This finding is parallel to the findings of Marcet and Sargent (1995) and Timmermann (1996). Our results show that this result remains true in the presence of a rational agent with dynamic predictor selection when learning conditions on the exogenous process, but is reversed when learning does not condition on m. The speed of convergence of an underparametrized learner who does not condition on m does not depend on the persistence of m; while in the presence of a rational agent and with endogenous weights, an underparametrized learning algorithm can even converge faster with a more persistent m.

1.3.2 Simulations

The previous section derived analytical results for speed of convergence, but analytical results are available only if the eigenvalues are smaller than $\frac{1}{2}$. Therefore in this section we show some numerical results on finite sample rate of convergence.

Following the methodology of Marcet and Sargent (1995) we use Monte Carlo simulations. We assume that there is a δ for which

$$t^{\delta}(\beta_t - \beta_f) \overset{\mathcal{D}}{\longrightarrow} F \tag{1.25}$$

for some non-degenerate well-defined distribution F with mean zero and variance σ_F^2. Then $t^{\overline{\delta}}(\beta_t - \beta_f) \to 0$ for $\overline{\delta} < \delta$ and we will call δ the rate of convergence of β_t. Intuitively, the faster β_t is converging to β_f, the higher t^{δ} has to be to ensure that the product on the left hand side of 1.25 converges to a proper distribution.

Equation 1.25 implies that $E[t^{\delta}(\beta_t - \beta_f)]^2 \to \sigma_F^2$ as $t \to \infty$. Therefore

$$\frac{E[t^{\delta}(\beta_t - \beta_f)]^2}{E[(tk)^{\delta}(\beta_{tk} - \beta_f)]^2} \to 1$$

which implies that

$$\frac{(E\beta_t - \beta_f)^2}{E(\beta_{tk} - \beta_f)^2} \to k^{2\delta} \qquad \text{as} \qquad t \to \infty$$

This justifies using

$$\delta = \frac{1}{\log(k)} \log \left[\frac{E(\beta_t - \beta_f)^2}{E(\beta_{tk} - \beta_f)^2)} \right]^{1/2}$$

for large t in finite samples, as an approximation for the rate of convergence.

Given t and k the expectations involved can be approximated by Monte Carlo simulation: calculating a large number of independent realizations of length t and tk and calculating the mean square difference from β_f.

For the simulations we calculated the rates of convergence with 1000 independent realizations. For initial conditions we set the least squares β and R equal to their limiting point $\beta_0 = \beta_f$ (or $a = a_f$), $R_0 = \frac{1}{1-\varrho^2}$. The initial weight on the least squares learner was set equal to 1. Several seeds of the random number generator were tried, rates of convergence were maximum 0.03 from each other for the long sample, and 0.06 in the short sample.

In the simulations we use $\varrho = 0.9$; with this ϱ the cutoff values of λ for root-t convergence are summarized in Table 1.3. For example, when the economy is populated with LS_2 and RE_2 agents and $\varrho = 0.9$, λ has to be smaller than 0.74 for root-t convergence.

Table 1.3: Conditions for $\sqrt{t}$ convergence with $\varrho = 0.9$

	Alone	with RE_1	with RE_2
LS_1	$\lambda < 0.5$	$\lambda < 0.685$	$\lambda < 0.82$
LS_2	$\lambda < 0.55$	$\lambda < 0.61$	$\lambda < 0.74$

Tables 1.7 and 1.6 summarize simulation results.The estimated δs are presented, for $\sqrt{t}$ convergence these have to be higher than 0.5.

The simulations in general coincide with the analytical results. Rate of convergence of least squares learning is higher with expert advice. Also convergence is faster the better is the information set of the rational agent [34]. Simulations also show that a rational agent speeds up convergence more when learning is misspecified compared to when learning is correctly specified.

From the simulation results one can observe that convergence is very slow for high values of λ, and the higher is λ the slower is the convergence. This confirms the theorem of Benveniste and P. (1990), the higher the eigenvalues of $h_\theta(\theta_f)$ the slower is the rate of convergence. Also rate of convergence can

[34]Finite sample rate of convergence of LS_1 is lower than the analytical results suggested.

be slower than $\frac{1}{2}$, when the conditions in the Benveniste et al. theorem do not hold.

Conclusion

Modelling expectations remains a controversy: rational expectations is criticized for assuming too much knowledge about agents; adaptive learning, which is an alternative, is criticized for being ad hoc and for having slow convergence to the rational equilibrium. Our paper addressed these criticisms of learning: we examined what happens when agents are learning but have access to the advice of somebody who has better forecasts. We also examined how the presence of such an expert increases speed of convergence of learning. We developed a self referential model with two types of agents: learners and rationals.

We proposed a new dynamic predictor selection, where weights on the two types evolved dynamically depending on their relative forecasting performance in the previous period. In other words, agents gradually use more and more the better predictive performance of the rational expert when the expert generates better forecasts. One can think about a central bank that does not have any credibility, but can gradually gain credibility by providing good forecasts.

Our paper showed that the coexistence of learners and rationals can be rationalized in an equilibrium framework: if agents forecast with a learning algorithm and have access to forecasts of a rational agent, they will not rush to abandon their ad hoc learning rule, even if rational forecasts are costless. Surprisingly, learning survives forecasting competition with a rational agent even if it is underparametrized. The equilibrium intrinsically depends on the type of agents in the economy. With an underparametrized learning algorithm the economy does not converge to a rational expectations equilibrium. These results coincide with recent surveys of inflation expectations, which find that expectations are well approximated as being a weighted average of forward-looking and backward-looking expectations. We believe our results strengthen the case for modelling expectations as a mixture of adaptive and forward looking expectations. Especially, since it is well documented that modelling expectations this way improves empirical performance of standard models.

Our second main result is: the presence of a rational expert 'helps' learning algorithms to converge faster. The more rationality the expert possesses the higher is the speed of convergence. This result is also confirmed by finite sample simulations. This result might further strengthen the case for using learning models enriched with rational agents to model expectations, because the criticisms against learning's slow convergence can be weakened.

Since learning and rational expectations imply different policy implications, we also find it important for future research to examine policy implications of modelling expectations as a mixture of them.

We have also shown that a rational agent might even induce instability: when rational agents have a misperceived model, and expectations have a negative feedback on the actual outcome, for example as in the cobweb model. In this sense making the forecasts of the rational 'expert' accessible for learners might even worsen the situation. When welfare loss is associated with being out of the steady state (like in the New Keynesian framework) a diverging economy clearly leads to huge welfare losses[35]. A practical example might be an economy with learners and a central bank with rational expectations. If the central bank were to hesitate between making its forecasts public or not, our paper suggests the answer depends on how confident the central bank is about the underlying economy. On the one hand publishing central bank forecasts can help agents learn the equilibrium faster, but when the central bank forecast is based on a misperceived model publishing it it might introduce instability.

[35] Ferrero (2003) shows that in the New Keynesian framework when the central bank target coincides with the steady state, higher speed of convergence is always beneficial.

Tables and Graphs

Table 1.4: **Derivative $\frac{\delta}{\delta\omega}E[F(\cdot)]$ evaluated at the fixed point, LS_1 with RE_1**

λ \ ϱ	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
0.1	0	0.000	0.001	0.004	0.006	0.008	0.012	0.016	0.0172	0.0146
0.235	0	0.001	0.003	0.008	0.014	0.018	0.026	0.033	0.0313	0.0225
0.37	0	0.001	0.005	0.012	0.021	0.027	0.039	0.045	0.0379	0.0244
0.505	0	0.001	0.007	0.016	0.026	0.034	0.049	0.053	0.0397	0.0254
0.64	0	0.001	0.008	0.020	0.032	0.041	0.058	0.058	0.0388	0.0267
0.775	0	0.002	0.010	0.024	0.036	0.047	0.065	0.060	0.0369	0.0283
0.91	0	0.002	0.011	0.027	0.040	0.053	0.071	0.060	0.0351	0.0299
0.955	0	0.002	0.012	0.028	0.041	0.055	0.072	0.060	0.0346	0.0304

$\sigma = \sigma_\varepsilon = \sigma_m = 1$.
For stability derivatives have to be smaller than 1.
For root-t convergence, smaller than 0.5.

Table 1.5: **Derivative $\frac{\delta}{\delta\omega}E[F(\cdot)]$ evaluated at the fixed point, LS_1 with RE_2**

λ \ ϱ	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
0.1	0	0.001	0.003	0.005	0.006	0.005	0.006	0.008	0.009	0.009
0.235	0	0.003	0.007	0.012	0.012	0.010	0.013	0.019	0.020	0.019
0.37	0	0.004	0.012	0.020	0.017	0.013	0.018	0.027	0.026	0.030
0.505	0	0.005	0.019	0.030	0.020	0.013	0.022	0.035	0.029	0.042
0.64	0	0.006	0.031	0.040	0.022	0.007	0.025	0.042	0.030	0.049
0.775	0	0.010	0.050	0.050	0.026	-0.001	0.033	0.045	0.031	0.045
0.91	0	0.025	0.073	0.056	0.032	-0.003	0.049	0.043	0.036	0.025
0.955	0	0.034	0.079	0.055	0.034	0.000	0.057	0.044	0.041	0.014

$\sigma = \sigma_\varepsilon = \sigma_m = 1$.
For stability derivatives have to be smaller than 1.
For root-t convergence, smaller than 0.5.

Table 1.6: Speed of convergence of ls learning in the MSV form (in short and long sample)

	LS_2		With RE_1		With RE_2	
λ	t=200 to 800	t=2000 to 10,000	t=200 to 800	t=2000 to 10,000	t=200 to 800	t=2000 to 10,000
0.1	0.6026	0.5190	0.6024	0.5171	0.6024	0.5171
0.145	0.6050	0.5225	0.6053	0.5191	0.6054	0.5189
0.19	0.6059	0.5272	0.6081	0.5216	0.6084	0.5213
0.235	0.6043	0.5329	0.6105	0.5249	0.6111	0.5242
0.28	0.5989	0.5389	0.6121	0.5290	0.6133	0.5277
0.325	0.5887	0.5436	0.6122	0.5338	0.6146	0.5318
0.37	0.5732	0.5447	0.6102	0.5391	0.6146	0.5361
0.415	0.5523	0.5397	0.6052	0.5443	0.6128	0.5404
0.46	0.5264	0.5272	0.5966	0.5481	0.6089	0.5441
0.505	0.4964	0.5069	0.5836	0.5488	0.6023	0.5463
0.55	0.4633	0.4800	0.5660	0.5442	0.5928	0.5458
0.595	0.4281	0.4481	0.5430	0.5331	0.5798	0.5412
0.64	0.3915	0.4128	0.5150	0.5142	0.5632	0.5304
0.685	0.3541	0.3755	0.4821	0.4880	0.5422	0.5133
0.73	0.3164	0.3369	0.4444	0.4555	0.5176	0.4889
0.775	0.2784	0.2976	0.4031	0.4175	0.4879	0.4571
0.82	0.2406	0.2580	0.3579	0.3745	0.4527	0.4188
0.865	0.2028	0.2182	0.3089	0.3252	0.4130	0.3746
0.91	0.1651	0.1782	0.2581	0.2695	0.3703	0.3255
0.955	0.1276	0.1383	0.2038	0.2133	0.3244	0.2703

$\varrho = 0.9, \sigma = \sigma_\varepsilon = \sigma_m = 1.$

.

1.4 Appendix

1.4.1 Convergence Proofs

Stochastic Approximation

Throughout our convergence proofs we build on technical results of stochastic approximation. So first let us briefly describe this method. Let us consider the following stochastic recursive algorithm (SRA)

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

where θ_t is a vector of parameter estimates, X_t is the state vector, and γ_t is a deterministic sequence of gains. The function $\mathcal{Q}$ expresses the way in which the estimates of θ are updated from period $t-1$ to t.

If $\mathcal{Q}$ and X_t satisfy some technical assumptions (See Evans and Honkapohja (2001b) Chapter 6 [36]) the stochastic approximation approach associates an ordinary differential equation (ODE) with the SRA form

[36]Ljung's Theorems 4 and 2. Ljung (1977)

Table 1.7: Speed of convergence of ls learning about a constant (in short and long sample)

	LS_1		With RE_1		With RE_2	
λ	t=200 to 800	t=2000 to 10,000	t=200 to 800	t=2000 to 10,000	t=200 to 800	t=2000 to 10,000
0.1	0.4661	0.5032	0.4713	0.5052	0.4716	0.5054
0.145	0.4610	0.5013	0.4697	0.5048	0.4704	0.5050
0.19	0.4546	0.4987	0.4676	0.5041	0.4690	0.5046
0.235	0.4467	0.4951	0.4650	0.5033	0.4673	0.5042
0.28	0.4371	0.4900	0.4616	0.5021	0.4653	0.5036
0.325	0.4255	0.4830	0.4573	0.5005	0.4629	0.5029
0.37	0.4117	0.4735	0.4519	0.4983	0.4601	0.5020
0.415	0.3956	0.4607	0.4450	0.4952	0.4567	0.5008
0.46	0.3769	0.4442	0.4363	0.4908	0.4524	0.4992
0.505	0.3558	0.4233	0.4252	0.4844	0.4472	0.4972
0.55	0.3320	0.3980	0.4112	0.4752	0.4407	0.4943
0.595	0.3057	0.3685	0.3937	0.4617	0.4324	0.4902
0.64	0.2769	0.3352	0.3720	0.4426	0.4219	0.4842
0.685	0.2459	0.2986	0.3452	0.4161	0.4081	0.4752
0.73	0.2128	0.2595	0.3127	0.3809	0.3899	0.4613
0.775	0.1778	0.2185	0.2739	0.3364	0.3654	0.4396
0.82	0.1412	0.1761	0.2283	0.2827	0.3320	0.4058
0.865	0.1032	0.1328	0.1759	0.2209	0.2850	0.3532
0.91	0.0639	0.0888	0.1167	0.1523	0.2182	0.2739
0.955	0.0237	0.0444	0.0512	0.0781	0.1202	0.1590

$\varrho = 0.9, \sigma = \sigma_\varepsilon = \sigma_m = 1.$

.

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

where $h(\theta)$ is obtained as

$$h(\theta) = \lim_{t\to\infty} E\mathcal{Q}(\theta, X_t(\theta))$$

for a fixed θ provided this limit exists. τ denotes "notional" or "artificial" time.

According to the results established by stochastic approximation theory, if this ODE has an equilibrium point[37] θ_f which is locally asymptotically stable[38] , then θ_f is a possible point of convergence of the algorithm. If θ_f is not a locally stable equilibrium point of the ODE, then θ_f is not a possible point of convergence of the SRA, i.e. $Pr(\theta \to \theta_f) = 0$.

[37] θ_f is an equilibrium point if $h(\theta_f) = 0$

[38] θ_f is locally stable if all eigenvalues of the derivative matrix (Jacobian) $Dh(\theta_f)$ have negative real parts.

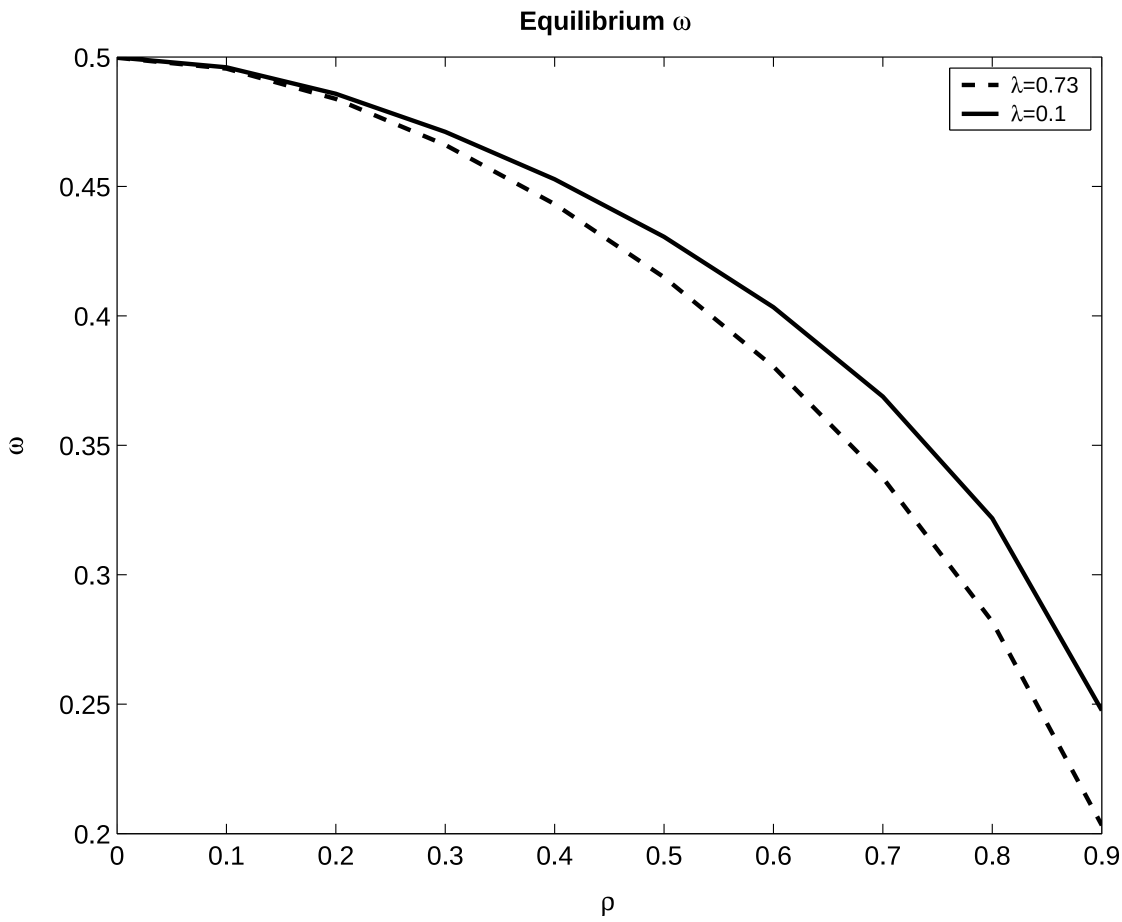

Figure 1.1: Equilibrium ω of LS_1 in the presence of rational agent RE_1

Convergence of α

Claim 1. *Let* α_t *evolve according to* (2.8) *in the model* (1.1)-(1.2). *Then* α_t *converges to 0 given* $\lambda < 1$

Proof. It can be easily verified that $h(\theta)$ is

$$\frac{da}{d\tau} = \lambda a - a \ , \tag{1.26}$$

thus the only possible fixed point is $a = 0$, which is stable if $\lambda < 1$. □

Proof of Proposition 1 part i.)

Let the economy (1.1)-(1.4) *be populated with two types of agents* LS_2 *and* RE_1. $[\beta_t, \omega_t]$ *converge to* $[\beta^{RE}, \frac{1}{2}]$ *if* $\lambda \in (-\frac{2}{\varrho}, \frac{1}{\varrho})$, *where* $\beta^{RE} = \frac{\varrho}{1-\lambda\varrho}$ *is the* β *corresponding to the rational expectations equilibrium.*

Proof. Both β and ω and R evolve over time, their joint dynamical system can be analyzed with the theory of stochastic approximation.

From equations (1.7), (1.11) and (1.16) in (1.4), and (1.8) our model in SRA form is:

$$\begin{aligned}\beta_t &= \beta_{t-1} + \frac{1}{t}\frac{1}{R_{t-1}} m_{t-2}\{[\ \lambda[\omega_{t-1}\beta_{t-1} + (1-\omega_{t-1})(\lambda\beta_{t-1}+1)\varrho] + 1\ ]\varrho m_{t-2} \qquad (1.27)\\ &- \beta_{t-1}m_{t-2} + [\lambda[\omega_{t-1}\beta_{t-1} + (1-\omega_{t-1})(\lambda\beta_{t-1}+1)\varrho]+1]\varepsilon_{t-1}\}\\ R_t &= R_{t-1} + \frac{1}{t}\left(m_{t-1}^2 - R_{t-1}\right) \qquad (1.28)\\ \omega_t &= \omega_{t-1} + \frac{1}{t}\left\{F(\ (E_{t-1}^{RE}p_t - p_t)^2 - (E_{t-1}^{LS}p_t - p_t)^2\ ) - w_{t-1}\right\}\ , \qquad (1.29)\end{aligned}$$

where $F(\cdot)$ denotes the c.d.f. of $N(0,\sigma^2)$ [39]. Using the notation of section 1.4.1: $\theta_{t-1} = [\beta_{t-1}, \omega_{t-1}, R_{t-1}]$. γ_t is $\frac{1}{t}$, $X_t = [m_{t-2}, \varepsilon_{t-1}]$ and $\mathcal{Q}(\theta_{t-1}, X_t)$ are the terms multiplying $\frac{1}{t}$[40][41].

Using the results of stochastic approximation we can consider the asymptotic behavior of the mean of $\mathcal{Q}(\theta, X_t)$ for a fixed θ, which gives a system of associated ordinary differential equations (ODE).

---

[39] For an interpretation of $F(\cdot)$ see page 7. The choice of $F(\cdot)$ does not alter the proof as long as it is differentiable.

[40] In equation (1.27) the original tracking parameter is $\gamma_t = \frac{1}{t-1}$. Rewriting it in the standard SRA form with $\gamma_t = \frac{1}{t}$ results a second-order complementary term which does not affect the associated ODE.

[41] The regularity conditions A1, A.2, A.3 are easy to verify (we use notation of Evans and Honkapohja (2001b)). To verify assumption A.2. we have to restrict R_{t-1} to be bounded away from 0 thus the domain for $\theta_t = (\beta_t, R_t, \omega_t)$ is $(\mathbb{R}, D = (R_L, \infty), [0,1])$ where $R_L \geq \epsilon > 0$. This is a natural restriction since R_t is the sample moment of m. Assumption A.3. is satisfied since $\mathcal{Q}(\theta, x)$ is twice continuously differentiable with bounded second derivatives. Assumptions B.1. and B.2. are satisfied since m is an AR(1) process and ε is white noise.

Claim 2. *The system of ODEs is*

$$\frac{\partial\beta}{\partial\tau} = \frac{1}{R}M_m[C(\theta)\varrho - \beta] \quad (1.30)$$

$$\frac{\partial R}{\partial\tau} = M_m - R \quad (1.31)$$

$$\frac{\partial\omega}{\partial\tau} = \iint F((A(\theta)^2 - B(\theta)^2)m^2 - 2C(\theta)(A(\theta)+B(\theta))m\varepsilon)\ d\phi_m\ d\phi_\varepsilon - \omega \quad (1.32)$$

where M_m is the variance of the limiting distribution of m. ϕ_m and ϕ_ε are the density functions of m and ε respectively. m and ε denote realizations of m_{t-1} and ε_t drawn from these distributions. $A(\theta) = \beta - C(\theta)\varrho$, $B(\theta) = (\lambda\beta+1)\varrho - C(\theta)\varrho$. $C(\theta) = [\lambda[\omega\beta + (1-\omega)(\lambda\beta+1)\varrho] + 1]$.

Proof. Let us fix $\theta = [\beta\ R\ \omega]$ and simplify (1.29) using the notation

$$(E^{RE}_{t-1}p_t - p_t)^2 - (E^{LS}_{t-1}p_t - p_t)^2 = (A(\theta)m_{t-1} - C(\theta)\varepsilon_t)^2 - (B(\theta)m_{t-1} - C(\theta)\varepsilon_t)^2$$
$$= (A(\theta)^2 - B(\theta)^2)m^2_{t-1} - 2C(\theta)(A(\theta)+B(\theta))m_{t-1}\varepsilon_t\ .$$

Then,

$$\lim_{t\to\infty} E\{F(\ (A(\theta)^2 - B(\theta)^2)m^2_{t-1} - 2C(\theta)(A(\theta)+B(\theta))m_{t-1}\varepsilon_t\ )\} = \quad (1.33)$$

$$\iint F((A(\theta)^2 - B(\theta)^2)m^2 - 2C(\theta)(A(\theta)+B(\theta))m\varepsilon)\ d\phi_m\ d\phi_\varepsilon\ . \quad (1.34)$$

This gives us equation (1.32).

To get equation (1.30)-(1.31) fix θ in equations (1.27)-(1.28) and use $E[m_{t-2}m_{t-2}] = \frac{\sigma^2_\varepsilon}{1-\varrho^2} = M_m$ and $E[m_{t-2}\varepsilon_{t-1}] = 0$. □

Equilibrium

Next let us examine the equilibrium points of the system of differential equations (1.30)-(1.32). From equation (1.31) the possible convergence point for $R_t = \frac{\sum_{i=1}^t m_i}{t}^2$ is $R_f = M_m$. Then from equation (1.30) the equilibrium β is the rational expectations equilibrium

$$\beta_f = \frac{\lambda(1-\omega)\varrho^2 + \varrho}{1 - \omega\lambda\varrho - (1-\omega)\lambda^2\varrho^2} = \frac{(\lambda(1-\omega)\varrho + 1)\varrho}{(\lambda(1-\omega)\varrho+1)(1-\lambda\varrho)} = \frac{\varrho}{1-\lambda\varrho}\ . \quad (1.35)$$

Finally, to find ω_f we need to calculate the fixed point of ω using $\beta_f = \frac{\varrho}{1-\lambda\varrho}$, $R_f = M_m$. Let us denote $\{\theta^* = [\beta_f, R_f, \omega] : \omega \in [0,1]\}$. Notice that $B(\theta^*) = A(\theta^*) = 0$. Thus (1.32) at θ^* is

$$\iint F(0)\ d\phi_m\ d\phi_\varepsilon - \omega = \iint \frac{1}{2}\ d\phi_m d\phi_\varepsilon - \omega = \frac{1}{2} - \omega\ .$$

The first equality comes from $F(0) = 0.5$, the second from $\int d\phi_m = \int d\phi_\varepsilon = 1$. Thus the fixed point of ω is $\omega_f = \frac{1}{2}$.

Our result is that the fixed point of (1.30)-(1.32) is[42].

$$\theta_f = [\beta_f, R_f, \omega_f] = \left[\frac{\varrho}{1-\lambda\varrho}, M_m, \frac{1}{2}\right]\ .$$

Stability of the fixed point

It remains to show that θ_f is a locally stable equilibrium point of (1.30)-(1.32). For this we have to find the eigenvalues of the derivative matrix (Jacobian) evaluated at θ_f, $Dh(\theta)|_{\theta_f}$. θ_f is locally stable if all eigenvalues have negative real parts.

$$Dh(\theta) = \begin{bmatrix} \frac{M_m}{R}[\lambda\varrho\omega + \lambda^2\varrho^2(1-\omega) - 1] & \frac{-M_m}{R^2}[\lambda\varrho(\omega\beta + (1-\omega)(\lambda\beta+1)\varrho) + \varrho - \beta] & \frac{M_m}{R}[\lambda\varrho\beta - \lambda\varrho^2(\lambda\beta+1)] \\ 0 & -1 & 0 \\ \frac{\partial \iint F(\cdot)\ d\phi_m\ d\phi_\varepsilon}{\partial\beta} & 0 & \frac{\partial \iint F(\cdot)\ d\phi_m\ d\phi_\varepsilon}{\partial\omega} - 1 \end{bmatrix}$$

Claim 3.

$$\left.\frac{\partial \iint F(\cdot)\ d\phi_m\ d\phi_\varepsilon}{\partial\omega}\right|_{\theta_f} = 0$$

Proof. Using (1.32) [43]

$$\begin{aligned} &\frac{\delta}{\delta\omega}\iint F((A(\theta)^2 - B(\theta)^2)m^2 - 2C(A(\theta)+B(\theta))m\varepsilon)\ d\phi_m\ d\phi_\varepsilon = \\ &\iint \frac{\delta}{\delta\omega}[(A(\theta)^2 - B(\theta)^2)m^2 - 2C(A(\theta)+B(\theta))m\varepsilon] \times \\ &f((A(\theta)^2 - B(\theta)^2)m^2 - 2C(\theta)(A(\theta)+B(\theta))m\varepsilon)\ d\phi_m\ d\phi_\varepsilon\ , \end{aligned} \tag{1.36}$$

where $\frac{\delta C(\theta)}{\delta\omega} = \lambda[\beta - \varrho(\lambda\beta+1)]$, $\frac{\delta A(\theta)}{\delta\omega} = \frac{\delta B}{\delta\omega} = -\varrho\frac{\delta C(\theta)}{\delta\omega}$. Next we evaluate (1.36) at the fixed point, θ_f. Note that $A(\theta_f) = B(\theta_f) = 0$ and $C(\theta_f) = \lambda\beta_f + 1$.

[42] Notice that to obtain this result only the symmetry of the distribution of ε is needed.

[43] We can interchange the integral and the derivative since $F(\cdot)$ is a continuous function $\mathbb{R} \to [0,1]$ (thus measurable).

Since at the fixed point $(\beta_f - (\lambda\beta_f + 1)\varrho) = 0$,(1.36) at θ_f is

$$\iint -2C(\theta_f)(-2\varrho\lambda[\beta_f - (\lambda\beta_f + 1)\varrho])]m\varepsilon\ f(0m - 0m\varepsilon)\ d\phi_m\ d\phi_\varepsilon = 0 \quad (1.37)$$

thus (1.36) at the fixed point is 0. □

Claim 4.

$$\left.\frac{\partial \iint F(\cdot)\mathrm{d}\phi_m \mathrm{d}\phi_\varepsilon}{\partial\beta}\right|_{\theta_f} = 0$$

Proof.

$$\frac{\delta}{\delta\beta}\iint F((A(\theta)^2 - B(\theta)^2)m^2 - 2C(\theta)(A(\theta) + B(\theta))m\varepsilon)\ d\phi_m\ d\phi_\varepsilon = \quad (1.38)$$
$$\iint \frac{\delta}{\delta\beta}[(A(\theta)^2 - B(\theta)^2)m^2 - 2C(\theta)(A(\theta) + B(\theta))m\varepsilon] \times$$
$$f((A(\theta)^2 - B(\theta)^2)m^2 - 2C(\theta)(A(\theta) + B(\theta))m\varepsilon)\ d\phi_m\ d\phi_\varepsilon$$

where $\frac{\delta C(\theta)}{\delta\beta} = \lambda[\omega - (1-\omega)\lambda\varrho]$, $\frac{\delta A(\theta)}{\delta\beta} = 1 - \varrho\frac{\delta C(\theta)}{\delta\beta}$ and $\frac{\delta B(\theta)}{\delta\beta} = \lambda\varrho - \varrho\frac{\delta C(\theta)}{\delta\beta}$.

At the fixed point (1.38) becomes

$$\iint -2C(\theta_f)[1 + \lambda\varrho - 2\varrho\lambda[\omega_f + (1-\omega_f)\lambda\varrho]]m\varepsilon\ f(0)\ d\phi_m\ d\phi_\varepsilon\ . \quad (1.39)$$

Using $\omega_f = 0.5$ and $\beta_f = \frac{\varrho}{1-\lambda\varrho}$ (1.39) can be further simplified to[44]

$$\iint -2(1+\lambda\varrho)m\varepsilon\ f(0)\ d\phi_m\ d\phi_\varepsilon = -2(1+\lambda\varrho)\frac{1}{\sigma\sqrt{2\pi}}\iint m\varepsilon\ d\phi_m\ d\phi_\varepsilon = 0$$

□

From Claim 3 and 4 it follows that the Jacobian at θ_f is

$$Dh(\theta)|_{\theta_f} = \begin{bmatrix} \lambda\varrho\frac{1}{2} + \lambda^2\varrho^2\frac{1}{2} - 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & -1 \end{bmatrix}.$$

θ_f is locally stable if all eigenvalues are negative. Clearly the last two eigenvalues are negative. The first one is negative if $\lambda\varrho\frac{1}{2} + \lambda^2\varrho^2\frac{1}{2} - 1 < 0$. This condition holds if $\lambda \in (-\frac{2}{\varrho}, \frac{1}{\varrho})$, where $\varrho \in [0, 1)$ by assumption. □

[44] Recall that m and ε are independent and $f(0) = \frac{1}{\sigma\sqrt{2\pi}}\ e^{\frac{0}{2\sigma^2}}$.

Proof of Proposition 1 part ii.)

Let the economy (1.1)-(1.4) be populated with two types of agents LS_2 and RE_2. $[\beta_t, \omega_t]$ converge to $[\beta^{RE}, \frac{1}{2}]$ if $\lambda < \frac{1}{\varrho}$, where $\beta^{RE} = \frac{\varrho}{1-\lambda\varrho}$ is the β corresponding to the rational expectations equilibrium.

Proof. The proof goes similarly to proof of Proposition (1) part . *i.)* We have to find the fixed point of $\theta = [\beta \; R \; \omega]$. Forecasts of p_t by the least squares learner and the rational agent are

$$E_{t-1}^{LS_2} p_t = \beta_{t-1} m_{t-1} \quad E_{t-1}^{RE_2} p_t = \frac{1 + \lambda\omega_{t-1}\beta_{t-1}}{1 - \lambda\varrho(1 - \omega_{t-1})} \varrho m_{t-1} \ ,$$

and the ALM of inflation is

$$p_t = \frac{1 + \lambda\omega_t\beta_t}{1 - \lambda\varrho(1 - \omega_t)} (\varrho m_{t-1} + \varepsilon_t) \ .$$

Introducing the notation $C = \frac{1+\lambda\omega\beta}{1-\lambda\varrho(1-\omega)}$

$$(E_{t-1}^{RE} p_t - p_t)^2 - (E_{t-1}^{LS} p_t - p_t)^2 = C(\theta)^2 \varepsilon^2 - ((\beta - \varrho C(\theta)) m_{t-1} - C(\theta)\varepsilon_t)^2 \ .$$

Substituting these into (1.4), (1.8) the associated system of ordinary differential equations, $\frac{d\theta}{d\tau} = h(\theta)$, is:

$$\frac{\partial \beta}{\partial \tau} = \frac{1}{R} M_m [\varrho \frac{1 + \lambda\omega\beta}{1 - \lambda\varrho(1 - \omega)} - \beta] \tag{1.40}$$

$$\frac{\partial R}{\partial \tau} = M_m - R \tag{1.41}$$

$$\frac{\partial \omega}{\partial \tau} = \iint F(2(\beta - \varrho C(\theta)) C(\theta) m\varepsilon - (\beta - \varrho C(\theta))^2 m^2) \, d\phi_m \, d\phi_\varepsilon - \omega \tag{1.42}$$

where M_m is the variance of the limiting distribution of m. ϕ_m and ϕ_ε are the density functions of m and ε respectively. m and ε denote realizations of m_{t-1} and ε_t drawn from these distributions.

Equilibrium

From equation (1.41) the possible convergence point for $R_t = \frac{\sum_{i=1}^{t} m_i}{t}^2$ is $R_f = M_m$. Then from equation (1.40) the equilibrium β is the rational expectations equilibrium $\frac{\varrho}{1-\lambda\varrho}$. Finally, to find ω_f let us denote $\{\theta^* = [\beta_f, R_f, \omega] : \omega \in [0, 1]\}$. Notice that $\beta - \varrho C(\theta^*) = 0$. Thus (1.42) at θ^* is

$$\iint F(0)\ d\phi_m\ d\phi_\varepsilon - \omega = \iint \frac{1}{2}\ d\phi_m d\phi_\varepsilon - \omega = \frac{1}{2} - \omega\ .$$

Thus, fixed point of ω is $\frac{1}{2}$. The fixed point of (1.40)-(1.42) is

$$\theta_f = [\beta_f, R_f, \omega_f] = \left[\frac{\varrho}{1-\lambda\varrho}, M_m, \frac{1}{2}\right]\ .$$

Stability

To obtain stability conditions of θ_f we will use the following two claims.

Claim 5.

$$\left.\frac{\partial \iint F(\cdot)\mathrm{d}\phi_m \mathrm{d}\phi_\varepsilon}{\partial\beta}\right|_{\theta_f} = 0$$

Proof.

$$\frac{\delta}{\delta\beta}\iint F[2(\beta-\varrho C(\theta))C(\theta)m\varepsilon - (\beta-\varrho C(\theta))^2 m^2]\ d\phi_m\ d\phi_\varepsilon = \tag{1.43}$$
$$\iint [2(1-\varrho\frac{\delta C(\theta)}{\delta\beta})C(\theta)m\varepsilon + 2(\beta-\varrho C(\theta))\frac{\delta C(\theta)}{\delta\beta}m\varepsilon - 2(\beta-\varrho C(\theta))(1-\frac{\delta C(\theta)}{\delta\beta})m^2]\times$$
$$f[2(\beta-\varrho C(\theta))C(\theta)m\varepsilon - (\beta-\varrho C(\theta))^2 m^2]\ d\phi_m\ d\phi_\varepsilon,$$

where $\frac{\delta C(\theta)}{\delta\beta} = \frac{\lambda\omega}{1-\lambda\varrho(1-\omega)}$. At the fixed point $\beta_f - \varrho C(\theta_f) = 0$, thus at the fixed point (1.43) becomes

$$\iint [2(1-\varrho\frac{\lambda\frac{1}{2}}{1-\lambda\varrho\frac{1}{2}})\frac{1+\frac{1}{2}\frac{\lambda\varrho}{1-\lambda\varrho}}{1-\frac{1}{2}\lambda\varrho}m\varepsilon f[0]\ d\phi_m\ d\phi_\varepsilon = \tag{1.44}$$
$$4\frac{1}{2-\lambda\varrho}\iint m\varepsilon\ d\phi_m\ d\phi_\varepsilon = 0,$$

where the last equality follows from the independence of m and ε[45]. □

Claim 6.

$$\left.\frac{\partial \iint F(\cdot)\mathrm{d}\phi_m \mathrm{d}\phi_\varepsilon}{\partial\omega}\right|_{\theta_f} = 0$$

[45]Recall that m and ε denote realizations of m_{t-1} and ε_t drawn from the distributions ϕ_m and ϕ_ε.

Proof.

$$\frac{\delta}{\delta\omega}\iint F[2(\beta-\varrho C(\theta))C(\theta)m\varepsilon-(\beta-\varrho C(\theta))^2m^2]\ d\phi_m\ d\phi_\varepsilon = \tag{1.45}$$
$$\iint[-2\varrho\frac{\delta C(\theta)}{\delta\omega})C(\theta)m\varepsilon+2(\beta-\varrho C(\theta))\frac{\delta C(\theta)}{\delta\omega}m\varepsilon-2(\beta-\varrho C(\theta))(-\frac{\delta C(\theta)}{\delta\omega})m^2]\times$$
$$f[2(\beta-\varrho C(\theta))C(\theta)m\varepsilon-(\beta-\varrho C(\theta))^2m^2]\ d\phi_m\ d\phi_\varepsilon,$$

where $\frac{\delta C(\theta)}{\delta\omega}=\frac{\lambda\beta-\lambda^2\varrho\beta-\lambda\varrho}{(1-\lambda\varrho(1-\omega))^2}$. At the fixed point $\beta_f-\varrho C(\theta_f)=0$ and (1.46) becomes

$$\iint[-2\varrho\frac{\frac{\lambda\varrho}{1-\lambda\varrho}-\frac{\lambda^2\varrho^2}{1-\lambda\varrho}-\lambda\varrho}{(1-\frac{1}{2}\lambda\varrho}})^2)C(\theta_f)m\varepsilon]f[0]\ d\phi_m\ d\phi_\varepsilon=\iint 0\ m\varepsilon f[0]\ d\phi_m\ d\phi_\varepsilon=0\ .$$

□

Using claims 5 and 6, the Jacobian of (1.40)-(1.42) evaluated at θ_f is

$$Dh(\theta)|_{\theta_f}=\begin{bmatrix}\frac{\lambda\varrho}{2-\lambda\varrho}-1 & 0 & 0\\ 0 & -1 & 0\\ 0 & 0 & -1\end{bmatrix}.$$

θ_f is a stable fixed point of (1.40)-(1.42) if the eigenvalues of the Jacobian evaluated at θ_f are eigenvalues. Clearly the last two eigenvalues are negative. The first eigenvalue is negative if $\frac{\lambda\varrho}{2-\lambda\varrho}<1$, which can be simplified to $\lambda\varrho<1$[46]. □

Proof of proposition 2 part i.)

Let the economy (1.1)-(1.4) *be populated with two types of agents* LS_1 *and* RE_1*.* $[\alpha_t,\omega_t]$ *converge to a fixed point* $[0,\omega_f]$*, if* $\lambda\in(\frac{-1}{1-\omega_f},1)$ *and* $\frac{2|\lambda|\varrho^2}{\sqrt{2\pi}\sigma}\frac{\sigma_\varepsilon^2}{1-\varrho^2}\left(\varrho+\frac{2\sqrt{1-\varrho^2}}{\pi}\right)<$ 1.

ω_f *is the solution to the equation* $\omega=\iint F(\ m^2\varrho^2-2m\varrho(m\varrho+\varepsilon)[\lambda\varrho(1-\omega)+1]\ )\ d\phi_m\ d\phi_\varepsilon$*. F is the c.d.f. of* $N(0,\sigma)$*,* ϕ_m,ϕ_ε *are distribution functions of m and* ε*.*

Proof. The proof goes similarly to proof of Proposition (1). We have to find the fixed point of $\theta=[a\ \omega]$. Using equation (1.5), (1.10) and (1.18) in (1.4) and (2.8), the associated system of ordinary differential equations, $\frac{d\theta}{d\tau}=h(\theta)$, is:

[46] $\frac{\lambda\varrho}{2-\lambda\varrho}<1$ holds also for $\lambda>\frac{2}{\varrho}$, however in this case Lipshitz conditions do not hold for (1.40) and stochastic approximation cannot be used.

$$\frac{\partial a}{\partial \tau} = (\lambda\omega + \lambda^2(1-\omega))a - a \tag{1.46}$$

$$\frac{\partial \omega}{\partial \tau} = \iint F(\cdot)\, d\phi_m\, d\phi_\varepsilon - \omega\ , \tag{1.47}$$

where $F(\cdot) = F(\ (a[\lambda - \lambda\omega - \lambda^2(1-\omega)] + m\varrho - (m\varrho+\varepsilon)[\lambda\varrho(1-\omega)+1])^2 - (a[1 - \lambda\omega - \lambda^2(1-\omega)] - (m\varrho+\varepsilon)[\lambda\varrho(1-\omega)+1])^2\ )$.

Equilibrium

From equation (1.46) it is easy to see that the possible convergence point of a is 0. Then, using $\alpha_f = 0$ from (1.47) it follows that ω_f is the solution of

$$\omega = \iint F(\ m^2\varrho^2 - 2m\varrho(m\varrho+\varepsilon)[\lambda\varrho(1-\omega)+1]\ )\ d\phi_m\ d\phi_\varepsilon\ . \tag{1.48}$$

There is no closed form solution. The existence of the solution follows from the fact that the right hand side of (1.48) is continuous function on a compact set , mapping from $\omega \in [0,1]$ to $[0,1]$.

Stability

$\theta_f = [0\ \omega_f]$ is a locally stable equilibrium point of (1.46)-(1.47) if the eigenvalues of the derivative matrix (Jacobian) evaluated at θ_f, $Dh(\theta)|_{\theta_f}$ are negative in modulus.

The Jacobian at θ_f is

$$Dh(\theta)|_{\theta_f} = \begin{bmatrix} \lambda\omega_f + \lambda^2(1-\omega_f) - 1 & 0 \\ \frac{\partial \iint F(\cdot)\mathrm{d}\phi_m \mathrm{d}\phi_\varepsilon}{\partial a}\Big|_{\theta_f} & \frac{\partial \iint F(\cdot)\mathrm{d}\phi_m \mathrm{d}\phi_\varepsilon}{\partial \omega}\Big|_{\theta_f} - 1 \end{bmatrix},$$

where

$$\frac{\partial \iint F(\cdot)\mathrm{d}\phi_m \mathrm{d}\phi_\varepsilon}{\partial a}\Big|_{\theta_f} = \tag{1.49}$$

$$\iint (2m\varrho[\lambda - \lambda\omega - \lambda^2(1-\omega)] - 2(m\varrho+\varepsilon)...$$

$$[\lambda\varrho(1-\omega)+1][\lambda-1]) f(\ m^2\varrho^2 - 2m\varrho(\varrho m+\varepsilon)[\lambda\varrho(1-\omega)+1]\ )\mathrm{d}\phi_m \mathrm{d}\phi_\varepsilon \tag{1.50}$$

$$\frac{\partial \iint F(\cdot)\mathrm{d}\phi_m \mathrm{d}\phi_\varepsilon}{\partial \omega}\Big|_{\theta_f} = \tag{1.51}$$

$$\iint 2\lambda\varrho^2(\varrho m+\varepsilon) m f(\ m^2\varrho^2 - 2m\varrho(\varrho m+\varepsilon)[\lambda\varrho(1-\omega_f)+1]\ )\mathrm{d}\phi_m \mathrm{d}\phi_\varepsilon\ .$$

This is a lower triangular matrix, the eigenvalues are the diagonal elements. The first eigenvalue is negative if $\lambda\omega_f + \lambda^2(1-\omega_f) < 1$, which is true if $\lambda \in (\frac{-1}{1-\omega_f}, 1)$. The following Claim shows that a sufficient condition for second eigenvalue to be negative is $\frac{2|\lambda|\varrho^2}{\sqrt{2\pi}\sigma}\frac{\sigma_\varepsilon^2}{1-\varrho^2}\left(\varrho + \frac{2\sqrt{1-\varrho^2}}{\pi}\right) < 1$.

Claim 7. $\left.\frac{\partial \iint F(\cdot)\mathrm{d}\phi_m \mathrm{d}\phi_\varepsilon}{\partial \omega}\right|_{\theta_f} < 1$ *if* $\frac{2|\lambda|\varrho^2}{\sqrt{2\pi}\sigma}\frac{\sigma_\varepsilon^2}{1-\varrho^2}\left(\varrho + \frac{2\sqrt{1-\varrho^2}}{\pi}\right) < 1$

Proof. From equation (1.51)

$$\begin{aligned}\iint 2\lambda\varrho^2(\varrho m+\varepsilon)mf(\ m^2\varrho^2 - 2m\varrho(\varrho m+\varepsilon)[\lambda\varrho(1-\omega_f)+1]\ )\mathrm{d}\phi_m\mathrm{d}\phi_\varepsilon < \\ 2\varrho^2\iint|\lambda(\varrho m+\varepsilon)m|f(0)\mathrm{d}\phi_m\mathrm{d}\phi_\varepsilon < \\ 2|\lambda|\varrho^2\iint(\ \varrho m^2+|m\varepsilon|\ )\frac{1}{\sqrt{2\pi}\sigma}d\phi_m\mathrm{d}\phi_\varepsilon \leq \\ \frac{2|\lambda|\varrho^3\sigma_m^2}{\sqrt{2\pi}\sigma}+\frac{2|\lambda|\varrho^2}{\sqrt{2\pi}\sigma}\iint|m||\varepsilon|\mathrm{d}\phi_m\mathrm{d}\phi_\varepsilon\ .\end{aligned} \tag{1.52}$$

$\int |m| d\phi_m$ has a closed form solution:

$$\int |m|d\phi_m = 2\int_0^\infty m\frac{1}{\sqrt{2\pi}\sigma_m}e^{-\frac{m^2}{2\sigma_m^2}} = \frac{2}{\sqrt{2\pi}\sigma_m}\left[-\sigma_m^2 e^{-\frac{m^2}{\sigma_m^2}}\right]_0^\infty = \frac{2\sigma_m}{\sqrt{2\pi}}\ . \tag{1.53}$$

Thus an upper bound for (1.52) is

$$\frac{2|\lambda|\varrho^2}{\sqrt{2\pi}\sigma}\left(\varrho\sigma_m^2 + \frac{2\sigma_m\sigma_\varepsilon}{\pi}\right)\ .$$

Using $\sigma_m^2 = \frac{\sigma_\varepsilon^2}{1-\varrho^2}$

$$\left.\frac{\partial\iint F(\cdot)\mathrm{d}\phi_m\mathrm{d}\phi_\varepsilon}{\partial\omega}\right|_{\theta_f} < \frac{2|\lambda|\varrho^2}{\sqrt{2\pi}\sigma}\frac{\sigma_\varepsilon^2}{1-\varrho^2}\left(\varrho + \frac{2\sqrt{1-\varrho^2}}{\pi}\right)\ .$$

From this, the claim follows. □

□

Proof of Proposition 2 part ii.)

Let the economy (1.1)-(1.4) be populated with two types of agents LS_1 and RE_2. $[a_t, \omega_t]$ converge to a fixed point $[0, \omega_f]$ if $\lambda < 1$ and
$\frac{2|\lambda|}{(1-\lambda(1-\omega_f))^3} \frac{\sigma_\varepsilon^2}{\sqrt{2\pi}\sigma(1-\varrho^2)} \left(\varrho^2 + \frac{2\sqrt{1-\varrho^2}}{\pi}\right) < 1.$
Where $\omega = \iint F\left(-\frac{\varrho^2 m^2 + 2\varrho m\varepsilon}{(1-\lambda(1-\omega))^2}\right) \mathrm{d}\phi_m \mathrm{d}\phi_\varepsilon$. F is the c.d.f. of $N(0,\sigma)$, ϕ_m, ϕ_ε are distribution functions of m and ε.

Proof. The proof goes similarly to proof of (1). We have to find the fixed point of $\theta = [a\ \omega]$. Substituting (1.5), (1.12) and (1.19) into (1.4) and (2.8), the associated system of ordinary differential equations, $\frac{d\theta}{d\tau} = h(\theta)$, is:

$$\frac{\partial a}{\partial \tau} = \frac{\lambda\omega}{1-\lambda(1-\omega)} a - a \tag{1.54}$$

$$\frac{\partial \omega}{\partial \tau} = \iint F\left(\left(\frac{\varepsilon}{1-\lambda(1-\omega)}\right)^2 - \left(\frac{(1-\lambda)a}{1-\lambda(1-\omega)} - \frac{\varrho m + \varepsilon}{1-\lambda(1-\omega)}\right)^2\right) \mathrm{d}\phi_m \mathrm{d}\phi_\varepsilon - \omega \tag{1.55}$$

Equilibrium

From equation (1.54) it is easy to see that the possible convergence point of a is 0. Then, using $\alpha_f = 0$ from (1.55) it follows that ω_f is the solution of

$$\omega = \iint F\left(-\frac{\varrho^2 m^2 + 2\varrho m\varepsilon}{(1-\lambda(1-\omega))^2}\right) \mathrm{d}\phi_m \mathrm{d}\phi_\varepsilon\ . \tag{1.56}$$

Stability

$\theta_f = [0\ \omega_f]$ is a locally stable equilibrium point of (1.54)-(1.55) if the eigenvalues of the derivative matrix (Jacobian) evaluated at θ_f, $Dh(\theta)|_{\theta_f}$ are negative in modulus.

The Jacobian at θ_f is

$$Dh(\theta)|_{\theta_f} = \begin{bmatrix} \frac{\lambda\omega_f}{1-\lambda(1-\omega_f)} - 1 & 0 \\ \left.\frac{\partial \iint F(\cdot)\mathrm{d}\phi_m \mathrm{d}\phi_\varepsilon}{\partial a}\right|_{\theta_f} & \left.\frac{\partial \iint F(\cdot)\mathrm{d}\phi_m \mathrm{d}\phi_\varepsilon}{\partial \omega}\right|_{\theta_f} - 1 \end{bmatrix},$$

where

$$\left.\frac{\partial \iint F(\cdot)\mathrm{d}\phi_m \mathrm{d}\phi_\varepsilon}{\partial a}\right|_{\theta_f} = \tag{1.57}$$

$$\iint 2\frac{\varrho m+\varepsilon}{1-\lambda(1-\omega_f)} f\left(-\frac{\varrho^2 m^2+2\varrho m\varepsilon}{(1-\lambda(1-\omega))^2}\right)\mathrm{d}\phi_m \mathrm{d}\phi_\varepsilon$$

$$\left.\frac{\partial \iint F(\cdot)\mathrm{d}\phi_m \mathrm{d}\phi_\varepsilon}{\partial \omega}\right|_{\theta_f} = \tag{1.58}$$

$$\iint 2\left[\frac{-\varepsilon^2\lambda}{(1-\lambda(1-\omega_f))^3}+\frac{(\varrho m+\varepsilon)^2\lambda}{(1-\lambda(1-\omega_f))^3}\right] f\left(-\frac{\varrho^2 m^2+2\varrho m\varepsilon}{(1-\lambda(1-\omega_f))^2}\right)\mathrm{d}\phi_m \mathrm{d}\phi_\varepsilon\ .$$

The first eigenvalue is negative if $\lambda < 1$[47]. The second eigenvalue is negative if $\left.\frac{\partial \iint F(\cdot)\mathrm{d}\phi_m \mathrm{d}\phi_\varepsilon}{\partial \omega}\right|_{\theta_f} < 1$. The following theorem shows a sufficient condition for this to hold.

Claim 8. $\left.\frac{\partial \iint F(\cdot)\mathrm{d}\phi_m \mathrm{d}\phi_\varepsilon}{\partial \omega}\right|_{\theta_f} < 1$ *if* $\frac{2|\lambda|}{(1-\lambda(1-\omega_f))^3}\frac{\sigma_\varepsilon^2}{\sqrt{2\pi}\sigma(1-\varrho^2)}\left(\varrho^2+\frac{2\sqrt{1-\varrho^2}}{\pi}\right) < 1$

Proof. From equation (1.58)

$$\iint 2\left[\frac{-\varepsilon^2\lambda}{(1-\lambda(1-\omega_f))^3}+\frac{(\varrho m+\varepsilon)^2\lambda}{(1-\lambda(1-\omega_f))^3}\right] f\left(-\frac{\varrho^2 m^2+2\varrho m\varepsilon}{(1-\lambda(1-\omega_f))^2}\right)\mathrm{d}\phi_m \mathrm{d}\phi_\varepsilon \le \tag{1.59}$$

$$\frac{2|\lambda|}{(1-\lambda(1-\omega_f))^3}\iint \varrho^2 m^2+2\varrho|\varepsilon m| f(0)\mathrm{d}\phi_m \mathrm{d}\phi_\varepsilon =$$

$$\frac{2|\lambda|}{(1-\lambda(1-\omega_f))^3}\frac{1}{\sqrt{2\pi}\sigma}[\varrho^2\sigma_m^2+\frac{2\sigma_m\sigma_\varepsilon}{\pi}]\ .$$

Using $\sigma_m^2 = \frac{\sigma_\varepsilon^2}{1-\varrho^2}$

$$\left.\frac{\partial \iint F(\cdot)\mathrm{d}\phi_m \mathrm{d}\phi_\varepsilon}{\partial \omega}\right|_{\theta_f} < \frac{2|\lambda|}{(1-\lambda(1-\omega_f))^3}\frac{\sigma_\varepsilon^2}{\sqrt{2\pi}\sigma(1-\varrho^2)}\left(\varrho^2+\frac{2\sqrt{1-\varrho^2}}{\pi}\right)\ .$$

From this, the claim follows. □

□

47 $\frac{\lambda\omega_f}{1-\lambda(1-\omega_f)} < 1$ holds also for $\lambda > \frac{1}{1-\omega_f}$, however in this case Lipshitz conditions do not hold for (1.54) and stochastic approximation cannot be applied.

1.4.2 Speed of convergence proofs

LS_2 **with** RE_1

Proposition 3. *Let the economy* (1.1)-(1.4) *be populated with two types of agents* LS_2 *and* RE_1*. Then* LS_2 *has* $\sqrt{t}$ *convergence if* $\frac{1}{2}\lambda\varrho + \frac{1}{2}\lambda^2\varrho^2 < \frac{1}{2}$.

Proof. **Speed of convergence of LS_2 with RE_1**
In the proof of Proposition 1 part *i.)* we have already derived:

$$Dh(\theta)|_{\theta_f} = \begin{bmatrix} \lambda\varrho\frac{1}{2} + \lambda^2\varrho^2\frac{1}{2} - 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & -1 \end{bmatrix}.$$

This is a diagonal matrix, the eigenvalues are the diagonal elements. Clearly the las two eigenvalues are smaller than $-\frac{1}{2}$. The first one is smaller than $-\frac{1}{2}$ if $\frac{1}{2}\lambda\varrho + \frac{1}{2}\lambda^2\varrho^2 < \frac{1}{2}$. □

LS_2 **with** RE_2

Proposition 4. *Let the economy* (1.1)-(1.4) *be populated with two types of agents* LS_2 *and* RE_2*. Then* LS_2 *has* $\sqrt{t}$ *convergence if* $\lambda\varrho < \frac{2}{3}$.

Proof. **Speed of convergence of LS_2 with RE_2**
In the proof of proposition 1 part *i.)* we derived

$$Dh(\theta)|_{\theta_f} = \begin{bmatrix} \frac{\lambda\varrho}{2-\lambda\varrho} - 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & -1 \end{bmatrix}.$$

For root-t convergence all eigenvalues have to be smaller than $-\frac{1}{2}$, which holds if $\lambda\varrho < \frac{2}{3}$. □

$LS1$ **with** $RE1$

Proposition 5. *Let the economy* (1.1)-(1.4) *be populated with two types of agents* LS_1 *and* RE_1*. Then sufficient conditions for* $\sqrt{t}$ *convergence of* LS_1 *are*

$$\lambda\omega_f + \lambda^2(1-\omega_f) < \frac{1}{2} \tag{1.60}$$

$$\frac{2|\lambda|\varrho^2}{\sqrt{2\pi}\sigma}\frac{\sigma_\varepsilon^2}{1-\varrho^2}\left(\varrho + \frac{2\sqrt{1-\varrho^2}}{\pi}\right) < \frac{1}{2}\ . \tag{1.61}$$

Proof. **Speed of convergence of LS_1 with RE_1**
Using derivations in the proof of Proposition 2 part *i.)*:

$$Dh(\theta)|_{\theta_f} = \begin{bmatrix} \lambda\omega_f + \lambda^2(1-\omega_f) - 1 & 0 \\ \frac{\partial \iint F(\cdot)\mathrm{d}\phi_m \mathrm{d}\phi_\varepsilon}{\partial a}\Big|_{\theta_f} & \frac{\partial \iint F(\cdot)\mathrm{d}\phi_m \mathrm{d}\phi_\varepsilon}{\partial \omega}\Big|_{\theta_f} - 1 \end{bmatrix}$$

With $\frac{\partial \iint F(\cdot)\mathrm{d}\phi_m \mathrm{d}\phi_\varepsilon}{\partial \omega}\Big|_{\theta_f} < \frac{2|\lambda|\varrho^2}{\sqrt{2\pi}\sigma}\frac{\sigma_\varepsilon^2}{1-\varrho^2}\left(\varrho + \frac{2\sqrt{1-\varrho^2}}{\pi}\right)$.

This is a lower triangular matrix, the eigenvalues are the diagonal elements. For root-t convergence the eigenvalues have to be smaller than $-\frac{1}{2}$, which holds if

$$\lambda\omega_f + \lambda^2(1-\omega_f) < \frac{1}{2} \tag{1.62}$$

$$\frac{2|\lambda|\varrho^2}{\sqrt{2\pi}\sigma}\frac{\sigma_\varepsilon^2}{1-\varrho^2}\left(\varrho + \frac{2\sqrt{1-\varrho^2}}{\pi}\right) < \frac{1}{2}\ . \tag{1.63}$$

□

$LS1$ with $RE2$

Proposition 6. *Let the economy* (1.1)-(1.4) *be populated with two types of agents LS_1 and RE_2. Then LS_1 has $\sqrt{t}$ convergence if*

$$\lambda(1+\omega_f) < 1 \tag{1.64}$$

$$\frac{2|\lambda|}{(1-\lambda(1-\omega_f))^3}\frac{\sigma_\varepsilon^2}{\sqrt{2\pi}\sigma(1-\varrho^2)}\left(\varrho^2 + \frac{2\sqrt{1-\varrho^2}}{\pi}\right) < \frac{1}{2} \tag{1.65}$$

Proof. **Speed of convergence of LS_1 with RE_2**
Using derivations in the proof of Proposition 2 part *ii.)*:

$$Dh(\theta)|_{\theta_f} = \begin{bmatrix} \frac{\lambda\omega_f}{1-\lambda(1-\omega_f)} - 1 & 0 \\ \left.\frac{\partial \iint F(\cdot)\mathrm{d}\phi_m \mathrm{d}\phi_\varepsilon}{\partial a}\right|_{\theta_f} & \left.\frac{\partial \iint F(\cdot)\mathrm{d}\phi_m \mathrm{d}\phi_\varepsilon}{\partial \omega}\right|_{\theta_f} - 1 \end{bmatrix},$$

with $\left.\frac{\partial \iint F(\cdot)\mathrm{d}\phi_m \mathrm{d}\phi_\varepsilon}{\partial \omega}\right|_{\theta_f} < \frac{2|\lambda|}{(1-\lambda(1-\omega_f))^3}\frac{\sigma_\varepsilon^2}{\sqrt{2\pi}\sigma(1-\varrho^2)}\left(\varrho^2 + \frac{2\sqrt{1-\varrho^2}}{\pi}\right)$.

This is a lower triangular matrix, the eigenvalues are the diagonal elements. For root-t convergence the eigenvalues have to be smaller than $-\frac{1}{2}$, which holds if

$$\lambda(1+\omega_f) < 1 \tag{1.66}$$

$$\frac{2|\lambda|}{(1-\lambda(1-\omega_f))^3}\frac{\sigma_\varepsilon^2}{\sqrt{2\pi}\sigma(1-\varrho^2)}\left(\varrho^2 + \frac{2\sqrt{1-\varrho^2}}{\pi}\right) < \frac{1}{2}\,. \tag{1.67}$$

□

Chapter 2

Optimal Monetary Policy When Agents Are Learning

2.1 Introduction

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

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

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

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

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

erature explores the robustness of the optimal policies, which were derived under rational expectations, when instead agents update their expectations according to a learning algorithm.[4] A typical result in this literature is that interest rate rules that are optimal under rational expectations may lead to instability under learning.

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

We are able to analytically derive optimal monetary policy in our theoretical model. One important feature of the optimal policy is that the Central Bank should act more aggressively towards inflation that what a rational expectations model suggests. Earlier work in the literature that uses ad hoc rules has shown similar results computationally (see Ferrero (2003), Orphanides and Williams (2005a), Orphanides and Williams (2005b)); here we establish that these results extend to the case when the central bank uses the optimal policy, and provide a formal proof. The intuition for the result is that aggressively driving inflation close to equilibrium helps private agents to learn the true equilibrium value of inflation at a faster pace. As is well-known, even with rational expectations the central bank cares about price stability due to nominal rigidities. When, in addition expectations of nominal variables are sluggish because of learning, our results show that monetary policy should be even more aggressive towards inflation. Being aggressive towards inflation generates a welfare cost in terms of an increased volatility of the output gap. We show analytically that the optimal policy involves a more volatile output gap then the rational expectations benchmark; this holds true even if the Central Bank puts a high weight on output gap stabilization.

A second important feature of the optimal policy is that it is time consistent, and qualitatively resembles the commitment solution under rational expectations in the sense that the optimal policy is unwilling to accommodate noisy shocks. As a consequence the impulse response of a cost push shock is also similar to the

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

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

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

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

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

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

The rest of the paper is organized as follows: in Section 2.2 we analyze optimal policy under constant gain learning where there is no exogenous cost-push shock; in Section 2.3 we study how the introduction of the cost-push shock affects our results; Section 2.4 relaxes the assumptions that expectations follow constant gain learning, and show that out main results remain valid under decreasing gain learning; Section 2.5 relaxes the assumption that the policy maker can perfectly observe the fundamental shocks and the beliefs of the agents; Section 2.6 concludes.

2.2 The Model without a cost push shock

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

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

$$x_t = E_t^* x_{t+1} - \sigma^{-1}(r_t - E_t^* \pi_{t+1} - \overline{rr}_t) + g_t \qquad (2.1)$$

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

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

$$\pi_t = \beta E_t^* \pi_{t+1} + \kappa x_t \qquad (2.2)$$

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

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

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

$$E_0 \sum_{t=0}^{\infty} \beta^t (\pi_t^2 + \alpha x_t^2) \qquad (2.3)$$

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

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

[7] In other words, in each period firm i can reset the price with a constant probability $1-\theta$, and with probability θ it keeps the same price as in the previous period. If firms take this structure into account when deciding the optimal price it can be shown (See Yun (1996)) that the aggregate inflation is given by (2.2). κ is decreasing in the level of stickiness, the longer are prices fixed in expectation the smaller the effect of the output gap is on inflation.

[8] As is shown in Rotemberg and Woodford (1998), equation (2.3) can be seen as a quadratic approximation to the expected household's utility function; in this case, α is a function of structural parameters.

2.2.1 Benchmark: discretionary solution under rational expectations and under learning

A key feature of this model is that, if expectations are rational (i.e., if $E_t^* = E_t$), there is *no trade-off between inflation and output gap stabilization*; in fact, following Gali (2003), we can solve forward equation (2.2) and impose a boundedness condition on π, obtaining:

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

Therefore, if the CB stabilizes output gap in every period, under RE inflation will also be equal to zero every period; moreover, this plan is time-consistent, in the sense that the optimal plan chosen by the CB if optimizing at period $t+1$ will be equal to the continuation of the optimal plan set when optimizing at t. The absence of inflation bias is due to the fact that, differently from Barro and Gordon (1983) and all the subsequent literature, the target for output chosen by the CB is the natural level of output, and not a higher level; in other words, the target for output gap is zero, as shown in (2.3). To restore an inflation stabilization/output gap stabilization trade-off it is necessary to modify the NKPC introducing a so-called cost-push shock[9].

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

Since the Central Bank can not credibly manipulate beliefs, in the optimization it takes expectations as given. The policy problem is to choose a time path for the nominal interest rate r_t [10] to engineer a time path of the target variables π_t and x_t such that the social welfare loss (2.3) is minimized, subject to the structural equations (2.1) and (2.2), and given the private sectors expectations.

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

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

$$\min_{\{\pi_t,x_t,r_t\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \beta^t(\pi_t^2 + \alpha x_t^2) \tag{2.4}$$

s.t. (2.1), (2.2)

$E_t\pi_{t+1}$, $E_t x_{t+1}$ given for $\forall\ t$

Because there are no endogenous state variables, problem (2.4) reduces to a sequence of static optimization problems. As shown in Clarida, Gali, and Gertler (1999), and the optimality condition to this problem (at time t) is

$$x_t = -\frac{\kappa}{\alpha}\pi_t\ . \tag{2.5}$$

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

Under non-rational expectations (E^*), using (2.5) the optimal allocations are:

$$\pi_t \;=\; \frac{\alpha\beta}{\alpha+\kappa^2}E_t^*\pi_{t+1} \tag{2.6a}$$

$$x_t \;=\; -\frac{\kappa\beta}{\alpha+\kappa^2}E_t^* x_{t+1}\ . \tag{2.6b}$$

and solving for r_t from the structural equations (2.1) and (2.2) yields

$$r_t = \overline{rr}_t + \delta_\pi^{EH} E_t^*\pi_{t+1} + \delta_x^{EH} E_t^* x_{t+1} + \delta_g^{EH} g_t \tag{2.7}$$

where:

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

and E_t^* denote non-rational expectations. Throughout the paper we denote the coefficients by EH referring to the paper Evans and Honkapohja (2003b), where the authors derive a rule analogous to (2.7)[11]. In the terminology introduced in Evans and Honkapohja (2003b), Evans and Honkapohja (2003a) (EH hereafter), this is an *expectations-based reaction function.* EH show that this rule guarantees not only determinacy under RE, but also convergence to the RE equilibrium when expectations E_t^* evolve according to least squares learning.

However, a rational Central Bank, knowing that private agents follow learning, could do even better. In other words the solution (2.7) under learning is not a full

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

optimum. In the next section we show how optimal monetary policy is modified when the CB optimizes taking into account its effect on private expectations.

2.2.2 Constant Gain Learning

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

$$E_t^* \pi_{t+1} \equiv a_t = a_{t-1} + \gamma(\pi_{t-1} - a_{t-1}) \tag{2.8}$$

$$E_t^* x_{t+1} \equiv b_t = b_{t-1} + \gamma(x_{t-1} - b_{t-1}) \tag{2.9}$$

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

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

In Section 2.4 we will relax this assumption, and examine how optimal policy changes when agents follow decreasing gain learning.

To analyze the optimal control problem faced by the CB, we use the standard Ramsey approach, namely we suppose that the policymakers take the structure of the economy (equations (2.1) and (2.2)) as given; moreover, we assume that the CB knows how private agents' expectations are formed, and takes into account its ability to influence the evolution of the beliefs. Hence, the CB problem can be stated as follows:

$$\min_{\{\pi_t, x_t, r_t, a_{t+1}, b_{t+1}\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \beta^t (\pi_t^2 + \alpha x_t^2) \tag{2.10}$$

$$\begin{aligned} &\text{s.t. } (2.1),\ (2.2),\ (2.8),\ (2.9) \\ &\quad a_0,\ b_0 \text{ given} \end{aligned}$$

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

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

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

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

$$\begin{aligned}
\lambda_{1t} &= 0 && (2.11)\\
2\pi_t - \lambda_{2t} + \gamma\lambda_{3t} &= 0 && (2.12)\\
2\alpha x_t + \kappa\lambda_{2t} - \lambda_{1t} + \gamma\lambda_{4t} &= 0 && (2.13)\\
E_t\left[\frac{\beta}{\sigma}\lambda_{1t+1} + \beta^2\lambda_{2t+1} + \beta(1-\gamma)\lambda_{3t+1}\right] &= \lambda_{3t} && (2.14)\\
E_t\left[\beta\lambda_{1t+1} + \beta(1-\gamma)\lambda_{4t+1}\right] &= \lambda_{4t} && (2.15)
\end{aligned}$$

where λ_{it}, $i = 1, ..., 4$ denote the Lagrange multipliers associated to (2.1), (2.2), (2.8) and (2.9), respectively. The necessary conditions for an optimum are the first order conditions, the structural equations (2.1)-(2.2) and the laws of motion of private agents' beliefs, (2.8)-(2.9). Combining equation (2.11) and (2.15), we get:

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

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

$$\lambda_{4t} = 0 \tag{2.16}$$

If we put together equations (2.11)-(2.14) and (2.16), we derive the following optimality condition:

$$\frac{\kappa}{\alpha}\pi_t + x_t = \beta E_t\left[\beta\gamma x_{t+1} + (1-\gamma)\left(\frac{\kappa}{\alpha}\pi_{t+1} + x_{t+1}\right)\right] \tag{2.17}$$

Inflation-Output Gap Tradeoff

A crucial difference from the rational expectations case is that under learning there is an inflation-output gap tradeoff even without a cost push shock. From equation (2.2) we can see that, if a_t is different from zero, inflation and output gap cannot be set contemporaneously equal to zero, as in the RE case. Hence, the fact that the expectations are not rational, introduces a trade-off between inflation and output gap stabilization that is not present under RE. In particular, we have

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

the contemporaneous presence of two trade-offs. There is an *intratemporal trade-off* between stabilization of inflation at t and output gap at t, determined by the presence of the nonzero term βa_t in the Phillips Curve (2.2). There is an additional *intertemporal trade-off* between optimal behavior at t and stabilization of output gap at $t+1$, which is generated by the ability of the CB to manipulate future values of a. This can be seen from iterating forward the optimality condition (2.17):

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

Hence, for a given positive value of x_t, the optimal disinflation is less harsh with respect to the one implied by (2.5), provided that future output gaps are also expected to be positive. A smaller deflation in turn guarantees that future inflationary expectations will be closer to the rational expectations equilibrium of inflation, zero.

Let us summarize our first result for later reference:

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

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

Optimal allocations

To derive the optimal allocations, we can use (2.2) to substitute out x_t in (2.17), then using the evolution of inflationary expectations (2.8)we get:

$$\pi_t + \frac{\alpha}{\kappa^2}[\pi_t - \beta a_t] - \beta E_t \left[\frac{\alpha}{\kappa^2} \left(1-\gamma \left(1-\beta \right) \right) \left[\pi_{t+1} - \beta a_{t+1} \right] + \left(1-\gamma \right) \pi_{t+1} \right] = 0 \tag{2.18}$$

Hence, at an optimum, the dynamics of the economy can be summarized by stacking equations (2.8), (2.9) and (2.18), obtaining the trivariate system[15]:

$$E_t y_{t+1} = A y_t \tag{2.19}$$

[15] Once we have the equilibrium laws of motion for $[\pi_t, a_t, b_t]$, we can use (2.1) and (2.2) to derive the equilibrium r_t and x_t.

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

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

The three boundary conditions of the above system are:

$$\begin{array}{c} a_0,\, b_0 \text{ given} \\ \lim_{t\to\infty} |E_t\pi_{t+1}| < \infty \end{array} \quad . \tag{2.20}$$

The last one is due to the fact that, if there exists a solution to the problem (2.10) when the possible sequences $\{\pi_t, x_t, r_t\}$ are restricted being bounded, then this would also be the minimizer in the unrestricted case[16].

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

$$A_1 \equiv \begin{pmatrix} \frac{\kappa^2+\alpha+\alpha\beta^2\gamma(1-\gamma(1-\beta))}{\alpha\beta(1-\gamma(1-\beta))+\kappa^2\beta(1-\gamma)} & -\frac{\alpha\beta(1-\beta(1-\gamma)(1-\gamma(1-\beta)))}{\alpha\beta(1-\gamma(1-\beta))+\kappa^2\beta(1-\gamma)} \\ \gamma & 1-\gamma \end{pmatrix} \tag{2.21}$$

In the Appendix we show that A_1 has one eigenvalue inside and one outside the unit circle, which implies (together with $(1-\gamma) \in (0,1)$) that we can invoke Proposition 1 of Blanchard and Kahn (1980) to conclude that the system (2.19)-(2.20) has one and only one solution. In other words, there exists one and only one stochastic process[17] for each of the three variables of y such that (2.20) is satisfied. Moreover, note that $y_{1t} \equiv [\pi_t, a_t]'$ does not depend on b_t; therefore, the processes for inflation and a that solve (together with the process for b) the system (2.19)-(2.20) are also a solution of the subsystem:

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

together with the boundary conditions:

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

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

[16]For a proof, see the Appendix.

[17]Since the system (2.19) does not depend on the only source of randomness in this economy (i.e., g), in equilibrium the process followed by the endogenous variables turns out to be deterministic, see below.

$$\pi_t = c_\pi^{cg} a_t \tag{2.22}$$

Following the adaptive learning terminology, we call (2.22) the Actual Law of Motion (ALM) of inflation.

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

Proposition 1. *Let c_π^{cg} be the feedback coefficient defined in (2.22); then, the following holds:*

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

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

Proof. See the Appendix. □

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

As is shown in the Appendix, c_π^{cg} depends on all the structural parameters; in particular, its dependence on the constant gain γ is not necessarily monotonic. In fact, a higher value of γ has two effects on c_π^{cg}: on one hand, it increases the effect of current inflation on future expectations, increasing the incentive for the CB to use this influence (i.e., it would determine a lower c_π^{cg}); on the other hand, it reduces the impact of current expectations on future expectations, thus reducing the benefits from a reduction of the expectations, so that there is an incentive to set a higher c_π^{cg}. In Figure 2.4 we show a numerical example with the calibration found in Woodford (1996), i.e. with $\beta = 0.99$, $\sigma = 0.157$, $\kappa = 0.024$ and $\alpha = 0.04$; in this case, the first effect dominates, so that c_π^{cg} is a monotonically decreasing function of γ.

Using the structural equation (2.2) we can derive the optimal allocation of the output gap:

$$x_t = c_x^{cg} a_t \tag{2.23}$$

where:

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

$c_\pi^{cg} < \frac{\alpha\beta}{\alpha+\kappa^2}$ (see Proposition 1) implies $c_x^{cg} < -\frac{\kappa\beta}{\alpha+\kappa^2}$; if the private sector expects inflation to be positive, the optimal CB response will imply a negative output gap, i.e. the policymaker will contract economic activity (using the interest rate instrument) in order to attain an actual inflation sufficiently smaller than the expected one. Using (2.22) and (2.23) in (2.1) we can derive the nominal interest rate:

$$r_t = \overline{rr}_t + \delta_\pi^{cg} a_t + \delta_x^{cg} b_t + \delta_g^{cg} g_t \tag{2.24}$$

where:

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

The interest rate rule (2.24) is an *expectations-based reaction function*, which is characterized by a coefficient on inflation expectations that is decreasing in c_π^{cg}: an optimal ALM for inflation that requires a more aggressive undercutting of inflation expectations (a lower c_π^{cg}) calls for a more aggressive behavior of the CB when it sets the interest rate (a higher coefficient on inflation expectations in the rule (2.24)). Moreover, the coefficients on b_t and g_t are such that their effects on the output gap in the IS curve are fully neutralized.

Since $c_{\pi,t}^{cg} < \beta$ (see Proposition 1) $\delta_{\pi,t}^{cg}$ is always bigger than 1. In response to a rise in expected inflation optimal policy should raise the nominal interest rate sufficiently to increase the real interest rate. An increase in the real rate has a negative effect on current output; this reflects the intertemporal substitution of consumption. Then a contraction in output will decrease current inflation through the Phillips Curve (2.2), and consequently through Equation (2.8) inflationary expectations in the next period will decrease. This criteria is also emphasized in Clarida, Gali, and Gertler (1999) under the discretionary rational expectations solution; since this holds both under RE and learning it provides a very simple criteria for evaluating monetary policy[18].

Asymptotically, the system will converge to the RE equilibrium, with inflation and output gap equal to zero, and so do the corresponding expectations; this can be seen from the autonomous, linear, homogeneous system of first-order difference equations (2.19). The asymptotic properties of this kind of systems are well-known[19], and with two eigenvalues inside and one outside the unit circle, and the set of boundary conditions (2.20), we have only one non-explosive solution, which is such that in the long run the system converges to the trivial solution $y_t = 0$.

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

[19]See for example Agarwal (2002).

2.3 Introduction of a cost-push shock

In this section we will change the model, introducing an additional term in the Phillips Curve, called a cost-push shock[20], so that equation (2.2) becomes:

$$\pi_t = \beta E_t^* \pi_{t+1} + \kappa x_t + u_t \tag{2.25}$$

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

2.3.1 Benchmark: discretionary solution under rational expectations and under learning

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

$$\pi_t^{RE} = \frac{\alpha}{\kappa^2 + \alpha} u_t \tag{2.26a}$$

$$x_t^{RE} = -\frac{\kappa}{\kappa^2 + \alpha} u_t \ . \tag{2.26b}$$

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

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

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

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

$$\pi_t = \frac{\alpha\beta}{\alpha+\kappa^2}a_t + \frac{\alpha}{\alpha+\kappa^2}u_t \quad (2.27a)$$

$$x_t = -\frac{\kappa\beta}{\alpha+\kappa^2}a_t - \frac{\kappa}{\alpha+\kappa^2}u_t \ . \quad (2.27b)$$

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

$$r_t = \overline{rr}_t + \delta_\pi^{EH} E_t^* \pi_{t+1} + \delta_x^{EH} E_t^* x_{t+1} + \delta_g^{EH} g_t + \delta_u^{EH} u_t \quad (2.28)$$

where:

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

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

2.3.2 Constant Gain Learning

At the presence of cost push shocks an additional problem arises in designing the optimal monetary policy when agents are learning: namely which PLM the agents are learning. As we explained above the actual law of motion of the discretion and the commitment solution have different functional forms. For analytical simplicity, in this paper we will restrict our attention to the discretionary case. In particular, we assume that agents believe that inflation and output gap are continuous invariant functions of the cost-push shock only, $\pi_t = \pi(u_t)$ and $x_t = x(u_t)$[24]; this hypothesis, together with the iid nature of the shock, implies that the conditional and unconditional expectations of inflation and output gap coincide, and are perceived by the agents as constants. Hence, it is natural to assume that agents will estimate them using their sample means: the stochastic recursive algorithms (2.8), (2.9)[25].

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

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

[25]To be precise, in the algorithms (2.8), (2.9) the observations are weighted geometrically, while in the normal sample average they all receive equal weight.

We can now follow a procedure analogous to the one used in the model without cost-push shock. The optimality condition we get is the same as before, Equation (2.17).

With cost push shocks, there is a well known intratemporal tradeoff between inflation and the output gap present under RE; the presence of learning introduces an additional intertemporal tradeoff (Result 1 holds). We can isolate the two tradeoffs from the optimality condition. When $\gamma = 0$ (which implies constant expectations) (2.17) implies:

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

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

$$\frac{\kappa}{\alpha}\pi_t + x_t = 0 \ , \tag{2.29}$$

which is identical to the optimality condition derived in the RE optimal monetary policy literature when the CB sets the optimal plan taking private sector's expectations as given (i.e., in the discretionary case). Clarida, Gali, and Gertler (1999) describe this relation as implying a 'lean against the wind' policy: in other words, if output gap (inflation) is above target, it is optimal to deflate the economy (contract demand below capacity). Because of the presence of the cost push shock in the Phillips Curve, the Central Bank cannot set π_t and x_t equal to zero every period; so an intratemporal tradeoff between inflation and output gap is present (even in the limit). When $\gamma > 0$ iterating forward (2.17) show the presence of an intertemporal tradeoff, just like in Section 2.2.2. When the current output gap is positive, the Central Bank will decrease inflation less then under RE; the Central Bank renounces to fully stabilize the current economy, in exchange of easing future inflation-output gap tradeoffs.

Stacking (2.17) with the Phillips Curve (2.25) and the algorithms (2.8)-(2.9), we can show that at the optimum the economy evolves according to:

$$\begin{pmatrix} E_t\pi_{t+1} \\ a_{t+1} \\ b_{t+1} \end{pmatrix} = A \begin{pmatrix} \pi_t \\ a_t \\ b_t \end{pmatrix} + \begin{pmatrix} -\frac{\alpha}{\alpha\beta(1-\gamma(1-\beta))+\kappa^2\beta(1-\gamma)} \\ 0 \\ -\frac{\gamma}{\kappa} \end{pmatrix} u_t \tag{2.30}$$

(where A is defined as in the previous section), plus the boundary conditions (2.20). The system (2.30)-(2.20) is in the form studied in Blanchard and Kahn (1980), so that we can use their results. In particular, since there are two predetermined variables and one non-predetermined, and A has one eigenvalue outside the unit circle and two inside, there exists one and only one solution. Moreover, also the system:

$$\left(\begin{array}{c} E_t\pi_{t+1} \\ a_{t+1} \end{array}\right) = A_1 \left(\begin{array}{c} \pi_t \\ a_t \end{array}\right) + \left(\begin{array}{c} -\frac{\alpha}{\alpha\beta(1-\gamma(1-\beta))+\kappa^2\beta(1-\gamma)} \\ 0 \end{array}\right) u_t \tag{2.31}$$

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

$$\pi_t = c_\pi^{cg} a_t + d_\pi^{cg} u_t \tag{2.32}$$

Combining $E_t\pi_{t+1} = c_\pi^{cg} a_{t+1}$ with the optimality condition (2.17) and the Phillips Curve (2.25), and making use of the law of motion of inflation expectations (2.8), we derive the values of the coefficients c_π^{cg} and d_π^{cg}, which are summarized in the next Proposition.

Proposition 2. *Let the economy evolve according to the system (2.30), (2.20); then the ALM for inflation is:*

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

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

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

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

$$x_t = c_x^{cg} a_t + d_x^{cg} u_t \tag{2.33}$$

$$r_t = \overline{rr}_t + \delta_\pi^{cg} a_t + \delta_x^{cg} b_t + \delta_g^{cg} g_t + \delta_u^{cg} u_t \tag{2.34}$$

where $c_x^{cg}, \delta_\pi, \delta_x^{cg}, \delta_g^{cg}$ are the same as in (2.24), and:

$$\begin{aligned} d_x^{cg} &= \tfrac{d_\pi^{cg}-1}{\kappa} \\ \delta_u^{cg} &= -\sigma\tfrac{d_\pi^{cg}-1}{\kappa} \end{aligned}$$

Plugging (2.32) into (2.8), we get:

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

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

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

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

2.3.3 Comparison with the myopic rule

In this section we state results regarding how optimal monetary policy under constant gain learning differs from myopic rules used earlier in the literature, where myopic means a rule that considers expectations as given in the optimization problem: in particular we refer to rule (2.28), derived in Evans and Honkapohja (2003b) (henceforth EH).

It is clear that the coefficients on the output gap expectations and on the demand shock are the same in rule (2.28) as in rule (2.34), while the other two coefficients are typically different. Proposition 1 implies $\delta_{\pi,t}^{cg} > \delta_{\pi}^{EH}$: the interest rate response of OP to out-of-equilibrium inflation expectations is more aggressive than the interest rate response of EH. This is due to the fact that when the CB takes into account its ability to influence agents' beliefs, it optimally chooses to undercut future inflation expectations more than what a myopic CB would do.

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

The inflation and output gap allocations implemented by the two different interest rate rules are also different. Under constant gain learning optimal allocations are characterized by (2.32) and (2.33). Under EH allocations are given by (2.27) with $E_t^*\pi_{t+1} = a_t$.

From Proposition 1 we know that the feedback coefficient under optimal policy c_{π}^{cg} is smaller than under the EH rule, in order to undercut inflation expectations more. Also the response to the cost push shock is of lesser magnitude when (2.34) is used instead of (2.28) (in fact, $c_{\pi}^{cg} < \frac{\alpha\beta}{\kappa^2+\alpha}$ implies that $d_{\pi}^{cg} < \frac{\alpha}{\kappa^2+\alpha}$), because the CB is less willing to accommodate noisy shocks, in order to make easier for the private sector to learn what is the long-term value of the conditional expectations of inflation.

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

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

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

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

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

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

Similarity to the commitment solution

From Result 2 and 3 it follows that the impact of a given nonzero cost push shock drives inflation (output gap) closer to (further from) target when agents are learning, relative to the discretionary RE case. Interestingly, this behavior *qualitatively resembles the optimal RE equilibrium under commitment* within a simple class of policy rules derived in Clarida, Gali, and Gertler (1999): if the CB can commit to a policy rule that is a linear function of u_t, the solution can be characterized, when compared to the discretionary equilibrium, by inequalities analogous to the ones summarized in the results stated above. However, the (constrained) commitment solution differs from the discretionary one only when the cost-push shock is an AR(1); if u -and consequently, the equilibrium processes for inflation and output gap- is iid, the two solutions coincide, since future (rational) expectations of the agents cannot be manipulated by the CB. Instead, if expectations are backward-looking, the future beliefs can be manipulated also when the shock is iid: the current actions of the CB influence future beliefs through (2.8) and (2.9) even if the shock is iid.

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

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

As a result of these two similarities, the impulse response function of inflation to a cost push shock will be also similar under OP and RE commitment. Figure 2.5 displays the impulse response function of inflation to a unit shock under OP and discretionary RE policy. In the optimal RE discretionary policy, inflation rises on impact and immediately reverts to the steady state once the shock dies out. Instead, under learning the policy maker engineers a smaller initial response of inflation; in subsequent periods inflation gradually converges back to the steady state value. Clarida, Gali, and Gertler (1999) and Gali (2003) show a *similar disinflation path for the Ramsey policy*: a smaller initial inflation compared to the discretionary case, in exchange for a more persistent deviation from the steady state later[28] This behavior of Ramsey policy leads to welfare gains over discretion due to the convexity of the loss function; this preference for slower but milder adjustment to shocks is at the heart of the stabilization bias.

The similarity to the RE commitment solution resembles the analysis carried out in Sargent (1999), Chapter 5, which shows that in the Phelps problem under adaptive expectations[29], the optimal monetary policy drives the economy close to the Ramsey optimum. Moreover, when the discount factor β equals 1, optimal policy under learning replicates the Ramsey equilibrium. In our case, optimal policy under learning cannot replicate the commitment solution even for β going to 1. This result follows from the particular nature of the gains from commitment; commitment calls for an ALM with a different functional form to the discretionary case[30]. In the Phelps problem, on the other hand, the Phillips Curve is such that the discretion and commitment outcome of inflation has the same functional form, but different coefficients. However, also in our case an increase in the discount factor makes the optimal disinflationary path under learning getting closer to the commitment solution. This can be seen in Table 2.1, where we summarize the behavior of inflation in response to a unit cost push shock when the model's parameters are calibrated as in Woodford (1996), apart from β which takes several values. As β goes to 1 the initial response of inflation is milder and the path back

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

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

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

to the steady state longer.

Table 2.1: Path of inflation for different βs after an initial cost push shock

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

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

Welfare Loss Analysis

To have a quantitative feeling of the welfare gains that the use of the optimal rule (2.34) instead of the EH rule (2.28) implies, we present a numerical welfare loss analysis.

Since welfare losses in utility terms are hard to interpret we report consumption equivalents: for a given monetary policy rule we calculate the cumulative utility losses resulting from deviations from the steady state allocation and then express what is the equivalent percentage decrease of the steady state consumption that results to the same cumulative utility loss (For details of the calculation see the Appendix.). We use the calibration of Woodford (1996): $\beta = 0.99$, $\kappa = 0.024$, $\alpha = 0.048$ and $\sigma = 0.157$. We perform Monte Carlo with a simulation length 10,000 and a cross sectional sample size of 1000, with the initial condition $a_0 = b_0 = 0$. Cost push shocks are drawn from a normal distribution with 0 mean and variance 0.1[31].

Table 2.2 reports consumption equivalents for a range of tracking parameters. For small tracking parameters the results are in the range of Lucas' original estimates[32]: consumption losses resulting from cyclical fluctuations are small. For higher tracking parameters the consumption equivalents are also higher, which results from the fact in the presence of a cost push shock, constant gain learning does not settle down to RE, but converges to a limiting distribution and the limiting variance of inflation expectation increases in γ (keeping other coefficients

[31]Note that the demand shock does not appear in the actual law of motion of the endogenous variables.

[32] See Lucas (1987).

constant). This is illustrated in Figure 2.7. A higher variance of inflationary expectations in turn results in higher variance of inflation and output gap both under OP (see equation (2.32) and (2.33)) and under EH (see equation (2.27)), thus a higher welfare loss equivalent permanent consumption reduction.

Since both inflation and output gap variance can be expressed as a linear function of the variance of the cost push shock, clearly the absolute value of consumption equivalents are also increasing with the variance of the cost push shock but the ratio of consumption equivalents under OP and EH should not be sensitive to the choice of σ_u^2. In Table 2.3 we report consumption equivalents for $\sigma_u^2 = 0.6$, an estimate of Milani (2005). We can observe that consumption equivalents are higher, but the ratio of them are not sensitive to the variance of the cost push shock.

Optimal policy decreases consumption equivalents relative to the myopic rule (2.28) (see the third column in Table 2.2 and 2.3). Even for tracking parameters bellow 0.05[33] the gain from using an optimal interest rate rule (2.34) compared to the EH rule (2.28) is around $1-3\%$. The gain in consumption equivalents is higher the higher is the gain parameter. For a very high tracking parameter $\gamma = 0.9$ the welfare loss in consumption terms of not using the optimal rule is twice as large as under OP. This follows from the fact that, optimal policy takes into account that expectations have a limiting variance while the EH policy is myopic in the sense that considers expectations to be fixed[34]. As a result, optimal policy aims to decrease the limiting variance of inflationary expectations while EH does not, and the higher is the tracking parameter the bigger is the decrease in the limiting variance OP engineers compared to EH (see Figure 2.7).

It is interesting to examine the composition of welfare losses coming from inflation variation and output gap variation. For this we calculate the equivalent permanent consumption decrease for welfare losses caused by only inflation variation or output gap variation respectively, and report the ratios of OP and EH in Table 2.4. The table demonstrates Results 2 and 3: optimal policy focuses on decreasing inflation variation even at the cost of higher output gap variation. The higher is the tracking parameter, the higher is the incentive of the Central Bank to focus on lowering inflation variance and allowing for an increase in output gap

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

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

Table 2.2: Consumption equivalents under optimal policy (OP) and myopic policy (EH) under constant gain learning, $\sigma_u^2 = 0.1$

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

Woodford (1996) calibration

Table 2.3: Consumption equivalents under optimal policy (OP) and myopic policy (EH) under constant gain learning, $\sigma_u^2 = 0.6$

Tracking parameter	p^{OP}	p^{EH}	p^{OP}/p^{EH}
0.0187	0.0787	0.0788	0.9990
0.05	0.0898	0.0920	0.9759
0.08	0.1041	0.1127	0.9243
0.1	0.1145	0.1297	0.8830
0.3	0.2256	0.3763	0.5996
0.5	0.3369	0.6824	0.4936
0.9	0.5537	1.3476	0.4108

Woodford (1996) calibration

deviation from the flexible price equilibrium. For $\gamma = 0.9$ compared to EH an optimally behaving Central Bank engineers a 78% lower welfare loss in inflation when it properly conditions on expectation formation, permitting at the same time 15 times more variation in output gap.

Moreover, it is worth noting that the use of a myopic rule under constant gain learning allows for the autocorrelation of inflation to rise, thus increasing the persistence of a shock's effect on inflation expectations. This problem arises from the relatively weak response to inflation expectations which feeds back to current inflation and, in turn, into subsequent expectations and inflations. The optimal rule's strong feedback to inflation expectations dampens this interaction between inflation and expectations[35].

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

Table 2.4: Ratio of consumption equivalents of losses due to inflation and output gap variations using OP and EH under constant gain learning

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

Woodford (1996) calibration

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

2.4 Decreasing Gain Learning

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

Using a constant gain parameter γ is appropriate when agents believe struc-

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

tural changes to occur. If instead the private sector confidently believes that the environment is stationary it is more reasonable to model their learning behavior with a decreasing gain rule, namely an algorithm of the form:

$$E_t^* \pi_{t+1} \equiv a_t = a_{t-1} + t^{-1}(\pi_{t-1} - a_{t-1}) \tag{2.35}$$

$$E_t^* x_{t+1} \equiv b_t = b_{t-1} + t^{-1}(x_{t-1} - b_{t-1}) \tag{2.36}$$

where the only difference with (2.8)-(2.9) is the substitution of γ with t^{-1}.

An updating scheme of this form is equivalent[36] to estimating inflation and output gap every period with OLS[37].

2.4.1 Without cost push shock

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

$$\min_{\{\pi_t, x_t, r_t, a_{t+1}, b_{t+1}\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \beta^t (\pi_t^2 + \alpha x_t^2) \tag{2.37}$$

$$\text{s.t. } (2.1),\ (2.2),\ (2.35),\ (2.36)$$
$$a_0,\ b_0 \text{ given}$$

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

$$\frac{\kappa}{\alpha}\pi_t + x_t = \beta E_t \left[\beta \frac{1}{t+1} x_{t+1} + \frac{\kappa}{\alpha}\pi_{t+1} + x_{t+1} \right] \tag{2.38}$$

Iterating it forward we get:

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

Similarly to Section 2.2 our result is that learning introduces an an intratemporal tradeoff between inflation and output that is not present under RE in an economy without a cost push shock and an additional intertemporal tradeoff that is not

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

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

present in general under rational expectations (Result 1). From the latter it follows that during the transition for a given positive value of x_t, the optimal disinflation is less harsh with respect to the one implied by (2.5) (optimizing taking expectations as given) provided that the series on the right hand side is expected to be positive. The intuition behind is that when the CB makes active use of the expectation formation, it renounces its ability to optimally stabilize the economy in period t, in exchange for a reduction in future inflation expectations (in absolute value) and this allows an ease in the future inflation-output gap trade-off embedded in the Phillips Curve.

To derive the optimal allocations, we can use (2.2) to substitute out x_t in (2.38), then using the evolution of inflationary expectations (2.35)we get:

$$E_t\left[\pi_{t+1}\right] = A_{11,t}\pi_t + A_{12,t}a_t, \tag{2.39}$$

where:

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

Hence, at an optimum, the dynamics of the economy can be summarized by stacking equations (2.35), (2.36) and (2.39), and obtaining the trivariate system:

$$E_t y_{t+1} = A_t y_t \tag{2.40}$$

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

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

The three boundary conditions of the above system are (2.20), the same as in section 2.2.2.

We can find the solution with the method of undetermined coefficients, with the guess[38]:

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

$$\pi_t = c^{dg}_{\pi,t} a_t. \tag{2.41}$$

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

$$c^{dg}_{\pi,t} = \frac{c^{dg}_{\pi,t+1}\left(1 - \frac{1}{t+1}\right) - A_{12,t}}{A_{11,t} - c^{dg}_{\pi,t+1}\frac{1}{t+1}} \tag{2.42}$$

Of course, there exist infinite sequences that satisfy equation (2.42), one for each initial value $c^{dg}_{\pi,0}$. However, since the boundary conditions require π_t to stay bounded, we will concentrate on the solutions that do not explode.

Proposition 3. *Let* $\left\{c^{dg}_{\pi,t}\right\}$ *be defined by* (2.42)*, and assume it is bounded; then,* $\lim_{t\to\infty} c^{dg}_{\pi,t}$ *exists, and is given by:*

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

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

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

Proof. See the Appendix. □

Thus Result 2 holds during the transition: when the CB takes into account its influence on expectations it is optimal to decrease the effect of out-of-equilibrium expectations on inflation compared to the myopic policy (see (2.6)), in order to undercut future inflation expectations by a larger amount. This relaxes the future inflation-output gap trade-off embedded in the Phillips Curve. The ALM for output gap is:

$$x_t = c^{dg}_{x,t} a_t \qquad c^{dg}_{x,t} = \frac{c^{dg}_{\pi,t} - \beta}{\kappa} \tag{2.43}$$

If the private sector expects inflation to be positive, the optimal CB will contract economic activity more than EH[39] (using the interest rate instrument); the CB is ready to pay a short-term cost represented by a wider current output gap in order to contain future inflationary expectations.

[39] From $c^{dg}_{\pi,t} < \frac{\alpha\beta}{\alpha+\kappa^2}$ it follows that $c^{dg}_{x,t} < -\frac{\kappa\beta}{\alpha+\kappa^2}$. Compare with the ALM under EH (2.6).

The nominal interest rate rule is as follows:

$$r_t = \overline{rr}_t + \delta^{dg}_{\pi,t} a_t + \delta^{dg}_x b_t + \delta^{dg}_g g_t \tag{2.44}$$

where:

$$\begin{aligned} \delta^{dg}_{\pi,t} &= 1 - \sigma \tfrac{c^{dg}_{\pi,t} - \beta}{\kappa} \\ \delta^{dg}_x &= \sigma \\ \delta^{dg}_g &= \sigma \end{aligned}$$

Since $c^{dg}_{\pi,t} < \beta$ (see Proposition 3) $\delta^{dg}_{\pi,t}$ is always bigger than 1. In response to a rise in expected inflation optimal policy should raise the nominal interest rate sufficiently to increase the real interest rate. The following proposition pertains to the characteristics of the optimal rule compared to the myopic EH rule (2.7):

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

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

An interesting result is that the coefficient on inflation expectations in the interest rate rule (2.44) is *time-varying*, reflecting the fact that the Central Bank's incentives to manipulate agents' beliefs evolve over time. This implies that during the transition optimal policy should be time varying even in a stationary environment.

In Figure 2.1, we show how this coefficient depends on time when the parameters are calibrated according to Clarida, Gali, and Gertler (2000): $\kappa = 0.075$, $\alpha = 0.3$, $\sigma = 1/4$. $\delta^{dg}_{\pi,t}$ is always above its limiting level (see analytical proof in Proposition 4), moreover, it decreases over time. Numerical analysis on the grid $\beta = 0.99$ and $\alpha \in [0.01, 2]$, $\kappa \in [0.01, 0.5]$ shows that this decreasing behavior of $\delta^{dg}_{\pi,t}$ is a robust feature of the model [40]. We find that after the 4th period (from the 4th to the 5th period and so on) $\delta^{dg}_{\pi,t}$ is always decreasing, while in the first 4 periods $\delta^{dg}_{\pi,t}$ might be increasing (hump-shaped) for a combination of low values of α and high values of κ (see Figure 2.3) [41]. We summarize our new results as:

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

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

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

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

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

Proposition 5. *Let* $\pi_t = c^{dg}_{\pi,t} a_t$*, where* $c^{dg}_{\pi t}$ *is given by (2.42); then,* $a_t \to 0$*.*

Proof. See the Appendix. □

Combining this result with the boundedness of $c^{dg}_{\pi,t}$, the ALM for inflation (2.41) and output gap (2.43) tell us that both these variables go to zero asymptotically, restoring the RE allocations. Optimal policy naturally chooses a non-explosive solution (it is E-stable), and drives expectations to the rational expectations equilibrium.

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

2.4.2 With Cost Push Shock

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

Proceeding with the same analysis as before, we get the same optimality condition (2.38). Substituting out x_t using the Phillips Curve (2.25), and using the evolution of inflation expectations (2.35) we can show that at the optimum the economy evolves according to:

$$E_t[\pi_{t+1}] = A_{11,t}\pi_t + A_{12,t}a_t + P_{1,t}u_t \tag{2.45}$$

where A_{11} and A_{12} are the same as in (2.39), and

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

Stacking together this condition, and the two learning algorithms (2.35) and (2.36), we again have a trivariate system. We can guess and verify that the ALM for inflation is of the form:

$$\pi_t = c_{\pi,t}^{dg} a_t + d_{\pi,t}^{dg} u_t \tag{2.46}$$

which implies that $E_t \pi_{t+1} = c_{\pi,t+1}^{dg} a_{t+1}$; substituting this expression in (2.45), and making use of the law of motion of inflation expectations (2.35), we obtain that the sequence $\left\{c_{\pi,t}^{dg}\right\}$ is identical to (2.42) and $\left\{d_{\pi,t}^{dg}\right\}$ must satisfy:

$$d_{\pi,t}^{dg} = \frac{P_{1,t}}{c_{\pi,t+1}^{dg}\frac{1}{t+1} - A_{11,t}}. \tag{2.47}$$

The solution of $\left\{c_{\pi,t}^{dg}\right\}$ is again characterized by Proposition 3. From Proposition 3 and (2.47) it follows that $0 < d_{\pi,t}^{dg} < \frac{\alpha}{\alpha+\kappa^2}$, a positive cost push shock increases inflation, but less than under myopic policy (compare with (2.27)). The ALM for output gap and the nominal interest rate rule are given by:

$$x_t = c_{x,t}^{dg} a_t + d_{x,t}^{dg} u_t \tag{2.48}$$

$$r_t = \overline{rr}_t + \delta_{\pi,t}^{dg} a_t + \delta_x^{dg} b_t + \delta_g^{dg} g_t + \delta_{ut}^{dg} u_t \tag{2.49}$$

where $c_{x,t}^{dg}, \delta_{\pi,t}^{dg}, \delta_x^{dg}, \delta_g^{dg}$are the same as in (2.44), and

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

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

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

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

δ_{ut}^{dg} is positive and decreasing over time (see Figure 2.2)[42]. Thus monetary policy should react to the cost push shock in a similar fashion as to out of equilibrium expectations (see Result 4):

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

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

The asymptotic properties of the ALM (2.46),(2.48) depend on the limiting behavior of a_t, which is given by the stochastic recursive algorithm:

$$a_{t+1} = a_t + (t+1)^{-1}\left((c_{\pi t}^{dg} - 1)a_t + d_{\pi,t}^{dg} u_t\right) \tag{2.50}$$

We study its properties in the Appendix, where we use the stochastic approximation techniques[43] to prove the following Proposition:

Proposition 7. *Let a_t evolve according to (2.50); 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 $d_{\pi,t}^{dg} \to \frac{\alpha}{\kappa^2+\alpha}$, so that we can conclude that $\pi_t \to \frac{\alpha}{\kappa^2+\alpha} v$ almost surely, where v is a random variable with the same probability distribution as u_t. The equilibrium corresponds to the discretionary rational expectations equilibrium. Optimal policy 'helps' private agents to learn the rational discretionary REE[44].

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

[42]Since $\delta_{u,t}^{dg} < 1$ from (2.49) it follows that the change of $\delta_{u,t}^{dg}$ through time is identical to that of $\delta_{\pi,t}^{dg}$ and the numerical analysis of Section 2.4.1 also applies here.

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

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

analysis on how the difference during the transition translates into welfare losses. Similarly to Section 2.3.3 we report consumption equivalents[45].

Tables 2.5 and 2.6 show that similarly to the constant gain case in the long run OP engineers a lower consumption equivalent then the EH policy, and OP engineers lower variation of inflation at the cost of allowing higher variation in output. The last row of the first column in Table 2.5 shows that if we start the economy from the steady state, $a_0 = 0$, in the long run the consumption equivalent of OP is about 10% lower then that of EH. Table 2.6 reports reports the composition of these losses: if we start the economy in the steady state optimal policy engineers an inflation variation 20 percent lower then EH and allows a 3-9 times higher welfare loss due to output gap variations.

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

Let us first examine the path of expectations. Both OP and EH are E-stable under learning, so guarantee that expectations converge to the discretionary REE, the difference is the speed of convergence. Figure 2.6 shows a typical realization of the evolution of expectations under OP and EH. We can observe that inflation expectations converge faster and output gap expectations converge more slowly with our rule than with the EH one. This is a consequence of the intertemporal tradeoff (Result 1): when the CB does take into account its influence on the learning algorithm, it has an incentive to undercut future inflation beliefs. The way the central bank can achieve this, is to keep inflation close to the steady state; since inflationary expectations are formed as averages of past inflation data, this policy undercuts future inflation expectations. Because of the intratemporal tradeoff between inflation and output, the cost of keeping inflation close to the steady state is a wider output gap and consequently a slower convergence of b to its RE value.

We report how the ratios of OP and EH consumption equivalents evolve during the transition in Table 2.5. In the first periods the optimal interest rate rule (2.49) yields *ex-post* higher cumulative welfare losses expressed in consumption terms than the EH rule; later, however, our rule starts generating smaller welfare losses. These findings are consistent with our finding that a CB that follows the optimal rule (2.49) reacts to out-of-equilibrium inflation expectations more aggressively than in the EH case, in order to undercut more future expectations, even if this means allowing a wider output gap in the short run. This implies that in the first periods, when this more aggressive behavior has not yet generated a pay-off in terms of a smaller $|a|$ sufficient to offset the costly output gap variability, our

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

rule performs worse than the EH one; as soon as inflation expectations become small enough, this initial disadvantage is more than compensated. This pattern is magnified by the time-varying behavior of $\delta_{\pi,t}^{dg}$ that we characterized above: the coefficient on inflation expectations in (2.49) is particularly large in the first periods, hence determining large output gap variations and large welfare losses in the short run, and large gains from the contraction of $|a|$ in the medium and long run.

Since the main advantage of OP is that it helps private agents' inflationary expectations to converge faster, the advantage of OP over EH increases the further away initial expectations are from the steady state. The different columns of Table 2.5 report ratios of consumption equivalents for different initial inflationary expectations. The higher is a_0 the bigger is the consumption cost of OP compared to EH in the first periods: OP allows for higher welfare losses in order to keep inflation closer to the SS in order to help inflationary expectations converge faster. As time goes on, inflation expectations converge closer to 0 under OP than under EH; the further away a_0 is from the SS, also the further away future inflation expectations remain from the SS under EH. Consequently the inflation output gap tradeoff remains worse under EH and consumption equivalents remain also higher then under OP. The bigger is a_0 the bigger is the gain in decreasing inflation variation of OP over EH, and the higher is the output gap variation OP allows compared to EH (See Table 2.6)[46].

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

2.5 Extensions

Up to now, we have supposed that the CB perfectly observes all the relevant state variables of the system, namely the exogenous shocks and the agents' beliefs. In

[46]Similarly to Section 2.3.3 ratios of consumption equivalents do not depend on the choice of σ_u^2.

Table 2.5: Path of cumulative consumption equivalent ratios under decreasing gain, using OP and EH

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

Woodford (1996) calibration

this section we show that our main results extend to a more general framework, where either the shocks or the expectations are not observable. In particular, to make the problem non-trivial, throughout this section we modify the structural equations (2.1) and (2.25) with the introduction of unobservable shocks, so that the model is now given by:

$$x_t = E_t^* x_{t+1} - \sigma^{-1}(r_t - E_t^* \pi_{t+1} - \overline{rr}_t) + g_t + e_{x,t} \tag{2.51}$$

and:

$$\pi_t = \beta E_t^* \pi_{t+1} + \kappa x_t + u_t + e_{\pi,t} \tag{2.52}$$

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

2.5.1 Measurement Error in the Shocks

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

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

Table 2.6: Ratio of welfare losses using OP and EH under decreasing gain learning due to inflation and output gap variations

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

Woodford (1996) calibration

To make the problem non-trivial, we also assume that the CB can observe π_t and x_t only with a lag. Note that the shocks do not depend on the policy followed by the CB; hence, the *separation principle* applies, namely, the optimization of the welfare criterion and the estimation of the realizations of the shocks can be solved as separate problems. As is well known, the above signal-extraction problem implies that the expected values of the shocks given the signals are[47]:

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

Moreover, the separation principle implies that certainty equivalence holds in designing the optimal interest rate rule, which turns out to be identical to (2.49), with g_t and u_t replaced by $E_t^{CB} g_t$ and $E_t^{CB} u_t$, respectively:

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

We can combine the above equation with (2.51) and (2.52) to obtain the ALM for inflation and output gap:

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

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

where:

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

As a consequence of the measurement error, inflation and output gap now depend on a wider set of state variables; however, it is easy to see that the main findings of the preceding section go through in this modified environment. First of all, the separation principle trivially implies that when the CB takes into account the effect of its decisions on future beliefs, the optimal policy is more aggressive against out-of-equilibrium inflation expectations, compared to the case in which the private sector's expectations are considered as exogenously given[48]; moreover, the analysis of convergence of learning algorithms to the optimal discretionary RE equilibrium[49] does not change in this modified environment.

2.5.2 Heterogenous Forecasts

As argued in Honkapohja and Mitra (2005) (HM hereafter), the hypothesis that the CB can perfectly observe private sector's expectations is subject to several criticisms[50]; it is therefore natural to verify the robustness of our results when this assumption is relaxed. In what follows, we assume that the optimal interest rate rule takes the same form as (2.49), but the agents' forecasts for inflation and output gap, a_t and b_t, are replaced by the CB internal forecasts, a_t^{CB} and b_t^{CB}[51]; in particular, we suppose that the CB and the private sector forecasts have the same form, and are updated according to the same algorithm, which is given by (2.35)-(2.36). The only difference is given by the initial beliefs. Note that this setup corresponds to a situation where the CB, in solving its optimization problem,

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

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

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

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

knows the adaptive algorithm used by the agents to form their expectations, but cannot observe the actual values of these expectations; instead, the CB has a tight prior on a_0 and b_0[52], and forms its internal forecasts accordingly. Plugging the interest rate rule into the structural equations (2.51) and (2.52), we get the ALM:

$$\begin{aligned} \pi_t &= \nu_a^1 a_t + \nu_{a^{CB}t}^1 a_t^{CB} + \nu_b^1 b_t + \nu_{b^{CB}}^1 b_t^{CB} + \nu_{ut}^1 u_t + \kappa e_{x,t} + e_{\pi,t} \\ x_t &= \nu_a^2 a_t + \nu_{a^{CB}t}^2 a_t^{CB} + \nu_b^2 b_t + \nu_{b^{CB}}^2 b_t^{CB} + \nu_{ut}^2 u_t + e_{x,t} \end{aligned} \tag{2.53}$$

where:

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

Again, our main results are unaffected by this change in the CB information set, both for $t < \infty$ and for $t \to \infty$. In fact, since the parameters in the optimal rule are the same as in rule (2.49), the results summarized in Propositions 4 and 6 are still valid. On the other hand, we can study E-stability of the system extending Proposition 2 in HM to a time-varying environment. In particular, it is easy to show[53]:

Corollary 1. *Consider the model (2.53); it is E-stable if and only if the corresponding model with homogenous expectations is E-stable.*

Since E-stability of the homogenous expectations model is ensured by Proposition 7, we conclude that also system (2.53) is E-stable, and it converges to the optimal discretionary RE equilibrium[54].

2.6 Conclusions

In this paper we analyzed the optimal monetary policy problem faced by a CB that tries to exploit its ability to influence future beliefs of the agents, when they follow adaptive learning to form their expectations. We have shown that monetary policy

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

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

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

should be aggressive on inflation, and the reason for this is that in this way private agents learn the true value of steady state inflation faster. We have shown that optimal policy can be implemented by an aggressive interest rate policy, and also that this behavior is optimal even at the cost of higher welfare losses from output gap volatility. We conclude by describing several areas where future research would be useful.

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

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

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

Our analysis was restricted to examining optimal policy given a certain learning algorithm. It would be interesting to examine how monetary should be conducted with endogenous expectation formation, in other words when private agents would change their expectation formation depending on their perception about the underlying economy. Endogenous expectation formation could be formulated for example along the lines of Marcet and Nicolini (2003) where agents dynamically switch between predictor use depending on the last forecast error. An alternative way would be to to model expectation formation as in Molnar (2005) where agents do not switch predictor use, but always a weighted average of predictor forecasts and adjust the weight on predictors dynamically depending on the relative forecasting performance.

2.7 Appendix

2.7.1 Constant Gain Learning

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

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

and let:

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

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

$$\begin{aligned} &\min_{\{\pi_t, x_t, r_t, a_{t+1}, b_{t+1}\} \in G} E_0 \sum_{t=0}^{\infty} \beta^t (\pi_t^2 + \alpha x_t^2) \\ &\text{s.t. (2.1), (2.2), (2.8), (2.9)} \\ &\quad a_0,\ b_0 \text{ given} \end{aligned} \tag{2.54}$$

then $\left\{\pi_t^*, x_t^*, r_t^*, a_{t+1}^*, b_{t+1}^*\right\}$ *solves also (2.10).*

Proof. Let $\left\{\widehat{\pi}_t, \widehat{x}_t, \widehat{r}_t, \widehat{a}_{t+1}, \widehat{b}_{t+1}\right\}$ be an arbitrary unbounded sequence that satisfies the constraints of (2.10), and such that:

$$\widehat{V} \equiv \sum_{t=0}^{\infty} \beta^t (\widehat{\pi}_t^2 + \alpha \widehat{x}_t^2) < \infty \tag{2.55}$$

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

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

and $\left\{\widehat{x}_t^n, \widehat{r}_t^n, \widehat{a}_{t+1}^n, \widehat{b}_{t+1}^n\right\}$ are defined accordingly to respect the constraints of (2.10); clearly, $\left\{\widehat{\pi}_t^n, \widehat{x}_t^n, \widehat{r}_t^n, \widehat{a}_{t+1}^n, \widehat{b}_{t+1}^n\right\}$ is bounded, so that:

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

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

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

if $\lim_{n\to\infty} \widehat{V}^n$ exists. However, it is easy to see that $\lim_{n\to\infty} \widehat{V}^n = \widehat{V}$; since $\left\{\widehat{\pi}_t, \widehat{x}_t, \widehat{r}_t, \widehat{a}_{t+1}, \widehat{b}_{t+1}\right\}$ was arbitrary, it proves the statement [55]. □

[55] Note that the condition (2.55) can be imposed without any loss of generality, since any $\left\{\widehat{\pi}_t, \widehat{x}_t, \widehat{r}_t, \widehat{a}_{t+1}, \widehat{b}_{t+1}\right\}$ that does not respect it, for sure cannot do better than $\left\{\pi_t^*, x_t^*, r_t^*, a_{t+1}^*, b_{t+1}^*\right\}$.

Lemma 2. *Let A_{11} be given by equation (2.21) in the text; then it has an eigenvalue inside and one outside the unit circle.*

Proof. First of all, we recall a result of linear algebra that we will use in the proof, i.e. that a necessary and sufficient condition for a 2 by 2 matrix to have an eigenvalue inside and one outside the unit circle, is that[56]:

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

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

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

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:

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

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, since β and γ are supposed smaller than one. □

We now prove Proposition 1. First of all, we can guess that inflation follows the ALM (2.22)[57] and use the optimality condition (2.18) and the method of undetermined coefficients to verify that c_π^{cg} must satisfy the following quadratic expression:

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

[56]LaSalle (1986).

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

where:

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

The above polynomial can be equivalently rewritten as follows:

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

We will prove that the function $f(\cdot)$, defined on the interval $[0,1]$, is a contraction, so that it admits one and only one fixed point; moreover, since the two roots of the quadratic expression have the same sign (it is due to the fact that both p_2 and p_0 are positive), it follows that the other candidate value for c_π^{cg} is greater than one, which is not compatible with the boundary conditions[58].

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

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

Proof. Note that:

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

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

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

After rearranging:

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

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

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

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

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

[58] Since it would imply an exploding inflation.

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

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

After rearranging, we get:

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

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

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

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

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

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

Proof. First of all, note that:

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

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

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

Since $\beta, \gamma \in (0, 1)$, 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). $\square$

Moreover, we prove the following result.

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

Proof. Note that:

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

if and only if:

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

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 $\gamma>0$, since the $\frac{\alpha\beta}{\alpha+\kappa^2}<\beta$, the LHS of the above inequality is smaller than:

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

which is equal to one; in fact:

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

is equivalent to:

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

But the LHS can simplified as:

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

which is equal to:

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

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

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

which implies that:

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

□

We are now ready to prove the Proposition.

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

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

2.7.2 Decreasing Gain Learning

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

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

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

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

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

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

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

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

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

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

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

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

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

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

We prove the second part of the statement by contradiction. Assume that there exists a $T < \infty$ such that $c^{dg}_{\pi,t} \geq \frac{\alpha\beta}{\alpha+\kappa^2}$; 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:

$$\left(1-\frac{1}{T+1}\left(1-\frac{\alpha\beta}{\alpha+\kappa^2}\right)\right)c^{dg}_{\pi,t+1} \geq \frac{\alpha\beta}{\alpha+\kappa^2}A_{11,T}+A_{12,T} \tag{2.56}$$

Note that the RHS is equal to:

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

where the last inequality is due to the fact that $\left(1+\beta\frac{1}{t+1}\right)^{-1} < 1$; putting together the last inequality and (2.56), we get:

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

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

□

Proof of Proposition 5. Recall that, as shown in Proposition 3, we have $\lim_{t\to\infty} c^{dg}_{\pi,t} = \frac{\alpha\beta}{\alpha+\kappa^2}$; since $0 < \frac{\alpha\beta}{\alpha+\kappa^2} < 1$, for any $\overline{C}$ with $\frac{\alpha\beta}{\alpha+\kappa^2} < \overline{C} < 1$, there exists a T such that, for any $t \geq T$ we will have $0 < c^{dg}_{\pi,t} < \overline{C}$; moreover, using the ALM for π_t, the law of motion of inflation expectations after T can be rewritten as[59]:

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

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

Finally, we prove Proposition 7. First of all, we will briefly describe some results of stochastic approximation[60] that we will exploit in the proof.

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

$$\theta_t = \theta_{t-1} + \gamma_t Q\left(t, \theta_{t-1}, X_t\right) \tag{2.57}$$

where X_t is a state vector with an invariant limiting distribution, and γ_t is a sequence of gains; the stochastic approximation literature shows how, provided certain technical conditions are met, the asymptotic behavior of the stochastic difference equation (2.57) can be analyzed using the associated deterministic ODE:

$$\frac{d\theta}{d\tau} = h\left(\theta(\tau)\right) \tag{2.58}$$

where:

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

E represents the expectations taken over the invariant limiting distribution of X_t, for any fixed θ. In particular, it can be shown that the set of limiting points of (2.57) is given by the stable resting points of the ODE (2.58).

Proof of Proposition 7. Note that our equation (2.50) is a special case of (2.57), where the technical conditions are easily shown to be satisfied; moreover, it is also easy to see that:

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

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

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

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

2.7.3 Comparison with EH Rule

Proof of Propositions 4 and 6. 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:

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

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

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

which implies, since

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

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

2.7.4 Derivations of Consumption Equivalents

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

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

$$U = -\bar{Y}U_c L^P \ , \tag{2.59}$$

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

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

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

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

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

where U_{cc} is the second derivative of utility of utility with respect to consumption evaluated at $\bar{Y}$. Equating this utility loss to (2.59), the welfare loss generated under policy rule P gives

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

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

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

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

$$p=100\sigma\left(-1+\sqrt{1+\frac{2(1-\beta)(L^P)\sigma}{100^2}}\right) . \tag{2.60}$$

2.8 Figures

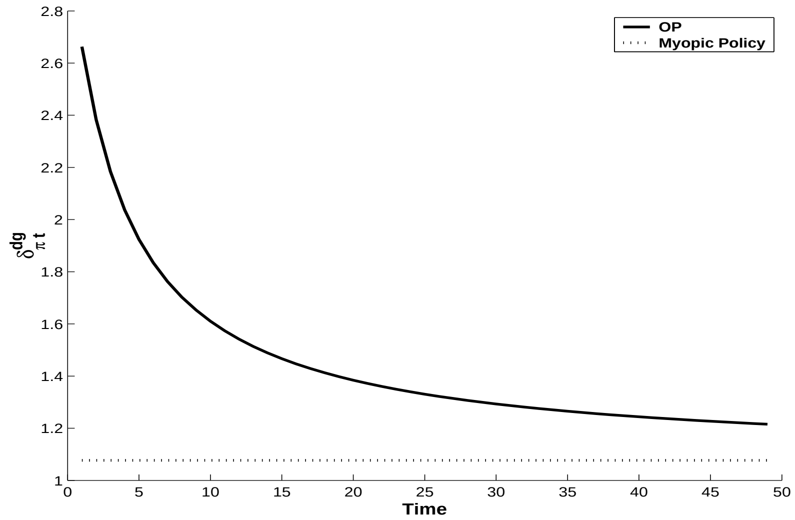

Figure 2.1: Interest rate rule coefficient on inflation expectations under decreasing gain learning.

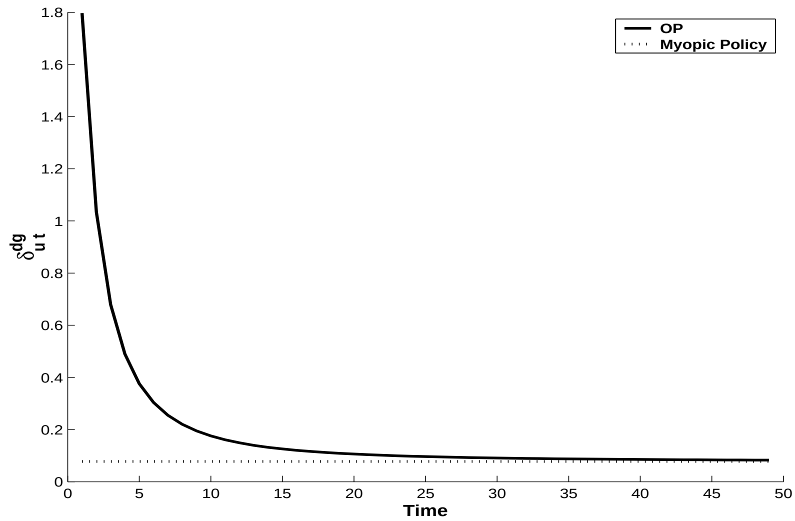

Figure 2.2: Optimal versus myopic interest rate rule: coefficient of the cost push shock

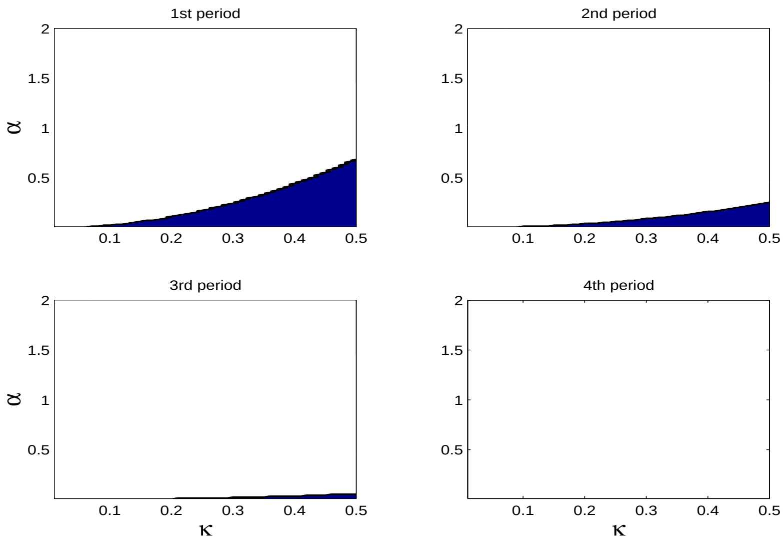

Figure 2.3: Values of α and κ for which δ_π^{dg} is increasing in the first 4 periods. From the 4th period on δ_π^{dg} is always decreasing. ($\beta = 0.99$)

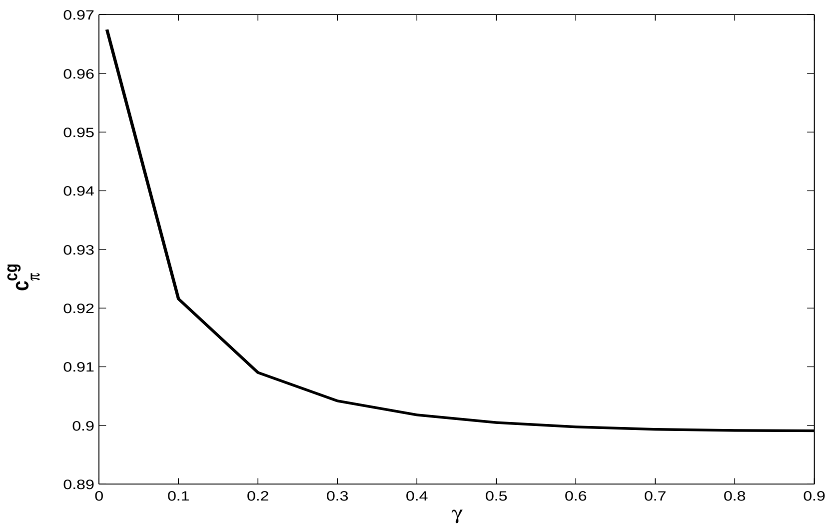

Figure 2.4: Feedback parameter in the ALM for inflation as a function of γ.

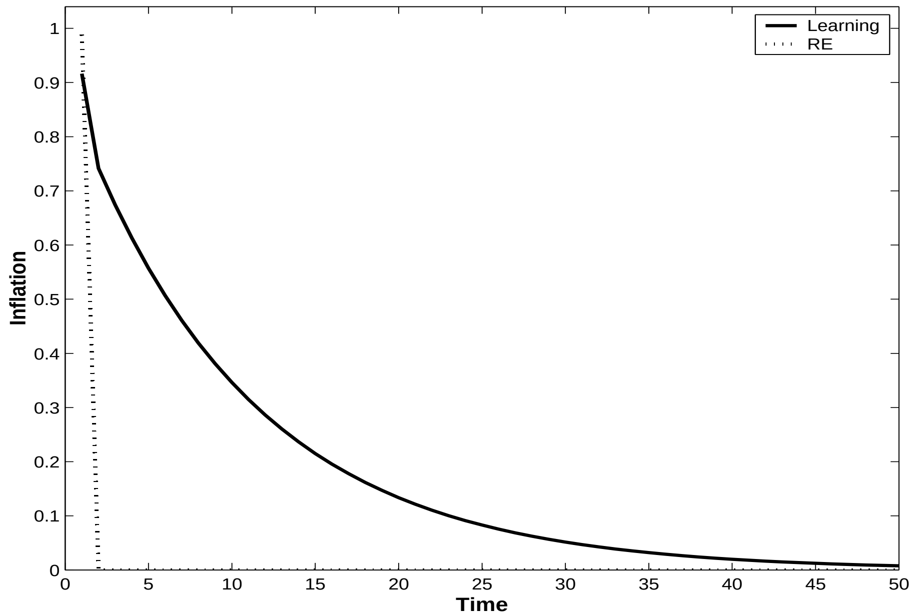

Figure 2.5: Impulse response of inflation for an initial cost-push shock $u = 1$. Solid line: optimal policy under learning and private agents following learning with $\gamma = 0.9$. Dashed line: optimal discretionary policy under RE with private agents have rational expectations. Initial conditions: $a_0 = 0$, $\pi_0 = 0$, $x_0 = 0$.

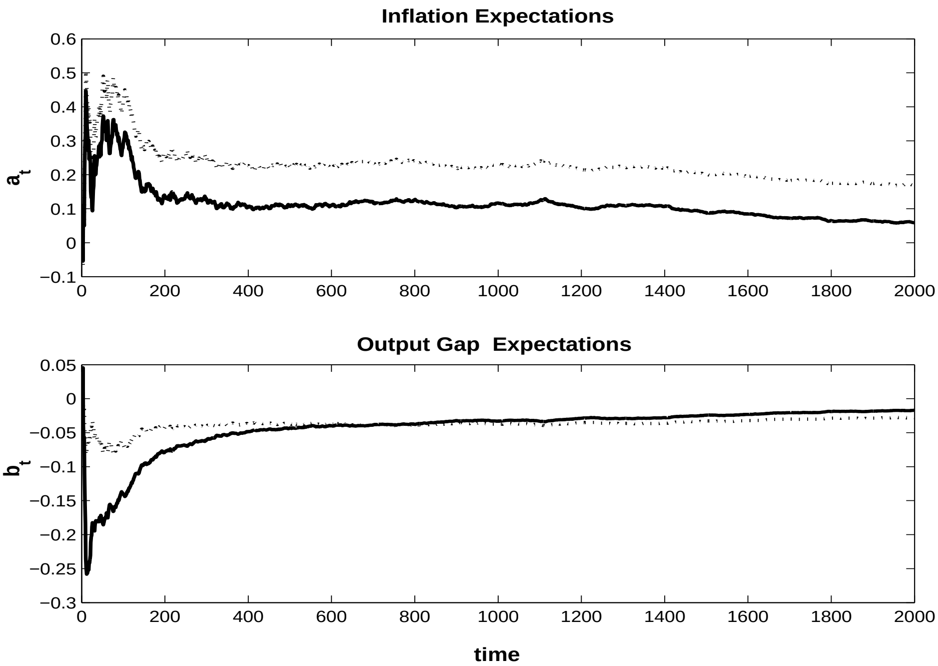

Figure 2.6: Evolution inflation and output gap expectations under the optimal (solid line) the myopic rule (dashed line), when private agents follow decreasing gain learning

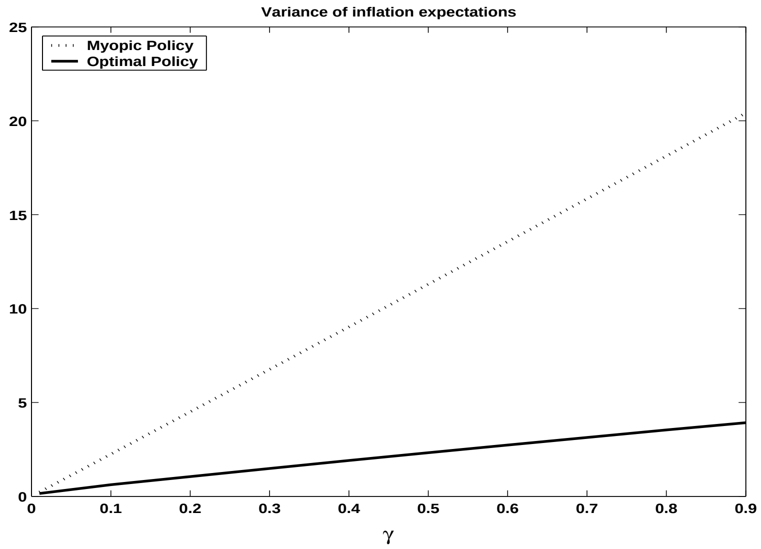

Figure 2.7: Variance of inflationary expectations

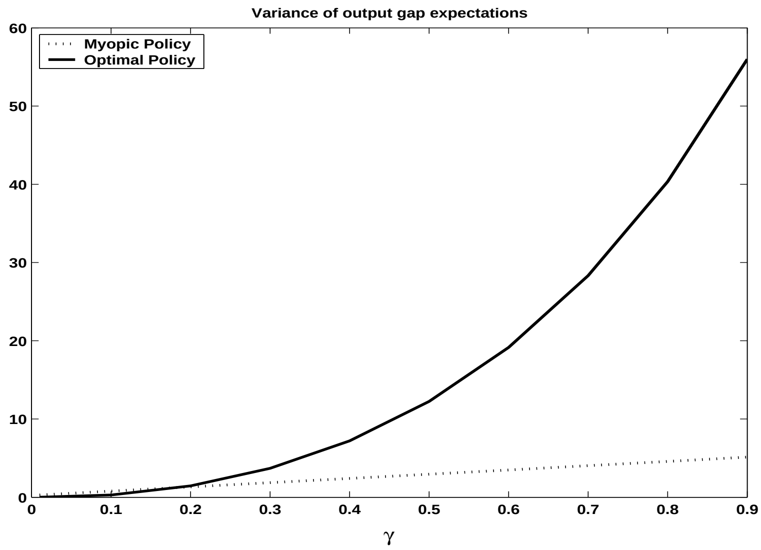

Figure 2.8: Variance of output gap expectations

Bibliography

ADAM, K. (2004): "Learning to Forecast and Cyclical Behavior of Output and Inflation," *forthcoming Macroeconomic Dynamics.*

ADAM, K., AND R. BILLI (2006): "Discretionary Monetary Policy and the Zero Lower Bound on Nominal Interest Rates," Discussion paper.

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

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

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

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

BAAK, S. J. (1999): "Tests for bounded rationality with a linear dynamic model distorted by heterogeneous expectations," *Journal of Economic Dynamics and Control*, 23.

BAGHESTANI, H. (1992): "Survey Evidence on the Muthian Rationality of the Inflation Forecasts of U.S. Consumers," *Oxford Bulletin of Economics and Statistics*, 54(2), 173–86.

BALL, L., AND D. CROUSHORE (1995): "Expectations and the effects of monetary policy," *NBER Working Paper*, (W5344).

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

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

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

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

Branch, W. A. (2002): "Local convergence properties of a cobweb model with rationally heterogeneous expectations," *The Economic Journal*, 27, 63–85.

——— (2004): "The Theory of Rationally Heterogeneous Expectations: Evidence from Survey Data on Inflation Expectations," *The Economic Journal*, 114.

Branch, W. A., and G. W. Evans (2004): "Intrinsic Heterogenity in Expectation Formation," .

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

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

Brock, W. A., and C. H. Hommes (1997): "A Rational Route to Randomness," *Econometrica*, 65(5), 1059–1096, available at http://ideas.repec.org/a/ecm/emetrp/v65y1997i5p1059-1096.html.

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

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

Caroll, C. D. (2003): "Macroeconomic expectations of households and professional forecasters," *Quarterly Journal of Economics*, 118(1), 269–298.

Cesa-Bianchi, N., and G. Lugosi (2006): *Prediction, Learning, and Games*. Cambridge University Press.

Chavas, J.-P. (2000): "On information and market dynamics: The case of the U.S. beef market," *Journal of Economic Dynamics and Control*, 24.

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

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

Englander, A. S., and G. Stone (1989): "Inflation expectations surveys as predictors of inflation and behavior in financial and labor markets," *Federal Reserve Bank of New York Quarterly Review*, pp. 20–32, roberts97 p178.

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

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

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

——— (2001a): "Expectations and the Stability Problem for Optimal Monetary Policies," *University of Oregon Economics Department Working Papers*, (2001-6).

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

LOVELL, M. C. (1986): "Tests of the Rational Expectations Hypothesis," *The American Economic Review*, 76,1, 110–124.

LUCAS, R. E. (1987): *Models of Business Cycles.* Oxford University Press.

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

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

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

——— (1995): *Speed of Convergence of Recursive Least Squares: Learning with Autoregressive Moving-Average Perceptions In Learning and Rationality in Economics.* Blackwell Publishers.

Marimon, R. (1996): "Learning from learning in economics," in *Advances in Economic Theory*, ed. by D. Kreps, and K.F.Wallis. Cambridge University Press.

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

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

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

Nunes, R. (2004): "Learning the Inflation Target," .

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

Orphanides, A., and J. C. Williams (2004): "The Decline of Activist Stabilization Policy: Natural Rate Misperceptions, Learning, and. Expectations," *Board of Governors of the Federal Reserve System (U.S.),International Finance Discussion Papers*, (804).

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

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

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

ROBERTS, J. M. (1995): "New Keynesian Economics and the Phillips Curve," *Journal of Money, Credit and Banking*, 27, 975–984.

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

——— (1998): "Inflation expectations and the transmission of monetary policy," Federal Reserve Board.

——— (2001): "How Well Does the New Keynesian Sticky-Price Model Fit the Data?," *The Federal Reserve Board Finance and Economics Discussion Series*, (13).

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

SARGENT, T. J. (1993): *Bounded Rationality in Macroeconomics*. Oxford University Press.

——— (1999): *The Conquest of American Inflation*. Princeton University Press.

SETHI, R., AND R. FRANKE (1995): "Behavioural Heterogeneity under Evolutionary Pressure: Macroeconomic Implications of Costly Optimisation," *Economic Journal*, 105(430), 583–600, available at http://ideas.repec.org/a/ecj/econjl/v105y1995i430p583-600.html.

TIMMERMANN, A. G. (1996): "Excess Volatility and Predictability of Stock Prices in Autoregressive Dividend Models with Learning," *Review of Economic Studies*, 63.

VIVES, X. (1993): "How Fast do Rational Agents Learn?," *Review of Economic Studies*, 60, 329–347.

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

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

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

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

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