Learning the Ramsey outcome in a Kydland & Prescott economy

Jasmina Arifovic and **Murat Yıldızoğlu***
Department of economics, Simon Fraser University, arifovic@sfu.ca
GREThA (UMR CNRS 5113), Bordeaux University, murat.yildizoglu@u-bordeaux.fr

November 16, 2017

Abstract

We study if adaptive learning by a Central Bank (CB) in the Kydland and Prescott environment can steer the economy to the Pareto-optimal outcome. Our CB evaluates its potential strategies regarding the announced and the actual inflation rate through expectations of the performance of these strategies, formed thanks to its *mental model* of the economy. This model is forward looking and adaptive at the same time. As a starting point, we follow Arifovic & al. (2010), and initially assume that there are two types of agents: Believers who set their inflation forecast equal to the announced inflation, and Non-believers who form static optimal forecast coupled with a forecast error correction mechanism. Our results show that the economy can reach near Ramsey outcomes most of the time. In the absence of Believers, the economies almost always converge to the Ramsey outcome.

In their experiments with human subjects, Arifovic and Sargent (2003) showed that experimental economies reach and stay close to the Ramsey outcome most of the time, giving support to the 'just do it' policy recommendation . In light of the experimental findings, our model is of particular interest as it is the only agent-based or adaptive learning model that consistently selects the Ramsey outcome in this context.

JEL categories: E50, C45, C72, D60
Keywords: Learning in the Kydland-Prescott environment; Expectations; Artificial neural networks; Credibility of economic policies; Convergence to Nash equilibrium; Ramsey outcome

*Corresponding author. GREThA (UMR CNRS 5113). Université de Bordeaux. Avenue Léon Duguit. 33608 Pessac. France.

1 Introduction

The objective of this paper is to study whether a boundedly rational policy maker that uses only past information and adaptive expectations, sophisticated enough not to be subject to the Lucas critique, could steer the economy towards the Pareto optimal solution. We study the issue in the framework of the original Kydland & Prescott model, where this question has initially been given a strong and negative answer in an economy composed of agents with substantive rationality and perfect foresight.

The economy considered in Kydland & Prescott (1977) is a simple dynamic game between the Central Bank (CB) and the private agents where the CB's objective is to minimize the social loss of inflation and unemployment. The CB sets the inflation rate while the unemployment depends on agents' expectational errors regarding the actual inflation rate and is determined through a standard Phillips curve. If agents have perfect foresight, the expected inflation rate is equal to the actual inflation rate, and the unemployment is equal to its natural level. The Nash equilibrium of this stage game is an outcome where the unemployment rate is equal to its natural level and the inflation rate is greater than zero. However, in the Pareto optimal, *Ramsey* outcome, while the unemployment is still equal to its natural rate, the rate of inflation is equal to zero. While socially desirable, the Ramsey outcome is not a Nash equilibrium of this model, and the Ramsey inflation rate is also *time inconsistent.*[1]

Basically, a solution of a game is time inconsistent if the policy maker has incentives to deviate ex post from decisions that were initially optimal. In our setting, ex ante, the CB would like to choose zero inflation rate. However, assuming that agents 'believe' it and form their expectations accordingly, the CB has an incentive to 'deviate and inflate'. This surprise inflation will lower the unemployment rate through the Phillips curve channel. Rational agents are aware of this temptation and will never form zero inflation forecast. As a result, the economy ends up in a Pareto inferior, Nash equilibrium.

Barro and Gordon (1983) found better reputational equilibria when the game is repeated over time. With repetition, the availability of history-dependent strategies multiplies the range of equilibrium outcomes, some better than the one-period time-consistent one, others worse. Following Kydland and Prescott (1977) and Barro and Gordon (1983), the issue has been extensively studied in the monetary policy literature. One strand is the reputational macroeconomics that studies subgame perfect equilibria with history dependent strategies and results in a set of equilibria. For a high enough discount factor, one of the equilibrium paths is the Ramsey outcome that repeats itself for ever. However there are others, some of which result in outcomes worse than Nash.[2]. Another strand investigated the possible use of non binding policy announcements to improve on the Nash solution. Most of them assume hidden information about the type of the CB or the state of the economy. Then, the CB can use a non binding policy

[1] Kydland & Prescott were the first to point out the problem of time inconsistency. Its potential consequences have been a recurrent issue in many areas of economics. It arises when a decision maker, especially a policy maker, prefers one policy in advance, but a different one when the time to implement arrives. Knowing this, others will not find the commitment to the first policy credible.

[2] For a short survey of this literature, see Stokey (1989)

announcement to provide a signal about her private information.[3] Thus, observing the announcement allows for a better prediction of the CB's decision about the actual inflation rate. Dawid and Deisssenberg (2005) study a different model to examine the role of cheap talk inflation announcements. In their continuous time framework there is a continuum of atomistic private agents (households) of two types, Believers and Non-believers and the CB that is able to correctly predict the impact of its inflation announcements on agents' forecasts. If the initial fraction of Believers is sufficiently high, the CB can get the economy to the steady state that is Pareto superior to the Nash solution. The existence of Believers helps economy achieve a better outcome than the Nash equilibrium.

There is also literature on how learning affects the outcomes of the Kydland & Prescott environments, see, for example, Sargent (1999), Cho et al. (2002), and Cho and Sargent (1997). This literature uses stochastic approximation algorithms such as recursive least squares, constant gain and stochastic gradient learning. In different setups, either policy maker is learning or both policy maker and private agents are learning using these algorithms. The learning outcomes narrow the set of equilibria down to the time consistent but Pareto inferior Nash equilibrium, and to the Pareto optimal but time inconsistent Ramsey outcome. However, Nash equilibrium is the most frequent outcome while the Ramsey outcome rarely occurs.

When theoretical and learning models result in a large number of equilibrium predictions and outcomes, experiments with human subjects might be successful in selecting among these. Arifovic and Sargent (2003) conducted experiments with human subjects in an infinitely (indefinitely) repeated Kydland & Prescott environment. One of the randomly selected subjects played the role of the CB and set a target inflation rate in each period. Other subjects played the role of the public, and each subject forecasted the inflation rate in each period. The unemployment was then generated using a standard Phillips curve. The results of their experiments showed that the Ramsey outcome emerged most of the time. The experimental economies only occasionally slid into the Nash equilibrium.

Our work is most closely related to Arifovic et al. (2010) who build upon Dawid and Deissenberg (2005), and develop a much richer model in which the fractions of two types of agents, Believers and Non-believers, change over time depending on how well they are performing in terms of their forecasting accuracy.[4] Non-believers use an adaptive learning scheme to revise their forecasts in each time period. The CB is boundedly rational. It has a collection of strategies that evolve over time based on the evaluation of their foregone payoffs.[5] In each time period, it chooses, probabilistically, the strategy that it actually implements. Thus, the CB does not have a priori knowledge of how its

[3]See for example, Stein (1989), Cukierman (1992), Walsh (1999), Persson and Tabellini (1993).

[4]In each period, *word of mouth* information exchange allows a fraction of the agents to compare their last-period payoffs with the ones obtained by agents who followed the other strategy. Each agent then adopts the strategy that provided the highest payoff, and uses it until a new comparison motivates it to switch strategies.

[5]They implement a version of the Individual Evolutionary Learning Algorithm (see Arifovic and Ledyard 2004).

behavior impacts agents' expectations. The CB learns how to sustain the economy with a positive fraction of Believers and 'better than Nash' outcomes. In addition, Arifovic et al.'s analysis shows that these Pareto superior outcomes also require a sufficient number of Non-believers that have to adjust quickly and not have too high of a cost of acquiring information. The economy has to have a right degree of heterogeneity to maintain 'better than Nash' outcomes and provide enough low cost information to those who are adjusting their forecast. We further discuss below the extensions of their idea.

In the current paper, we develop a model (an agent–based model) inspired by Arifovic et al (2010), but in which the CB chooses now strategies on the basis of its adaptive expectations about the future consequences of these strategies. For the purpose of comparison, we also incorporate two types of private agents, Believers and Non-believers in the baseline version of the model.

As in Arifovic et al., the CB updates a collection of strategies in each period. A strategy consists of two elements, the value of the inflation rate that they announce to the public, and the value of the inflation rate to be actually implemented. The collection of strategies evolves following a simple individual learning process based on replication and combination of the already discovered strategies, as well as random experimentation that introduces new strategies (Arifovic and Ledyard, 2004). However, we depart from Arifovic et al. in a very important way in how we evaluate and choose the CB's strategies. Our CB learns about the future consequences of its decisions through the adaptation of its representation of how the economy works, i.e. its *mental model*. In order to model this CB's decision making process and updating, we follow the approach developed in Yildizoglu (2001) and Yildizoglu et al. (2014): the *mental model* (Holland et al., 1989) of the CB is represented as an artificial neural network (**ANN**) that evolves in response to the experience that the CB gathers about the economy. On the basis of the expectations formed using this mental model, the CB selects in each period, from its collection, the strategy with the highest expected future performance. Using this mental model, the CB can also form expectations about the performance of new potential strategies that have never been used by it before.

Our approach hence extends the evolutionary search process in the strategy space, by a forward looking dimension that involves evaluation of strategies based on their expected performance. This is the only type of a setup that allows a possibility of 'learning' of the future consequences of the current decisions in each period. This learning plays a crucial role in what strategy the CB chooses in each time period. At the beginning, expectations maybe quite different from the actual outcomes. However, observed errors between the expected and actual performance are used to correct the CB's representation of the economy. Its learning results in emergence of more accurate expectations over time.

We use this framework to test the ability of the CB to take the economy towards the Pareto optimal outcome. We also consider learning environments more complex than the initial setup, including one where the expectations of the CB and the households can co-evolve.

Our main result shows that this modeling approach where the CB is actively involved in trying to learn about the future consequences of its current decisions results in the

convergence of the economies to the neighborhood of the Ramsey outcome. Using its mental model it learns what the Non-believers' complete reaction function is and thus, can take into account the future reaction of all agents (Believers and Non-believers) to its current behavior. More remarkably, it takes into account what the future response of Non-believers (including correction of their errors) will be if it deviates from the announced inflation rate. Finally, the economy performs 'better' in the absence of Believers, i.e. it reaches Ramsey outcome with a much higher frequency.

In comparison to Arifovic et al., our results differ in two important ways: While Arifovic et al. observe improvements upon Nash, our economies spend most of the time in the neighborhood of the Ramsey outcome. Moreover, we show that our results do not depend on the existence of two types of agents. In fact, our economies do better in the absence of Believers in terms of spending more time close to the Ramsey outcome.

Regarding the broader impact, this is the first model that we are aware of that clearly results in the outcomes very close to Ramsey, and shows that adaptive economic policies, driven with enough forward-looking, can be effective.

The rest of the paper is organized as follows. In section 2, we describe the structure of the model, with a focus on the new learning scheme with expectations that we adopt for the behavior of the CB. We provide a description of our simulation protocol in section 3 and we describe and discuss our main results in section 4, where we show that the CB is able to steer the economy towards the Ramsey outcome. Section 5 explores the abilities of the CB in more complex learning environments, and also studies the dynamics of the economy when the households can also use dynamic adaptive expectations. Last section concludes the article.

2 The model

2.1 A very simple economy

We model a simple dynamic game between the Central Bank (CB) and the private agents, *à la* Kydland and Prescott (1977) and Barro and Gordon (1983), where the CB aims to maximize the social welfare, or minimize the social loss resulting from the unemployment (u) and actual inflation (y) rates of the economy:

$$J^G = -\frac{1}{2}\left(u^2 + y^2\right) \tag{1}$$

The CB evaluates its performance in terms of this social welfare.

The CB is assumed to have the means to control the inflation level of the economy, but the unemployment level results from the expectation errors of the households. Two potential outcomes are then possible in this economy: with non-cooperative rational agents and perfect foresight the outcome corresponds to a Nash equilibrium where the unemployment rate is the natural one $\left(u^N = u^*\right)$, but the CB is not able to eliminate inflation from the economy $\left(y^N = u^*/2 > 0\right)$. An outcome with superior social performance, with the same rate of unemployment $\left(u^R = u^*\right)$, but without any inflation $\left(y^R = 0\right)$ is also possible: the Ramsey outcome. Unfortunately the latter is not time-consistent, (it is

not a non-cooperative equilibrium of this game), and the ability of the economic policy to drive the economy towards the socially preferable (Pareto optimal) outcome becomes very weak in this setting. This strong conclusion had important consequences on our vision of the credibility of economic policies, but it has also been questioned by a vast literature. The main question is on the dependance of this conclusion on the assumptions adopted in these simple models, especially concerning the rationality of behavior and expectations of the agents.

Our model starts from the most recent contribution to this debate that we are aware of: Arifovic et al. (2010), which introduce bounded rationality in the behavior of the CB and the households, in a Kydland & Prescott type economy where the CB can use two policy instruments to influence the economic performance: the inflation rate announced at the beginning of each time period, and the inflation rate actually implemented at the end of the period. In Arifovic et al. (2010) households use one of the two simple strategies to form their inflation expectations: either they can just follow the announcement of the CB on the inflation rate when they form their expectations (such households are called "Believers"–B), or they can base their expectations on their best reply to the CB's inflation announcement in the static game, and correct it when they make expectation errors (such households are called "Non-believers"– NB). The actual economic performance depends on the inflation set by the CB, and the unemployment rate of the economy that results from the expectation errors of the households. The fraction of Believers in the population, denoted as ϕ, evolves as a result of an imitation process between households.

Following the notations of Arifovic et al. (2010), we can describe this setting in the following way. Households form inflation expectations (x) after having observed the inflation rate announced by the CB (y^a).

Each *Believer* sets their expectations, x^B, simply equal to the announced inflation y^a:

$$x^B = y^a$$

and each *Non-believer* (**NB**) forms adaptive expectations by correcting its forecast errors around the optimal solution of the static game:

$$x_t^{NB} = \frac{\theta^2 \phi y_t^a + \theta u^*}{1 + \theta^2 \phi} + d_t \tag{2}$$

where d_t is the correction term that takes into account the forecast error of the previous period:

$$d_{t+1} = d_t + \gamma \left(y_t - x_t \right), \; d_0 = d_0 = 0 \tag{3}$$

and $\gamma > 0$ represents the speed of learning of the households. Parameters θ and ϕ are introduced respectively in equations (5) and (6) below.

The utility of each agent depends on its expectational error, the actual inflation rate and the cost of forming a forecast, in the case of the Non-believers

$$J^i = -\frac{1}{2} \left[\left(y - x^i \right)^2 + y^2 \right] - c^i \tag{4}$$

$c^i \geq 0$ if $i \in NB$, and $c^i = 0$ if $i \in B$.

The unemployment rate of the economy results from expectation errors of the households, following an augmented Philips curve

$$u = u^* - \theta\,(y - \overline{x}) \tag{5}$$

where u^* is the unemployment rate that would arise if agents had correct inflation expectations, $\theta > 0$ is the Phillip's curve parameter, and $\overline{x}$ represents the average expectation of the households

$$\bar{x} = \phi x^B + (1 - \phi) x^{NB} \tag{6}$$

when ϕ represents the share of the Believers in the population, and NB households share the same forecasts.

The unemployment rate hence depends on the composition of the agent population.

2.2 The CB's problem

The aim of the CB is to choose a strategy (y^a, y) that would minimize the welfare loss that could result from inflation and unemployment (see equation 1). But it cannot observe or compute in advance the expectation errors of the households. It must hence form expectations on the economic performance that would result from different available strategies. In Arifovic et al. (2010), the CB explores the strategy space using an evolutionary algorithm, where it computes the potential fitness of each strategy using the strategies that have been used by the households in the last period. Consequently, its adaptive expectations are prone to the Lucas critique.

We consider here that the CB could try to form expectations in a more forward–looking way on the economic consequences of its strategies, trying to anticipate the reaction of the agents to these strategies. The CB can do it if it has a *mental model* representing its economic environment, where it can simulate the potential outcome resulting from each available strategy. In this case, its would choose for each period the strategy corresponding to the highest expected performance. It could consequently observe the actual economic performance that results from this strategy and the reactions of the agents to it. Any discrepancy between the initial expected performance and the actual performance would correspond to expectations errors, and the CB could use them to adjust its mental model with the aim of reducing such errors in the future, by better adapting its representation of the economy to the actual state of the latter. In this case, even if the expectations are only based on past observations (and can consequently be called *adaptive*), they would stand up to the Lucas critique, since if the World changes, so would do the mental model of the CB, and the expectations could catch with the state of the economy, instead of running behind it. More generally, if the behavior of the agents is guided by such dynamic forward–looking expectations, any policy changes would be integrated to their mental model, and hence to their behavior.

Obviously, in this context, the concept of a "model" must be understood in a very loose sense. More than a mathematical construction, it is a representation of the CB's perception of the environment: "In (. . .) situations that are not sufficiently simple as to

be transparent to human mind, we must expect that the mind will use such imperfect information as it has, will simplify and represent the situation as it can, and make such calculations as are within its powers" (Simon, 1976, p. 144). These calculations are "As if" experiments that enable the CB to evaluate the possible consequences of its decisions. In other words, before making a decision, the CB simulates the potential outcomes of different available strategies by using its mental model. The output of these simulations provides the expectations of the CB. The latter makes a decision on the basis of these expectations. This decision yields an actual outcome, which can be compared with the expected one. Discrepancies between those outcomes may lead to an update of the mental model. Hence, we have a dynamic structure which evolves as depicted by Figure 1.

How could we take into account such dynamic adaptive expectations through a mental model?

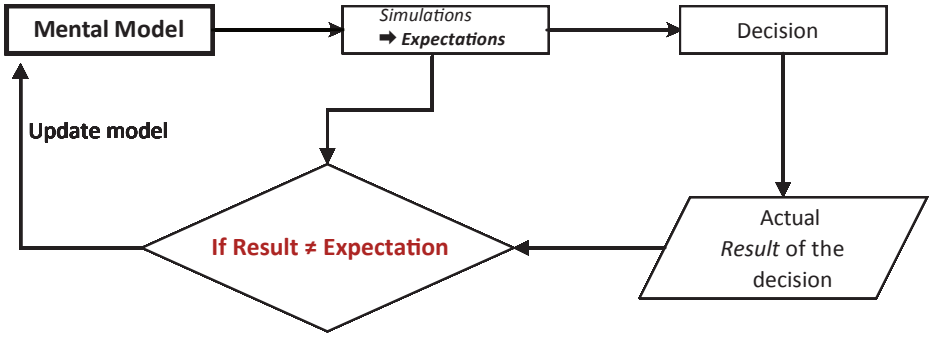

Figure 1: Dynamics of the mental model of the agents (Yildizoglu 2001)

2.3 Mental model, adaptive expectations and forward–looking learning of the CB

In order to introduce such dynamic adaptive expectations, we follow the approach developed in Yildizoglu (2001), and Yildizoglu et al. (2014).

In the framework of the Kydland & Prescott economy, if the CB considers a strategy $y_t = (y^a, y)$, to implement for the following period $t+1$, it would like to have an idea about its expected performance, $J_{t+1}^{G,e}$ in order to compare it with other available strategies. Even if it does not have any information about the exact structure of the economy, and the future behavior of the agents, it can try to infer from its past decisions and experience in this economy, how the latter could react to y_t, and what would be the performance resulting from this reaction. In order to make this inference, it must be able to *generalize* from the past experience to potential future outcomes, and form expectations using some personal representation of the workings of this economy: its *mental model* of the economy that connects the current state of the economy (including past decisions of the CB, and the corresponding economic performance) and the strategy

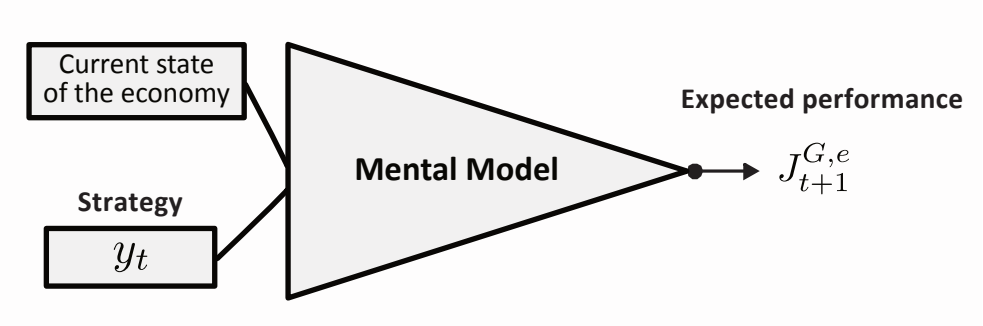

Figure 2: Mental model and expectations

y_t under consideration, to the expected future performance , $J_{t+1}^{G,e}$ (see Figure 2).

On could imagine that this inference could be based on an econometric model of the economy that the agents could build and correct following their past experience in the economy. Recursive least square estimations have indeed been used, in this perspective, albeit at the aggregate level, by the macroeconomic learning literature (see Evans and Honkapohja, 2001, for a modern synthesis on this literature). However, this method makes strong assumptions, since it relies upon a specific (and linear) functional structure for the mental model.

We adopt, here, a more flexible tool, better adapted for a more general context and modeling. Our approach is independent of the structure and the parametrization of the mental model, in order to incorporate only its most primitive dimensions: its existence and its influence on the decisions of the CB. We propose to use artificial neural networks (ANN), because we think that an ANN is a good candidate for representing the role of the mental model, and its adaptive nature. An ANN can indeed represent the connections that the agent can have in mind between the state of the economy and its strategies (represented together as the inputs of the mental model), to the expected performance of the strategy (represented as the output of the mental model). With only minimal structural assumptions, namely the list of endogenous (outputs) and explicative variables (inputs), and the structure of the *hidden layer* (see below), the ANN can also easily represent the fact that the CB adjusts its mental model to the flow of experience. The specific structure is not very important in our high-level approach, because for many practical problems even a very simple feed forward ANN with one hidden layer of few hidden nodes gives quite robust representations (*see* Masters (1993), for the discussion of properties of ANNs).

More particularly, an ANN here is composed of *input nodes*, corresponding to variables indicating the state of the economy and the strategies to connect to the expectations of performance, an *output node* (here we consider only one output: the expected performance of each available strategy), and intermediate nodes, grouped on a *hidden level* of the network, representing other variables that agent can have in mind on the path from the inputs to the output. All nodes but the input ones aggregate the signal they receive and send a signal to the following nodes. This aggregation is non-linear because

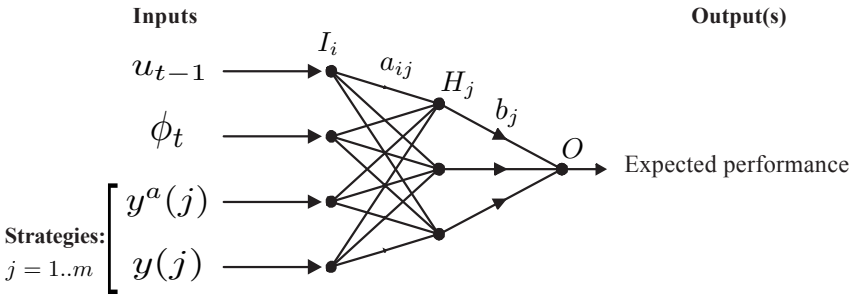

Figure 3: A feed forward ANN with one hidden layer (and with only two inputs for illustration)

the linear combination of the signals received by the node is submitted to a non-linear *activation function*, $\psi(\bullet)$ (generally a sigmoid, thanks to its simple derivative) to compute the signal that will be emitted by the node itself. For example, in the case of a hidden node H_2, connected to two input nodes I_1 and I_2 , its output signal o_j could be given by:

$$o_j = \psi\left(a_{12}I_1 + a_{22}I_2\right), \text{ where } \psi(x) = \frac{1}{1+e^{-x}} \tag{7}$$

a_{ij} correspond to the weight of each incoming signal in the activation of the receiving node. These weights are randomly initialized and they are continuously corrected by the ANN using the local expectation error it observes at the level of this node, between the output given by it, as a function of its current state, and the output it should have given if it was a perfect representation of the connection between its inputs and its output. This correction process is called the *training* of the ANN.

Consequently, a complete ANN, composed of such nodes, and trained by a continuous flow of observations, provides a time varying flexible functional form that delivers an approximation of the connections between the inputs and the output of the mental model. This approximation is obtained by the calibration of the parameters of the ANN (a_{ij} and b_j in Figure 3) according to the series of input and output data, submitted to the ANN in successive training periods. To train the ANN, a past history of inputs and outputs (a data matrix, containing in each line the values of all inputs and of the corresponding output for each past period) is given to the ANN, which confronts the output that its would have expected given the current weights in the network, and the provided inputs (its expected output), with the output effectively observed in the corresponding period and provided in the matrix. Any discrepancy between these expected and observed outputs generates an expectation error and the weights are increased or decreased given the derivative of the activation function in respect to each input of the node, starting with the output node, and propagating back the corrections towards the initial input nodes of the ANN (back-propagation error correction). This correction process is adaptive and the expectation error cannot be corrected in one pass: several passes ($numEpoch$) are necessary, and each of them only corrects a fraction ($learnRate$) of the

error. This progressive correction method provides a smooth error correction process. While training the ANN the CB can use all past observations (they would be perfectly relevant in a static environment), or only observations for a given number of the most recent periods (a $windowSize$ corresponding to the last 50 periods for example), in a changing environment, like here, since the reactions of the households are also adaptive.

As an aside, we could observe that, hidden nodes represent unobserved state variables, and they may well correspond to the variables of the mental model of the agent that are not observable by the modeler. Using such nodes in an ANN, we can thus cover a wide range of mental model structures from the simplest linear one when there is no hidden layers, to the increasingly sophisticated ones when the number of the hidden nodes ($numHidden$) increases. This number can even be used to represent the complexity of the agent's mental model.

At the end, the goal of this training process is to try to minimize the ANN's expectation errors and this process should result in an ANN progressively better adapted to its evolving environment. How the CB can perceive the state of this environment, and base its expectations on these inputs?

2.4 Information and the perception of the state of the economy by the CB

In our case the CB is placed in a very simple economy, where it can only observe a very limited set of variables on the state of the economy. In its mental model, the CB could use past observations on the following variables for composing its perception of the state of the economy, and use them as inputs of its mental model. For each of them, we propose a simple coding method that makes them a possible input for a simple ANN.

1. Past rate of unemployment: u_{t-1};
2. Current composition of the agents population: the proportion of Believers ϕ_t (only relevant if there are Believers in the model);
3. Direction of this composition (1 : increase, 0 : decrease/stable):

$$\Delta\phi_t = \begin{cases} 1 & if\ \phi_{t-1} < \phi_t \\ 0 & otherwise \end{cases}$$

4. Direction of inflation:

$$\Delta y_t = \begin{cases} 1 & if\ y_{t-2} < y_{t-1} \\ 0 & otherwise \end{cases}$$

5. Direction of unemployment:

$$\Delta u_t = \begin{cases} 1 & if\ u_{t-2} < u_{t-1} \\ 0 & otherwise \end{cases}$$

The level of past unemployment is quite natural as an input, and it is directly connected with the objective of the CB. The second element (ϕ) is introduced in accordance with assumption adopted in the reference literature, since it should normally influence the dynamics of the model that the CB is learning to anticipate. Later, we will show that Believers are not necessary in our model, and we will eliminate them. Direction indicators are introduced in order to give a very rough idea about the dynamics of the economy and the consequences of the strategies of the CB on these dynamics. Moreover, the relevance of all these variables in the ANN is endogenous, since their weight may become zero, if the CB ends by discovering that they do not play a real role.

2.5 CB's decision process: Using expectations to compare available strategies

In order to connect its strategies, the CB must supplement these economic indicators as inputs, with the two components of its strategies: $j = (y^a(j), y(j))$. These seven elements hence constitute the inputs of its mental model. For a given state of the economy, characterized by the values of the above five inputs, the CB can now complete them with any strategy $\left(y_j^a, y_j\right)$, to form an expectation on the potential performance that may result from that strategy, given the state of the economy at period t, and the state of its mental model resulting from its past experience (weights of the ANN resulting from the training with the most up-to-date data matrix):

$$\left(I_1, I_2, I_3, I_4, I_5, y_j^a, y_j\right) \longrightarrow Mental\ model \longrightarrow J_{t+1}^{G,e}\left(y_j^a, y_j\right)$$

CB hence forms expectations on the expected performance that could result from each strategy currently available, and after having compared these expected performances, it implements the strategy that yields the highest expected performance. Households react to y_t^a and form their expectations x_t. CB implements y_t, and the unemployment results from the expectation errors of the households as in equation (5).

At the end of period t, the CB acquires a new set of observations $\left(u_t, y_t^a, y_t; J_t^G\right)$, and it can adjust its mental model by training it, using data for the last period for which it now has a complete set of observations. Consequently, all previous observations (including the last period's set) are used to train the ANN before its use in the following period, and to update its mental model. See Figure 4 for the description of the sequence of decisions and operations of the agents in each period.

2.6 CB's decision process: Building and curating the set of available strategies

The decision process described so far, shows how the CB can choose a strategy in a given set of available strategies. But how should it proceed to build such a strategy set? We proceed in a way very similar to Arifovic et al. (2010) for this component of the model, and assume a genetic-algorithm-inspired exploration process.

CB	Agents
Observe: $(y^a_{t-1}, y_{t-1}, u_{t-1}, \phi_t, \Delta\phi_t, \Delta y_t, \Delta u_t)$ **Update:** expected fitness for $Y_t \leftarrow$ ANN **Set:** (y^a_t, y_t)	
Implement: y_t **Observe:** u_t, J^G_t, ϕ_{t+1}	$i = B, NB$ **Observe:** y^a_t, y_{t-1} **Set:** d_t, x^i_t **Observe:** y_t, u^i_t, J^i_t
	Set: $i = B, NB \rightarrow \phi$
Train ANN: using the new data $(Inputs, Output)$ **If GARate:** Modify population of strategies $\rightarrow Y_{t+1}$ **Update:** expected fitness in Y_{t+1} if GARate	

Figure 4: Sequence of actions in each period

In each period t, the CB has a collection of m strategies $Y_t = \{(y^a_t(j), y_t(j))\}_{j=1...m}$. Each strategy consists of two elements: the announced inflation rate (y^a), and the actual rate of inflation (y). The initial strategy set of the CB is generated randomly. Then the CB computes an expected value for each available strategy, and chooses to use the strategy with the highest expected value in period t.

We rescale the expected performance of strategy j, $J^{G,e}_{t+1}$ in order to help the learning of the CB, and use a value function, $V(\bullet) \in]0,1[$, to normalize it:

$$V^e_t = f\left(J^{G,e}_{t+1}\right) \in [0,1], f' > 0 \tag{8}$$

This value function[6] guides the exploration by the CB of its strategy space, through the evolution of its population of strategies.

In some periods, the CB reconsiders the strategies present in Y, and makes them evolve. More precisely, every $GARate$ periods[7], the CB revises its population of strategies Y_t through the following three steps:

1. Replicating the strategies through a roulette-wheel based on the expected relative performance of the strategies: this selection operator creates a new population of strategies, where the probability of each strategy to be reproduced is proportional to its relative performance $\left(V^e_l / \sum_j V^e_j\right)$.

[6] In our simulations we will simply use $f(J^G) = exp(J^G/10)$, see below. This is a simple scaling function, and any scaling function giving positive values for the fitness, and enough differentiating performances resulting from different policy decisions would allow the learning of the CB.

[7] $GARate$ is hence the frequency of modification of the strategy population of the CB. A limit case is $GARate = 1$, corresponding to a modification of the population of strategies every period, but we also consider the possibility of letting more time to the CB for evaluating the current population, before making it evolve again. We know in the evolutionary learning literature that a too frequent modification of the strategy population can perturb the learning of the value of each present strategy. Vallée & Yildizoglu (2009) show, for example, that this frequency can play an important role in the ability of firms of learning the Cournot equilibrium.

2. Combining the already discovered strategies (crossover): each strategy in the population can be combined, with a probability p^C, with another strategy. If strategies $j = (y^a(j), y(j))$ and $k = (y^a(k), y(k))$ are chosen for crossover, they are replaced by two new strategies: $l, m = ((y^a(j) + y^a(k))/2, (y(j) + y(k))/2)$.

3. Random experimentation (mutation): with a probability p^m one element of each strategy ($y^a(j)$ or $y(j)$) is modified with a gaussian noise

$$new\,value = old\,value + \varepsilon, \quad \varepsilon \rightsquigarrow \aleph(0, 1).$$

These tree steps (Replication–Combination–Mutation) are quite common in genetic-algorithm inspired evolutionary learning schemes that are used in the literature, and they are similar to the one used by Arifovic et al. (2010).

2.7 Evolution of the types of agents

In Arifovic et al. (2010), the types of agents evolve over time: in each period a proportion β of the agents are drawn randomly and they meet other households. If their utility level is lower than the observed utility of the agent they meet, they adopt its behavior as a Believer or Non-believer. The switching agent starts with no error to correct ($d_{it} = 0$). This imitation process determines the dynamics of ϕ. This will only be relevant in our model in the initial configurations where we keep Believers only for allowing the comparison of our results with Arifovic et al. (2010).

We have now completely specified the mechanisms of our model. Figure 4 describes the decisions and observations made by all agents, as they take place sequentially in each period, and the pseudo code of the model, with the reference configuration, is given in Figure 5.

Can such a learning scheme favor the ability of the CB to drive the economy towards the Ramsey outcome, and reinstate the credibility and effectiveness of economic policy? We will explore this question in our simulation experiments, starting with the baseline configuration of Arifovic et al. (2010) and under more complex environments, and learning schemes of the households.

3 Simulation protocol

We use Arifovic et al. (2010) as a reference for the baseline setup for the economy:

- Number of private agents: $N = 100$
- Natural unemployment rate: $u^* = 5.5$;
- Proportion of agents participating to the tournament for learning on types: $\beta = 1\%$;
- Learning speed of the Non-believers: $\gamma = 0.1$;

1. Read the data corresponding to the current experiment, and initialize the parameters accordingly.
2. Using the initial value of ϕ, create n^B Believers, n^{NB} Non-believers, and the Central Bank (CB) with randomly initialized strategies (decisions of the CB and the forecasts of the agents).
3. Randomly initialize the population of strategies (Y) of the CB and the weights of her neural network.
4. CB chooses a strategy (y^a, y).
5. for $t \leq T$ (T is the length of each run),
   (a) Compute the new forecasts of the agents.
   (b) Compute the current unemployment rates and utility of the agents given the strategy used by the CB.
   (c) Compute individual and global indicators.
   (d) Agents imitate each other's type and a new value is computed for ϕ.
   (e) Train the ANN with the new observations.
   (f) If $t \ mod \ GARate = 0$, the population of strategies of the CB is updated using selection, crossover, mutation operators, and the expected fitness of the elements of the new population is updated using the actual state of the ANN, and the relevant input data.
   (g) The CB chooses a new strategy for the next period, from her population of strategies using her expectations given by her mental model and the current input data.

Figure 5: The pseudo code of the model.

- Parameter of the Philips curve: $\theta = 1$;
- Non-believers' forecasting cost: $c^{NB} = 3.3$
- Initialization and mutations: we draw y^a and $y \in [-10, 15]$.

These parameter values correspond to the following values of the CB's optima:

Ramsey solution	Nash solution
$y^R = 0$	$y^N = 5.5$
$JG^R = -15.125$	$JG^N = -30.25$
$V(JG^R) = 0.224$	$V(JG^N) = 0.0486$

with the following simple value function used in our model (see equation (8)):

$$f\left(J^{G,e}\right) = exp\left(J^{G,e}/10\right)$$

For parameters specific to our setup (concerning mainly the characteristics of the CB's ANN, and exploration process), we adopt:

Size of the strategy collection, $m = 400$; frequency of the GA, $GARate = 4$; number of hidden nodes, $nbHidden = 4$; probability of crossover, $p^C = 0.3$; probability of mutation, $p^m = 0.05$. These are quite standard values for these parameters and our simulations confirm what we observed in our previous models using the same approach (Yildizoglu 2001, and Yildizoglu et al. 2014): except with extreme values, the learning process based on such a mental model is robust.

For each different configuration of the model we ran 50 simulations where each simulation lasts for $5,000$ periods.[8]

For each configuration, we represent the resulting distribution of variables and indicators using density plots of the performance of the CB and of its inflation strategy. The density plots give one-dimensional kernel density estimation of the distribution of the considered variable. The x-axis represents the values of the variable and the y-axis, the estimated density at these values. In our case, these plots quite nicely show the distributions of the observed values of J^G under different assumptions, and around which values the economy spends its time.

As our simulations are quite long and contain a large number of observations, we present the data of a random subsample in each figure (with a sampling rate of 5%, corresponding to $\mathbf{10,000}$ observations coming from different time periods and simulation runs).We discard the data generated in the first 1000 simulation periods in order to get an idea about the distribution we observe once the initial randomness is absorbed by the learning processes of the agents. Consequently, we present the data on the last 4000 periods only. When we analyze the dynamics of the learning, we represent the evolution of variables over time, and, in this case, we plot the per-period average of observations over all runs for each variable.

4 Results

Original literature has underlined the inability of the CB to drive the economy to the Ramsey outcome, Dawid and Deissenberg (2005) showed that the presence of naive Believers could help to push the economy above the Nash equilibrium and Arifovic et al. (2010) showed that a boundedly rational CB could learn to keep the Believers in the population, and keep again the economy above the Nash welfare. These are the most positive results we have in the theoretical literature, but none shows a convergence towards the Ramsey outcome as it has been observed in the experiments of Arifovic and Sargent (2003). Our first two questions concern the dynamics of this economy under the baseline setup adopted by Arifovic et al. (2010), except for the forward-looking learning and strategy selection used by the CB: *i)* Does the economy converge on a specific outcome and, if yes, *ii)* is it able to attain the Ramsey solution?

[8] We analyze the results of these simulations using R-project (R Development Core Team 2003), and `ggplot2` library of R-project.

4.1 Convergence? To the Ramsey outcome?

We give in Figure 6–(a) the plot of the distribution of the CB's payoff over all relevant periods and simulation runs, and compare it with the performance that we would have in Nash and Ramsey outcomes. We clearly observe that the economy spends a dominant share of its history close to the Ramsey outcome, and nearly never close to the Nash outcome. We observe quite a strong convergence onto the Ramsey outcome in terms of CB's payoff: a very significant fraction of the $10,000$ observations are grouped around the Ramsey payoff. Figure 6–(b) clearly shows the CB is indeed able to use the Ramsey inflation rate (the zero–inflation) in a very large proportion of observations.

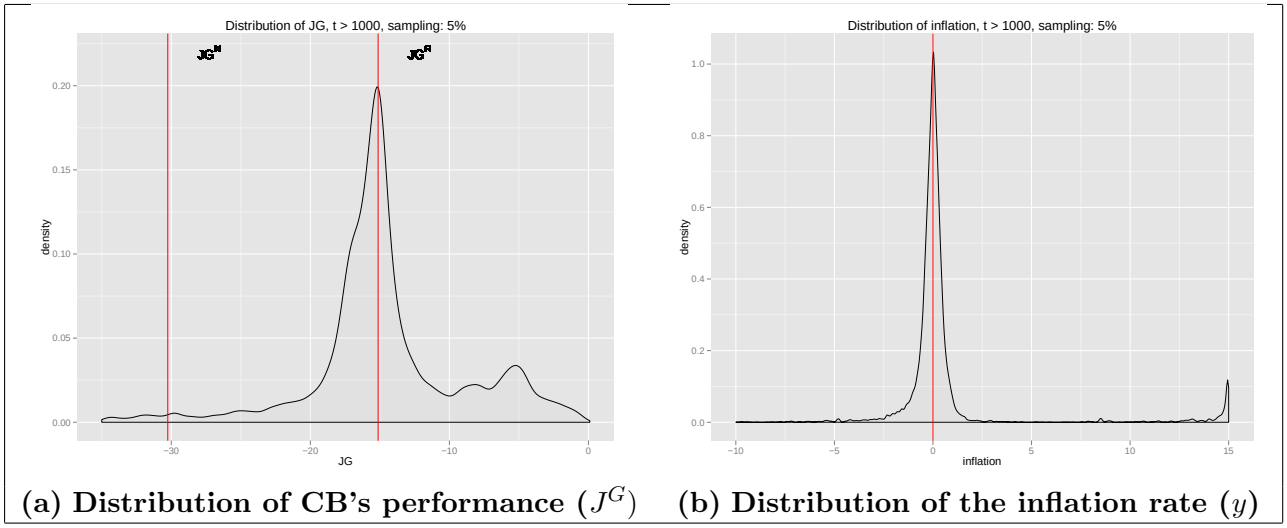

(a) Distribution of CB's performance (J^G) **(b) Distribution of the inflation rate (y)**

Figure 6: Convergence of J^G and inflation: Nash $\left(JG^N\right)$ or Ramsey $\left(JG^R\right)$?

Proposition 1. *The forward-looking CB is able to steer the economy towards the Ramsey outcome.*

Proposition 1 is the main result of our model, and it contrasts our model with all previous work on this topic. Remember that we use here exactly the same setup as Arifovic et al. (2010), so the dynamic adaptive expectations of the CB, based on its mental model, is necessarily the determining factor that drives the ability of the CB to implement the Ramsey outcome. And as we will show below, the CB does not even need the naive Believers for this ability, and the effectiveness of the economic policy does not depend on them.

4.2 Does the CB need the Believers for going to Ramsey?

Are Believers necessary for getting at the Ramsey outcome in this economy? To check this, we do a batch of runs without any Believers.

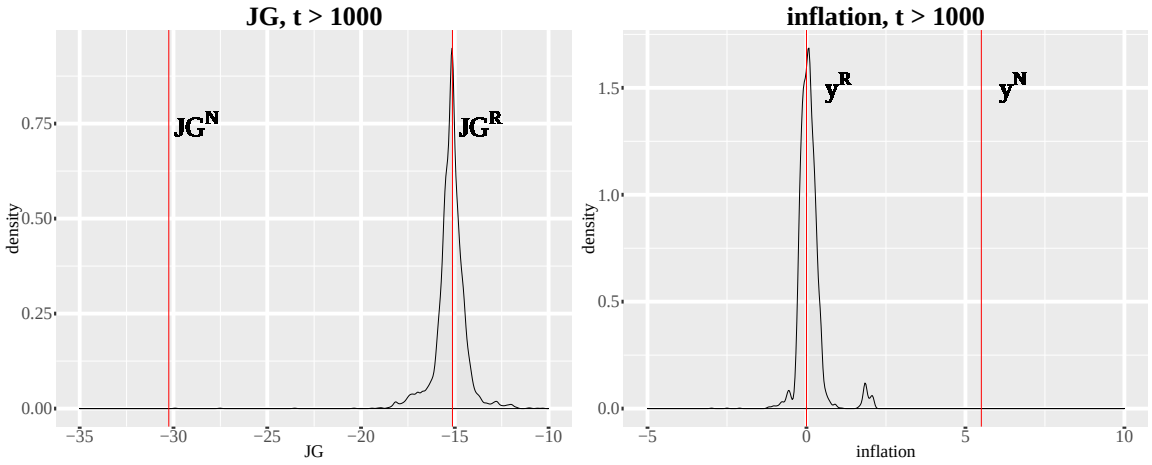

Figure 7: Convergence to the Ramsey outcome without Believers ($\phi = 0$)?

We plot in Figure 7 the distributions of the J^G and the inflation rate in this specific configuration with no Believers. We observe that the convergence onto the Ramsey outcome is definitely possible, and even stronger, in this case.

Proposition 2. *Even in the absence of naive Believers, the CB is able to drive the economy towards the Ramsey outcome. Hence, Believers are not necessary for the effectiveness of the economic policy.*

If we compare the distributions of this J^G in the left panel of this plot and in Figure 6–(a), we observe that the hump on the right that can be observed in the latter is absent from the new plot: the Believers are mainly necessary for the CB to attain performances higher than the Ramsey outcomes (since they can be easily *cheated*), but they are not necessary for attaining the Ramsey outcome itself. This result is quite remarkable and specific to our model, in comparison with other models we have discussed in the introduction. Consequently, since they are not necessary at all for our results, **we drop the Believers from the model from this point on** by setting $\phi = 0$. Thus, the announcement of inflation, y^a, also becomes irrelevant in the basic model (we will reintroduce it again later, when it will potentially become relevant).

4.3 How the CB can steer the economy towards to Ramsey outcome?

The previous results are quite strong, and they may look somewhat miraculous, given that they contrast strongly with the existing literature on this question. To check that they are robust, we must verify that **(a)** the learning of the CB is indeed behind these results, and **(b)** that it is not using the zero-inflation strategy just by chance, or because of some quirk of the evolutionary algorithm that evolves its strategy population (for example, if all strategies of the CB converge to $y = 0$ for a strange reason, we would observe the Ramsey outcome by "accident" of early convergence of the exploration scheme).

4.3.1 Is the CB learning how the economy functions?

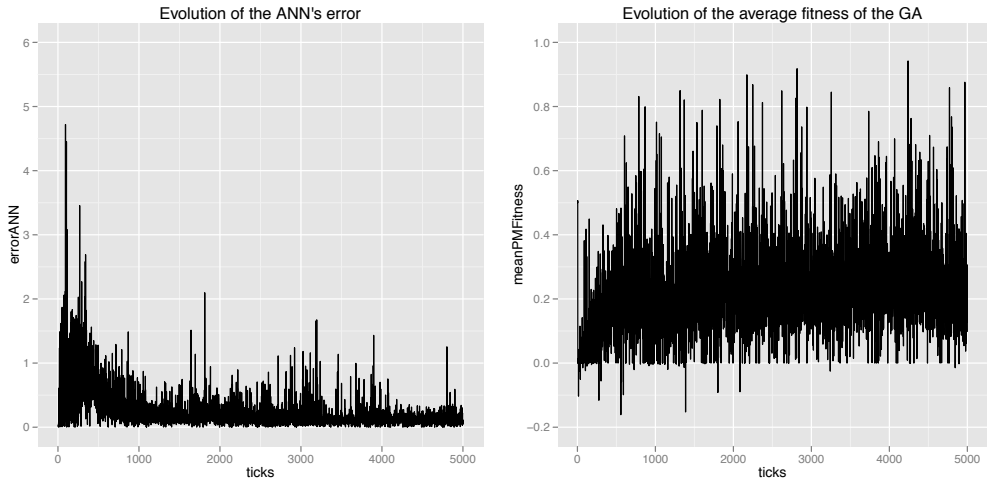

Figure 8: Learning of the CB: Decreasing total errors in the ANN (on the left) and increasing average fitness of the CB's strategy population (on the right)

How is this economy able to discover Ramsey outcome and to remain in its neighborhood with such a high frequency? Is the CB learning this outcome, or is it implementing the zero–inflation even if this strategy implies important errors in its representation of the economy?

The plot on the left side of Figure 8 represents the evolution of the total error in the CB's ANN (the error observed during the training phase in each period, averaged over all runs). We observe that, after an initial phase of increase, this error decreases over time and the CB learns to correctly predict the consequences of its policies, including the reactions of households. In parallel, the plot on the right side of Figure 8 shows that the average fitness of the population of strategies of the CB (again averaged over all runs) increases over time, confirming the learning hypothesis.

Moreover, we can observe in parallel with the evolution of the ANN error, the inflation strategy used by the CB, and the forecast error of the Non-believers. Figure 9 clearly shows the co-evolution of these variables: the CB discovers the favorable role of zero inflation, and the Non-believers progressively adapt to this strategy, their errors decrease (in absolute value) in parallel with the error in the mental model of the CB. These medians, as well as data from individual runs show that the convergence to zero inflation is never final, and the CB uses transient deviations from zero-inflation, and the corrective reactions of the NB (see below) to keep the economy in the vicinity of the Ramsey outcome.

Proposition 3. *The convergence to Ramsey is made possible by the forward-looking learning of the CB, and the CB uses transient deviations from zero-inflation to keep the economy in the vicinity of the Ramsey outcome.*

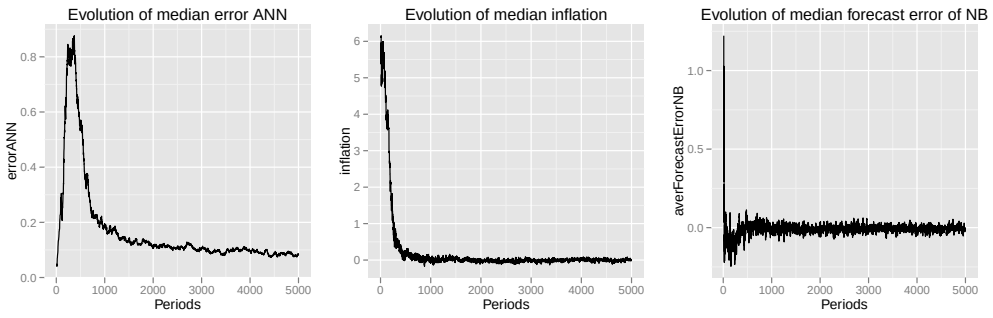

Figure 9: Evolution of expectation errors, and inflation strategy

Observing this learning is already comforting on the significance of the convergence result but we can also check that this zero-inflation policy is not an artifact of the mechanisms of the genetic algorithm either (like an early convergence to a population with only zero inflation as strategy).

To check that the CB uses the zero-inflation strategy non by chance, but because it expects to get a higher fitness using it, in comparison with other inflation levels, we can check the expected values given by the ANN, for different inflation rates, when the CB has driven the economy in the vicinity of the Ramsey solution.

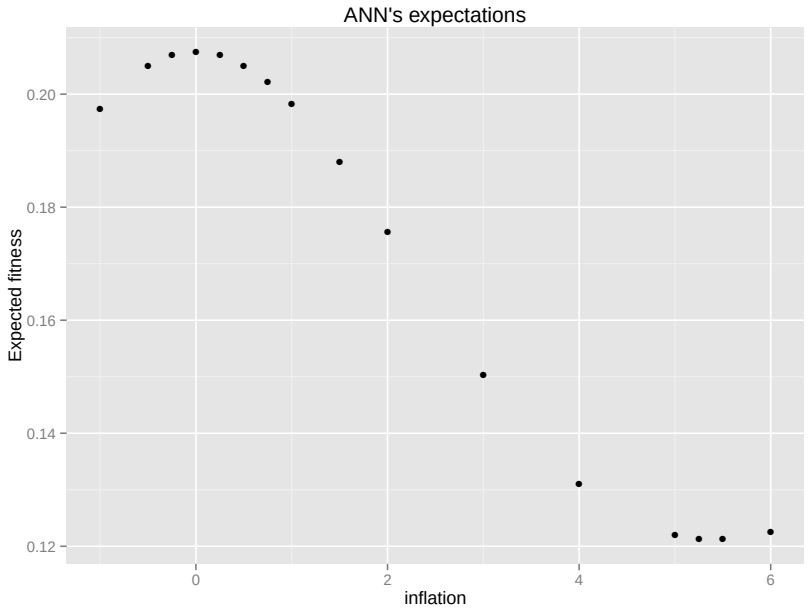

Figure 10: Output values of the ANN (expected value of CB's performance, for different values of y, when the economy is close to the Ramsey outcomes

Figure 10 represents the value surface expected by the ANN when the economy

is in the neighborhood of the Ramsey outcome, for different values of inflation that we manually check, independently from the actual configuration of the population (to eliminate, as a cause, any strange convergence of the strategy population to the zero-inflation strategy):

$$y \in \{-1, -0.5, -0.25, \mathbf{0}, 0.25, 0.5, 0.75, 1, 1.5, 2, 3, 4, 5, 5.25, \mathbf{5.5}, 6\}$$

These values cover a large spectrum that contains both the Ramsey inflation (0) and the Nash one (5.5). We clearly observe in Figure 10 that the maximal expected value for the output of the ANN corresponds to a strategy with zero inflation. Moreover, this maximal value is quite close to the one corresponding to the Ramsey outcome in this case $\big(f(JG^R) = 0.22\big)$. When we check the configuration of the CB's strategy population, we observe that the average inflation strategy is -0.34 and the variance of the inflation strategies is $0.055 > 0$, so some diversity is clearly present, and we do not observe a strange convergence. We can hence conclude that the choice of zero inflation by the CB is perfectly optimal here, and it is not due to some idiosyncratic reason.

Proposition 4. *The CB adopts the zero-inflation Ramsey strategy, because it learns that it would indeed yield an optimal policy performance.*

4.4 Discussion: What does the CB learn in order to be able to drive the economy to the Ramsey outcome?

The preceding results show that the convergence on the Ramsey outcome is quite robust. This result is very different from the convergence to the Nash outcome that is observed in other models. Even in Arifovic and al. (2010), which is very close to our model, outcomes better than Nash can be observed, but not the convergence to Ramsey outcome. Given that the Non-believers mainly play their static best reply, convergence to Nash can indeed seem quite natural. Moreover, we obtain convergence to Ramsey even in the absence of the Believers (while they are necessary for moving away from the Nash in the previous models). What is the mechanism that, in our case with adaptive learning and expectations, drives the economy towards the Ramsey outcome?

This mechanism is the expectations based decision process of the CB, combined with the ability of the Non-believers to correct their errors. Because of this mechanisms, the CB ends by learning to expect the *full reaction function* of the Non-believers (which consists of their best reply augmented by the error correction process), and it takes into account this reaction function when choosing the inflation level. As a consequence of these adjustment processes (coevolution of behaviors), its expectation errors, as well as Non-believers forecast errors decrease over time and the economy converges to zero–inflation Ramsey outcome (see Figure 9).

Figure 11 displays the distribution of J^G and inflation (y) in a series of simulations run with the baseline configuration, but without the error correction ability of the Non-believers: they can only play their best reply function (2) with $d = Cste = 0$. These figures exhibit convergence on the Nash equilibrium (panel a), with the persistence of

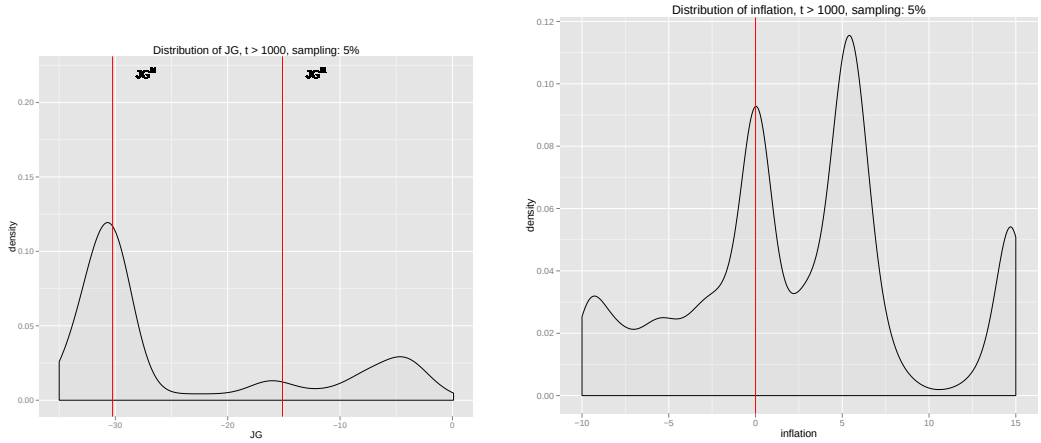

(a) (b)

Figure 11: Distribution of J^G and y when Non-believers do not correct their errors

positive rates of inflation (panel b). Consequently, the ability of the Non-believers to correct their forecast errors, and of the CB to incorporate this correction in its expectations and decisions is the mechanism that drives the convergence to the Ramsey outcome in our model. The flexibility provided by the error correction process of the Non-believers is necessary for the emergence of the coevolutionary dynamics that drive the economy towards the Ramsey outcome, thus escaping the Nash equilibrium.

In fact, in the absence of the error correction process, the CB learns to expect the best reply of the Non-believers, and the payoff it can get by moving after their choice, and by increasing the inflation. We observe, in this case, that it frequently deviates as illustrated in Figure 11–(b). The mode of the distribution of the inflation rate is strictly positive and very close to $u^* = 5.5$, which corresponds to the Nash equilibrium inflation rate of the static game).

Proposition 5. *The CB is able to drive the economy to the Ramsey outcome, because it is able to learn to expect the full reaction function of the Non-believers (which consists of their best reply augmented by the error correction process), and it takes into account this reaction function when choosing the inflation level. Under this constraint, zero-inflation Ramsey strategy becomes optimal (see above).*

Moreover, as a corollary of the discussion in paragraph 4.2, we can also observe that the presence of the Believers perturbs the learning process of the CB: it has more difficulty to take into account the error–correction feed-back of the Non-believers, and to bring the economy to the Ramsey outcome.

When there is the Non-believers' error correction mechanism in place, Nash outcomes cannot be supported any more by the learning dynamics as Non-believers *respond* to the observed, realized inflation rate. When the CB learns to *expect* this correction, it is constrained to take into account the full reaction of the Non-believers, and, consequently,

learns to take decisions under this constraint when choosing the inflation rate. The Non-believers' error correction mechanism disciplines the CB in its tendency to exploit the agents' expectations, and, as a consequence, it starts to act credibly. This results in the convergence to the Ramsey outcome, instead of the Nash equilibrium. The adaptive dimension of the Non-believers' behavior helps the CB in pulling the economy towards the Ramsey outcome.

5 Learning in more complex environments

5.1 Learning with less information: What is the necessary information level for the CB to learn the Ramsey outcome?

Another important dimension that could explain our favorable results would be the quantity of information processed by the CB. Could the CB attain the Ramsey outcome with less information? In particular, to be inline with the relevant literature, we have chosen to let the CB observe the composition of the agent population (through the proportion of Believers, ϕ). Was this information necessary? One could consider that it is not necessarily easily available to the CB. Since we observed that the Believers are in fact not necessary for the convergence to the Ramsey solution, this information was not playing an important role anyway. We also check the role played by other information elements in characterizing the state of the economy, by comparing the results obtained under alternative assumptions about the CB's information level, represented by the variable $infoLevel$. Different values of this variable correspond to different information elements observed by the CB:

1. Past inflation strategy: y_{t-1} ;
2. (1) + past unemployment: (y_{t-1}, u_{t-1})
3. (1) + direction of inflation (see above) as for the CB: $(y_{t-1}, \Delta y_t)$
4. (1) + (2) + (3): $(y_{t-1}, u_{t-1}, \Delta y_t,)$
5. (4) + direction of unemployment: $(y_{t-1}, u_{t-1}, \Delta y_t, \Delta u_t)$

Figure 12[9] shows that the most important information concerns the state variable of the model, even with a delay (u_{t-1}). As soon as this information is used, the convergence to Ramsey is assured, and more strongly when this signal is not disturbed by others (for the distribution of J^G, the case $infoLevel = 2$ dominates strongly and clearly other cases in terms of convergence to the Ramsey outcome). More information is not always favorable to economic performance. In any case, we never observe convergence to the Nash equilibrium, and the economy very frequently stays close to Ramsey.

[9]We use here simple box plots because it is not easy to distinguish different distributions in the corresponding density plot for J^G. Box plots give the minimum, 3 quartiles, and the maximum of the distribution in each case. The center line gives the median of the distribution.

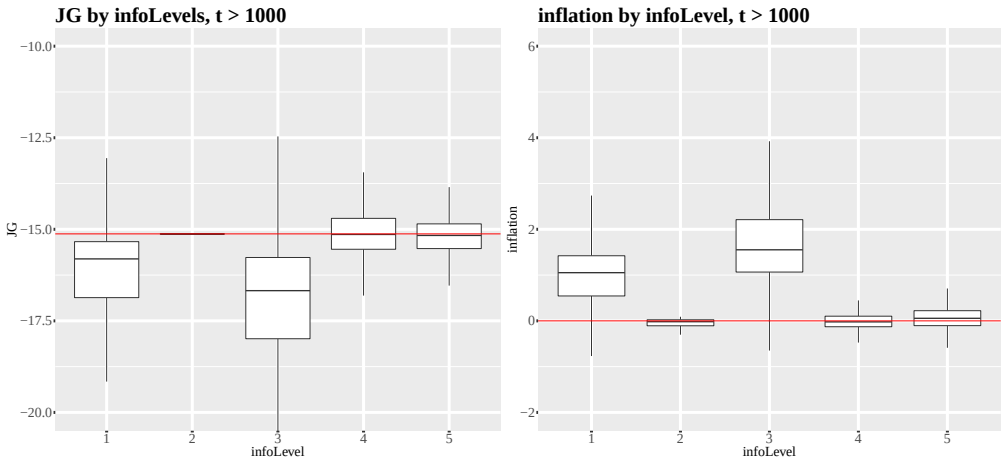

Figure 12: Learning with different information levels

Proposition 6. *The main information necessary for learning the Ramsey outcome is the past level of the state variable unemployment, u. In its absence, the CB is nevertheless able to keep the economy above the Nash equilibrium, but the convergence on Ramsey is less frequent.*

5.2 Learning the Ramsey outcome in the presence of error-prone households?

Does the ability of the CB of driving the economy towards the Ramsey solution depend on the deterministic error correction process of the households (see (3)) ? To check the robustness of this ability, we next consider that households can only imperfectly observe the current inflation when they compute the current correction of their expectation error:

$$d^i_{t+1} = d^i_t + \gamma\left((y_t + \varepsilon_t) - x^i_t\right), \tag{9}$$

$$\text{with } \varepsilon_t \rightsquigarrow \mathcal{N}(0, \sigma) \tag{10}$$

where we represent the error rate of the households through σ: the higher the σ, the more imprecise the inflation observations of the households, and the less reliable their error correction process.

The box plots[10] given in Figure 13 show that getting the economy to Ramsey is clearly more difficult with increasing error rates of the households. In terms of the performance of the CB, the best convergence is observed when the households can perfectly observe the inflation, and correct their expectation error ($\sigma = 0$). But the existence of errors, when σ is not too high, can help the convergence to the Ramsey inflation ($\sigma = 5\%$), by giving more flexibility to the error correction mechanism of the households. In any

[10]We use again here simple boxplots because it is not easy to distinguish different distributions in the corresponding density plot for J^G.

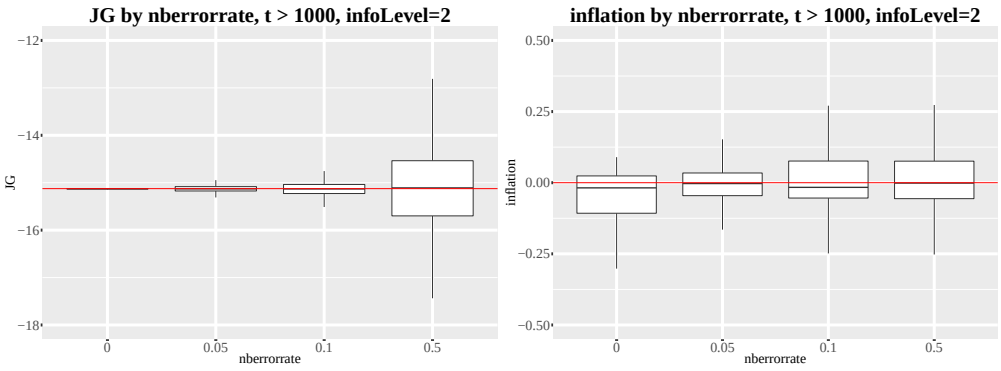

Figure 13: The convergence to Ramsey when households are error-prone (infoLevel = 2)

case, we never observe convergence towards the Nash equilibrium, in terms of both performance and inflation rates, even with noisy error correction by the households.

Proposition 7. *A less precise error correction by the households, with noisy observation of the inflation rates does not impede the ability of the CB to drive the economy away from the Nash equilibrium, and converge to the Ramsey solution. The convergence is relatively less frequent when the observations are subject to high levels of noise.*

5.3 When the households also use dynamic adaptive expectations: Co-evolution of mental models

We consider now an economy where the households also use a learning scheme based on dynamic adaptive expectations, in a similar way as for the CB. Each household has a mental model (ANN) on the connection between the state of the economy, the announcement of inflation by the CB (y_t^a), and the expected fitness of each available strategy to the household h (different potential inflation expectations, x_j^h). So, we consider that the CB can potentially again influence households' inflation expectations with its announces.

In a similar way to the CB, given the expected performance of each of the strategies in its strategy population, $X_t^h = \left\{ \left(x_j^h, V_j^e \right)_{j=1..m^H} \right\}$, each household h selects the expected inflation level that corresponds to the highest expected value.

When trying to characterize the state of the economy, in this simple setup, we assume by symmetry that, the households will use the same type of information as the CB (but their information level ($infoLevelH$) may differ from the CB's one). For the learning scheme parameters, we use the same configuration as for the CB, and we run simulations for 1000 periods.

CB	**Households** (h)
Observe: $\left(y_{t-1}^a, y_{t-1}, u_{t-1}, \Delta y_t, \Delta u_t\right)$	**Observe:** $\left(y_{t-1}^a, y_{t-1}, u_{t-1}, \Delta y_t, \Delta u_t\right)$
Update: expected fitness for $Y_t \leftarrow$ ANN	
Set: $\left(y_t^a, y_t\right)$, and announce y_t^a	
	Observe: y_t^a
	Update: expected fitness for $X_t^h \leftarrow$ ANN
	Set: inflation expectation x_t
Implement: y_t	
Observe: u_t, J_t^G	**Observe:** y_t, u_t^h, J_t^h
Train ANN: using the new data $(Inputs, Output)$	
If GARate: Modify strategy population $\rightarrow Y_{t+1}$	**If GARateH:** Modify strategy population $\rightarrow X_{t+1}^h$
Update: expected fitness in Y_{t+1} if GARate	**Update:** expected fitness in X_{t+1}^h, if GARateH

Figure 14: New sequence of actions with coevolving mental models

The new sequence of events in each period of the model is given in Figure 14, which clearly shows the symmetry between these two types of agents now. Instead of using the ad hoc Nash best reply and error correction mechanism, the households use now their mental model for forming dynamic adaptive expectations.

Can the CB drive again the economy towards the Ramsey outcome with these households who are normally able to adapt their representation of the economy when the functioning of the latter changes following the monetary policy used by the Central Bank ? We can observe that we have here, also for the households, adaptive expectations that are not subject to the Lucas critique, because the households adapt not only their expectations, but also the representation of the economy, when the latter changes.

We consider the following potential information levels for both types of agents (we distinguish now $infoLevel$ for the CB and $infoLevelH$ for the households):

1. Past announced and actual inflation rates: $\left(y_{t-1}^a, y_{t-1}\right)$
2. (1) + past unemployment rate: $\left(y_{t-1}^a, y_{t-1}, u_{t-1}\right)$
3. (1) + direction of inflation (see above): $\left(y_{t-1}^a, y_{t-1}, \Delta y_t\right)$
4. (2) + direction of unemployment: $\left(y_{t-1}^a, y_{t-1}, u_{t-1}, \Delta u_t\right)$
5. (2) + direction of inflation:$\left(y_{t-1}^a, y_{t-1}, u_{t-1}, \Delta y_t\right)$
6. (4) + (5): $\left(y_{t-1}^a, y_{t-1}, u_{t-1}, \Delta y_t, \Delta u_t\right)$

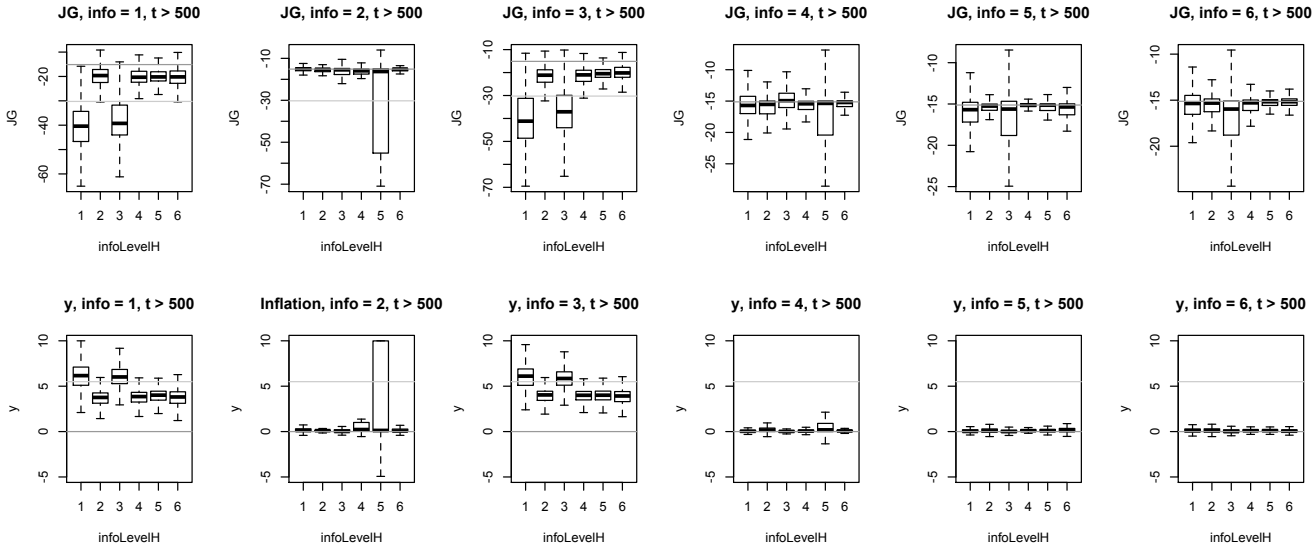

Figure 15: Distribution of J^G (top line) and of y (bottom line) with coevolving expectations for different information levels of the CB (each column) and the households (horizontal axis of each figure)

We show in Figure 15 the results corresponding to different combinations of these information levels for the CB and the households. We observe in this figure that the CB is able to bring the economy towards the Ramsey outcome if it uses information on past unemployment level (all cases but $infoLevel = 1, 3$). We observe a clear convergence to the Ramsey outcome on all columns except columns 1 and 3. But, even in these unfavorable cases ($infoLevel = 1, 3$), if the households use information on unemployment, the CB can use this information to pull the economy above the Nash equilibrium, and we observe lower inflation rates, and higher economic performance when $infoLevelH \neq 1, 3$.

Proposition 8. *As long as at least one type of agents, CB or households, use information on unemployment, the CB is again able to steer the economy towards the Ramsey outcome (and converge to it when the CB itself is using this information) with coevolving dynamic adaptive expectations.*

Consequently, the dynamic adaptive expectations based learning process of the CB is quite robust in its ability to drive the economy to the Ramsey outcome, under quite different assumptions about the learning of the households. We should remember here that in all cases, agents only use information on their past experience in the economy, or on their current decisions. They build instead an adaptive representation of this economy (their mental model), and form their expectations using that potentially imperfect, but adaptive model.

6 Conclusion

We study an extension of the Kydland-Prescott environment with two types of agents, Believers and Non-believers as in Arifovic et al. (2010). While Believers set their inflation expectations equal to the inflation announced by the Central Bank (CB), Non-believers update their inflation expectation every period. The fraction of Believers and Non-believers changes over time in response to their relative performance in terms of forecasting accuracy. The CB has a collection of strategies that evolve over time. Each strategy consists of two elements, the announced and the actual inflation. The CB evaluates ex ante each strategy using its mental model of the economy that is in our framework represented by an artificial neural network. In each time period, the CB selects the strategy that receives the highest expected value when evaluated through the artificial neural network.

Using this decision making and updating process in which alternative strategies (policies) evolve over time and are evaluated through the artificial neural network that is also updated every period, the CB brings the economy to the neighborhood of the Ramsey outcome. The frequency of Ramsey outcome is even higher in the environment populated with Non-believers only. While rational expectations theories predict a possibility of Ramsey outcomes alongside with many other equilibria, there is no selection mechanism that can be used to choose this Pareto optimal outcome. Modeling of learning in this type of environment usually results in Nash outcomes with a few exceptions, cited in the introduction, that can achieve better than Nash outcomes. However, to the best of our knowledge, the research in this area has not demonstrated the possibility of achieving the Pareto optimal, Ramsey outcome in the Kydland-Prescott environment.

Our result is novel and demonstrates that the Ramsey outcome can be achieved in this type of environments. This result is due to the fact that we model our Central Bank as 'truly' forward looking, but still adaptive. By updating its mental model, the CB becomes able to fully take into account the reaction of Non-believers, including their ability to correct their errors, and adapt to the economic environment. Using this ability of Non-believers, and expecting the future payoff that may result from it, the CB is able to discover the Ramsey outcome, and steer the Non-believers (and the economy) towards it. Our results show that without this adaptive dimension of the Non-believers behavior, the only outcome is the Nash equilibrium, with lower welfare. Consequently, adaptive behavior on both sides of the economy, and forward looking on the side of the policy maker both seem necessary for the emergence of the Ramsey outcome. It is interesting to note that, in addition to its collection of strategies, the only other information that the CB needs to bring the economy to the Ramsey outcome is the information on the previous period's unemployment.

When we add noise to the households observation of inflation, our results remain robust, i.e. the outcomes of the adaptation of the CB and the households results in Ramsey. Finally, when we endow each of our households with the expectations formation scheme symmetrical to the one used by the Central Bank, our results remain robust to this change in expectations scheme providing that, at least, either the CB or

the households use information on past unemployment for characterizing the economic environment in which they are learning.

In general, it is very difficult to model forward-looking, yet adaptive, behavior. We succeeded in capturing the behavior that is more realistic in terms of the actual CB behavior. Standard models of learning are only 'backward looking', and rational expectations models do not provide guidance as to how to reach Pareto optimal outcomes. Combining the adaptation, in a changing environment, with evaluation of forward looking strategies, we were able to describe how this process can take place.

References

Arifovic, J., H. Dawid, C. Deissenberg, and O. Kostyshyna. 2010, Learning benevolent leadership in a heterogeneous agents economy. *Journal of Economic Dynamics and Control*, 34:1768–1790.

Arifovic, J., and J. Ledyard, 2004, Scaling Up Learning Models in Public Good Games, *Journal of Public Economic Theory*, 6: 205-238.

Arifovic, J., T.J. Sargent, 2003, Experiments with human subjects in a Kydland–Prescott Phillips curve economy. In: Altig, D., Smith, B. (Eds.), *The Origins and Evolution of Central Banking: Volume to inaugurate the Institute on Central Banking of the Federal Reserve Bank of Cleveland*. Cambridge University Press, pp. 23–56.

Barro, R.J., D.B. Gordon, 1983. Rules discretion and reputation in a model of monetary policy. *Journal of Monetary Economics*, 12, 101–121.

Cho, I.-K., Sargent, T.J., 1997. Learning to be credible. Manuscript prepared for presentation at the conference to celebrate the Bank of Portugal's 150th birthday.

Cho, I.-K., Williams, N., Sargent, T.J., 2002. Escaping Nash inflation. *Review of Economic Studies*, 69, 1–40.

Cukierman, A., 1992. *Central Bank Strategy, Credibility, and Independence: Theory and Evidence*. The MIT Press.

Dawid, H., C. Deissenberg, 2005. On the efficiency effects of private (Dis)-trust in the policy maker. *Journal of Economic Behavior and Organization* 57, 530–550.

Evans, G. W. and S. Honkapohja, 2001, *Learning and expectations in macroeconomics*. Princeton University Press.

Holland, J. H., K. J. Holyoak, and P. R. Thagard, 1989, *Induction. Processes of Inference, Learning, and Discovery*. MIT Press, Cambridge:MA, 1989.

Kydland, F.E., Prescott, E.C., 1977. Rules rather than discretion: the inconsistency of optimal plans. *Journal of Political Economy* 85, 473–491.

Lettau, M. and H. Uhlig. Rules of thumb versus dynamic program- ming. *American Economic Review*, 89(1):148–174, 1999.

Masters, T., 1993, *Practical Neural Network recipes in C++*. Academic Press, New York.

Persson, T., Tabellini, G., 1993. Designing Institutions for Monetary Stability. In: Meltzer, A.H., Plosser, C.I. (Eds.), *Carnegie-Rochester Conference Series on Public Policy*, vol. 39, pp. 53–84.

R Development Core Team. 2003. *R: A Language and Environment for Statistical Computing*. http://www.r-project.org/. Vienna: R Foundation for Statistical Computing.

Salmon, M., 1995, Bounded rationality and learning: Procedural learning. In Mark Kirman, Alan; Salmon, editor, *Learning and Rationality in Economics*, pages 236–275. Blackwell, Oxford.

Sargent, T. J., 1999, *The Conquest of American Inflation*, Princeton University Press, Princeton: NJ.

Simon, H.A., 1976, From substantial to procedural rationality. In S. J. Latsis, editor, *Method and Appraisal in Economics*, pages 129–148. Cambridge University Press, Cambridge.

Stein, J., 1989, Cheap talk and the fed: a theory of imprecise policy announcement. *The American Economic Review*, 79, 32–42.

Stokey, N., 1989, Reputation and Time Consistency, *American Economic Review*, 79, 134-39.

Vallée, T., 1999, Jeux de Stackelberg dynamiques : incohérence temporelle, triche et apprentissage. PhD thesis, Université de Nantes.

Vallée T., M. Yildizoglu, "Convergence in the Finite Cournot Oligopoly with Social and Individual Learning", *Journal of Economic Behaviour and Organization*, 72, 670-690, 2009.

Walsh, C.E., 1999, Announcements, inflation targeting, and central bank incentives. *Economica*, 66, 255–269.

Wilson, S. W., 1995, Classifier fitness based on accuracy. *Evolutionary Computation*, 3(2): 149–175.

Yıldızoğlu, M., 2001, Connecting adaptive behaviour and expectations in models of innovation: The potential role of artificial neural networks. *European Journal of Economics and Social Systems*, 15(3):203–220.

Yıldızoğlu, M. M.-A. Sénégas, I. Salle, and M. Zumpe, Learning the optimal buffer-stock consumption rule of Carroll, in *Macroeconomic Dynamics*, 18(4), 727-75, 2014.