Evolution of Beliefs, Policy Implications in Agent-Based and Experimental Economies: 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)
Department of economics
Bordeaux University
yildi@u-bordeaux4.fr

December 2013

Abstract

We study learning in the Kydland and Prescott environment. Our policy maker evaluates its potential strategies regarding the announced and the actual inflation rate using its *mental model* . This model is forward looking and adaptive at the same time. There are two types of agents: Believers who set their inflation forecast equal to the announced inflation, and nonbelievers 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.

JEL categories: E50, C45, C72, D60

Keywords: learning in the Kydland-Prescott environment, artificial neural networks, Ramsey outcome

1 Introduction

Following Kydland and Prescott (1977) and Barro and Gordon (1983), the issue of time inconsistency and credibility building has been extensively studied. A number of studies investigate the possible use of non-binding policy announcements which can improve upon the time-inconsistent Nash solution. Most of them assume hidden information about the type of the central bank (**CB**) or the state of the economy. In this case, the CB can indeed use a non-binding policy announcement to provide a signal about its private information (e.g. Stein, 1989; Cukierman, 1992; Walsh, 1999; and Persson and Tabellini, 1993). Thus, the observed announcement allows for a better prediction of the CB's decision about the actual inflation rate.

In addition to models that studied reputation building in a world of rational policy makers and rational agents, there have been a number of papers that studied how learning affects the outcomes of the Kydland and Prescott environments, see, for example, Sargent (1999), Cho et al. (2002), and Cho and Sargent (1997). This literature uses recursive least squares, constant gain and stochastic gradient learning and 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. On the other hand, evidence from the experiments with human subjects (Arifovic and Sargent, 2003) subjects in Kydland and Prescott environments shows that Ramsey outcomes emerge most of the time and that the experimental economies only occasionally slide into the Nash equilibrium.

In a continuous time framework, Dawid and Deissenberg (2005) study outcomes of the model in which there is a continuum of two types of atomistic private agents, *believers* and *nonbelievers,* whose relative proportion in the population is constant over time. The model also includes a rational CB that is able to correctly predict the impact of its actions on agents' forecasts. Providing that there is sufficiently high stock of believers, the economy can reach a steady state that is Pareto superior to the Nash equilibrium.

While the fraction of believers and nonbelievers is fixed in Dawid and Deissenberg (2005), Arifovic et al. (2010) assume that each private agent can choose between two strategies, to believe or not to the CB's inflation announcement. Believers take the inflation announcement at the face value, and set their inflation forecast equal to the announced inflation. Nonbelievers use an adaptive learning scheme to revise their forecasts in each time period. The fraction of believers and nonbelievers changes over time in response to their relative performance regarding their forecast errors. In this more complex world, the CB can no longer make its decisions based on a solution to a dynamic programming problem. Instead, it uses individual evolutionary learning (Arifovic and Ledyard, 2004) to update its strategy set. Each strategy is evaluated using a *pseudo* value function which, in addition to current inflation and unemployment takes into account the strategy's intertemporal effect.

The CB learns to sustain an economy with a positive fraction of believers. This outcome is Pareto superior to the outcome predicted by standard theory. In addition, the economy has to have 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 develop an agent–based model inspired by Arifovic et al (2010) in which the CB adaptively chooses strategies on the basis of its (adaptive) expectations about the future consequences of these strategies. In order to keep our results comparable with the results of the original model, we do not change the type of private agents, believers and nonbelievers, and their behavior.

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. This collection of strategies evolves over time. However, we depart from Arifovic et al. in a very important way in how we evaluate and choose CB's strategies, as we do not use the pseudo-value function that contains an element borrowed from a solution to the dynamic programming problem.

Instead, our CB learns about the future consequences of its decisions through the adaptation of their 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. (2013): 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, the strategy with the highest expected future performance. 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. The adoption of the strategy that is used in each period, and also the replication of the strategies in the collection, is based on their expected performance.

Our approach hence complements 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 provide sensitivity analysis for other parameters of the model, using a standard Monte Carlo simulations.

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: the CB is capable of learning the Stackelberg-like behavior. Using its mental model it learns what the nonbelievers reaction function is, and thus, can take into account the future reaction of all agents (believers and nonbelievers) to its current behavior. More remarkably, it takes into account what the future response of nonbelievers will be if it deviates from the announced inflation rate.

Furthermore, the economy performs 'better' in the absence of believers, i.e. it reaches Ramsey outcome with a much higher frequency. In response to nonbelievers reaction to its policies, the CB learns to act in a credible manner. The presence of believers reduces the occurrence of Ramsey outcomes somewhat even though it is still the most frequent outcome. The presence of believers is tempting for the CB and it does deviate from Ramsey from time to time. These deviations can result in the outcomes where there is inflation as well as the outcomes where the CB deflates the economy.

The issue of CB's credibility, how to achieve it and what the best way of communicating its policies to the public is have attracted a lot of interest among the policy makers as well as in academia. How does credibility evolve and can it evolve? These are very difficult questions that have been hard to address in economic models . In variants of Barro-Gordon type of environments there is multiplicity of equilibria, some of which can result in Ramsey outcomes. Models with non-binding policy announcements and rational agents can result in improvement upon Nash. As mentioned above, dynamics of the models of learning and adaptation are characterized by 'escape dynamics' where the economies occasionally (infrequently) escape from Nash.

In addition to ours, the paper that addresses the issue of evolution of credibility is Arifovic et al. (2010). Our paper addresses this question and demonstrates how the CB can successfully attain credibility. From the results of Arifovic et al., as well as our own, it is clear that there has to be the right type of interaction between the policy maker and the private agents in order to give credibility a chance.

Recently, the New Keynesian types of environments have been used to discuss issues related to the effectiveness of monetary policies (this is the only type of environment where policy can matter because of the monopolistic competition and price stickiness). However, in all of these studies, the issue of credibility has not been brought out explicitly. There has been discussion of what type of communication might lead to better outcomes. But, the issue of credibility and communication go hand in hand, as effective communication cannot be achieved unless there is credibility in CB's actions. From this perspective, we see our study as the first one to demonstrate the dynamics that can lead to the evolution of credibility that can bring about better outcomes.

Overall, the qualitative features of our results match those of the experiments with human subjects conducted by Arifovic and Sargent. They report that most of the experimental economies gradually reach the Ramsey outcome with occasional backsliding to the Nash equilibrium.

Macroeconomic experiments with policy implications have recently recently attracted a lot interest among researchers (e.g. Pfajfar and Zakelj, 2010; Assenza et al, 2011; Petersen and Kryvstov, 2013) However, most of these experiments study the behavior of private agents while the role of a policy maker is automated and a robot implements a particular type of a rule. The objective is to study how human subjects expectations respond and adjust in response to these rules. Of course, unlike the models that are based on the rational expectations hypothesis, these experimental economies exhibit heterogeneity of expectations that are suboptimal. While experiments with human subjects are a great methodological tool that we can use to explore the behavior of economic agents, this research agenda should be accompanied by agent-based models as these represent an appropriate theoretical framework capable of handling the issues of suboptimal, heterogeneous behavior. [1]

The rest of the paper is organized as follows. In section 2, we describe the agents' and the CB's behavior and learning. We provide a description of our simulation protocol in section 3 and we describe and discuss our results in sections 4 and 5. Concluding remarks are given in section 6.

2 Agents' behavior and learning

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). Our model directly follows Arifovic et al. (2010)[2], but reconsiders the potential role of an expectations based learning of the CB.

In this framework, 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. The agents compute their inflation expectations after having observed the inflation rate announced by the CB. Then the latter implements the actual inflation rate, and the actual unemployment rate of the period depends positively on the expectational errors of the private agents.

Our representation of the agents' behavior directly follows Arifovic et all (2010): we introduce two types of private agents (*believers* and *unbelievers*), and the specific learning and expectations rules that they use.

However, we significantly modify the behavior and the learning process of the CB, since, instead of introducing a component trying to approximate the standard dynamic optimization solution, we introduce a behavior guided by expectations formed using a *mental model*. The learning of the CB takes place both through the exploration of the strategy space like in Arifovic et al. (2010), and through the updating of its expectations on the outcomes of these strategies. In each period, the CB adopts the strategy/policy with the highest expected performance. Then, following the observed expectational errors (the gap between the expected and observed performance), the CB updates its mental model.

[1] The exception is Assenza et al (2011) that uses a simple model of individual learning, with performance based evolutionary selection to explain coordination of individual expectations and aggregate macro behavior observed in the laboratory experiments.

[2] *See* also Vallée (1999) for a version of the Barro-Gordon model with a simpler learning process.

2.1 Central Bank (CB)

At the beginning of each period, the CB communicates an inflation rate (y^a), hoping that this communication will be able to influence the expectations of the private agents in a direction of the CB's policy objective $\left(J^G\right)$ that takes into account the unemployment (u) and actual inflation (y) rates of the period:

$$J_t^G = -\frac{1}{2}\left(u_t^2 + y_t^2\right) \tag{1}$$

The standard approach to this game where agents are rational takes the form of a dynamic optimization by the CB to determine the strategies that would be used over time. We adopt another approach based on bounded rationality *à la* Simon, taking also into account the forward looking behavior of the CB. But, differently than Arifovic et all (2010), we do not model this forward looking behavior in direct similarity with dynamic optimization, but using a *mental model*, that allows the CB to form explicit expectations about the potential consequences of different available strategies. Thus, this mental model allows the CB to *ex ante* evaluate all already discovered strategies, and it evolves following the learning process of the CB in response to its expectational errors (see section 2.2 for a detailed presentation of this process). The approach adopted here is hence very close to the one studied in Yildizoglu et al. (2013).

In each period t, the CB has a population of strategies $Y_t = \{(y_t^a(j), y_t(j))\}_{j=1\ldots N}$. Each strategy consists of two elements: the announced inflation rate (y^a), and the actual rate of inflation (y). The CB computes an expected value $V_t^e(j)$ (*see* below and the next section) for each available strategy, and the CB chooses to use the strategy with the highest expected value in period t.

The expected performance of strategy j, $J_{t+1}^{G,e}$, is defined as a measure of the performance that the CB expects to achieve during the next period if it uses strategy j. For technical reasons, we do not directly use this performance in the learning process of the CB, and use a fitness function, $f(\bullet) \in]0,1[$, to normalize it:

$$V_t^e = f\left(J_{t+1}^{G,e}\right) \tag{2}$$

This fitness guides the exploration by the CB of its strategy space, through the evolution of its population of strategies.

Every *GARate* periods[3], the CB revises its population of strategies Y_t through the following three steps:

1. Replicating the strategies for the next experimentation period, 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 $\left(V_l^e / \sum_j V_j^e\right)$.

[3] 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.

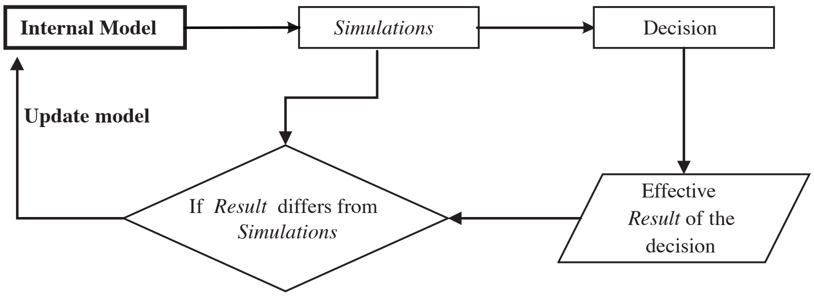

Figure 1: Dynamics of the mental model of the agents

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, 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 by perturbating it with a gaussian noise

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

How can we represent the formulation of the expectations that guide the learning process and the decisions of the CB?

2.2 Learning of the CB with adaptive expectations

In a dynamic setting, we can expect that the CB's decisions are based on the projection of past performances on the future circumstances. Such a projection would require a capacity of the CB to *generalize* or form *expectations*: an ability of attributing an expected performance even to strategies that had not been previously used by the CB.

This generalization would in turn entail that the CB develops a representation of her environment, i.e. her *mental model*. We closely follow here Yildizoglu (2001) and Yildizoglu et al. (2013) in representation of this mental model in the following way:

The mental model of the CB summarizes the state of its knowledge and evolves as a consequence of evolution of this knowledge. It guides the decision process since it enables the CB to test the connections between the alternatives of choice and their consequences. The presence of such a mental model can reflect the intentionality of decisions. 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 (Yildizoglu 2001).

While this line of thought is quite natural, the economic literature has not found a definitive way of incorporating it in models. This is the reason why purely adaptive models (see the preceding section) generally neglect the dynamic process of expectation formation. Evolutionary learning may result in better decision rules, through trial and error, and evaluation that is solely based on past performance. In this case, the agent can only judge decisions which have been used before. On the contrary, the approach based on the dynamics of the mental model takes into account that agents can have a relatively precise (if not perfect) perception of the value of their decisions, even if they have never been used before. This is made possible by means of simulations using the mental model. The ability of the mental model to "generalize" from past observations will determine the quality of these expectations. If the model is "overfit" for these observations, it will give precise expectations for the situations already met in the past, but it will frequently fail for completely new situations. A more general model should be able to also give meaningful expectations for the latter.

The standard way of formalizing such a model is to rely upon the subjective probabilities approach of Savage. In this case, the mental model of the agent corresponds to a set of conditional probability distributions. The update of this model can be imagined through successive least square estimations or applications of Bayes' rule. The Bayesian approach has the advantage of not assuming any particular structure for the mental model. But it is very demanding in terms of agents' rationality. Moreover, "there is substantial evidence that Bayes' theorem lacks empirical relevance and hence its procedural justification is weak" (Salmon 1995, p.245).

Recursive least square estimations have been used, in this perspective, albeit at the aggregate level, by the recent macroeconomic learning literature (Evans and Honkapohja, 2001). However, this method relies upon a specific functional structure for the mental model. We adopt, here, a more flexible tool. 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. In this respect, an artificial neural network (ANN) is a good candidate for representing the role of the mental model, and its adaptive nature. With only minimal structural assumptions, namely the list of endogenous and explicative variables, and the structure of the hidden layer, it can represent the fact that the agent adjusts her mental model to the flow of experience. For many practical problems, even a very simple feed forward ANN with one hidden layer of few hidden nodes gives quite robust results (*see* Masters (1993), for the discussion of properties of ANNs).

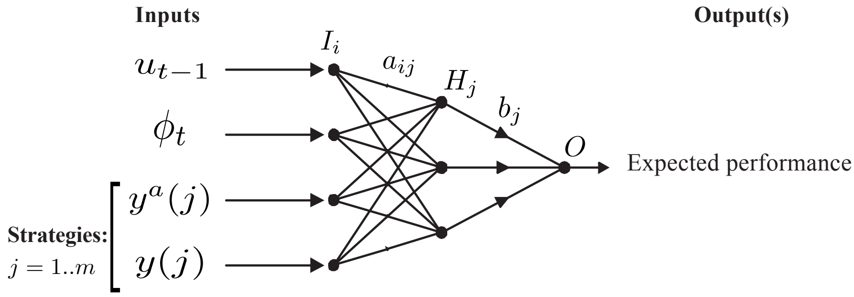

Figure 2: A feed forward ANN with one hidden layer (and with only two inputs for illustration)

Another potentially interesting modeling approach is the learning classifier systems (LCS).[4] A complete LCS, combining a generalization capability with a reinforcement learning mechanism (like the XCS developed by Wilson (1995)) could model context dependent choice of strategies by the agents. The main limit of this approach, in the context of our discussion, is the fact that expectations included in the rules of the XCS are necessarily implicit, and it is impossible to separate them from the actions. The use of a mental model represented by an ANN allows for such a separation. Moreover, the behaviors modeled using this representation seem more realistic (exhibiting some inertia, but also, performance increasing in time.

More particularly, an ANN 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 2) according to the series of input and output data, submitted to the ANN in successive training periods. To train the ANN, the complete past history of inputs and outputs can be used, or only observations for a given number of past periods ($windowSize$). In each training period (an *epoch*), a number of passes ($numEpoch$) are executed through the ANN in order to correct the error observed between the observed outputs and the predicted ones.

With each pass, the parameters a_{ij}, b_j are adjusted in order to correct a fraction $learnRate$ of the residual error. The goal of this adjustment process is to try to minimize the ANN's prediction errors. This process should result in the ANN's better adaptation to the environment.

Parameters a_{ij}, b_j reflect the intensity of the connections in the network. A better approximation can be achieved through the introduction of hidden nodes in the network, that is nodes that represent unobserved state variables or, more particularly, unobserved variables of the mental model of the agent. ANN thus covers a wide range of models 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

[4]See, for example, Lettau and Uhlig (1999) for an application of a very simplified LCS to the standard intertemporal consumption decision problem.

even be used to represent the complexity of the agent's mental model.

In our case the CB is placed in a very simple economic context, and can only observe a very limited set of variables on the state of the economy. We assume that the CB takes into account the following variables as inputs of its mental model:

1. Level of unemployment: u_{t-1};

2. Composition of the agents population ϕ_t;

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 literature, since it should normally influence the dynamics of the model that the CB is learning to anticipate. 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 these variables in the ANN is endogeneous, since their weight may become zero, if the CB ends by discovering that they do not play a real role.

The CB supplements these economic indicators with the two components of its strategies: $j = (y^a(j), y(j))$. These seven elements then constitute the inputs of its mental model. We will also consider different assumptions about the information level of the CB, starting with a very limited information only based on the strategy of the CB, and progressively incorporating more elements from the above indicators list. The strategy component of the inputs is used by it to compare different potential strategies on the basis of the expected performance they can yield (hence the single output of the mental model in Figure 2). This comparison serves as a basis for selecting the strategy that will be used in the current period, after the observation of the corresponding performance.

More particularly, at each period t, the CB uses the ANN as follows. At the beginning of the period, it compares strategies on the basis of the expected performance resulting from them. It *feeds* the ANN with the state of the environment, and the components of each strategy, and observes the performance predicted by the ANN (*see* (2)). After having compared the expected performance corresponding all strategies, it adopts the strategy that yields the highest expected performance.

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. For this training, at the end of the period t, the CB computes the difference (error) between, on the one hand, the expectations formed and used in period t, and on the other hand, the observation of the performance observed at the end of the same period. Then, it updates its mental model (*see* also Figure 3).

2.3 Private agents

Agents form inflation expectations (x) after having observed the inflation rate announced by the CB.

As in Arifovic et al. (2010), our model includes two types of private agents: the *believers* who set their expectations, x^B, equal to the announced inflation y^a:

$$x^i = y^a, i \in B$$

and *non-believers* ($i \in NB$) who form adaptive expectations by correcting their forecast errors around the optimal solution of the static game:

$$x_t^{NB,i} = \frac{\theta^2 \phi y_t^a + \theta u^*}{1+\theta^2\phi} + d_t^i \tag{3}$$

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

$$d_{t+1}^i = d_t^i + \gamma\left(y_t - x_t^i\right),\ d_0^i = d_0 = 0 \tag{4}$$

and $\gamma > 0$ represents the speed of learning of the agents.

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

$$J^i = -\frac{1}{2}\left[\left(y - x^i\right)^2 + y^2\right] - c^i \tag{5}$$

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

The unemployment rate of each type of agent is determined by the following augmented Philips curve

$$u^i = u^* - \theta\left(y - \overline{x}^i\right), i \in B, NB \tag{6}$$

where u^* is the unemployment rate that would arise if agents had correct inflation expectations, $\theta > 0$ is a parameter, and $\overline{x}^i$ represents the average expectation in the corresponding set of agents. The fraction of believers in the population is denoted by

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)	
	$i = B, NB$ **Observe:** y^a_t, y_{t-1} **Set:** d_t, x^i_t
Observe: $u_t, J^G_t, \phi_{t+1},$	**Observe:** y_t, u^i_t, J^i_t
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 3: Sequence of decisions in each period

ϕ. Thus, the unemployment rate that determines the performance of the CB (equation (1)) depends on the composition of the agent population

$$J^G_t = -\frac{1}{2}\left(\phi\left(u^B\right)^2 + (1-\phi)\left(u^{NB}\right)^2 + y^2_t\right) \tag{7}$$

2.4 Evolution of the population of agents

The population of agents evolves in time: in each period a proportion β of the agents are drawn randomly and they meet other agents. If their utility level is lower than the agent they meet, they adopt its behavior as a believer or nonbeliever. The switching agent starts from scratch with $d_{it} = 0$. This imitation process determines the dynamics of ϕ.

Figure 3 describes the decisions and observations made by all agents, as they take place sequentially in each period, and the pseudo code of the model is given in Figure 11, in the Appendix.

3 Simulation protocol

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

- $N = 100$
- $u^* = 5.5$;
- $\beta = 1\%$;
- $\gamma = 0.1$;
- $\theta = 1$;

- $c^{NB} = 3.3$
- Initialization and mutations: y^a and $y \in [-10, 15]$.

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

$$m = 400\,;\, gaRate = 4\,;\, nbHidden = 4\,;\, p^C = 0.3\,;\, p^m = 0.05.$$

We vary the amount and type of information, *level of information*, that the CB's mental model uses at its inputs. These include:

For each level of information we ran 50 simulations where each simulation lasted for 5000 periods. [5]

For each level of information, we represent the distribution of all variables and indicators using density figures: the performance of the CB and its inflation strategy. The density plots give one-dimensional kernel density estimation of the distribution of the concerned 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.

Because our simulations are quite long and contain a large number of observations, we present the data of a random subsample (with a sampling rate of 5%) in each figure.

We discard the data generated in the first $1,000$ 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,

we present data of the last 4000 periods only. When we analyze the dynamics of the learning, we represent the evolution of strategies over time, and, in this case, we plot the per-period average of observations over all runs for each variable.

4 Results

The first issue we address concerns the behavior of this economy under the baseline setup adopted by Arifovic et al. (2010). Does the economy converge on a specific outcome, if yes, is it able to attain the Ramsey solution?

4.1 Convergence to the Ramsey outcome?

The Figure 4–(a) plots the distribution of the CB's payoff and compares it to the Nash and Ramsey values. This figure clearly shows that the economy spends a dominant share of its history close to the Ramsey outcome, and very rarely close to the Nash outcome. We observe quite a strong convergence onto the Ramsey outcome, at least from the point of view of the CB's payoff. We can also already note the small bump that appears on the right of the Ramsey outcome. We will discuss later the signification of these observations.

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

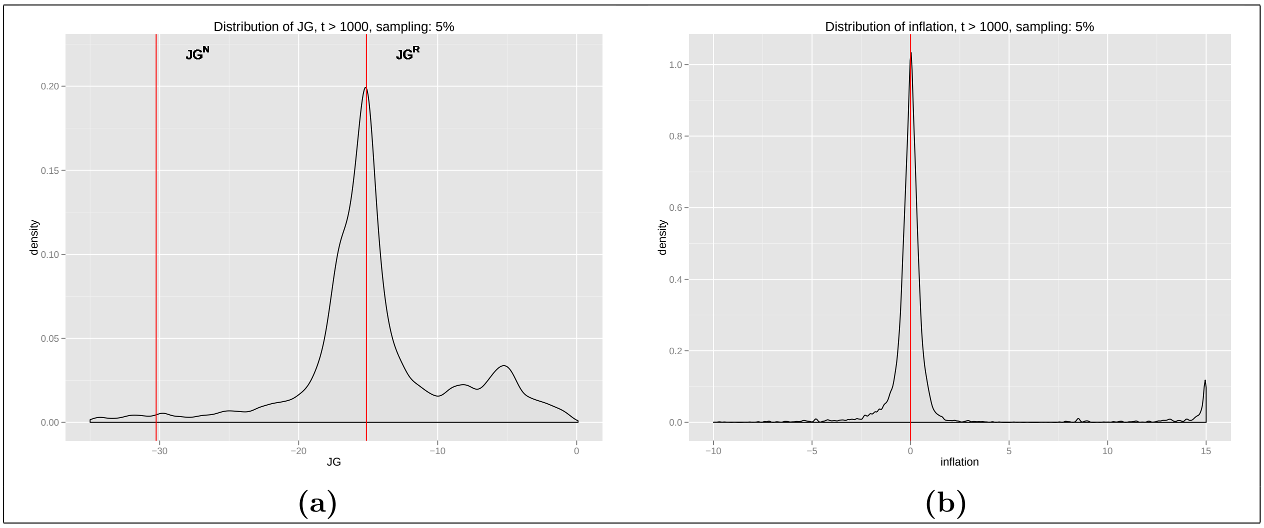

Figure 4: Convergence of J^G and inflation: Nash or Ramsey?

Is the result on J^G is confirmed by the distribution of the inflation rates implemented by the CB? Figure 4–(b) clearly shows the convergence onto the zero–inflation outcome.

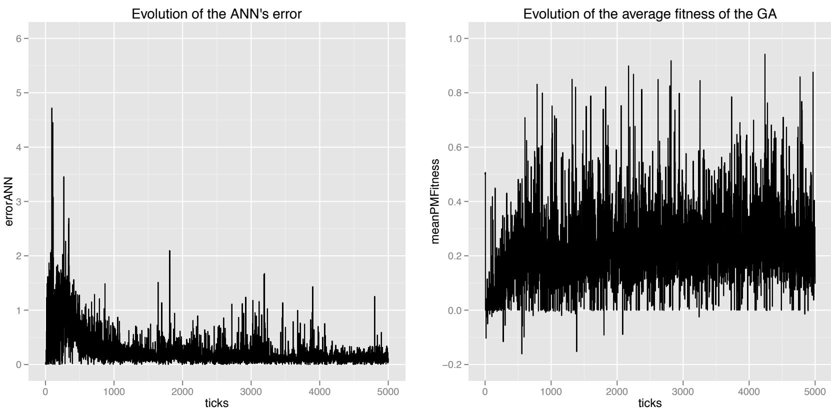

Figure 5: Learning of the CB: Evolution of the errors and the average fitnes of strategies

4.2 Learning of the CB?

How is this economy able to discover Ramsey outcome and to stay 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 5 represents the evolution of the CB's total prediction error (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 agents of both type. In parallel, the plot on the right side of Figure 5 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 confront the evolution of the ANN error, with the forecast error of the nonbelievers and the inflation strategy used by the CB. Figure 6 clearly shows the co-evolution of these variables: the CB discovers the favorable role of zero inflation, and the nonbelievers progressively adapt to this strategy, their errors decreases in parallel with the error of the mental model of the CB. These medians, as well as data from individual runs show that the convergence to zero inflation is not complete, and the CB continues to cheat in a transitory way, but always come back to zero inflation, and keeps the economy in the vicinity of the Ramsey outcome.

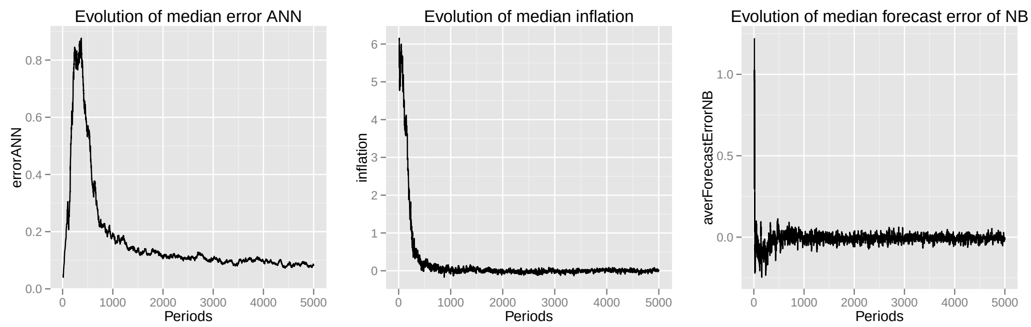

Figure 6: Evolution of expectation errors, and inflation strategy

Consequently these results show that the CB learns to implement the Stackelberg solution of the game (with full reaction on the side of the nonbelievers, including their error correction), and pull the economy towards the Ramsey outcome. This result is very new in the literature and this systematic convergence to the Ramsey outcome is not common in models with optimal control, or adaptive learning. The expectation forming ability of the CB, and the flexibility allowed in these expectations through the ANN is the main novelty in this model, and can be considered as the source of this surprising result.

But are these outcomes really desirable for the agents? What can be said about their welfare when the CB can obtain a payoff above the Ramsey J^G? Figure 7 plots the distribution of the global average welfare of the agents in two types of situation: when the CB's payoff is above the Ramsey one ($overRamsey = 1$), and when the CB is below the Ramsey payoff. We observe that agents average welfare maybe higher when the CB is above the Ramsey outcome, and in a significant number of cases, the nonbelievers can even attain their highest utility, which corresponds to $-c^{NB}$, when their expectation error is zero.[6] Consequently,the CB is able to guide the economy to a

[6]We should note that this result does not imply that these cases Pareto-dominate Ramsey, since the believers are necessarily worse off than in the Ramsey outcome.

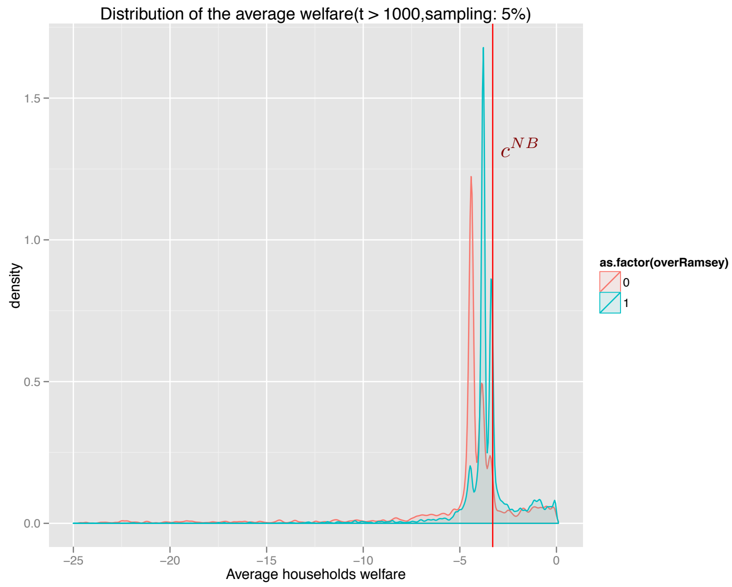

Figure 7: Agents welfare above the Ramsey outcome for the CB

Stackelberg/Ramsey solution that is indeed the first best for all agents.

This result is much favorable than the ones obtained in the literature, including the models where the believers are included to motivate the CB to adapt strategies closer to the Ramsey outcome.

4.3 Do we need the believers for Ramsey?

Are believers necessary for getting at the Ramsey outcome in this economy? To check this, we have done a batch of runs in which we have fixed the number of believers to the lowest number possible in our setting: 1/100.

Figure 8 plots the distribution of the J^G in this particular situation with the minimal number of believers. We observe that the convergence onto the Ramsey outcome is perfectly possible and very frequent in this case too. If we compare this plot with Figure 4–(a), we observe that the hump at the right of 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.

4.4 Role of the CB's information level

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 our reference literature, we have chosen to let the CB observe the composition of the agent population (through the

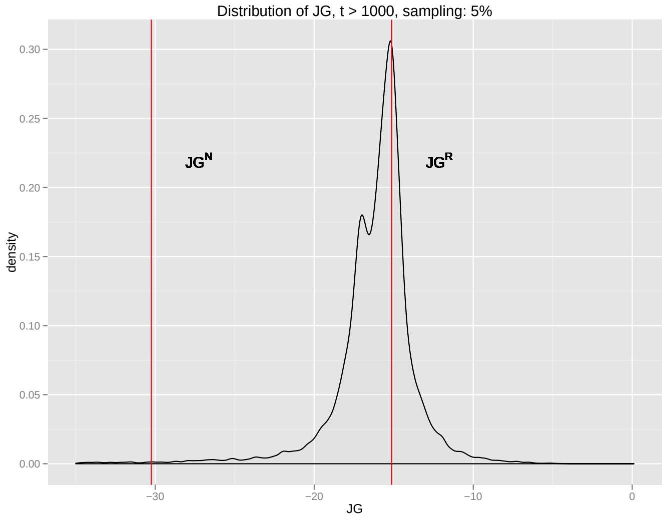

Figure 8: Are the believers necessary for the Ramsey outcome ($\phi = 1\%$)?

proportion of believers, ϕ). Is this information necessary? One could consider that it is not necessarily easily available to the CB. To check the robustness of our results, and the role played by the level of information, we compare the distribution of the J^G resulting from three different levels of information, corresponding to increasingly richer sets of variables that can be observed by the CB, represented by the increasing values of the variable $infoLevel$:

1. (y^a, y)
2. $(y^a, y, u_{t-1}, \phi_t)$
3. $(y^a, y, u_{t-1}, \phi_t, \Delta\phi_t, \Delta y_t, \Delta u_t)$

Figure 9 shows that even a minimal level of information (when the CB only observes its decision variables and their consequences in terms of payoff) allows the attainment of the Ramsey outcome with a remarkable frequency. Taking into account the level of unemployment and ϕ makes this convergence much stronger, while the inclusion of the change direction of these variables secures an even stronger convergence. We can conclude that higher levels of information (including ϕ) can favor convergence, without being indispensable for it.

5 Discussion: Why Ramsey?

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) to which our model is very close, outcomes better than Nash can be observed, but not the convergence to Ramsey outcome. Given

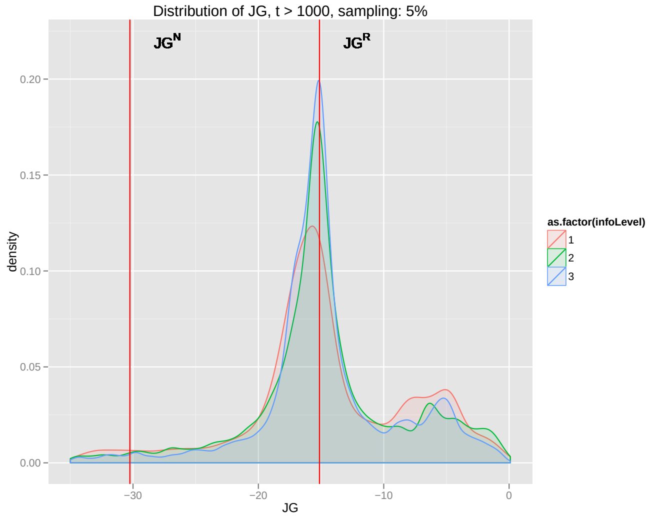

Figure 9: Role of the CB's information level

that the nonbelievers mainly play their static best reply, convergence to Nash can indeed seem quite natural. Moreover, we obtain convergence on Ramsey even in the absence of the believers (while they are necessary for beating 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 nonbelievers to correct their errors. Because of this mechanisms, the CB ends by learning to expect the full reaction function of the nonbelievers (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 process (coevolution of behaviors), its expectation errors, as well as nonbelievers forecast errors decrease over time and the economy converges to zero–inflation Ramsey–Stackelberg equilibrium (see Figure 6).

Figure 10 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 nonbelievers: they can only play their best reply function (3) with $d = Cste = 0$. These figures exhibit convergence on the Nash equilibrium (panel a), with the persistence of positive rates of inflation (panel b). Consequently, the ability of the nonbelievers 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 nonbelievers 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 nonbelievers, and the payoff it can get by moving after their choice,

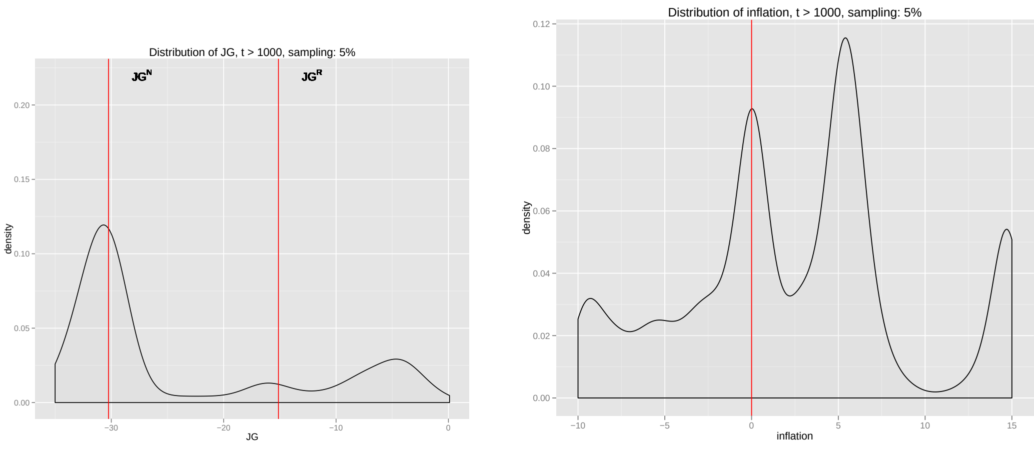

(a) (b)

Figure 10: Distribution of J^G and y when nonbelievers do not correct their errors

and by increasing the inflation. We observe, in this case, that it frequently deviates as illustrated in Figure 10–(b). The mode 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).

Moreover, as a cororally of the discussion in paragraph 4.3, 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 nonbelievers, and pull the economy on the Ramsey outcome.

When there is the nonbelievers' error correction mechanism in place, Nash outcomes cannot be supported any more by the learning dynamics as nonbelievers *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 nonbelievers, and, consequently, learns to take decisions under this constraint when choosing the inflation rate. The nonbelievers' 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–Stackelberg equilibrium, instead of the Nash equilibrium.

6 Conclusion

We study an extension of the Kydland-Prescott environment with two types of agents, believers and nonbelievers. While believers set their inflation expectations equal to the inflation announced by the central bank, nonbelievers update their inflation expectation every period. The fraction of believers and nonbelievers changes over time in response to their relative performance in terms of forecasting accuracy. The central bank has a collection of strategies that evolve over time. Each strategy consists of two elements, the announced and the actual inflation. The central bank evaluates 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 central bank 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 central bank brings the economy to the neighbourhood of the Ramsey outcome. The frequency of Ramsey outcomes 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 results in Nash outcomes. The exception is Arifovic et al. (2010) that obtain better than Nash outcomes. However, to 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 policy maker as 'truly' forward looking, but still adaptive. 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 '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

Assenza, T., P. Heemijer, C. Hommes, and D. Massari, 2013, Individual Expectations and Aggregate Macro Behavior, manuscript, University of Amsterdam.

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., 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.

Kryvstov, O. and L. Petersen, 2013, Expectations and Monetary Policy: Experimental Evidence, working paper.

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.

Pfajfer, D. and B. Zakelj, 2010, Inflation Expectations and Monetary Policy Design: Evidence from Laboratory, manuscript, Tilburg University.

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.

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

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

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, forthcoming, Learning the optimal buffer-stock consumption rule of Carroll, forthcoming in *Macroeconomic Dynamics*, 2013.

A Appendix: Pseudo code of the model

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} nonbelievers, 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 11: Pseudo code of the model.