# GREQAM

**Groupement de Recherche en Economie
Quantitative d'Aix-Marseille - UMR-CNRS 6579
Ecole des Hautes Etudes en Sciences Sociales
Universités d'Aix-Marseille II et III**

# Learning Benevolent Leadership in a Heterogenous Agents Economy

**Jasmina ARIFOVIC
Herbert DAWID
Christophe DEISSENBERG
Olena KOSTYSHYNA**

**22 Octobre 2008**

# Learning Benevolent Leadership in a Heterogenous Agents Economy

**Jasmina Arifovic**[*]     **Herbert Dawid**[†]     **Christophe Deissenberg**[‡]

**Olena Kostyshyna**[§]

October 22, 2008

## Abstract

This paper studies the potential commitment value of cheap talk inflation announcements in an agent-based dynamic extension of the Kydland-Prescott model. In every period, the policy maker makes a non-binding inflation announcement before setting the actual inflation rate. It updates its decisions using individual evolutionary learning. The private agents can choose between two different forecasting strategies: They can either set their forecast equal to the announcement or compute it, at a cost, using an adaptive learning scheme. They switch between these two strategies as a function of information about the associated payoffs they obtain through word-of-mouth, choosing always the currently most favorable one. We show that the policy maker is able to sustain a situation with a positive but fluctuating fraction of believers. This equilibrium is Pareto superior to the outcome predicted by standard theory. The influence of changes in key parameters and the impact of transmission of information among nonbelievers on the dynamics are studied.

*Keywords*: time inconsistency, bounded rationality, forecast and agent heterogeneity, cheap talk, evolutionary learning

*JEL codes*: C69, D83, E5

---

[*]Corresponding Author: Department of Economics, Simon Fraser University, Burnaby, BC V5A 1S6, Canada. Phone: + 1604 291-5603, Fax : +1604 291-5944, Email: arifovic@sfu.ca.

[†]Bielefeld University

[‡]Université de la Méditerranée (Aix-Marseille II) and GREQAM

[§]Portland State University

# 1 Introduction

In the last three decades or so, time inconsistency and its potential consequences have been a recurrent issue in economics. In a nutshell, a solution is time inconsistent if the policy maker has incentives to deviate ex post from the decisions that were initially optimal. The time inconsistent solution is often Pareto optimal. However, it is not credible since rational agents recognize ex ante that it will not be implemented as announced. Thus, the economy is likely to end up at a Pareto inferior but time consistent equilibrium. This equilibrium will often prove to be the Nash solution of a game between the policy maker and the private agents.

The issue has been exhaustively studied in the monetary policy literature following Kydland and Prescott (1977) and Barro and Gordon (1983), giving rise to a vast literature on credibility building and other means to practically mitigate the negative consequences of time inconsistency. In particular, diverse studies investigate 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 central banker or the state of the economy. In this case, the policy maker can indeed use a non binding policy announcement that provides a signal about policy maker's private information, see Stein (1989), Cukierman (1992), Walsh (1999), Persson and Tabellini (1993)). Hence, observing the announcement allows for a better prediction of the policy maker's decision concerning the actual inflation rate.

In this paper, we also investigate the role of cheap talk announcements by the policy maker as a means to sustain a Pareto superior solution in economy closely related to the model of Kydland-Prescott (1977). The mechanisms at play, however, are not related to asymmetric information in the conventional sense. Following Deissenberg and Alvarez (2002), Dawid and Deissenberg (2005), we assume that each private agent can choose in any period between two strategies: believe, that is, act as if the policy announcement was true; or not believe and compute a potentially exact forecast of the policy maker's next action. Word of mouth information exchange allows, in each period, a fraction of the agents to compare the past payoffs obtained by following the one or the other strategy. The agent then adopts the strategy that fared best in this comparison, and uses it until new information motivates it to switch strategies anew. Thus, the proportions of agents using the believe strategy may vary over time and can be interpreted as a measure of the policy maker's credibility.

In the two papers just mentioned, the policy maker and nonbelievers are fully rational, but do not learn. In this paper, they are boundedly rational but learn, addressing another issue of prime importance in the monetary policy/time inconsistency context, see e.g. Phelan (2005), Sargent (1999) and Cho, Williams and Sargent (2002), Cho and Sargent (1997). The learning mechanisms used in the related literature include principally Bayesian updating, recursive least squares, and stochastic gradient learning. Learning either by the policy maker or the agents, or by both, often narrows the choice of equilibria down to the Nash equilibrium and to the time inconsistent but Pareto optimal Ramsey outcome. However, simulations of the models show that the Nash equilibrium occurs most frequently. This contradicts Arifovic and Sargent (2003), who find from experiments with human agents using Kydland-Prescott model that the Ramsey outcome emerges most of the time, and that the economy only occasionally slides to the Nash equilibrium.

Departing from the existing literature, we assume that the nonbelievers use a simple adaptive learning mechanism to update their forecasts. The policy maker uses individual evolutionary learning, Arifovic and Ledyard (2004). As the authors have shown, and contrary to most evolutionary learning procedures suggested in the literature, individual evolutionary learning handles environments with large strategy spaces well and mimics the behavior observed in experiments with human subjects. Moreover, contrary to other arguably better known learning procedures, it does not model only the refinement of a given behavior as additional information become available. Rather, it captures a process of purposive exploration of a a priori loosely determined space of possible rules of action. In that way, it will allow us to show that a policy maker can find out how to efficiently use cheap talk announcements to increase its payoffs, rather than only demonstrating that learning can fine-tune the parameters of a given decision-making rule. That is, it will permit us to, in a limited way, document on the emergence of cheap talk announcements as an effective policy tool. As an additional benefit, our results will prove closely related with, or at least not contradictory to, the experimental ones.

Still another central issue addressed in the paper is the importance of the heterogeneity in agents, forecasts, and forecasting rules. At the empirical level diverse studies, surveyed by Branch (2004a), find persistent

heterogeneity in survey data on inflationary expectations. But heterogeneity has become also an increasingly influential strand of research in macroeconomics, following among others Evans and Ramey, (1992, 1998), Brock and Hommes (1997, 1998), Branch and Evans (2006b, 2007), Berardi (2007). In these models, the agents have different forecasting models. Equilibria with heterogeneity of beliefs, i.e. with positive fractions of agents using different forecasting models, can arise. Interesting phenomena take place that are not present in the standard models with homogeneous expectations. See Branch (2004) on the importance of explicitly modeling heterogeneity, and Giannitsarou (2003) and Honkapohja and Mitra (2005, 2006) on different possible forms of heterogeneity.

The inclusion of two forecasting models (believe/not believe) allows us to contribute to the discussion concerning the importance of an explicit representation of varying degrees of agents heterogeneity in macroeconomic models. So does also the comparison we conduct between two different forecast formation schemes (common and private) for the nonbelievers. In the first case, all nonbelievers share a common forecast based on a record of all previous actions of the policy maker. In the second, each agent starts learning anew, without using prior information, each time it switches to nonbelieving. This makes nonbelieving less attractive for the private agents, but also complicates the task of the policy makers. We will see that this has important implications for the outcome, and raises difficult questions on whether or not the policy maker should facilitate information exchange among agents.

In a nutshell, the main research questions addressed in the paper are: Can the policy maker, although it has very little a priori and online information and is only boundedly rational, build up and nurture a positive stock of believers? If so, does it lead to Pareto superior outcomes, and under which conditions? What is the impact of variations of the main model parameters and of the forecasting schemes on the outcome? Are there preferable (ranges of) values for these parameters? In particular, how much concern for the future should the policy maker have? What is the preferable speed of learning for the believers? Should the policy maker favor information exchange among the nonbelievers, and how?

The paper is organized as follows. In section 2, we present the static game underlying the agent-based model, discuss briefly its dynamic extension in Dawid and Deissenberg (2005) and present the agent based model. Sections 3 and 4 are devoted to the presentation and discussion of the simulation results. The final section 5 briefly summarizes the framework and the main results, and succinctly concludes on the importance of taking heterogeneity explicitly into account in economic analysis.

## 2 The Model

### 2.1 The static game

Before starting with the agent-based dynamic analysis, consider first the underlying one-shot game, see Dawid and Deissenberg (2005). The economy, a la Kydland-Prescott (1977), consists of a policy maker $G$ and of a continuum of (private) agents $i$, $i \in [0, 1]$. Agent's $i$ unemployment rate $u^i$ is given by an expectation augmented Phillips curve:

$$u^i = u^* - \theta(y - x^i), \tag{1}$$

where $u^*$ is the natural rate of inflation, $u^i$ is agent's $i$ unemployment rate, $x^i$ is the inflation forecasted by agent $i$, $y$ is the true inflation, and $\theta > 0$ is a parameter. The payoff of agent $i$, $J^i$, is given by:

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

where, as discussed below, $c^i \geq 0$ is the cost of forming a forecast. The policy maker's payoff, $J^G$, will be defined at a later place.

The policy maker and the agents play the following three stages, complete information game:

1. The policy maker announces a level of inflation $y^a$. The announcement is non-binding.

2. Each agent $i$ forms its own forecast $x^i$ of the true inflation.

3

3. The policy maker sets the true inflation level, $y$.

Each agent can choose between two strategies: to believe, $B$, or not to believe, $NB$. With the first strategy, the agent sets its inflation forecast equal to the announced inflation:

$$x^i = y^a.$$

We say that the agent is a believer and denote its forecast by $x^B$. The believers do not occur any forecasting costs: $c^B = 0$.

An agent that uses the second strategy is a nonbeliever. The nonbelievers are traditional rational atomistic agents that play Nash against the policy-maker. Doing so, they take into account both the proportion $\phi$ of believers in the economy and their forecasts $x^B$. We consider only symmetric equilibria where all nonbelievers make the same forecast $x^{NB}$. In the Brock-Hommes (1997) tradition , we assume that the forecasting costs of the nonbelievers may be positive, $c^{NB} \geq 0$, due to the computational, informational and other efforts they possibly need to make their own forecasts.

The policy maker's payoff $J^G$ is the sum of the $\phi-$weighted average squared rate of unemployment, and of the squared rate of inflation, i.e.

$$J^G = -\frac{1}{2}[\phi(u^B)^2 + (1 - \phi)(u^{NB})^2 + y^2]. \tag{3}$$

As usual, the game is solved through backwards induction. At the last stage 3, the policy maker knows $\phi$ and the forecasts $x^B$ and $x^{NB}$ of the believers and nonbelievers. Maximizing (3) with respect to $y$ given $\phi$, $x^B$ and $x^{NB}$ leads to the policy maker's optimal reaction function:

$$y = \frac{\theta}{1 + \theta^2}\left[u^* + \theta\phi x^B + \theta(1 - \phi)x^{NB}\right]. \tag{4}$$

At stage 2, as previously stated, the believers form their expectations according to:

$$x^B = y^a. \tag{5}$$

The nonbelievers, on the other hand, know $\phi$, $y^a$, and $x^B$, and the optimal reaction function (4). Each of them is rightly aware that it is too small to influence the policy maker's actions. Thus, the nonbeliever's inflation forecasts are given by

$$x^{NB} = \frac{\theta^2\phi y^a + \theta u^*}{1 + \theta^2\phi}. \tag{6}$$

At stage 1, given the optimal forecast functions of the believers and nonbelievers (5, 6), the policy maker determinates the optimal announcement of inflation, $y^a$, and the optimal actual inflation, $y^*$,

$$y^{a*} = -\frac{u^*}{\theta}, \tag{7}$$

$$y^* = \frac{\theta(1 - \phi)}{1 + \theta^2\phi}u^*. \tag{8}$$

Accordingly, the optimal choices of believers and nonbelievers at equilibrium are

$$x^{B*} = y^{a*},$$
$$x^{NB*} = y^*.$$

Notice that the nonbelievers perfectly forecast the true inflation. Notice also that the true inflation, $y^*$, decreases with $\phi$. So does the difference $y^* - y^{a*}$ between the true and the announced inflation. Since $x^{NB*} = y^{a*}$, this difference is equal to the difference between the forecasts of the nonbelievers and believers, $x^{NB*} - x^{B*}$.

4

Most important for us is the following. At the equilibrium, the policy maker's payoff is

$$J^{G*} = -\frac{1+\theta^2}{1+\theta^2\phi}(1-\phi)u^*, \tag{9}$$

the believers' payoff is

$$J^{B*} = -\frac{1}{2}\frac{(1+2\theta^2+2\theta^4-2\phi\theta^4+\phi^2+2\theta^4)}{\theta^2(1+\phi+2\theta^2)^2}u^*, \tag{10}$$

and the nonbelievers' payoff is

$$J^{NB*} = -\frac{1}{2}y^{*2} = -\frac{1}{2}\left[\frac{\theta(1-\phi)}{1+\theta^2\phi}u^*\right] - c^{NB}. \tag{11}$$

Thus, for all $\phi \in (0,1)$ and sufficiently small forecasting costs $c^{NB}$, the payoff of the nonbelievers is always higher than the payoff of the believers: not-believing is individually rational. However, all payoffs increase with $\phi$: the policy maker, the believers and the non-believers are all better off when the fraction of believers increases. Indeed, as $\phi$ increases, the nonbelievers' unemployment stays at $u^*$. However, the true inflation decreases. This is beneficial for the policy maker, the believers, and the nonbelievers. Furthermore, the believers' unemployment decreases, which is good for them and for the policy maker.

Thus, believing is socially efficient. The higher the proportion of believers, the better every actor is. A situation with a positive fraction of believers Pareto-dominates the one with where all private agents act rationally. But such a situation is not an equilibrium since any individual agent can improve its payoff by switching from believing to nonbelieving.

This framework naturally suggest searching for mechanisms that could sustain a Pareto superior equilibrium with $\phi > 0$. As mentioned in the introduction, we have shown in a previous paper that, in a dynamic setting, cheap talk announcements can provide such a mechanism. Let us have a closer look at this dynamic version, as it provides useful insights on the mechanisms at work in the agent-based model.

## 2.2 An analytical dynamic extension

In a nutshell, assume that (a) the static game just described is played at each instant of time $t$; (b) the proportion of believers, $\phi$, increases over time whenever the instantaneous payoffs received by the believers are larger than the ones received by the nonbelievers, decreases if they are smaller; and (c) the policy maker is aware of the $\phi-$dynamics and solves a standard optimal control problem to maximize with respect to the time paths of the two decision variables $y^a$ and $y$ its cumulated discounted stream of payoffs. The nonbelievers build their expectations according to the equilibrium forecast of the static game given by (6) and do not learn.

The dynamic model always has a stable equilibrium at $\phi^0 = 0$. However, if the policy maker is sufficiently patient (i.e., uses a sufficiently low rate of time preference), the model also admits an equilibrium at some interior value of $\phi$, $\phi^E$, with $0 < \phi^E < 1$. The instantaneous payoffs of the policy maker and of the private agents are larger at $\phi^E$ than at $\phi^0$.

A threshold $\phi^S$, $0 < \phi^S < \phi^E$, separates the basins of attraction of the two equilibria: If the initial value of $\phi$ is less than $\phi^S$, it is optimal for the policy maker to let the economy converge towards $\phi^0 = 0$, that is, to let the stock of believers go to zero. This may appear counterintuitive, since the instantaneous payoffs of the policy maker are smaller at $\phi^0$ than at $\phi^S$. However, the smaller payoff at equilibrium is more than compensated by the gains the policy maker makes by exploiting the stock of believers in the transitory phase towards the equilibrium $\phi^0$. If the initial value of $\phi$ is greater than $\phi^S$, it is optimal to build up a stock of believers until $\phi^E$ is reached.

In the dynamic setting, the policy maker never implements the actions that are optimal in the static game unless $\phi = 0$. That is, it never optimizes its instantaneous payoff. Rather, it implements an actual inflation that lies closer to the optimal static announcement (7) than the statically optimal reaction function (8) suggests. Doing so, it increases the payoffs of the believers compared to those of the nonbelievers. This, in turn, increases the rate of change of $\phi$, with positive consequences for the policy maker's cumulated payoffs. Similar considerations also drive the dynamic of the policy maker's actions in the agent based model we consider here, but we shall see that the arising dynamic patterns and implications are richer and quite different.

5

## 2.3 The agent-based formulation – common forecasts

The agent-based model is closely related to the static game and its dynamic extension presented above. The variables are the same as before. In every period $t$ the same actors carry the same sequence of actions. The economy (1) and the instantaneous payoffs (2) and (3) are unchanged. The proportion of believers changes depending upon the difference in the payoffs received by the believers and nonbelievers.

Minor differences from the previous framework are that time is now discrete, and that the number of agents is finite (but typically large.) The crucial departure is the hypothesis that the policy maker and the nonbelievers are boundedly rational and use individual evolutionary respectively adaptive learning to improve their decisions over time. This might lead to heterogeneity of expectations not only between believers and nonbelievers but also among nonbelievers. Agents know the economy (1) and observe without delay the actions of the other actors, but do not know precisely how these actions are taken. That is, they act under a substantial veil of ignorance. The policy maker, in particular, no longer optimizes but attempts to improve its actions using a genetic-algorithm-inspired learning procedure. The details of these modifications are presented in the next sub-sections. The pseudo-code of the model is presented in Appendix A.

### 2.3.1 The private agents

At the beginning of any period $t$, there are $\mu_t$ believers and $\nu_t$ nonbelievers, with $\mu_t + \nu_t = \Theta$, where $\Theta$ is the (constant) total number of agents in the population. Accordingly, the fraction of beleivers at $t$ reads $\phi_t = \frac{\mu_t}{\Theta}$. As in the static game, the policy maker announces $y_t^a$, and the believers set their inflation forecasts equal to the announcement, $x_t^B = y_t^a$. Contrary to the static case, however, the nonbelievers cannot perfectly predict the actions of the policy maker in the current period. They know that these actions may change over time in a for them unpredictable manner. They use an adaptive learning scheme to revise from period to period the way they form their forecasts. Specifically, we assume that they add an error correction term, $d_t$, to the forecast that would had been optimal in the static game,

$$x_t^{NB} = x_t^{NB}(static) + d_t,$$

where $x_t^{NB}(static)$ is given by (6). In other words, they use the static game as a reference situation. However, they know that the policy maker takes into account intertemporal effects and adapts its policy in an unpredictable manner and try to adaptively improve their forecasts over time. At the end of each period, the error correction term $d_t$ is updated as a function of the discrepancy between predicted and realized inflation,

$$
\begin{aligned}
d_{t+1} &= d_t + \gamma(y_t - x_t^{NB}), & (12) \\
d_0 &= 0, & (13)
\end{aligned}
$$

where $\gamma > 0$ is a parameter capturing the speed of learning. It should be noted that due to this dynamic adaptation algorithm the forecasts of nonbelievers converge over time to the actual value of the inflation rate if this rate stays constant. Put differently, in case the actions of the policy maker settle at a steady state all nonbelievers in the population eventually attain correct inflation forecasts.

Moreover, $\phi$ is now a dynamic variable. It changes, following a word-of-mouth information exchange among the private agents, as a function of the payoff difference between believers and nonbelievers. This is modelled in the following way. In each period, a fraction $\beta$ of the private agents is chosen randomly. The chosen agents are then randomly paired, e.g., agent $i$ with agent $k$. Each agent $i, k$ observes its partner's current strategy, to believe or not to believe. It also *imperfectly* observe the payoff of the other agent. Agent $i$ e.g. observes the payoff of its partner, agent $k$, as

$$J_{observed}^k = J^k + \epsilon,$$

where $\epsilon$ is a random noise. If $J^i < J_{observed}^k$, that is, if agent $i$'s payoff proves smaller than $k$'s observed payoff, agent $i$ adopts the strategy of agent $k$. Thus, a believer may become a nonbeliever, and vice versa.

The resulting dynamics of $\phi_t$ is stochastic. Assuming for analytical simplicity that $\epsilon$ is drawn from a distribution qualitatively similar to a Gaussian distribution[1], the expected change in the proportion of believers

---

[1]To obtain the analytically convenient expression (14), we assume that $\epsilon$ is drawn from the unimodal distribution with mean zero: $\epsilon = 2 \tan(\pi * (\text{rand} - 0.5))/\pi$, with rand drawn from the uniform distribution on [0,1].

6

is given by

$$\Delta\hat{\phi}_t := \mathbb{E}\phi_{t+1} - \phi_t = \beta\phi_t(1-\phi_t)arctan(J_t^B - J_t^{NB}). \tag{14}$$

In the case of common forecasts we suppose that, when an agent switches from $B$ to $NB$, it observes the common value of $d$ used by the others nonbelievers. Thus, all nonbelievers have the same $d$ in any $t$, independently of their history of strategy choices.

### 2.3.2 The policy maker

At the beginning of each period $t$, the policy maker randomly chooses the announcement $y_t^a$ it will make and the inflation rate $y_t$ it will implement from a given pool $Y_t = \{y_t^a(j), y_t(j)\}$ of $N$ possible actions or *rules* $[y_t^a(j), y_t(j)]$, $j = 1, 2, ..., N$. The initial pool of rules, $Y_0$, is assumed to be inherited from past history (In practice, this means that it will be randomly generated). From period to period, the current pool $Y_t$ is revised in a way that allows a systematic exploration of the space of all potentially possible rules, and not only of those present in $Y_0$. One speaks of *experimentation*. Parallely, *replication* attempts to improve over time the quality or *fitness* of the rules in the current pool. Finally, in every period, the policy maker randomly chooses the rule to be actually used with a probability proportional to its fitness (*rule selection*). This genetic-algorithm-inspired procedure known as individual evolutionary learning works as follows.

**Experimentation** Experimentation is used to generate new rules that are not included in the current pool $Y_t$. Each element of each rule $[y_t^a(j), y_t^a(j)] \in Y_t$, $j \in \{1, .., N\}$, is changed independently with some probability *mexp*. The new value after experimentation is computed as:

$$new \quad value = old \quad value + \xi,$$

where $\xi$ is a random number drawn from a standard Normal distribution.

**Computation of foregone outcomes, fitness and the pseudo-value function** After having observed the actions taken by the private agents, the policy maker $G$ is able to calculate both the payoff it would had received if it had used in period $t$ any other rule in $Y_t$, and the corresponding expected change in $\phi_t$. Let the values that would had been obtained by using the rule $[y_t^a(j), y_t(j)] \in Y_t$ be noted $\tilde{J}_t^G(j)$, $\tilde{u}_t^B(j)$, etc (see Appendix B for details). We have

$$\tilde{J}_t^G(j) = -\frac{1}{2}[\phi_t(\tilde{u}_t^B(j))^2 + (1-\phi_t)(\tilde{u}_t^{NB}(j))^2 + y_t^2(j)]. \tag{15}$$

Now, if $G$ was solving a standard dynamic optimization problem (maximizing its cumulated payoffs subject to the relevant dynamic constraints), the value for $G$ of using a given rule $j$ would not be limited to the resulting instantaneous payoff $\tilde{J}_t^G(j)$. It would also include the changes in the state variables, weighted by the proper dynamic multipliers, in order to capture the consequence of these changes on the future optimal payoffs stream. We do not assume that the policy maker acts based on infinite horizon dynamic optimization, but nevertheless takes into account the intertemporal effects of its policy. Thus motivated, we assume in this paper that $G$ values the quality or *fitness* of the different rules it might had used in $t$ in terms not of $\tilde{J}_t^G(j)$ but in terms of the pseudo-value function

$$V_t^G(j) = \tilde{J}_t^G(j) + \Omega \ \Delta\tilde{\phi}_t(j),$$

where $\Omega > 0$ is a parameter and $\Delta\tilde{\phi}_t(j)$ the expected change of $\phi_t$ if rule $j$ had been applied (see Appendix B). In that way, it assigns a positive value to an increase in $\phi$, that is, takes into account that a higher $\phi$ now should allow higher payoff in the future. Loosely speaking, the parameter $\Omega$ measures the concern of $G$ for the long term consequences of its actions.

**Replication** Replication increases the frequency of high fitness rules. It allows rules that are likely to generate high payoffs to replace inferior ones. In the model, replication occurs through so-called *tournament selection.* Pairs of rules are drawn randomly with replacement from the existing pool. The rule with the higher replaces in the pool the one with the lower fitness. The procedure is repeated $N$ times, leading to a new pool with an increased proportion of high fitness rules.

**Rule selection** Experimentation and replication allow the policy maker to construct starting with $Y_t$ a potentially improved pool of rules, $Y_{t+1}$. Compared to $Y_t$, the new pool $Y_{t+1}$ contains both new rules and a higher proportion of rules with high fitness. The rule effectively used in $t + 1$ is chosen randomly from $Y_{t+1}$, with a probability that is increasing with the rule fitness. Specifically, we assume that the probability for rule $j$ to be used in period $t + 1$ is given by:

$$P_{t+1}(j) = \frac{V_t^G(j)}{\sum_{i=1}^{N} V_t^G(j)}. \tag{16}$$

Note, however, that the fitness is evaluated on the basis of the actions of the agents in period $t$. Since these agents learn between $t$ and $t + 1$, it is not warranted that a rule that is highly fit in the sense defined above will indeed perform well when applied in $t + 1$.

## 2.4 The agent-based formulation – private forecasts

The previous version of the agent-based model captures an extreme situation: In every period all nonbelievers share the same error correction term $d_t$ and thus make the same forecast $x_t^{NB}$. In other words, this version is compatible with the existence of a representative nonbeliever. To investigate the impact of additional heterogeneity in a well defined benchmark case , we also consider the following variant. Whenever an agent $i$ starts using the $NB$ strategy (say, in $t$), it has no information additional to the one describing the static game and therefore uses the statically optimal reaction function (6) to make its prediction for period $t$. That is, it uses 0 as value of the error correction term in period $t$. In the subsequent periods, and as long as it uses the $NB$ strategy, it updates the error correction term according to (12). In this way, we introduce heterogeneity among the nonbelievers: Each nonbeliever uses in each $t$ its own, private error correction term $d_t^i$ (Notice that all agents that became believers in the same period nonetheless use the same error correction term). Accordingly, the forecast of a nonbeliever $i$ in period $t$ is given by

$$x_t^{NB,i} = \frac{\theta^2 \phi_t y^a + \theta u^*}{1 + \theta^2 \phi_t} + d_t^i \tag{17}$$

and its payoff by

$$J_t^{NB,i} = -\frac{1}{2}[(y_t - x_t^{NB,i})^2 + y^2] - c^{NB}. \tag{18}$$

Since $d_t^i$ is individually determined, these forecasts and payoffs may differ from nonbeliever to nonbeliever.

The unemployment of the nonbelievers is now

$$u_t^{NB} = u^* - \theta \left( y_t - \frac{1}{\nu_t} \sum_{i=1}^{\nu_t} x_t^{NB,i} \right),$$

where $\nu_t$ is, as previously stated, the current number of nonbelievers. Otherwise, the model is unchanged.

## 3 Simulation Results

Unless stated otherwise, we use in the simulations the following baseline parameter values:

- $u^* = 5.5$, $\theta = 1$, see Sargent (1999)

- $\beta = 0.05$, $c = 0.1$, $\gamma = 0.1$, $\Omega = 1000$

- $mexp = 0.2$, $\Theta = 100$, $N = 100$

In period 0, each agent is initialized as a believer or a nonbeliever with probability $1/2$. The $N$ initial rules in $Y_0$ are randomly generated from the uniform distribution with support [-10,15]. So is the rule $[y_0^a, y_0]$ used in $t = 0$. As previously indicated, $d_0 = 0$.

Simulations are run for $T = 300$ periods. All data presented in the tables and figures are averages over 100 runs. The results are robust – qualitatively similar outcomes are obtained for wide ranges of parameters.

For future comparison, note that for the parameters values $u^* = 5.5$ and $\theta = 1$ used in the simulations the instantaneous payoffs are $J^{G*} = -30.25$ and $J^{NB*} = -30.25$ at the equilibrium with $\phi = 0$ (no believers).

## 3.1 Dynamics under common forecasts

Figure 1 illustrates the typical dynamics of $\phi_t$, $J_t^G$, $y_t^a$, $y_t$, $x_t^{NB} - y_t$, and $x_t^B - y_t$ in our agent-based economy for a sufficiently large $\Omega$. After some initial transients, a clear pattern arises between periods 50 and 200. The true inflation $y_t$ oscillates around zero. The announced inflation $y_t^a$ is smaller than $y_t$, but to a large extend exhibits fairly parallel oscillations. This oscillatory behavior makes it impossible for the nonbelievers to effectively adjust their learning parameter $d_t$. Compared to those of the believers, their forecast errors remain large and their payoffs low. Accordingly, the stock of believers keeps increasing. Once the policy maker has built up a high proportion of believers it starts (around period 150) to exploit their gullibility by increasing the difference between announced and true inflation. The payoffs of the believers fall below the payoffs of the nonbelievers, and $\phi$ starts to decrease. At the same time, the payoffs of the policy maker increase due to a decrease in the unemployment of the believers. However, since now $J^B < J^{NB}$, the change $\Delta\phi$ in the proportion of believers becomes negative, with adverse consequences for the future payoffs $J^G$. Therefore, around period 220, the policy maker tries to reverse this downward trend by reducing the discrepancy between $y^a$ and $y$. The payoffs of the believers increase, leading to an increase of $\phi$ to 0.7 at $t \approx 260$. Then, the policy maker starts again exploiting the large stock of believers. It announces an inflation of about -5.5 but sets $y$ close to 0 (these are the optimal values for the policy maker in the static game). This leads to a new reduction of the proportion of believers.

Thus, one observes irregular cycles consisting in a phase of trust building (where the policy maker makes the believers better off than the nonbelievers) followed by a phase of trust exploitation (where the policy maker uses the high proportion of believers to increase its own payoffs at their expense). During the fluctuations described here the average payoff of the private agents and the payoff of the policy maker are higher during the periods with a large proportion of believers. This does not preclude the possibility that the payoffs of all or some actors are very low in some transitory phase just after initialization, reflecting the fact that the actors did not learn and still have a very arbitrary behavior at that time. For that reason, we will typically disregard outcome of the 20 first periods when analyzing the results.

Thus, cheap talk can sustain in the long run a situation with a positive, fluctuating proportion of believers. Most importantly, the first row of Table 2 shows that this cyclical behavior is Pareto-superior compared to the standard Nash solution with $\phi = 0$. The average payoff of the policy maker is higher, $-16$ instead of $-30.25$. So are the average payoffs of the private agents ($-11.07$ for the believers and $-4.40$ for the nonbelievers, compared to $-30.25$ for all agents at the Nash equilibrium).

In other words, in spite of the fluctuations and of the forecast errors they generate, cheap talk can indeed sustain not only a positive fraction of believers, but more crucially: it can lead to Pareto-superior outcomes. However, although true for large ranges of parameters, this may not be valid for other. In particular, as mentioned earlier, the results presented in Figure 1 require sufficiently large values of $\Omega$, that is, sufficient concern of the policy maker for the future. For lower values of $\Omega$ it is possible to generate simulations where $\phi$ becomes zero in finite time. Such an extinction of the believers' population will also occur if, e.g., the policy maker learns too slowly, due among others to too small a value of the rate of experimentation $mexp$.

9

## 3.2  Is there an optimal value of $\Omega$?

By construction, a high value of $\Omega$ implies that the policy maker gives much weight to an increase over time in the proportion of believers and thus, favors a solution with a high average value of $\phi$. Our discussion of the dynamics in the agent-based model and of the underlying static game suggests that a higher $\phi$ should make it easier the policy maker to achieve higher instantaneous payoffs. However, increasing the stock of believers is not without cost. To achieve and sustain a high value of $\Delta\phi$, the policy maker must favor the believers and, therefore, accept lower current payoffs. Therefore, it is not clear whether or not an increase of $\Omega$ implies a higher cumulated payoff for the policy maker.

In order to shed light on this issue, let define the discounted cumulated payoff of the policy maker as

$$\Gamma^G = \sum_{t=0}^{T} \rho^t \tilde{J}_t^G,$$

where $\rho \in (0,1)$ is the discount factor. Figure 2 shows how $\Gamma^G$ varies with $\Omega$ for $\Omega \in [800, 2000]$ and $\rho = 0.98$[2].

One recognizes that $\Gamma^G$ is first increasing, then decreasing in $\Omega$. To understand this result remember that a policy maker with a high value of $\Omega$ is strongly concerned with the stock of believers and, therefore, with the believers' payoffs. These payoffs depend on the deviation between the believers' inflation forecast and the actual inflation, that is, between $y^a$ and $y$. Thus, the higher $\Omega$ is, the lower is, ceteris paribus, the difference $y^a - y$. However, a low value of $y^a - y$ implies a high current unemployment and a low value of the current payoff $J^G$. If the value of $\Omega$ is excessive, the short term losses needed to obtain a high value of $\phi$ dominate the gains this high value generates in the long term, and $\Gamma^G$ decreases.

## 3.3  Sensitivity analysis: The impact of $c^{NB}$ and $\gamma$

From the previous discussion it should be clear that, ceteris paribus, the policy maker $G$ would welcome any parameter change that lowers the instantaneous payoffs of the nonbelievers compared to those of the believers. Such a change indeed allows $G$ to invest less effort (that is, to forego a smaller fraction of its instantaneous payoffs) to insure relatively high payoffs to the believers and, thus, to insure a high value either of $\Delta\phi$ or of $\phi$.

Two parameter changes are of particular interest in that regard: an increase of the nonbelievers' forecasting costs, $c^{NB}$; and a decrease of their learning speed $\gamma$. As we shall see, both changes indeed increase $G$'s average payoff $\bar{J}^G$ and cumulated discounted payoff $\Gamma^G$. However, they have markedly different impacts on the key economic variables, as the simulations presented below will show. In these simulations, we use as before $\rho = 0.98$ to compute $\Gamma^G$. The parameter $\Omega$ is given the value $\Omega = 1000$, that is, the value that roughly maximizes $\Gamma^G$ for the reference parameter constellation. We use a Wilcoxon test to check whether or not the changes in $c^{NB}$ and $\gamma$ generate statistically significant differences in the average values of the main variables[3].

**Impact of an increase in $c^{NB}$**   Table 2 shows the discounted payoff $\Gamma^G$ as well as the averages over periods 20 to 300 of the main variables of interest for $c^{NB} = 0.1$ (its value in the basis scenario) and $c^{NB} = 1$. A bar - above a variable indicates a time average. Row 3 of the table gives the significance levels of the Wilcoxon test. It indicates that the differences between the two sets of runs are almost always statistically significant at the 95% level.

One recognizes that $\bar{y}^a$ and $\bar{y}$ both decrease when $c^{NB}$ increases. Furthermore, $|\bar{y}^a - \bar{y}|$ becomes smaller. This tends to increase both $\bar{\phi}$ and $\bar{J}^G$. However, the average unemployment rates of both believers and nonbelievers also increase, with adverse effects on both $\Gamma^G$ and $\bar{J}^G$. Since $\Delta\phi > 0$ and $u^B < u^{NB}$, however, the increase in average unemployment is smaller at the population level than the increase observed for each agent type. In toto, the policy maker average payoffs increase.

---

[2]If one interprets each period as two weeks, which seems reasonable given $\beta = 0.05$ and the adjustment speed of $\phi_t$, this corresponds to a yearly discounting of about 10%.

[3]To carry out the statistical analysis of the effects of changes of $c^{NB}$ we generate 100 pairs of $(\gamma, \Omega)$ values using a uniform distribution on $[0.01, 0.1] \times [800, 1300]$. For each of these $(\gamma, \Omega)$ pairs one simulation is run with $c^{NB} = 0.1$ and one with $c^{NB} = 1$, where all remaining parameters are set to their default values. Differences of the variable under consideration between the two runs are taken for each of the 100 parameter profiles and the Wilcoxon test is used to check whether the differences are (statistically) significantly above or below zero. Analogous for changes of $\gamma$.

Likewise, both nonbelievers and believers are better off after an increase in $c^{NB}$, due to the decrease in the difference $|\bar{y}^a - \bar{y}|$ and to the lowering of the rate of inflation. Moreover, since the believers obtain a higher payoff than the nonbelievers, and since $\bar{\phi}$ increases, the average payoff of the private agents also increases.

**Impact of a decrease in** $\gamma$ Quite a different picture emerges when $\gamma$ is reduced from $\gamma = 0.1$ to $\gamma = 0.01$, see Table 3. Here again, he differences between the two sets of runs are almost always statistically significant at the 95% level, see Row 3 of the Table. Now, the policy maker reacts to the parameter change not by decreasing, but by strongly *increasing* $|\bar{y}^a - \bar{y}|$. The announcement $\bar{y}^a$ decreases, the average inflation rate goes up. As a consequence, one observes in the long run a decrease in the proportion of believers. However, the policy maker is able to quickly reduce unemployment among the believers to zero and, therefore, increases its payoffs in the short run. Since the benefits occur early, this results in an increase of its discounted cumulated payoff $\Gamma^G$. But the loss of credibility in later periods implies a lower average payoff $\bar{J}^G$. What is more, both believers and nonbelievers have lower average payoffs after a decrease of $\gamma$.

We thus have the striking result that the policy maker, although best off when confronted with a population consisting only of believers, obtains a higher average payoff when the nonbelievers are efficient learners, able to adjust quickly their forecasts to a changing realty. Society as a whole profits, too, from more efficient nonbelievers.

**Is there an optimal** $c^{NB}$? Additional simulations indicate that the described above hold for all $c^{NB} \in [0.1, 1]$ and $\gamma \in [0.01, 0.1]$. A decrease in the learning speed $\gamma$ induces the policy maker to lower unemployment. This results in a worse long-run economic outcome for all agents. By contrast, an increase in the forecasting costs $c^{NB}$ triggers a policy that primarily aims at insuring a low inflation rate. This leads to a larger proportion of believers, supported by small differences between the announced and the true inflation, and to higher payoffs for all actors.

However, one may wonder whether increasing in $c^{NB}$ always improves the payoffs. Is there rather an optimal level of $c^{NB}$ behind which the one or the other payoff starts declining? To address this question, we plot in Figure 3 the average payoffs of $G$ and of the $B$s and $NB$s as $c^{NB}$ varies from 0 to 4 in increments of 0.1. One recognizes that the payoff of the believers reaches a maximum for $c^{NB} = 1.3$ while the payoff of the nonbelievers is highest for $c^{NB} = 0.8$. The average, $\phi-$weighted payoff of all private agents also reaches it maximum at $c^{NB} = 0.8$. By contrast, $\Gamma^G$ and $\bar{J}^G$ increase monotonically with $c^{NB}$.

An examination of the time series shows a sharp qualitative change in the policy maker behavior around $c^{NB} = 1.3$, that is, at the value that maximizes the payoff of the nonbelievers. The policy maker stops implementing a low $y$ and starts increasing the gap between $y^a$ and $y$. As a consequence, the believers' unemployment rate decreases. This behavior is similar to the one observed when $\gamma$ is low.

# 4 Dynamics under private forecasts

So far we have assumed common forecasts, i.e., we supposed that all nonbelievers always use the same error correction term $d$. In other words, we hypothesized that a private agent who just started using the $NB$ strategy is able to forecast the true inflation rate as accurately as other nonbelievers with a possibly long learning history. As an other extreme case, we now consider the private forecast situation described in section 2.4. Here an agent, when becoming a believer, cannot rely on any knowledge accumulated by the other nonbelievers but must start from the onset a new individual learning process, using 0 as initial value of its error correction term. Thus, an agent that just decided to use the strategy $NB$ is likely to make larger forecast errors under the private than under the common forecast regime. This, in turn, may lower its payoffs and make it more prone to switch back to the $B$ strategy. In other words, private forecasts make switching to the $NB$ strategy potentially more costly and less stable.

## 4.1 Comparison of the baseline simulations under common and private forecasts

Table 4 summarizes the outcome of siumulations under common and private forecasts, using the baseline values of the parameters. Here again, the values for the first 20 periods have been discarded while constructing the

averages. One recognizes that $\bar{\phi}$, $\Gamma^B$, $\bar{y}$, $\bar{y}^a$, and $\bar{u}^B$ are lower, and $\bar{u}^{NB}$ is higher in the private than in the common forecast case. The differences between the two sets of runs are almost always statistically significant at the 95% level, as indicated by the significance levels of the Wilcoxon test, see row 3.

The economy's time behavior under the private forecast regime is illustrated in Figure 4. Compared to the common forecast case (Figure 1), $\phi$ at first increases faster but starts to decrease earlier. As noted before, this results in a lower average $\bar{\phi}$. The dynamics can be explained in the following way. Initially, it is easier for the policy maker to build up the proportion of believers because, as previously explained, switching to nonbelieving is less attractive. Once the policy maker builds up a large stock of believers, it starts exploiting it by increasing the gap between actual and announced inflation. This leads to a decrease in the proportion of believers.

Compared to the common forecast case, the greater heterogeneity of the private forecasts makes it more difficult for the policy maker to determine the actual inflation rate that maximizes its payoff. As a result, even in the long run, the average forecast of the nonbelievers overstate the actual inflation rate, and the average unemployment of the nonbelievers is larger than the natural rate of $u^* = 5.5$ (see table 4). This leads to lower payoffs for both $G$ and the $NB$s. The existence of a systematic forecasting bias by the nonbelievers is a strong illustration of the coordination problems that may arise in a system with a high level of agents' heterogeneity.

## 4.2 Sensitivity analysis: The impact of $c^{NB}$ and $\gamma$

As noted earlier, in case of common forecasts, the policy maker responds in qualitatively different ways to changes in the parameters $c^{NB}$ and $\gamma$. An increase in $c^{NB}$ leads to less inflation, a decrease in $\gamma$ to a lower unemployment. Up to some threshold value, an increase in $c^{NB}$ improves the average payoffs of all actors, while a decrease in $\gamma$ increases $\Gamma^G$ but reduces $\bar{J}^G$, $\bar{J}^B$, and $\bar{J}^{NB}$.

**The impact of $c^{NB}$ in the private forecast case**   Table 5 shows the effect of an increase of $c^{NB}$ from 0.01 to 0.1 on the key variables in the private forecast case. Row 3 of the table gives the significance levels of the Wilcoxon test already applied in the common forecast case.

One recognizes that an increase in $c^{NB}$ no longer leads to a reduction in inflation, but to a lowering of unemployment. The true inflation goes up, the announced inflation goes down, and the gap between these two values widens. This reduces the unemployment rate of the believers, with favorable consequences for the policy maker. However, $\bar{\phi}$ decreases, leading to a lower average payoff $\bar{J}^G$. The timing of the different effects is such that $\Gamma^G$ nonetheless increases. Contrary to the common forecast case, the average payoff of both the believers and the nonbelievers decreases.

Thus, an increase of $c^{NB}$ is no longer Pareto-improving. However, remember that in the common forecast case the impact of such an increase was not monotonic. Behind a certain threshold, increases in $c^{NB}$ lead to lower payoffs for all agents. We have argued that using individual forecasts imposes, in a sense, additional costs on the nonbelievers. Thus, it might lower this threshold. Could our last results be simply explained by the fact that the high value we used for $c^{NB}$, 0.1, was superior to this lower threshold?

Figure 5 confirms the dual hypothesis of a lower threshold that is inferior to 0.1. The figure was produced analogously to figure 3 by letting $c^{NB}$ vary from 0 to 0.4 by increments of 0.01. One recognizes that $\bar{J}^B$ is highest for $c^{NB} = 0.02$, while $\bar{J}^{NB}$ reaches a maximum for $c^{NB} = 0.01$. The average $\phi-$weighted payoff of the whole population is highest for $c^{NB} = 0.01$. Thus, contrary to the common forecast case, increasing $c^{NB}$ above 0.1 is not Pareto-improving. As before, $\Gamma^G$ and $\bar{J}^G$ keep increasing with $c^{NB}$, up to some volatility.

**The impact of $\gamma$ in the private forecast case**   The result of a lowering of $\gamma$ from 0.1 to 0.01 and those of the associated Wilcoxon test are presented in Table 6. When the speed of learning, $\gamma$, is reduced, $\bar{y}$ and $\bar{y}^a$ both decrease while $|\bar{y}^a - \bar{y}|$ increases. The higher value of $|\bar{y}^a - \bar{y}|$ leads to a reduction of $\bar{u}^B$. Since the speed of learning $\gamma$ is lower, the nonbelievers make larger forecast errors and therefore experience a higher rate of unemployment. Contrary to the common forecast case, $\bar{\phi}$ increases. Since $\bar{y}$ and $\bar{u}^B$ decrease but $\bar{\phi}$ increases, $\bar{J}^G$ and $\Gamma^B$ are lower. So are $\bar{J}^B$ and $\bar{J}^{NB}$. Overall, we observe that the direction of movement of

|  | $J^{G,disc}$ | | $\bar{\phi}$ | | $\bar{y}$ | | $U^B$ | | $U^{NB}$ | |
|---|---|---|---|---|---|---|---|---|---|---|
| $\gamma$ ↓ | ↑ | ↑ | ↓ | ↑ | ↑ | ↓ | ↓ | ↓ | ↑ | ↑ |
| $c$ ↑ | ↑ | ↑ | ↑ | ↓ | ↓ | ↑ | ↑ | ↓ | ↑ | ↓ |

Table 1: Effect of changes of the speed of error correction learning ($\gamma$) and the cost of forecasting for non-believers ($c$). The first arrow in each cell gives the effect under common forecasts of non-believers, the second under private ones.

the unemployment rates triggered by a decrease of $\gamma$ is the same for the common and private forecast cases. The direction of movement of the inflation rate differs however between the two cases.

**The consequences of heterogeneity** As we have seen, the impact of changes in $c^{NB}$ and $\gamma$ are very different in the case of common and private forecasts, see Table 2 and Table 5, Table 3 and Table 6. In particular, the policy response (acting primarily on the unemployment or the true inflation) depends crucially from the modalities of forecast formation. Table 1 summarizes these differences. The first arrow in each cell indicates the direction of change of the time averaged variable in the common, the second the direction of change in the private forecast case. For changes in the speed of learning ($\gamma$) the direction, in which the fraction of believers and the inflation rate move, differs between the common and private forecasts case. If the cost of forecast formation of non-believers ($c$) varies, then all listed variables other than the discounted payoff of the policy maker move in opposite directions under common and private forecasts!

All this illustrates the crucial importance of explicitly taking into account potential agent heterogeneity when devising and analyzing policy measures. The only difference between the scenarios considered here lies in the homogeneity or the heterogeneity of the nonbelievers' forecasts. Nevertheless two qualitatively very different policy reactions to a change of $c^{NB}$ or $\gamma$ are observed. Thus, in the present setting but presumably also in many others, using an heterogenous agent model appears unavoidable if one wants to exhaustively examine the qualitative features of the problem of interest.

In our model, the policy maker obtains a higher payoff when the nonbeliever's forecasts are homogenous. Thus, it has an incitation to facilitate the information flow between nonbelievers, thus helping those agents that just adopted the $NB$ strategy to build on the experience of the other nonbelievers. *Making too much data publicly available might, however, reduces the forecasting costs $c^{NB}$. This, as we have seen above, is not necessarily desirable.* Thus, the policy maker faces the non-trivial problem to keep the informational costs of the nonbelievers as high as possible while avoiding that they build wildly diverging and inexact forecasts.

## 4.3 A more differentiated look at the effects of parameter changes

In order to shed additional light on the mechanisms that underlie the reported consequences of parameter changes under common and private forecasts, we consider in more detail the dynamic patterns of the key variables in the model. We concentrate on changes of $\gamma$, since our main general points can be made for that case. As a first step we consider (in addition to the time averages for the periods 21 to 300 already presented in 3 and 6) the time averages over the periods 21 to 100 and 21 to 300. Once again, the Wilcoxon test shows a statistically significant impact of the studied parameters changes at the 95% level.

Tables 7 and 8 show the results obtained for $\gamma = 0.1$ and $\gamma = 0.01$ in the common and private forecast cases. One recognizes that in both cases a lower $\gamma$ leads to a higher $\bar{\phi}$ over the first 100 periods. Over 300 periods, $\bar{\phi}$ decreases in the common forecast case, but increases if the forecasts are private. In both cases, $\bar{\phi}$ is lower over 300 periods than over 100 periods. That is summarizing, the proportion of believers always declines over time, but slower in the private than in the common forecast case.

For a given $\gamma$, the true inflation rate $\bar{y}$ is lower under private forecasts, independently of the time span considered. In the common forecast case, $\bar{y}$ decreases over the 100 initial periods, but increases over the 300 periods. The spread $|\bar{y}^a - \bar{y}|$ increases over all time intervals in both the common and the private forecast cases. Thus, a lower $\gamma$ implies a reduced unemployment rate for the believers, but a higher one for the nonbelievers

due to higher forecast errors. Altogether these effects lead to a decrease in $\bar{J}^G$. However, $\Gamma^G$ increases since $\bar{y}$ is lower, $\bar{\phi}$ higher, and $\bar{u}^B$ lower over the initial 100 periods.

The previous observations allow to summarize the impact of a decrease of the speed of learning $\gamma$. In the case of common as in the case of private expectations, this decrease implies that the nonbelievers take longer to improve their forecasts. Consequently, the policy maker is able to build up faster a large mass of believers. In both cases, once $\phi$ has reached a sufficiently high value, the policy maker starts to exploit the existing believers by increasing the spread $|y^a - y|$. This leads to a decline of $\phi$. This decline starts earlier when $\gamma$ is lower, as illustrated in Figures 1 and 6 for the common forecasts case, and in Figures 4 and 7 for the private forecasts case. As can be seen in these figures, in cases of a low speed of adaptation of nonbelievers forecasts the exploitation of the stock of believers by the government eventually leads to the extinction of believers. This is not in the (long-run) interest of the policy maker, but in the given stochastic and heterogeneous environment it is not able to fully predict the implications of its policy choices. These observations of the dynamics of the relevant variables reinforce an additional important conclusion: To sustain a positive mass of believers, with all the associated economic advantages, it is necessary that the nonbelievers update sufficiently quickly and accurately their forecasts, that is, $\gamma$ must be sufficiently high. This hinders the policy maker from exploiting believers excessively and thereby avoids the eventual extinction of believers.

Using private instead of common forecasts has similar consequences than a lowering of $\gamma$. The number of believers starts to decline earlier in the private (Figure 4) than in the common forecast case (Figure 1). The logic behind this was already discussed. The learning of the nonbelievers is slower in the private forecast case, since a new nonbeliever ignores the experience accumulated by the others.

## 5 Conclusions

In this paper, we investigated an agent-based dynamic version of the seminal Kydland-Prescott (1977) model. In our framework, the policy maker makes cheap talk announcements of future inflation before implementing the true one. The private agents have two possible strategies: Believing, i.e. using the announcement as forecast for the future inflation; or Nonbelieving, i.e, trying to improve adaptively their inflation forecasts. They tend to adopt the strategy that provided the highest payoff in the last period.

The policy maker uses individual evolutionary learning to determine how it should set the inflation announcements and the true inflation in order to gain high payoffs. Doing so, it gives attention to the future evolution of the stock of believers. Indeed, ceteris paribus, the policy maker receives higher instantaneous payoffs if the proportion of believers is high. However, since the private agents tend to adopt the strategy that gives them the highest payoff, building and maintaining such a high proportion implies that the policy maker must forego possible current gains in order to insure that the believers receive higher payoffs than statically optimal. In other words, to obtain satisfactory payoffs in every period, the policymaker must solve an complex intertemporal tradeoff between short and long term gains.

The simulations show that the policy maker is able to learn how to reach an outcome that is Pareto-superior compared to the one that would be attained if it did not make adequate cheap talk announcements. This outcome is characterized by a succession of trust building phases, where the announcemnts and the true inflation are chosen in order to increase the proportion of believers; and of trust exploitation phases, where the policy maker uses the existence of a large stock of believers to achieve for itself high payoffs at the cost of a decrease in this stock. Accordingly, the number of believers fluctuates over time, but always remain strictly positive.
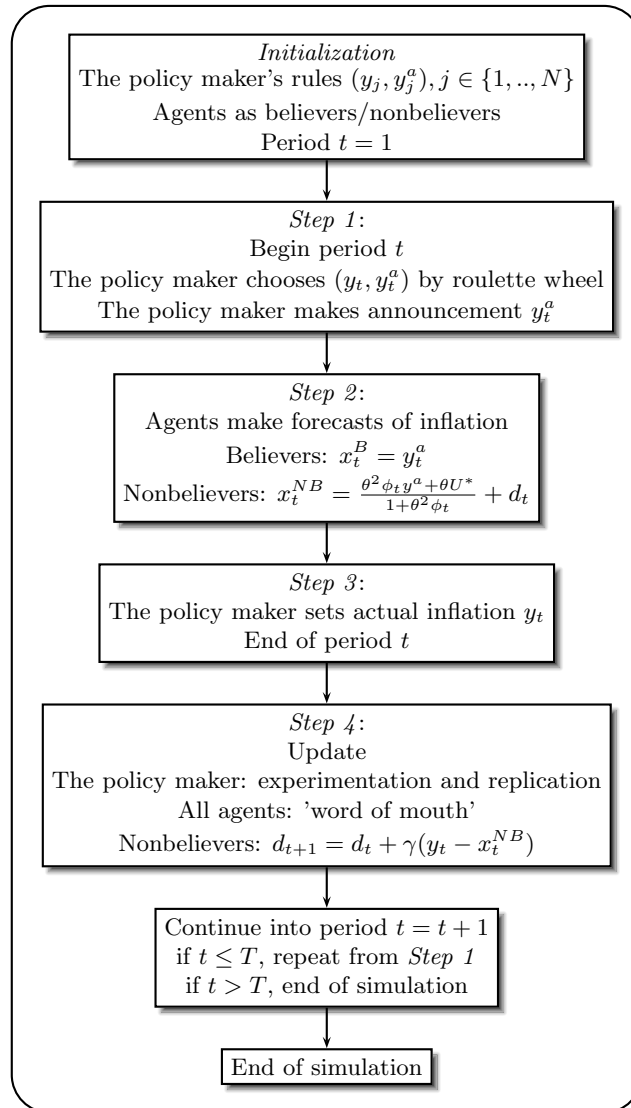
We furthermore showed that the speed with which the nonbelievers learn, and the costs they occur to compute their inflation forecast are of crucial importance for the qualitative properties of the outcome. Most importantly, the consequences of a modification of the value of one of these two parameters are not monotonic in the change. Depending on the starting value and the magnitude of the change, they may e.g. lead the decision maker to favor either a lower unemployement of the believers or a lower inflation. Thus, in some numerical exercices, we show that changes in either one of the two parameters can generate strikingly different policy responses, although they conceptually may appear largely equivalent. Interestingly enough, we show that the policy maker is better off when the nonbelievers learn relatively fast.

14

In the paper, we also investigated the consequences of giving up the assumption of identical nonbelievers, replacing it by the following hypothesis: Everytime an agent switches from believing to nonbelieving, it builds independently its inflation forecast without using the information accumulated in the past by other nonbelievers. This assumption creates an additional heterogeneity among the nonbelievers, tends to lower the quality of the nonbelievers forecasts, makes the agents less inclined to become nonbelievers, and increases the probability that they switch back to believing. It also complicates the task of the policy maker, since it is confronted to many different inflation forecasts (instead of only two). We show that this additional heterogeneity is in many respects equivalent to a lowering of the nonbelievers learning speed or to an increase of their forecasting costs. In any event, it can strongly impact on the decision makers actions and the outcome.

An interesting dilemna arises here. The policy maker is better off if all nonbelievers use the same, accurate inflation forecast. Thus, it is in its interest to facilitate the information flow between nonbelievers so that each of them can always use the past experience of the others. Doing so may, however, reduce the forecasting costs of the nonbelievers and therefore dimish the attractivity of being a believer. One is therefore interested in mechanisms that insure homogenous and accurate nonbeliever forecasts while maintening the incentives for being a believer. Finding them is not a trivial task.

But not least, all our results strongly suggests that explicitly taking into account the ubiquitous heterogeneity that characterizes a real economy is primordial if one wants to reach a proper understanding of its working and to make appropriate policy recommendations.

APPENDIX A. The algorithm of the simulation.



*Initialization*
The policy maker's rules $(y_j, y_j^a), j \in \{1, .., N\}$
Agents as believers/nonbelievers
Period $t = 1$

*Step 1*:
Begin period $t$
The policy maker chooses $(y_t, y_t^a)$ by roulette wheel
The policy maker makes announcement $y_t^a$

*Step 2*:
Agents make forecasts of inflation
Believers: $x_t^B = y_t^a$
Nonbelievers: $x_t^{NB} = \frac{\theta^2 \phi_t y^a + \theta U^*}{1 + \theta^2 \phi_t} + d_t$

*Step 3*:
The policy maker sets actual inflation $y_t$
End of period $t$

*Step 4*:
Update
The policy maker: experimentation and replication
All agents: 'word of mouth'
Nonbelievers: $d_{t+1} = d_t + \gamma(y_t - x_t^{NB})$

Continue into period $t = t + 1$
if $t \leq T$, repeat from *Step 1*
if $t > T$, end of simulation

End of simulation

APPENDIX B. Computation of hypothetical payoffs.

The hypothetical payoff for each rule $[y_t^a(j), y_t(j)] \in Y_t$ in period $t$ is calculated as:

$$\tilde{J}_t^G(j) = -\frac{1}{2}[\phi_t(\tilde{u}_t^B(j))^2 + (1 - \phi_t)(\tilde{u}_t^{NB}(j))^2 + y_t^2(j)] + \Omega \, \Delta\tilde{\phi}_t(j) \tag{19}$$

The hypothetical values of inflation forecasts for believers are calculated for each rule $j$'s announcement of inflation, $y_t^a(j)$, as:

$$\tilde{x}_t^B(j) = y_t^a(j) \tag{20}$$

When computing hypothetical forecasts of nonbelievers, we assume that the policy maker can observe the average inflation forecasts of nonbelievers in the previous period and that he knows how nonbelievers forecast are affected by changes in the inflation announcement. The policy maker computes the hypothetical value of the average nonbelievers' forecast as:

$$\tilde{x}_t^{NB}(j) = x_t^{NB,aver} + \frac{\theta^2 \phi_t}{1 + \theta^2 \phi_t}(y_t^a(j) - y_t^a) \tag{21}$$

where $x_t^{NB,aver} = \frac{1}{\nu_t} \sum_{i=1}^{\nu_t} x_t^i$.

The hypothetical unemployment rates are computed as

$$\tilde{u}_t^i(j) = u^* - \theta(y_t(j) - \tilde{x}_t^i(j)), \quad i = B, NB \tag{22}$$

The policy maker computes the expected change in the proportion of believers, knowing that it is determined by (14). The necessary computations include the following:

$$
\begin{aligned}
\tilde{J}_t^B(j)(\tilde{x}_t^B(j), y_t(j)) &= -\frac{1}{2}((y_t(j) - \tilde{x}_t^B(j))^2 + y_t^2(j)) \\
\tilde{J}_t^{NB}(j)(\tilde{x}_t^{NB}(j), y_t(j)) &= -\frac{1}{2}((y_t(j) - \tilde{x}_t^{NB}(j))^2 + y_t^2(j)) \\
\Delta\tilde{\phi}_t(j) &= \beta\phi_t(1 - \phi_t)arctan(\tilde{J}_t^B(j) - \tilde{J}_t^{NB}(j)) \\
\tilde{\phi}_{t+1}(j) &= \phi_t + \Delta\tilde{\phi}_t(j)
\end{aligned} \tag{23}
$$

If according to these calculations $\tilde{\phi}_{t+1}(j) > 1$, then we set $\Delta\tilde{\phi}_t(j) = 1 - \phi_t$. If $\tilde{\phi}_{t+1}(j) < 0$, then $\Delta\tilde{\phi}_t(j) = -\phi_t$.

APPENDIX C. Tables.

| cost | $J^{G,disc}$ | $\bar{J}^G$ | $\bar{\phi}$ | $\bar{y}$ | $\bar{y}^a$ | $\bar{U}^B$ | $\bar{U}^{NB}$ | $\bar{J}^{P,B}$ | $\bar{J}^{P,NB}$ |
|---|---|---|---|---|---|---|---|---|---|
| $c = 0.1$ | -745.84 | -16.00 | 0.51 | 1.82 | -0.50 | 3.18 | 5.47 | -11.07 | -4.40 |
| $c = 1$ | -625.99 | -11.42 | 0.69 | 0.94 | -0.83 | 3.73 | 5.48 | -3.89 | -2.30 |
| $\alpha$ | 0.000005 | 0.3859 | 0.0129 | 0.00005 | 0.0885 | 0.4602 | 0.0007 | 0.000005 | 0.000005 |

Table 2: Effect of an increase in $c$ with homogeneous expectations of non-believers.

| $\gamma$ | $J^{G,disc}$ | $\bar{J}^G$ | $\bar{\phi}$ | $\bar{y}$ | $\bar{y}^a$ | $\bar{U}^B$ | $\bar{U}^{NB}$ | $\bar{J}^{P,B}$ | $\bar{J}^{P,NB}$ |
|---|---|---|---|---|---|---|---|---|---|
| $\gamma = 0.1$ | -745.84 | -16.00 | 0.51 | 1.82 | -0.50 | 3.18 | 5.47 | -11.07 | -4.40 |
| $\gamma = 0.01$ | -689.80 | -17.42 | 0.35 | 2.33 | -4.04 | -0.87 | 5.63 | -33.87 | -6.55 |
| $\alpha$ | 0.000005 | 0.3859 | 0.1492 | 0.0005 | 0.000005 | 0.000005 | 0.000005 | 0.000005 | 0.000005 |

Table 3: Effect of a decrease of $\gamma$ with homogeneous expectations of non-believers.

| type | $J^{G,disc}$ | $\bar{J}^G$ | $\phi$ | $\bar{y}$ | $\bar{y}^a$ | $\bar{U}^B$ | $\bar{U}^{NB}$ | $\bar{J}^{P,B}$ | $\bar{J}^{P,NB}$ |
|---|---|---|---|---|---|---|---|---|---|
| 1 | -745.84 | -16.00 | 0.51 | 1.82 | -0.50 | 3.18 | 5.47 | -11.07 | -4.40 |
| 2 | -750.71 | -16.51 | 0.48 | 1.54 | -1.38 | 2.58 | 5.70 | -14.00 | -4.51 |
| $\alpha$ | 0.0005 | 0.0007 | 0.0009 | 0.242 | 0.000005 | 0.000005 | 0.000005 | 0.000005 | 0.000005 |

Table 4: The difference between simulations with homogeneous and heterogeneous nonbelievers. Row 1 is for the common forecasts case, row 2 is for the private forecasts case.

| cost | $J^{G,disc}$ | $\bar{J}^G$ | $\bar{\phi}$ | $\bar{y}$ | $\bar{y}^a$ | $\bar{U}^B$ | $\bar{U}^{NB}$ | $\bar{J}^{P,B}$ | $\bar{J}^{P,NB}$ |
|---|---|---|---|---|---|---|---|---|---|
| $c = 0.1$ | -750.71 | -16.51 | 0.48 | 1.54 | -1.38 | 2.58 | 5.70 | -14.00 | -4.51 |
| $c = 1$ | -659.30 | -16.59 | 0.44 | 1.85 | -3.04 | 0.61 | 5.63 | -25.01 | -6.15 |
| $\alpha$ | 0.000005 | 0.00005 | 0.0015 | 0.0019 | 0.000005 | 0.000005 | 0.000005 | 0.000005 | 0.000005 |

Table 5: Effect of an increase of $c$ with heterogenous expectations of nonbelievers.

| $\gamma$ | $J^{G,disc}$ | $\bar{J}^G$ | $\bar{\phi}$ | $\bar{y}$ | $\bar{y}^a$ | $\bar{U}^B$ | $\bar{U}^{NB}$ | $\bar{J}^{P,B}$ | $\bar{J}^{P,NB}$ |
|---|---|---|---|---|---|---|---|---|---|
| $\gamma = 0.1$ | -750.71 | -16.51 | 0.48 | 1.54 | -1.38 | 2.58 | 5.70 | -14.00 | -4.51 |
| $\gamma = 0.01$ | -697.85 | -15.59 | 0.54 | 1.03 | -3.78 | 0.69 | 6.46 | -22.34 | -4.98 |
| $\alpha$ | 0.00005 | 0.00005 | 0.4286 | 0.1562 | 0.000005 | 0.0005 | 0.000005 | 0.00005 | 0.00005 |

Table 6: Effect of a decrease of $\gamma$ with heterogenous expectations of nonbelievers.

| periods | | $J^{G,disc}$ | $J^G$ | $\bar{\phi}$ | $\bar{y}$ | $\bar{y}^a$ | $U^B$ | $U^{NB}$ | $J^{P,B}$ | $\bar{J}^{P,NB}$ | $y - x^{NB}$ |
|---|---|---|---|---|---|---|---|---|---|---|---|
| [21, 100] | $\gamma = 0.1$ | -597.41 | -15.41 | 0.53 | 1.69 | 0.38 | 4.19 | 5.43 | -4.86 | -3.79 | 0.07 |
| | $\gamma = 0.01$ | -552.65 | -12.40 | 0.69 | -0.24 | -2.77 | 2.97 | 6.72 | -5.82 | -2.22 | -1.22 |
| [21, 300] | $\gamma = 0.1$ | -745.84 | -16.00 | 0.51 | 1.82 | -0.50 | 3.18 | 5.47 | -11.07 | -4.40 | 0.03 |
| | $\gamma = 0.01$ | -689.80 | -17.42 | 0.35 | 2.33 | -4.04 | -0.87 | 5.63 | -33.87 | -6.55 | -0.13 |

Table 7: Effect of a decrease in $\gamma$ with homogeneous expectations of nonbelievers.

| periods | | $J^{G,disc}$ | $\bar{J}^G$ | $\bar{\phi}$ | $\bar{y}$ | $\bar{y}^a$ | $\bar{U}^B$ | $\bar{U}^{NB}$ | $\bar{J}^{P,B}$ | $\bar{J}^{P,NB}$ | $y - x^{NB}$ |
|---|---|---|---|---|---|---|---|---|---|---|---|
| [21, 100] | $\gamma = 0.1$ | -604.71 | -15.37 | 0.53 | 1.18 | -0.17 | 4.15 | 5.69 | -4.38 | -3.36 | -0.19 |
| | $\gamma = 0.01$ | -577.28 | -13.28 | 0.70 | -0.39 | -2.62 | 3.27 | 7.02 | -4.84 | -2.59 | -1.52 |
| [21, 300] | $\gamma = 0.1$ | -750.71 | -16.51 | 0.48 | 1.54 | -1.38 | 2.58 | 5.70 | -14.00 | -4.51 | -0.20 |
| | $\gamma = 0.01$ | -697.85 | -15.59 | 0.54 | 1.03 | -3.78 | 0.69 | 6.46 | -22.34 | -4.98 | -0.96 |

Table 8: Effect of a decrease in $\gamma$ with heterogenous expectations of nonbelievers.

# References

[1] Arifovic, J., Ledyard, J. 2004. Computer Testbeds: The Dynamics of Groves-Ledyard Mechanisms. Manuscript, April.

[2] Arifovic, J., Sargent, T.J. 2003. Experiments with human subjects in a Kydland-Prescott Phillips curve economy. in: D. Altig and B. Smith (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.

[3] Barro R.J., Gordon D.B. 1983. Rules, Discretion, and Reputation in a Model of Monetary Policy. *Journal of Monetary Economics*, 12, no.1: 101-121.

[4] Berardi,M. 2007. Heterogeneity and misspecifications in learning. *Journal of Economic Dynamics and Control*, 31

[5] Branch, W. 2004. The theory of rationally heterogeneous expectations: evidence from survery data of inflation expectations. *The Economic Journal*, 114: 592-621

[6] Branch, W., and G.Evans. 2006. Intrinsic heterogeneity in expectation formation. *Journal of Economic Theory*, 127: 264-295.

[7] Branch, W., and G. Evans. 2007. Model uncertainty and endogenous volatility. *Review of Economic Dynamics*, 19: 207-237.

[8] Brock, W.A., and C.H.Hommes. 1997. A rational route to randomness. *Econometrica*, 65: 1059-1160

[9] Brock, W.A., and C.H.Hommes. 1998. Heterogeneous beliefs and routed to chaos in a simple asset pricing model. *Journal of Economic Dynamics and Control*, 22: 1235-1274

[10] Cho, In-Koo, Sargent, Thomas J. 1997. Learning to be Credible.*Mimeo*

[11] Cho, In-Koo, Williams, N., Sargent, T.J. 2002. Escaping Nash Inflation. *Review of Economic Studies*, 69:1-40

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

[13] Dawid, H., Deissenberg, C. 2005. On the Efficiency Effects of Private (Dis)-Trust in the Government. *Journal of Economic Behavior and Organization*, 57: 530-550.

19

[14] Deissenberg C., Alvarez, G. F. 2002. Cheating for the common good in a macroeconomic policy game. *Journal of Economic Dynamics and Control*, 26, Issues 9-10 (August): 1353-1785.

[15] Evans, G., Ramey, G. 1992. Expectation calculation and macroeconomic dynamics. *American Economics Review*, March, 82(1) : 207-224

[16] Evans, G., Ramey, G. 1998. Calculation, Adaptation and Rational Expectations. *Macroeconomic dynamics*, 2(2): 156-182

[17] Giannitsarou, C. 2003. Heterogeneous learning. *Review of Economic Dynamics* 6(4): 885-906.

[18] Honkapohja, S., and K. Mitra. 2005. Performance of monetary policy with internal central bank forecasting. *Journal of Economic Dynamics and Control* (29): 627-658.

[19] Honkapohja, S., and K. Mitra. 2006. Learning stability in economies with heterogeneous agents. *Review of Economic Dynamics*, forthcoming.

[20] Kydland F.E., Prescott E.C. 1977. Rules Rather Than Discretion: The Inconsistency of Optimal Plans. *Journal of Political Economy*, 85, No.3 (June): 473-491.

[21] Persson, T. and Tabellini, G. 1993. Designing institutions for monetary stability. In A. H. Meltzer and C. I. Plosser (eds.), CarnegieRochester Conference Series on Public Policy, 39, 5384.

[22] Phelan, Ch. 2005. Public Trust and Government Betrayal. *Federal Reserve Bank of Minneapolis.*

[23] Sargent, T.J., 1999. *The Conquest of American Inflation* Princeton University ress, Princeton.

[24] Siegel, S., and N.J. Castellan. 1956. *Nonparametric Statistics for the Behavioral Sciences.* McGraw-Hill Book Company.

[25] Stein J. 1989. Cheap Talk and the Fed: A Theory of Imprecise Policy Announcement. *The American Economic Review*, 79, no.1 (March): 32-42.

[26] Stokey, N.L. 1989. Reputation and Time Consistency. *The American Economic Review.* 79, (May) 2 : 134-139.

[27] Walsh, C.E. 1995. Optimal Contracts for Central Bankers. *American Economic Review*, 85, no.1 (March): 150-167

[28] Walsh, C.E. 1999. Announcements, Inflation Targeting, and Central Bank Incentives. *Economica*, 66: 255-269
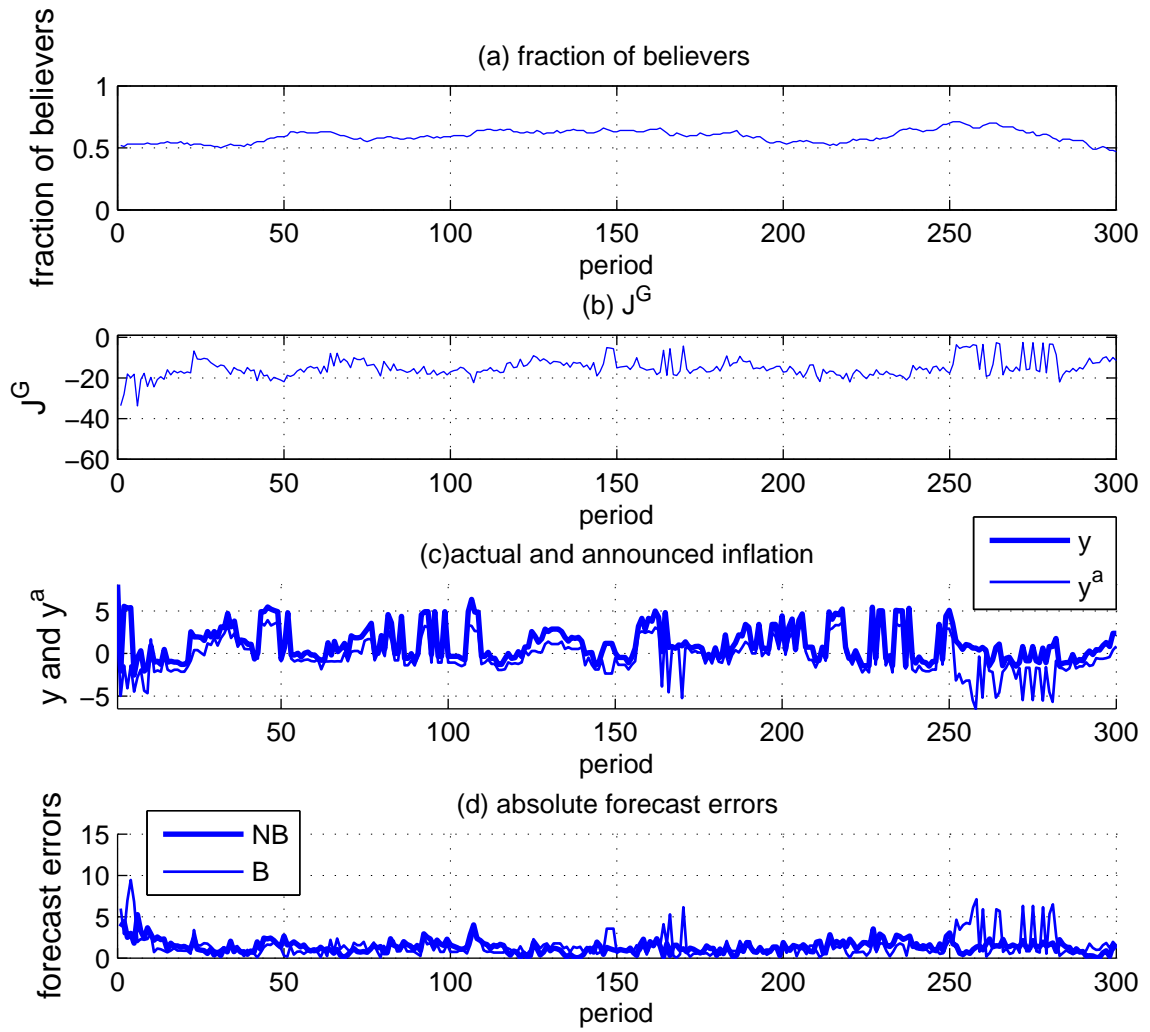
Figure 1: Evolution of the fraction of believers (a), policy maker payoffs (b), inflation rate and announced inflation rate (c), absolute forecast errors (d) for $c^{NB} = 0.1, \gamma = 0.1$ in the economy with common forecasts of nonbelievers.

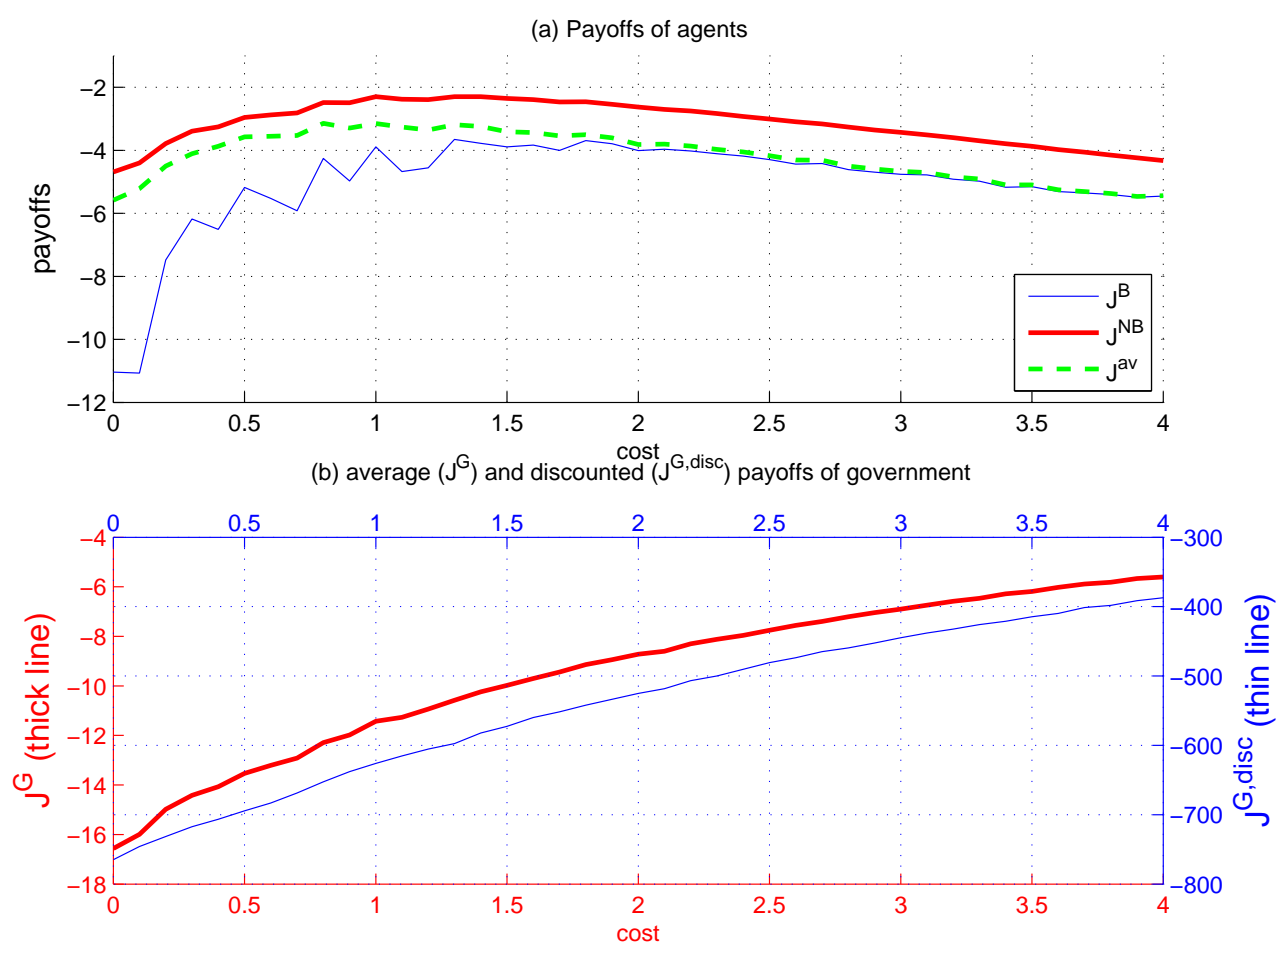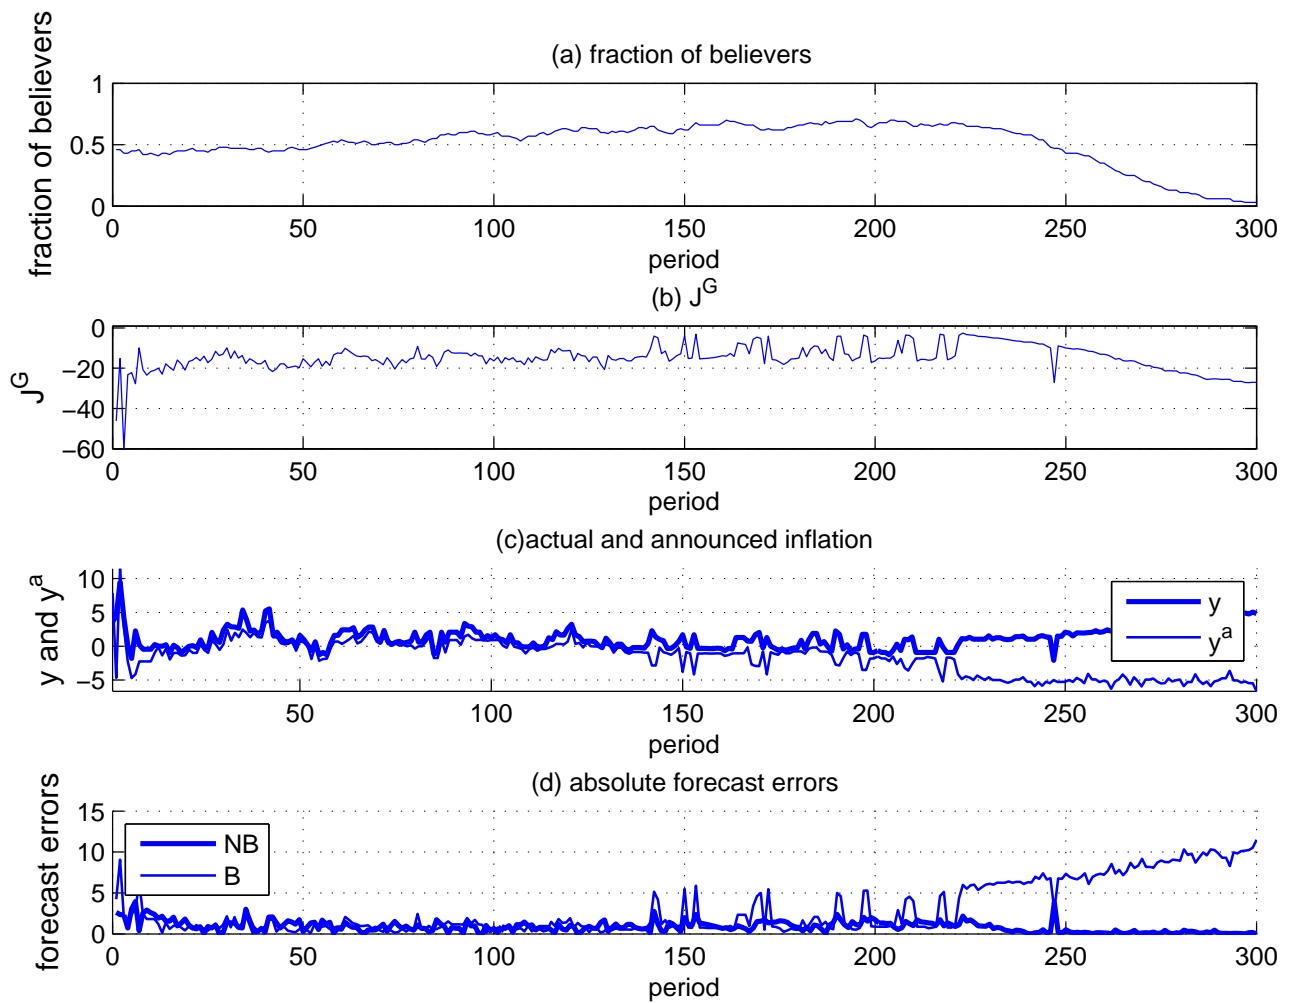Figure 2: Discounted payoff of the policy maker for different values $\Omega$.

22

Figure 3: Payoffs of believers and nonbelievers and average agents' payoffs (a) and discounted and average payoff of the policy maker (b) for different cost $c^{NB}$ for the case of homogeneous nonbelievers.
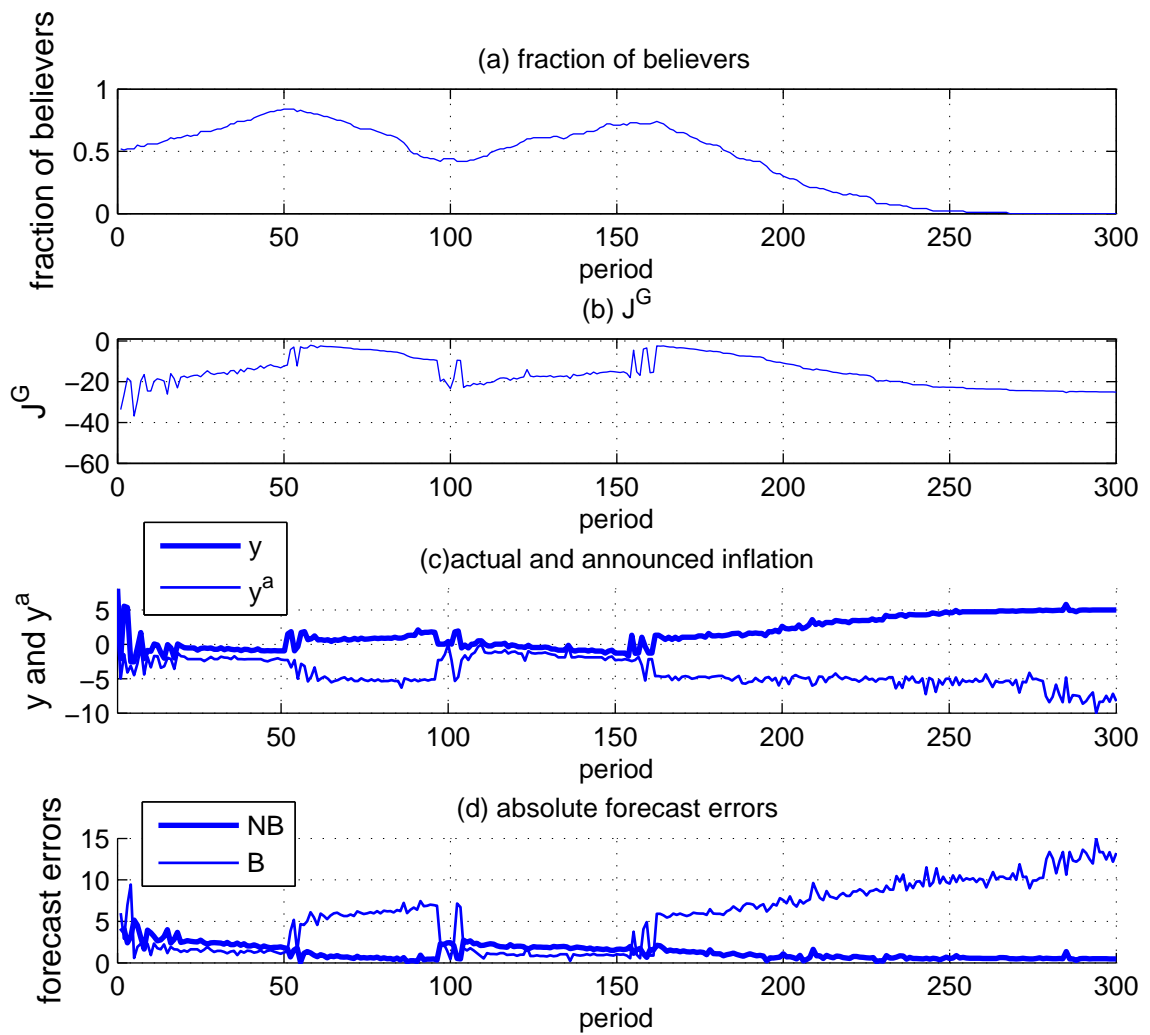
Figure 4: Evolution of the fraction of believers (a), payoff of the policy maker (b), inflation rate and announced inflation rate (c), absolute forecast errors (d) for $c^{NB} = 0.1, \gamma = 0.1$ in the economy with private forecasts of nonbelievers.

24

Figure 5: Payoffs of believers and nonbelievers and average agents' payoffs (a) and discounted and average payoff of the policy maker (b) for different cost $c^{NB}$ for the case of private forecasts of nonbelievers.

Figure 6: Evolution of the fraction of believers (a), payoff of the policy maker (b), inflation rate and announced inflation rate (c), absolute forecast errors (d) for $c^{NB} = 0.1, \gamma = 0.01$ in the economy with common forecasts of nonbelievers.

26

Figure 7: Evolution of the fraction of believers (a), payoff of the policy maker (b), inflation rate and announced inflation rate (c), absolute forecast errors (d) for $c^{NB} = 0.1, \gamma = 0.01$ in the economy with private forecasts of nonbelievers.