

Learning Benevolent Leadership in a Heterogenous Agents Economy *

Jasmina Arifovic[†] Herbert Dawid[‡] Christophe Deissenberg[§]
Olena Kostyshyna[¶]

January 15, 2010

Abstract

This paper studies the potential commitment value of cheap talk 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 use an adaptive learning scheme to (potentially) forecast the true inflation. 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. While all agents using the first strategy make the same forecast, those using the second strategy may generate different individual forecasts.

In spite of the complexity of the environment, the boundedly rational policy maker learns to sustain a situation with a positive but fluctuating fraction of believers. This outcome is Pareto superior to the outcome predicted by standard theory. Interestingly enough, the actions taken by the policy maker undergo marked qualitative changes as a function of the prevailing heterogeneity among and learning characteristics of the private agents.

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

JEL codes: C69, D83, E5

*The authors are grateful for helpful comments of two anonymous referees.

[†]Corresponding Author: Department of Economics, Simon Fraser University, Burnaby, BC V5A 1S6, Canada. Phone: + 1778 782-5603, Fax : +1778 782-5944, 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 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 to provide a signal about his 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 likewise investigate, in an economy closely related to the model of Kydland-Prescott (1977), the role of cheap talk announcements by the policy maker as a means to sustain a Pareto superior solution. 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 the best possible forecast of the policy maker's next action. An agent who uses the (not) believe strategy will be called a (*non*)*believer*. In each period, word of mouth information exchange allows a fraction of the agents to compare their last-period payoffs with the ones obtained by agents who followed the other strategy. Each agent then adopts the strategy that provided the highest payoff, and uses it until a new comparison motivates it to switch strategies anew. Thus, the proportion of believers may change over time and can be interpreted as a measure of the policy maker's credibility.

The consequences of the coexistence of believers and non believers were investigated in Dawid and Deissenberg (2005) in a continuous time framework, assuming a continuum of atomistic private agents and a perfectly rational policy maker able to correctly predict the impact of its actions on the agent's forecasts. In this earlier model, the policy maker and the nonbelievers do not learn. Moreover, all believers and all nonbelievers are symmetric. It was shown that in this highly stylized setting the policy maker will steer the economy towards a steady state that is Pareto superior to the Nash solution if the initial fraction of believers is sufficiently large. Thus, the existence of believers in the population proved to be Pareto improving.

Strongly reinforcing our confidence in the real value of this earlier work, similar results are obtained in the present paper under distinctively less demanding, and arguably more realistic assumptions: the policy maker has now limited abilities for dynamic optimization and forecasting, both the policy maker and the nonbelievers learn (so that there is coevolution of learning between the policy maker and the private agents), the population and time are discrete. Moreover, by contrast with the earlier work, the present paper main focus is on the implications of *learning* and *heterogeneity*. The very relevant policy framework in which we conduct our analysis allows us to highlight several issues of importance for policy design in a complex dynamic environment with boundedly rational and heterogeneous agents.

As indicated before, this paper extends an earlier framework in two main directions, *learning*, and *heterogeneity*. The learning issue is of prime importance in the monetary policy/time inconsistency context, see among others Sargent (1999), Cho, Williams and Sargent (2002), and Cho and Sargent (1997). The learning mechanisms used in the relevant literature include principally recursive least squares, constant gain and stochastic gradient learning. Learning either by the policy maker or by both the policy maker and the agents often narrows the choice of equilibria down to the time consistent but Pareto inferior Nash equilibrium and to the Pareto optimal but time inconsistent Ramsey outcome. However, simulations of the models show that the Nash equilibrium is most frequent occurrence, while the Ramsey outcome is observed only for brief periods of time. This however, contradicts Arifovic and Sargent (2003), who find from experiments with human agents

using the Kydland-Prescott model that the Ramsey outcome emerges most of the time and that the economy only occasionally slides to the Nash equilibrium.

In this paper, we depart from the existing learning literature by assuming that the nonbelievers use a simple adaptive learning mechanism to update their forecasts. The policy maker uses individual evolutionary learning to improve his strategy, see 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 an *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 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.

Likewise of prime importance is the issue of heterogeneity of agents, forecasts, and forecasting rules. At the empirical level, Branch (2004) finds evidence of 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. Also, the conditions for stability of learning are different under heterogenous expectations than with homogeneous expectations, see Giannitsarou (2003) and Honkapohja and Mitra (2005, 2006).

Here, the inclusion of two forecasting models (believe/not believe) allows us to contribute to the discussion on the importance of an explicit representation of varying degrees of agents heterogeneity in macroeconomic models. So does 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 maker. This has important implications for the outcome, and raises intricate 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 a boundedly rational policy maker with very little *a priori* and online information build up and nurture a positive stock of believers? If so, under which conditions do Pareto superior outcomes arise compared to scenarios without believers? What is the impact of variations of the main model parameters and of the forecasting schemes? Assuming that a policy maker attempts to maximize its cumulated discounted payoffs over an infinite time horizon using the short-horizon decision making rule considered in this paper, what are the preferable (ranges of) values for the main model parameters? In particular, how much concern for the intertemporal effects of its actions should the policy maker have? Should the policy maker favor speedy learning and extensive information exchange among the nonbelievers, and how?

Our answer to these questions provides important insights at two different levels. First, and as previously mentioned, the paper shows that the results obtained in e.g. Dawid and Deissenberg (2005) in a highly stylized model remain qualitatively valid in a more complex and arguably more realistic framework that includes bounded rationality, noise, coevolution of learning, agents' heterogeneity, and the emergence of efficient behavior (rather than a simple refinement of pre-defined decision rules). Thus, the paper gives a crucial (and typically missing) test of robustness with respect to real-world imperfection for the theoretical framework developed in earlier papers. Such a test, although important for assessing the practical relevance of a theoretical hypothesis, is very rarely conducted. Second, the paper reveals that variations in the difficulty of learning and degree of agents' heterogeneity lead at equilibrium to qualitatively different responses from the policy maker. In other words, it adds one more piece to the growing body of evidence indicating that learning and

heterogeneity are not subsidiary concepts that can be innocently superposed to a given model, but play a fundamental role in economics. Besides describing the consequences of agent's heterogeneity on the aggregate dynamics, the paper also highlights the different mechanisms whose interplay at the micro-level is responsible for such aggregate behavior.

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, respectively under the common and the private information cases. 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. Three Appendices provide respectively a pseudo-code of the simulation algorithm used in the paper, a detailed description of the computation of the hypothetical payoffs used by the policy maker to assess the fitness of alternative decision rules, and a synopsis of the main variables and their definition.

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 consists of a policy maker G and of a continuum of (private) agents i , $i \in [0, 1]$. Each agent i builds some forecast x^i about inflation. Its payoff, J^i , is given by

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

where y denotes actual inflation and, 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 a three stage, complete information game with the following structure:

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

Each agent can form its forecast in one of two ways. On one hand, it can set its inflation forecast equal to the announced inflation,

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

We say then that the agent acts as (or simply: is) a *believer*. Obviously, all believers share the same forecast $x^B = y^a$. We assume that the believers do not incur any forecasting costs: $c^B = 0$.

On the other hand, the agent can make a rational forecast of the future inflation, taking into account the fact that it is atomistic, the proportion of believers in the economy, and their forecasts. We then say that the agent acts as (is) a *nonbeliever*. Assume here (this will be relaxed at a later place) that all nonbelievers make the same forecast x^{NB} . In the Brock-Hommes (1997) tradition, the forecasting costs of the nonbelievers c^{NB} are assumed to be non-negative due to the computational, informational and other efforts they possibly need to build their own forecasts.

In the static game the fraction of believers in the population is an exogenous parameter denoted by ϕ . However, in the dynamic extension considered in the main part of this paper, each agent repeatedly decides which of the two forecasting strategies (to believe or not to believe) to follow. In that case, the proportion of believers in the economy, ϕ_t , is endogenously determined in each period t .

The unemployment rate of the believers respectively of the nonbelievers is given by an expectation augmented Phillips curve

$$u^i = u^* - \theta(y - \bar{x}^i), \quad i \in B, NB, \quad (2)$$

where u^* is the unemployment rate that would arise if agents had correct inflation expectations¹, \bar{x}^i is the average inflation forecast of the believers respectively of the nonbelievers, and $\theta > 0$ is a parameter.

The policy maker's payoff J^G is the ϕ -weighted sum of the squared rate of unemployment of believers and nonbelievers and of the squared rate of inflation, i.e.

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

This model would be a standard variant of the Kydland-Prescott (1977) seminal contribution if all agents rationally anticipated the same x^i – see e.g. Sargent (1999). In this paper, however, we will allow for heterogeneous, imperfect anticipations.

The static game for given ϕ can be solved through backward induction. At the last stage 3, the policy maker knows ϕ and the average forecasts \bar{x}^B and \bar{x}^{NB} of the believers and nonbelievers. Maximizing (3) with respect to y given ϕ , \bar{x}^B and \bar{x}^{NB} leads to the policy maker's optimal reaction function:

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

At stage 2, as previously stated, the forecasts of the believers is given by

$$\bar{x}^B = y^a. \quad (5)$$

The nonbelievers know ϕ , y^a , \bar{x}^B , and the optimal reaction function (4) of the policy maker. Each of them is rightly aware that it is too small to influence \bar{x}^{NB} . Thus, the inflation forecast of each nonbeliever i is given by $x^{NB,i} = y$, which implies $\bar{x}^{NB} := x^{NB} = y$ (Exploiting symmetry, we no longer differentiate between x and \bar{x} unless necessary). One has

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

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

$$y^{a*} = -\frac{u^*}{\theta}, \quad (7)$$

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

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

$$\begin{aligned} x^{B*} &= y^{a*}, \\ x^{NB*} &= y^*. \end{aligned}$$

Notice that the nonbelievers perfectly forecast the true inflation. Notice also that the true inflation, y^* , decreases with ϕ . 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, $x^{NB*} - x^{B*}$,

$$y^* - y^{a*} = x^{NB*} - x^{B*}. \quad (9)$$

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

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

¹Equation (2) implicitly assumes that the unemployment rate u instantaneously moves to the level u^* if the forecast error is 0. Although a standard assumption in neoclassical models, this might seem at odds with the general approach followed in this paper, which focusses on policy design under conditions of bounded rationality and adaptive learning. We nonetheless use this formulation to keep the treatment of the real effects of incorrect expectations as simple as possible, thus allowing us to focus on the core issues of expectation formation dynamics and policy formation.

the believers' payoff is

$$J^{B*} = -\frac{1}{2} \frac{(1 + 2\theta^2 + 2\theta^4 - 2\phi\theta^4 + \phi^2\theta^4)}{\theta^2(1 + \phi\theta^2)^2} u^{*2}, \quad (11)$$

and the nonbelievers' payoff is

$$J^{NB*} = -\frac{1}{2} y^{*2} - c^{NB} = -\frac{1}{2} \left[\frac{\theta(1 - \phi)}{1 + \theta^2\phi} u^* \right]^2 - c^{NB}. \quad (12)$$

One recognizes that, 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 ϕ : the policy maker, the believers and the nonbelievers are all better off when the proportion of believers increases. Indeed, as ϕ 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 the policy maker.

The higher the proportion of believers, the better every type of actor is. Actually, for $\theta > 1$ the equilibrium outcome of the static game with a sufficiently large proportion of believers even Pareto dominates the equilibrium where all private agents act rationally (i.e. there are no believers). But, if private agents can choose whether to believe the government announcement, having a positive fraction of believers is not an equilibrium, since then any individual agent can improve its payoff by switching from believing to nonbelieving.

This framework naturally suggests 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, ϕ , 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 ϕ -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 thus defined 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 ϕ , ϕ^E , with $0 < \phi^E < 1$. The instantaneous payoffs of the policy maker and of the private agents are larger at ϕ^E than at ϕ^0 .

A threshold ϕ^S , $0 < \phi^S < \phi^E$, separates the basins of attraction of the two equilibria: If the initial value of ϕ is less than ϕ^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 payoff of the policy maker is smaller at ϕ^0 than at ϕ^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 ϕ^0 . If the initial value of ϕ is greater than ϕ^S , it is optimal to build up a stock of believers until ϕ^E is reached. Since the non-believers do not adapt their forecasting rule over time (as previously mentioned, they use the reaction function of the static game (6) to form their forecasts), they keep making systematic forecasting mistakes even after the equilibrium ϕ^E has been 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 ϕ , with positive consequences for the policy maker's cumulated payoffs. Similar considerations also drive the dynamics 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.

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 out the same sequence of actions. The economy (2) and the instantaneous payoffs (1) 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. All actors know the economy (2), but do not know precisely how the other actors take their decisions. 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. There is coevolution of learning among nonbelievers and the policy maker. Nonetheless, and contrary to the dynamic model presented in the previous section, the believers change their forecasting rules in a way that guarantees that they would ultimately form correct forecasts, should the policy maker's actions remain constant. 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 μ_t believers and ν_t nonbelievers, with $\mu_t + \nu_t = \Theta$, where Θ is the (constant) total number of agents in the population. Accordingly, the fraction of believers at t is $\phi_t = \frac{\mu_t}{\Theta}$. 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 an 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 have 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, but know that the policy maker changes 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,

$$d_{t+1} = d_t + \gamma(y_t - x_t^{NB}), \quad (13)$$

$$d_0 = 0, \quad (14)$$

where $\gamma > 0$ is a parameter capturing the speed of learning. This mechanism insures that all nonbelievers will in the long run make correct inflation forecasts if the actions of the policy maker are constant.

Moreover, ϕ 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 β of the private agents is chosen randomly. The chosen agents are then randomly paired, e.g., agent i with agent k . Agent i observes agent's k current strategy, to believe or not to believe, and *vice-versa*. Each agent also *imperfectly* observes the payoff of the other. Agent i observes the payoff of agent k , as

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

where ϵ 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 ϕ_t is stochastic. Assuming for analytical simplicity that ϵ is drawn from a distribution qualitatively similar to a Gaussian distribution², the expected change in the proportion of believers is given by

$$\Delta \hat{\phi}_t := \mathbb{E} \phi_{t+1} - \phi_t = \beta \phi_t (1 - \phi_t) \arctan(J_t^B - J_t^{NB}). \quad (15)$$

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 other nonbelievers. Thus, all nonbelievers have the same d in any t , independently of their past choices of strategy.

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* $[y_t^a(j), y_t(j)]$, $j = 1, 2, \dots, N$. The initial pool of rules, Y_0 , is 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 (*experimentation*). In parallel, *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 $[y_t^a(j), y_t^b(j)] \in Y_t$, $j \in \{1, \dots, N\}$, is changed independently with some probability *maxp*. The new value after experimentation is computed as:

$$\text{new value} = \text{old value} + \xi,$$

where ξ 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 he would have received if he had used in period t any other rule in Y_t , and the corresponding expected change in ϕ_t . Let the values that would have 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

$$\tilde{J}_t^G(j) = -\frac{1}{2}[\phi_t(\tilde{u}_t^B(j))^2 + (1 - \phi_t)(\tilde{u}_t^{NB}(j))^2 + y_t^2(j)]. \quad (16)$$

Now, if G was solving a standard dynamic optimization problem (i.e., did maximize its cumulated discounted 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. Here, we do not assume that the policy maker explicitly solves an infinite horizon dynamic optimization problem, but that it nonetheless takes into account, albeit in a simplified form, 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 have used in t in terms not of $\tilde{J}_t^G(j)$ but in terms of the *pseudo-value function*

$$V_t^G(j) = \tilde{J}_t^G(j) + \Omega \Delta \tilde{\phi}_t(j), \quad (17)$$

where $\Omega > 0$ is a parameter and $\Delta \tilde{\phi}_t(j)$ the expected change of ϕ_t if rule j had been applied (see Appendix B). In that way, it assigns a positive value to an increase in ϕ , that is, takes into account that a higher ϕ now should allow a higher payoff in the future. Loosely speaking, the parameter Ω measures how strongly G takes into account the expected intertemporal consequences of the different actions he considers.

²To obtain the analytically convenient expression (15), we assume that ϵ 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 $[0, 1]$.

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 Given Y_t , experimentation and replication allow the policy maker to construct 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's fitness. Specifically, we assume that the probability for rule j to be used in period $t + 1$ is given by:

$$P_{t+1}(j) = \frac{V_t^G(j)}{\sum_{i=1}^N V_t^G(j)}. \quad (18)$$

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 (i.e., that is associated with a high value of (17)) will indeed perform well when applied in $t + 1$.

On the meaning of Ω The individual learning procedure outlined above describes the behavioral rule that the policy maker is assumed to follow. Within this rule, G evaluates the different rules in its pool Y_t based on a function which only takes into account the intertemporal effects for the next period. This does not mean that G is not concerned about further periods. Rather, this pragmatically reflects the fact that the complexity of the environment does not allow G to correctly predict the future, even statistically.³ Within the learning procedure, G approximates the trade-off between the current and the future effects of an action by the value of the pseudo-value function (17). If G was able to perfectly predict the future, then Ω would be the derivative of the standard value function with respect to the state ϕ_t and thus, would typically depend upon ϕ_t . But, since G is unable to compute the value of the true co-state, we interpret Ω in (17) as a constant control parameter of the behavioral rule. In order to evaluate the outcomes induced by the decision procedure of G from the perspective of the policy maker, we compute the cumulated discounted payoff of G over an infinite time horizon for a given discount factor ρ . In spite of the complexity of our environment, this allows us to investigate which values of the control parameter Ω a policy maker with given time preferences should choose in order to maximize its cumulated discounted payoffs.

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 besides the one describing the static game. It therefore uses the statically optimal reaction function (6) to make its prediction for period t , i.e., it sets $d = 0$. In the subsequent periods, and as long as it uses the *NB* strategy, it updates the error correction term according to (13). In this way, we introduce heterogeneity among the nonbelievers: Each nonbeliever uses in each t its own private error correction term d_t^i (Nonetheless, all agents that became nonbelievers in the same period use the same error correction term). Accordingly, the forecast of a nonbeliever i in period t is given by

$$x_t^{NB,i} = \frac{\theta^2 \phi_t y^a + \theta u^*}{1 + \theta^2 \phi_t} + d_t^i, \quad (19)$$

³In Dawid (2005) it was shown that in environments where a policy maker is not able to make correct predictions about the future marginal returns on investments, a heuristic with a planning horizon of one period may generate higher cumulated discounted payoffs over an infinite horizon than infinite horizon dynamic programming.

and its payoff by

$$J_t^{NB,i} = -\frac{1}{2}[(y_t - x_t^{NB,i})^2 + y^2] - c^{NB}. \quad (20)$$

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 ν_t is, as previously stated, the current number of nonbelievers. Otherwise, the model is unchanged.

3 Simulation Results in the Common Forecasts Case

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

- $u^* = 5.5$, $\theta = 1$, see Sargent (1999)
- $\beta = 0.05$, $c^{NB} = 0.1$, $\gamma = 0.1$, $\Omega = 1000$
- $mexp = 0.2$, $\Theta = 100$, $N = 100$

A summary of the main notations used in the paper is provided in

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 and initialization.

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). From the onset, let us emphasize that our results *do not* depend upon the assumption of positive forecasting costs, $c^{NB} > 0$. They are also valid when $c^{NB} = 0$. We use $c^{NB} = 0.1$ in our baseline simulation only to respect the assumption of positive forecasting costs usually made in the heterogeneous agents literature.

3.1 Dynamics under common forecasts

Figure 1 illustrates the typical dynamics of ϕ_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 when Ω is sufficiently large. A clear pattern arises between periods 50 and 200, after some initial transients. The actual inflation y_t oscillates around zero. The announced inflation y_t^a is smaller than y_t , but both time series exhibit 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 ϕ 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 ϕ 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. Qualitatively similar results are obtained for $c^{NB} = 0$, see Figure 2.⁴

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). The proportion of believers fluctuates in response to the gap between the announced and the actual inflation, which determines the relative payoffs of the private agents. The policy maker enjoys a high credibility when its announcement is close to the actual inflation, a low credibility when the announced inflation differs much from the actual one. In the first case the believers are better off than the nonbelievers, and worse off in the second case. A high proportion of believers (a high credibility) can be sustained even when the actual inflation is not precisely equal to the announced one – sufficient therefore is just a reasonably small discrepancy between the two values.

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 the outcome of the 20 initial 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 1 shows that this cyclical behavior *Pareto dominates* 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).

This Pareto improvement occurs not only *in spite* of the fluctuations and of the forecast errors that characterize our economy. Most importantly, it also occurs *thanks to* these fluctuations and errors. Indeed, were there no *explicit* forecasting costs (i.e., $c^{NB} = 0$) and no fluctuations, the *NBs* would learn to perfectly forecast the true inflation and to systematically fare better than the *Bs*. Only *NBs* would survive in the long run. However, since the actions of the policy maker fluctuate over time in a highly unpredictable manner, the nonbelievers' forecasts are not necessarily more accurate than the believers' forecasts. Thus, it may be advantageous to be a believer even if $c^{NB} = 0$. This, in turn, increases the economic welfare – providing an interesting argument (that will, however, be qualified at a later place) to mitigate the usual call for greater predictability of economic policy.

In other words, the forecasting errors generate *implicit* forecasting costs, since they make the forecasts inexact and the decision of the *NBs* incorrect. The relevant cost measure, the total forecasting costs incurred by a *NB's*, is the sum of the explicit and implicit costs. *Thus, although forecasting costs are essential for sustaining a population of Bs, they may stem both from limits to the exactitude of human forecasting and/or from the direct costs of the forecasting activities.*

Extensive numerical analyses showed that cheap talk insures the long-run sustainability of positive fraction of *Bs* and a Pareto-improvement for *large compact ranges* of parameters. However, it will not occur if some parameters are too large or too small. In particular, as mentioned earlier, the results presented in Figure 1 require sufficiently large values of Ω , that is, sufficient concern of the policy maker for the future. For lower values of Ω it is possible to generate simulations where ϕ 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 *merp*.

⁴Our economy exhibits fluctuations similar to those in Sargent (1999). However, the underlying mechanisms are different. In Sargent's model, the policy maker learns from a misspecified Phillips curve and the agents are rational. The dynamics consist of mean dynamics that pushes the economy towards a self-confirming equilibrium with Nash inflation, and of escape dynamics that pushes the economy away from the self-confirming Nash towards the Ramsey equilibrium. These dynamics are characterized analytically in Cho, Williams, Sargent (2002). The escapes are possible because the policy maker discounts past data in its estimation of the Phillips curve by using constant gain recursive least squares. The economy spends most of the time near the Nash equilibrium. Escapes towards the Ramsey outcome are rare events.

3.2 Is there an optimal value of Ω ?

In the previous simulations, the control parameter Ω was chosen arbitrarily. We only required it to be large enough to insure the existence of a solution with a sustainable population of believers. However, remember that it is through Ω that the policy maker takes intertemporal effects into account. Therefore, it should choose the value of Ω that maximizes its true, intertemporal preferences, if such a value exists.

Assume that the policy maker's underlying preferences are given by

$$\Gamma^G = \sum_{t=0}^T \rho^t J_t^G,$$

where $\rho \in (0, 1)$ is the discount factor. Figure 3 shows how Γ^G varies with Ω for $\Omega \in [800, 2000]$ and $\rho = 0.98^5$.

One recognizes that Γ^G is first increasing, then decreasing in Ω . Thus, there exists indeed an optimal value of Ω . To understand this result remember that a policy maker with a high value of Ω 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 Ω , 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 Ω is excessive, the short term losses needed to obtain a high proportion of believers dominate the gains this high proportion generates in the long term, and Γ^G decreases.

3.3 Sensitivity analysis: The impact of c^{NB} and γ

The previous discussion suggests 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 reduces the decrease in instantaneous payoffs G must accept in order to improve the relative position of the believers, and thus, to ensure a high value of $\Delta\phi$ and ϕ . 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, γ .

As we shall see, both changes indeed increase G 's average payoff \bar{J}^G and cumulated discounted payoff Γ^G . However, they have markedly different impacts on the key economic variables. In the simulations, we use as before $\rho = 0.98$ to compute Γ^G . The parameter Ω is given the value $\Omega = 1000$, that is, the value that roughly maximizes Γ^G for the reference parameter constellation. We use a Wilcoxon test to check whether or not the changes in c^{NB} and γ lead to statistically significant differences in the average values of the main variables⁶. Note that we investigate here the impact of a *local* change of c^{NB} or γ at the baseline parameter values. At later places, we will show that the policy response to variations of c^{NB} respectively γ is not monotonic: It qualitatively changes at a threshold, whose value depends upon the values of all other parameters. In other words, depending upon where we are situated in the parameter space, a change in the value of c^{NB} or γ may cause one of two possible qualitative policy answers.

Impact of a local increase in c^{NB} Table 1 shows the discounted payoff Γ^G as well as the averages over the periods 20 to 300 of the main variables of interest when $c^{NB} = 0.1$ (its value in the basis scenario) and when $c^{NB} = 1$. 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 . But the average unemployment rates of both believers and nonbelievers

⁵If one interprets each period as two weeks, which seems reasonable given $\beta = 0.05$ and the adjustment speed of ϕ_t , this corresponds to a yearly discounting of about 10%.

⁶To carry out the statistical analysis of the effects of changes of c^{NB} we generate 100 pairs of (γ, Ω) values using a uniform distribution on $[0.01, 0.1] \times [800, 1300]$. For each of these (γ, Ω) 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. For each of the 100 parameter profile, we take the differences of the variable under consideration between the two runs and use the Wilcoxon test to check for statistically significant differences. We proceed in a similar manner in the cases of changes of γ .

also increase, with adverse effects on Γ^G and \bar{J}^G . Since $\Delta\phi > 0$ and $u^B < u^{NB}$, however, the increase in average unemployment is smaller at the aggregate population level than the increase observed for each agent type. *In toto*, the policy maker's average payoffs increase.

Likewise, both nonbelievers and believers are better off after an increase in c^{NB} , due to a smaller value of $|\bar{y}^a - \bar{y}|$ and a lower 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.

The mechanisms underlying these observations are quite straight forward. As c^{NB} increases, a nonbeliever needs to forecast y more and more accurately in order to perform better than a believer. Initially, the *NBs* do not make sufficiently accurate forecasts, and their number decreases. As time goes, the remaining nonbelievers learn to make better forecasts until they obtain, in spite of the high value of c^{NB} , approximately the same payoffs as the *Bs*. As their forecasts are more accurate, their unemployment is lower and very close to the natural rate of unemployment. The increase in the stock of believers reduces the incentives for the policy maker to choose high inflation, since unemployment among believers can be reduced by lowering y^a as well as by increasing y , whereas increasing inflation is the only instrument to reduce unemployment among nonbelievers. As discussed below, this reasoning only applies as long as the increase in c^{NB} does not push the level of believers to such a high level that it becomes optimal for the policy maker to exploit that stock by choosing (y^a, y) in a way that low inflation is combined with low unemployment among believers. Positive forecasting costs $c^{NB} > 0$ are helpful but not necessary for the sustainability of credibility. In the simulations presented in Table 1 for instance, see also Figure 2, a positive fraction of believers is maintained although $c^{NB} = 0$.

Summarizing, positive explicit forecasting costs for the nonbelievers ($c^{NB} > 0$) help establish the credibility of cheap talk announcements, but are not necessary for it. However, persistent fluctuations of the inflation rate are necessary for this credibility when $c^{NB} = 0$. Without such fluctuations and the implicit costs they generate, the nonbelievers would, in the long run, learn the true inflation rate. If they learned it, they would always realize a higher payoff than the believers if $c^{NB} = 0$, leading to the extinction of the latter.

Impact of a local decrease in γ Quite a different picture emerges when γ is reduced from $\gamma = 0.1$ to $\gamma = 0.01$, see Table 2. Here again, the differences between the two sets of runs are almost always statistically significant at the 95% level, see Row 3 of the Table. But 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 his discounted cumulated payoff Γ^G . But the loss of credibility in later periods implies a lower average payoff \bar{J}^G . Both believers and nonbelievers also have a lower average payoffs after a decrease of γ .

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 reality. Society as a whole profits, too, from efficient nonbelievers.

The intuition for this finding is as follows. If the *NB* learn slowly, their forecasts remain inaccurate and their unemployment high for a long time. That is, there is no or little incitation for a *B* to switch from believing to not believing. The policy maker uses this opportunity to lower the unemployment rate of the believers by increasing the discrepancy between announced and actual inflation. Even when they learn slowly, however, the nonbelievers eventually start making more accurate forecasts than the believers, and the latter start to switch to not believing. The policy maker reacts with some inertia to this new situation, by moving to a more believers-friendly policy. (This inertia is due to the fact that the policy maker's pool of rules still consists mostly of 'exploitative' rules selected during the period when the nonbelievers were inaccurate forecasters, and that it takes time until rules aiming at nurturing the stock of believers establish themselves through experimentation.) With some probability, the proportion of believers becomes zero before this policy switch is realized, making all agents worse off. By contrast, if the nonbelievers are fast learners, the believers are not exploited over long periods, the inertia mentioned above does not take serious proportions, and believers do not disappear.

These observations raise the question why the response to an increase in costs of forecasts differs from that

to a decrease of the speed of update. To understand this different responses it should be noted that higher cost affects the NBs ' payoffs uniformly throughout the simulation. The policy maker is confronted in every period with its consequences, and can consistently learn how to react to it. The lower learning speed, by contrast, impacts the NBs ' most when their forecasts are poor. In particular, when the policy maker switches from a policy aiming at the build-up of the stock of believers to one exploiting the Bs , initially also the NBs have relatively low payoffs due to the slow adaptation of their beliefs. A low learning speed prolongs the time during which the NBs have low payoffs and the policy maker can exploit the Bs . When the NBs have learned to make accurate forecasts, the learning speed is no longer of importance to them. But it takes time for the policy maker to realize that new state of affairs and to switch to appropriate policies. If the policy maker changes its policy too slowly and keeps exploiting the Bs , they will switch to nonbelieving since this strategy now insures the highest payoffs. The inertia in the response of the policy maker may even lead to a complete extinction of the Bs .

Global analysis: The optimal value of c^{NB} , from the perspective of the different actors Additional simulations indicate that the results described above hold for all $c^{NB} \in [0.1, 1]$ and $\gamma \in [0.01, 0.1]$. A decrease in the learning speed γ 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, increasing c^{NB} does not always improve all payoffs. Rather, for any given value of γ , there is an optimal level of c^{NB} beyond which the one or the other payoff starts declining. Figure 4 shows the average payoffs of G and of the Bs and NBs as c^{NB} varies from 0 to 4 in increments of 0.1, for $\gamma = 0.1$. 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, ϕ -weighted payoff of all private agents also reaches its maximum at $c^{NB} = 0.8$. By contrast, Γ^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's 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 γ is low.

The fact that (different) interior values of c^{NB} are optimal for the believers respectively the nonbelievers is by no means trivial. Intuitively, one might have expected that an increase in c^{NB} would harm the nonbelievers but would not affect the believers. However, our findings show that these direct effects of an increase in c^{NB} on the NBs and Bs are dominated by the policy changes this increase induces, i.e., by a lowering of the inflation rate when c^{NB} is sufficiently small. The total effect of an increase of c^{NB} is then positive for believers and nonbelievers alike. Likewise, it is not attractive to switch to not believing when γ is low and the nonbelievers learn slowly, thus making inaccurate forecasts.

The existence of an optimal c^{NB} helps us understand the difference in the policy response to changes in c^{NB} and in γ . When c^{NB} increases from 0.1 to 1, it becomes easier for the policy maker to sustain a population of believers. However, c^{NB} is below the critical level of 1.3 and thus not high enough to hinder the Bs from easily switching to not believing. Therefore, the policy maker cannot use an overly aggressive policy to lower the believers' unemployment. When c^{NB} is above the critical value, it is not attractive for the believers to change their strategy. The policy maker uses this to aggressively lower their unemployment. Likewise, a sufficiently low γ makes not believing unattractive during the long time needed for the nonbelievers' forecasts to become sufficiently accurate. The policy maker uses this time to reduce the believers' unemployment.

4 Simulation Results in the Private Forecasts Case

So far we 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 nonbelievers with a possibly long learning history. As another extreme case, we now consider the private forecast situation described in section 2.4. Here

an agent, when becoming a nonbeliever, 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. Its total forecasting costs (defined as the sum of the explicit costs c^{NB} and of the implicit costs resulting from inaccurate forecasts) are now higher. This lowers its payoffs and make it more prone to switch back to the B strategy. That is, 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 3 summarizes the outcome of simulations 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 $\bar{\phi}$, Γ^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 according to the Wilcoxon test, see row 3.

The economy's time behavior under the private forecast regime is illustrated in Figure 5. Compared to the common forecast case (Figure 1), ϕ 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 because of the increased inexactitude of the initial forecasts. 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. Note that the mechanisms that lead to an initially higher proportion of believers and to its possible decline over time are similar to those described in section 3.3 for the case a lower speed of adjustment. We explore these dynamics further in section 4.2.

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 overstates the actual inflation rate, and the average unemployment of the nonbelievers is larger than the natural rate of $u^* = 5.5$ (see table 3). This leads to lower payoffs for both G and the NBs . 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 γ

As noted earlier, in case of common forecasts, the policy maker responds in qualitatively different ways to changes in the parameters c^{NB} and γ . An increase in c^{NB} leads to less inflation, a decrease in γ 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 γ increases Γ^G but reduces \bar{J}^G , \bar{J}^B , and \bar{J}^{NB} .

The impact of c^{NB} in the private forecast case Table 4 shows, in the private forecast case, the effect on the key variables of an increase of c^{NB} from 0.1 to 1. 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 Γ^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. To understand why, remember that in the common forecast case the impact of such an increase is not monotonic. For reasons discussed earlier, an increase in c^{NB} beyond a given threshold leads to lower payoffs for all agents. Intuitively, this threshold corresponds to a value of c^{NB} where the level of believers after the initial build-up of believers becomes so large that the policy

maker has incentives to exploit that stock. As discussed in section 3.3 the profitability of an exploitation strategy for the policy maker depends on the way the nonbelievers adapt their beliefs. In particular, it has been shown that even for $c^{NB} = 0.1$ the policy maker uses a strategy that exploits (and thus depletes) the stock of believers when the speed of learning is decreased to $\gamma = 0.01$. Table 4 is based on simulations where the *individual* speed of learning for each nonbeliever is $\gamma = 0.1$. However, due to the facts that forecasts are now private and that after a switch to nonbelieving an agent starts with $d = 0$, the *average* speed by which the performance of the *NBs* improves after the policy maker implements an exploitation strategy is much lower than in the common forecasts case with $\gamma = 0.1$. Hence, contrary to the common forecasts case, a higher c^{NB} now makes it more attractive for the policy maker to exploit the *Bs*. This causes a decrease in the average level of believers, a higher inflation, and a lower payoff for the *Bs*, see Table 4.

As discussed above, the existence of private forecasts also implies that the payoff that a *B* *observes* when paired with a *NB* is typically higher than the payoff this *B* *obtains* immediately after switching to nonbelieving. Therefore, the probability that an agent that recently became a nonbeliever returns to believing is much higher than under common forecasts. Thus, the number of believers tends to decrease slower in the private than in the common forecasts case.

Figure 6 provides additional evidence on these phenomena. It was generated analogously to Figure 4 by letting c^{NB} vary from 0 to 0.4, in 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 ϕ -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, Γ^G and \bar{J}^G keep increasing with c^{NB} , (with some volatility).

Additional simulations not reported here show that, under private forecasts too, the fraction of believers remains positive in the long run even if $c^{NB} = 0$. The underlying mechanisms are the same as those described earlier for the common forecasts case.

The impact of γ in the private forecast case The impact of lowering γ from 0.1 to 0.01 and the associated Wilcoxon test are presented in Table 5. When the speed of learning, γ , 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 γ is lower, the nonbelievers make larger forecast errors and experience a higher rate of unemployment that exceeds the natural rate). As in the common forecasts case, this reflects the fact that the policy makers tends to exploit more the believers for $\gamma = 0.01$. Contrary to the common forecast case, however $\bar{\phi}$ increases. Since \bar{y} and \bar{u}^B decrease and $\bar{\phi}$ increases, \bar{J}^G and Γ^B are higher. \bar{J}^B and \bar{J}^{NB} are lower. Overall, the direction of the changes of the unemployment rates triggered by a decrease of γ is the same for the common and private forecast cases, but the direction of the change of the inflation rate differs. As discussed above, this is due to the fact that the stock of believers increases in one case, and decreases in the other.

In order to more closely examine the differences in the directions of change of the stock of believers and other variables, let investigate in more detail the dynamic patterns of the key variables in the model for $\gamma = 0.1$ and $\gamma = 0.01$. As a first step consider (in addition to the time averages for the periods 21 to 300 already presented in 2 and 5) the time averages over the periods 21 to 100 and 21 to 300. Once again, the Wilcoxon test shows a statistically significant impact at the 95% level.

Tables 6 and 7 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 γ leads to a higher $\bar{\phi}$ over the first 100 periods. Over 300 periods, $\bar{\phi}$ decreases in the common, increases in the private forecasts case. In both cases, $\bar{\phi}$ is lower over 300 periods than over 100 periods. This illustrates the point made above: The proportion of believers declines over time when the policy maker uses an exploitation strategy, but much slower in the private than in the common forecast case.

For a given γ , 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, but increases over the 300 initial 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 γ implies a reduced unemployment rate for the believers but, due to higher forecast errors, a higher one for the nonbelievers. Altogether these effects lead to a decrease in \bar{J}^G . However, Γ^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 γ . In the

case of common as in the case of private expectations, this decrease implies that the nonbelievers take longer to improve their forecasts and to reduce their total forecasting costs. Consequently, the policy maker is able to build up faster a large mass of believers. In both cases, once ϕ 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 ϕ . This decline starts earlier when γ is lower, as illustrated in Figures 1 and 7 for the common forecasts case, and in Figures 5 and 8 for the private forecasts case. For a given γ this decline is slower under private than under common forecasts. As can be seen in these figures, when the nonbelievers learn slowly, the exploitation of the believers by the policy maker eventually leads to their extinction. This is not in the long-run interest of the policy maker G , but happens because in the given complex environment G is not able to fully predict the implications of his 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, γ must be sufficiently high. This hinders the policy maker from exploiting the believers excessively and thereby avoids their eventual extinction.

The consequences of heterogeneity As we have seen, the impact of changes in c^{NB} and γ are for the same parameter values very different in the case of common and private forecasts, see Table 1 and Table 4, Table 2 and Table 5. This is in particular the case in what concerns the policy response, i.e., acting primarily on the unemployment or on the actual inflation. Table 8 summarizes these differences. The first arrow in each cell indicates the direction of change of the time averaged variable in the common, the second arrow the direction of change in the private forecast case. For changes in the speed of learning (γ) the direction of change of the fraction of believers and the inflation rate differs between the common and private forecasts case. If c^{NB} varies, then all listed variables other than the discounted payoff of the policy maker move in opposite directions under common and private forecasts! In our discussion above we have shown that the mechanisms driving these differences can be clearly identified.

All this illustrates the crucial importance of explicitly taking into account any 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 γ are observed. Thus, in the present setting but presumably also in many others, *using a 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. This creates 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, reduce both the explicit forecasting costs c^{NB} and the implicit ones. This, as we have seen above, is not necessarily desirable.* Thus, the policy maker faces the non-trivial problem of keeping the total forecasting costs of the nonbelievers as high as possible while avoiding that they build wildly diverging and inexact forecasts.

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 not believing, 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 policy maker must solve an complex intertemporal trade-off 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 announcements 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 remains strictly positive. The irregular cycles of trust building and trust exploitation are shown to be sufficient for the sustainability of a positive proportion of believers, with all the associated positive economic effects. Without them, the nonbelievers would learn to make exact forecasts. If the forecasting errors are zero, they would have always higher payoffs than the believers, who would be thus driven to extinction. The exercise shows 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.

We furthermore showed that the speed with which the nonbelievers learn, and the costs they must pay to compute their inflation forecast are of crucial importance for the qualitative properties of the outcome. These costs are not limited to the explicit forecasting costs c^{NB} . They also include the implicit costs resulting from inaccurate forecasts. Most importantly, the consequences of a modification of the value of one of these two parameters are not monotonic in the change. We have analyzed in details why in different circumstances the decision maker may favor a lower unemployment of the believers or a lower inflation. Thus, in some numerical exercises, 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. To maintain a positive proportion of believers, it is necessary that the nonbelievers update quickly and are accurate in their forecasts: this prevents the policy maker from exploiting believers too much.

In the paper, we also investigated the consequences of giving up the assumption of identical nonbelievers, replacing it by the following hypothesis: Every time 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 similar to a lowering of the nonbelievers learning speed, but also has distinctive implications. Heterogeneity amounts to an increase of private agents' total forecasting costs and has a strong impact on the decision makers' actions and the outcomes.

An interesting dilemma 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 *NBs* and therefore diminish the attractiveness of being a *B*. One is therefore interested in mechanisms that insure homogenous and accurate nonbeliever forecasts while maintaining the incentives for being a believer. Finding them is not a trivial task.

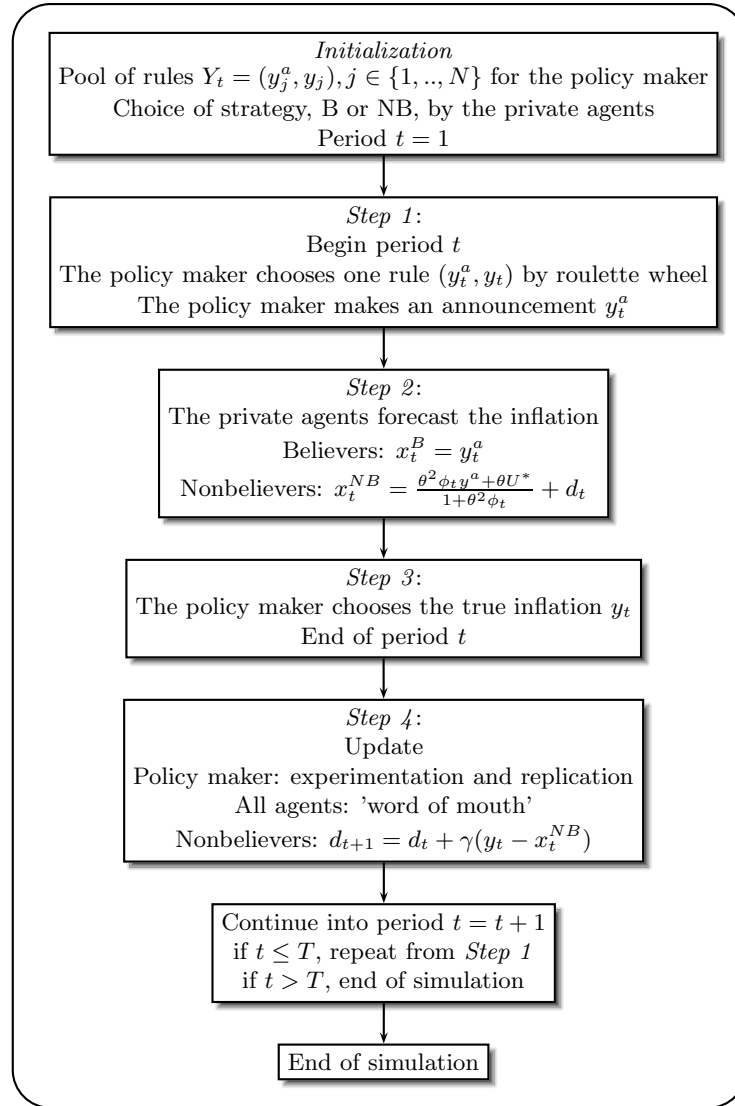
Overall our results strongly suggests that explicitly taking into account the ubiquitous heterogeneity and imperfection that characterizes a real economy is primordial if one wants to reach a proper understanding of its working and to make appropriate policy recommendations.

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., and D.B. Gordon. 1983. Rules, Discretion, and Reputation in a Model of Monetary Policy. *Journal of Monetary Economics*, 12 (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 Survey 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 Route to Chaos in a Simple Asset Pricing Model. *Journal of Economic Dynamics and Control*, 22: 1235-1274
- [10] Cho, In-Koo, and T.J. Sargent. 1997. Learning to be Credible. Manuscript.
- [11] Cho, In-Koo, Williams, N., and T.J. Sargent. 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. 2005. Long Horizon versus Short Horizon Planning in Dynamic Optimization Problems with Incomplete Information. *Economic Theory*, 25: 575-597.
- [14] Dawid, H., and C. Deissenberg. 2005. On the Efficiency Effects of Private (Dis)-Trust in the policy maker. *Journal of Economic Behavior and Organization*, 57: 530-550.
- [15] Deissenberg C., and G.F. Alvarez. 2002. Cheating for the Common Good in a Macroeconomic Policy Game. *Journal of Economic Dynamics and Control*, 26, Issues 9-10 (August): 1353-1785.
- [16] Evans, G., and G. Ramey. 1992. Expectation Calculation and Macroeconomic Dynamics. *American Economic Review*, March, 82 (1) : 207-224
- [17] Evans, G., and G. Ramey. 1998. Calculation, Adaptation and Rational Expectations. *Macroeconomic Dynamics*, 2(2): 156-182
- [18] Giannitsarou, C. 2003. Heterogeneous Learning. *Review of Economic Dynamics* 6(4): 885-906.
- [19] Honkapohja, S., and K. Mitra. 2005. Performance of Monetary Policy with Internal Central Bank Forecasting. *Journal of Economic Dynamics and Control* (29): 627-658.

- [20] Honkapohja, S., and K. Mitra. 2006. Learning Stability in Economies with Heterogeneous Agents. *Review of Economic Dynamics*, 9 (2): 284-309.
- [21] Kydland F.E., and E.C.Prescott. 1977. Rules Rather Than Discretion: The Inconsistency of Optimal Plans. *Journal of Political Economy*, 85, No.3 (June): 473-491.
- [22] Persson, T. and G. Tabellini. 1993. Designing Institutions for Monetary Stability. In A. H. Meltzer and C. I. Plosser (eds.), *Carnegie-Rochester Conference Series on Public Policy*, 39, 53-84.
- [23] Sargent, T.J., 1999. *The Conquest of American Inflation* Princeton University Press, Princeton.
- [24] Stein J. 1989. Cheap Talk and the Fed: A Theory of Imprecise Policy Announcement. *The American Economic Review*, 79, no.1 (March): 32-42.
- [25] Walsh, C.E. 1999. Announcements, Inflation Targeting, and Central Bank Incentives. *Economica*, 66: 255-269

APPENDIX A. The pseudo-algorithm of the 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

$$\tilde{J}_t^G(j) = -\frac{1}{2}[\phi_t(\tilde{u}_t^B(j))^2 + (1 - \phi_t)(\tilde{u}_t^{NB}(j))^2 + y_t^2(j)] + \Omega \Delta \tilde{\phi}_t(j). \quad (21)$$

The hypothetical values of the believers' inflation forecasts are calculated for each rule j as

$$\tilde{x}_t^B(j) = y_t^a(j). \quad (22)$$

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

$$\tilde{x}_t^{NB}(j) = x_t^{NB,aver} + \frac{\theta^2 \phi_t}{1 + \theta^2 \phi_t} (y_t^a(j) - y_t^a), \quad (23)$$

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

The hypothetical unemployment rates are computed as

$$\tilde{u}_t^i(j) = u^* - \theta(y_t(j) - \tilde{x}_t^i(j)), \quad i = B, NB. \quad (24)$$

The policy maker computes the expected change in the proportion of believers, knowing that it is determined by (15). This requires to calculate

$$\begin{aligned} \tilde{J}_t^B(j)(\tilde{x}_t^B(j), y_t(j)) &= -\frac{1}{2}((y_t(j) - \tilde{x}_t^B(j))^2 + y_t^2(j)), \\ \tilde{J}_t^{NB}(j)(\tilde{x}_t^{NB}(j), y_t(j)) &= -\frac{1}{2}((y_t(j) - \tilde{x}_t^{NB}(j))^2 + y_t^2(j)), \\ \Delta \tilde{\phi}_t(j) &= \beta \phi_t (1 - \phi_t) \arctan(\tilde{J}_t^B(j) - \tilde{J}_t^{NB}(j)), \\ \tilde{\phi}_{t+1}(j) &= \phi_t + \Delta \tilde{\phi}_t(j). \end{aligned} \quad (25)$$

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 Notation

Notation	Meaning
μ	number of believers
ν	number of nonbelievers
Θ	total number of private agents
ϕ	fraction of believers
y^a	announced inflation
x^j	inflation forecast of group j , $j = B, NB$
y	realized inflation
u^j	unemployment of group j , $j = B, NB$
θ	parameter in the expectations-augmented Phillips curve
J^G	instantaneous payoff of the policy maker
ρ	discount factor
Γ^G	cumulated payoff of the policy maker
V^G	fitness function of the policy maker
Ω	parameter in the policy maker's fitness function
$J^{P,j}$	instantaneous payoff of group j , $j = B, NB$
c^{NB}	(explicit) forecasting costs of the nonbelievers
d	error correction term
γ	updating speed of the error correction term
β	fraction of private agents participating in 'word-of-mouth' information exchange
N	number of rules in the policy maker's pool of rules
$mexp$	probability of experimentation

cost	Γ^G	J^G	ϕ	\bar{y}	\bar{y}^a	U^B	U^{NB}	$J^{P,B}$	$J^{P,NB}$
$c^{NB} = 0.1$	-745.84	-16.00	0.51	1.82	-0.50	3.18	5.47	-11.07	-4.40
$c^{NB} = 1$	-625.99	-11.42	0.69	0.94	-0.83	3.73	5.48	-3.89	-2.30
α	0.000005	0.3859	0.0129	0.00005	0.0885	0.4602	0.0007	0.000005	0.000005
$c^{NB} = 0$	-764.75	-16.56	0.47	1.96	-0.32	3.23	5.46	-11.04	-4.68

Table 1: Effect of an increase in c^{NB} under common forecasts. This table reports average values of the main variables in the simulations with cost, c^{NB} , equal 0.1, 1 and 0. α is the significance level of Wilcoxon test applied to simulations with cost, c^{NB} , equal 0.1 and 1.

γ	Γ^G	J^G	ϕ	\bar{y}	\bar{y}^a	U^B	U^{NB}	$J^{P,B}$	$J^{P,NB}$
$\gamma = 0.1$	-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
α	0.000005	0.3859	0.1492	0.0005	0.000005	0.000005	0.000005	0.000005	0.000005

Table 2: Effect of a decrease of γ under common forecasts. This table reports average values of the main variables in the simulations with updating speed, γ , equal 0.1 and 0.01. The last row reports the significance level of Wilcoxon test, α .

type	Γ^G	J^G	ϕ	\bar{y}	\bar{y}^a	U^B	U^{NB}	$J^{P,B}$	$J^{P,NB}$
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
α	0.0005	0.0007	0.0009	0.242	0.000005	0.000005	0.000005	0.000005	0.000005

Table 3: This table reports average values of the main variables in the simulations with common forecasts (row 1) and private forecasts (row 2). The last row reports the significance level of Wilcoxon test, α .

cost	Γ^G	J^G	ϕ	\bar{y}	\bar{y}^a	U^B	U^{NB}	$J^{P,B}$	$J^{P,NB}$
$c^{NB} = 0.1$	-750.71	-16.51	0.48	1.54	-1.38	2.58	5.70	-14.00	-4.51
$c^{NB} = 1$	-659.30	-16.59	0.44	1.85	-3.04	0.61	5.63	-25.01	-6.15
α	0.000005	0.00005	0.0015	0.0019	0.000005	0.000005	0.000005	0.000005	0.000005
$c^{NB} = 0$	-766.63	-16.29	0.49	1.43	-1.03	3.04	5.71	-11.49	-4.08

Table 4: Effect of an increase of c^{NB} under private forecasts. This table reports average values of the main variables in the simulations with cost, c^{NB} , equal 0.1, 1 and 0. α is the significance level of Wilcoxon test applied to the simulations with cost, c^{NB} , equal 0.1 and 1.

γ	Γ^G	J^G	ϕ	\bar{y}	\bar{y}^a	U^B	U^{NB}	$J^{P,B}$	$J^{P,NB}$
$\gamma = 0.1$	-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
α	0.00005	0.00005	0.4286	0.1562	0.000005	0.0005	0.000005	0.00005	0.00005

Table 5: Effect of a decrease of γ under private forecasts. This table reports average values of the main variables in the simulations with updating speed, γ , equal 0.1 and 0.01. The last row reports the significance level of Wilcoxon test, α .

periods		Γ^G	J^G	ϕ	\bar{y}	\bar{y}^a	U^B	U^{NB}	$J^{P,B}$	$J^{P,NB}$	$y - x^{NB}$
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
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 6: Effect of a decrease in γ under common forecasts. This table reports average values of the main variables during initial 100 periods and during 300 periods of the simulations with updating speed, γ , equal 0.1 and 0.01.

periods		Γ^G	J^G	ϕ	\bar{y}	\bar{y}^a	U^B	U^{NB}	$J^{P,B}$	$J^{P,NB}$	$y - x^{NB}$
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
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 7: Effect of a decrease in γ under private forecasts. This table reports average values of the main variables during initial 100 periods and during 300 periods of the simulations with updating speed, γ , equal 0.1 and 0.01.

	Γ^G		ϕ		\bar{y}		U^B		U^{NB}	
$\gamma \downarrow$	\uparrow	\uparrow	\downarrow	\uparrow	\uparrow	\downarrow	\downarrow	\downarrow	\uparrow	\uparrow
$c^{NB} \uparrow$	\uparrow	\uparrow	\uparrow	\downarrow	\downarrow	\uparrow	\uparrow	\downarrow	\uparrow	\downarrow

Table 8: Effect of changes of γ and c^{NB} . The first arrow in each cell gives the effect under common forecasts, the second under private ones.

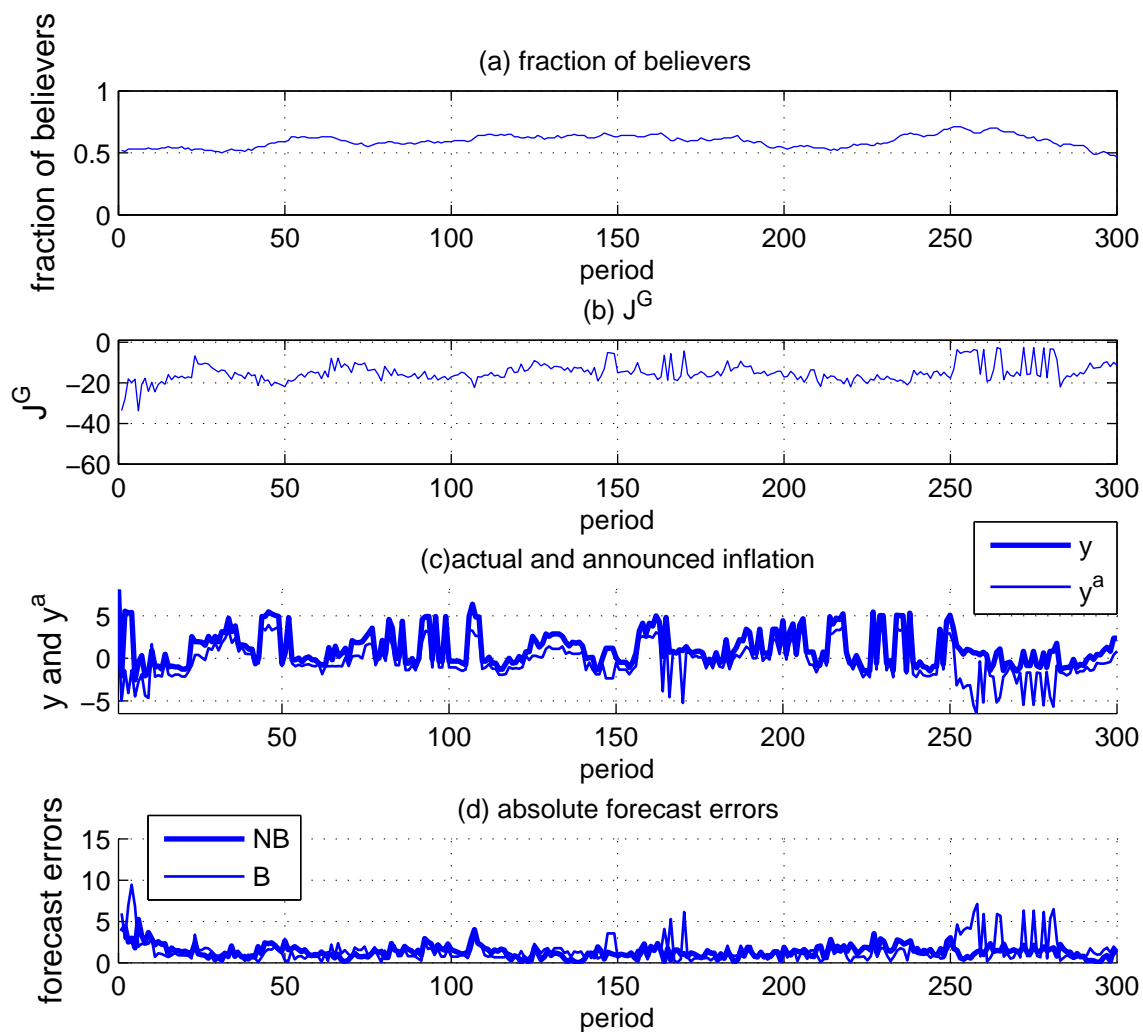


Figure 1: (a) Evolution of the proportion of believers, (b) policy maker payoffs, (c) announced and realized inflation rate, (d) absolute forecast errors, for $c^{NB} = 0.1, \gamma = 0.1$ under common forecasts.

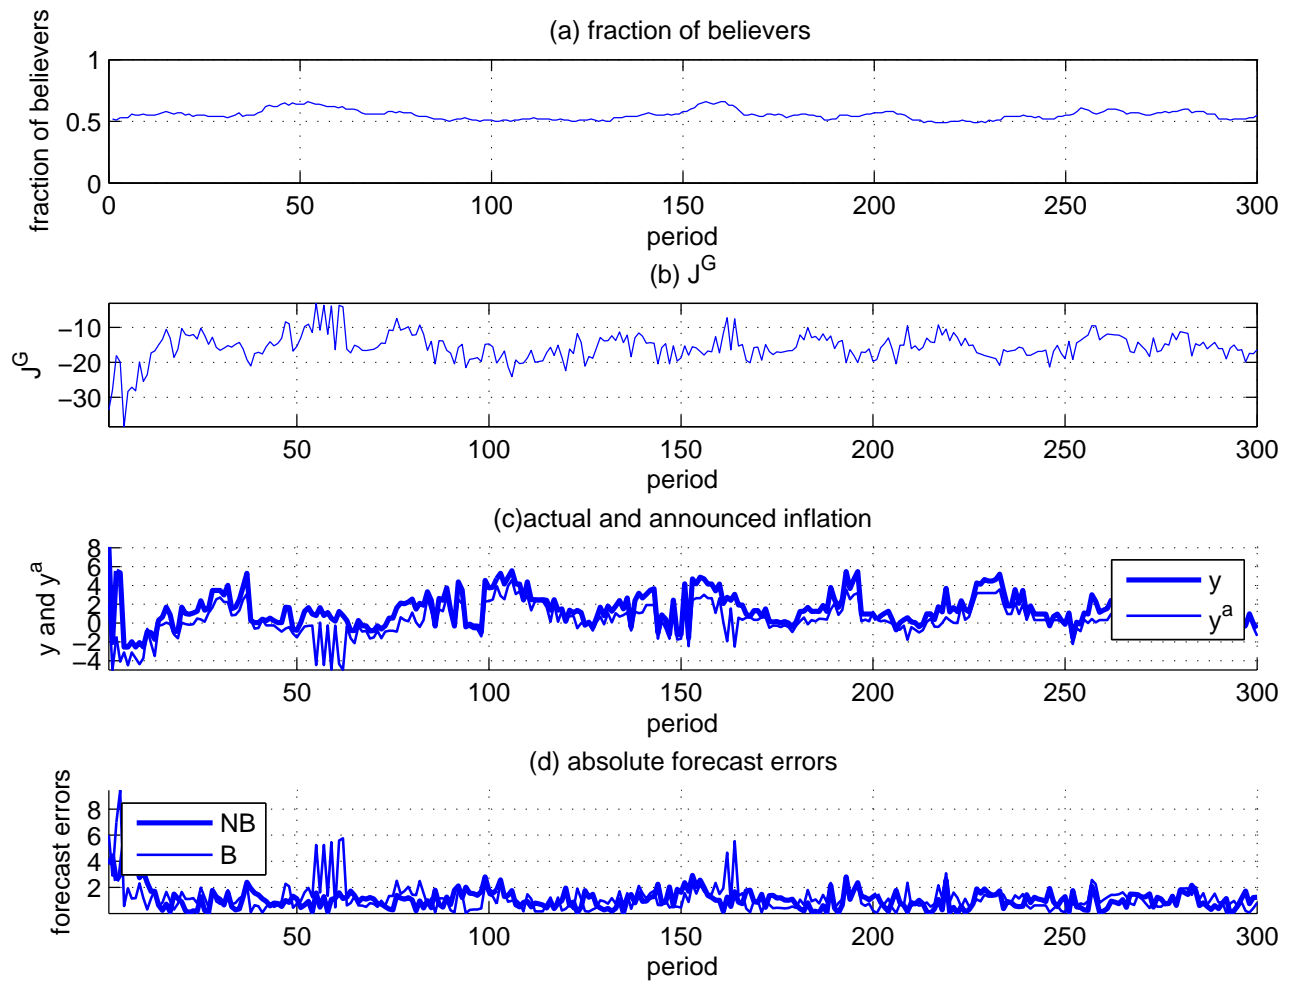


Figure 2: (a) Evolution of the proportion of believers, (b) policy maker payoffs, (c) announced and realized inflation rate, (d) absolute forecast errors, for $c^{NB} = 0, \gamma = 0.1$ under common forecasts.




Figure 3: Discounted cumulated payoff of the policy maker for different values Ω .


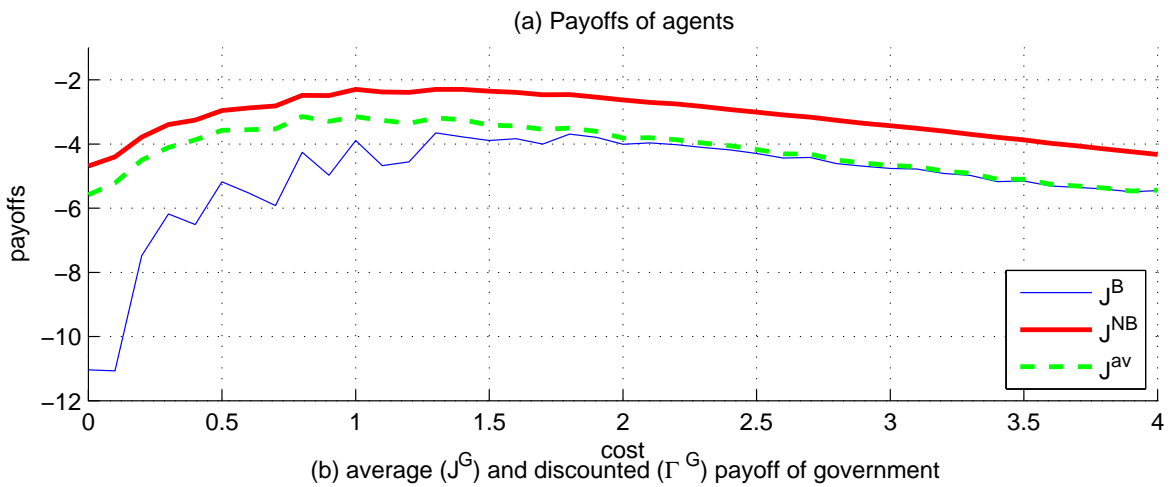


Figure 4: (a) Payoffs of believers and nonbelievers and average private payoffs; (b) cumulated discounted and average payoff of the policy maker, for different costs c^{NB} under common forecasts.

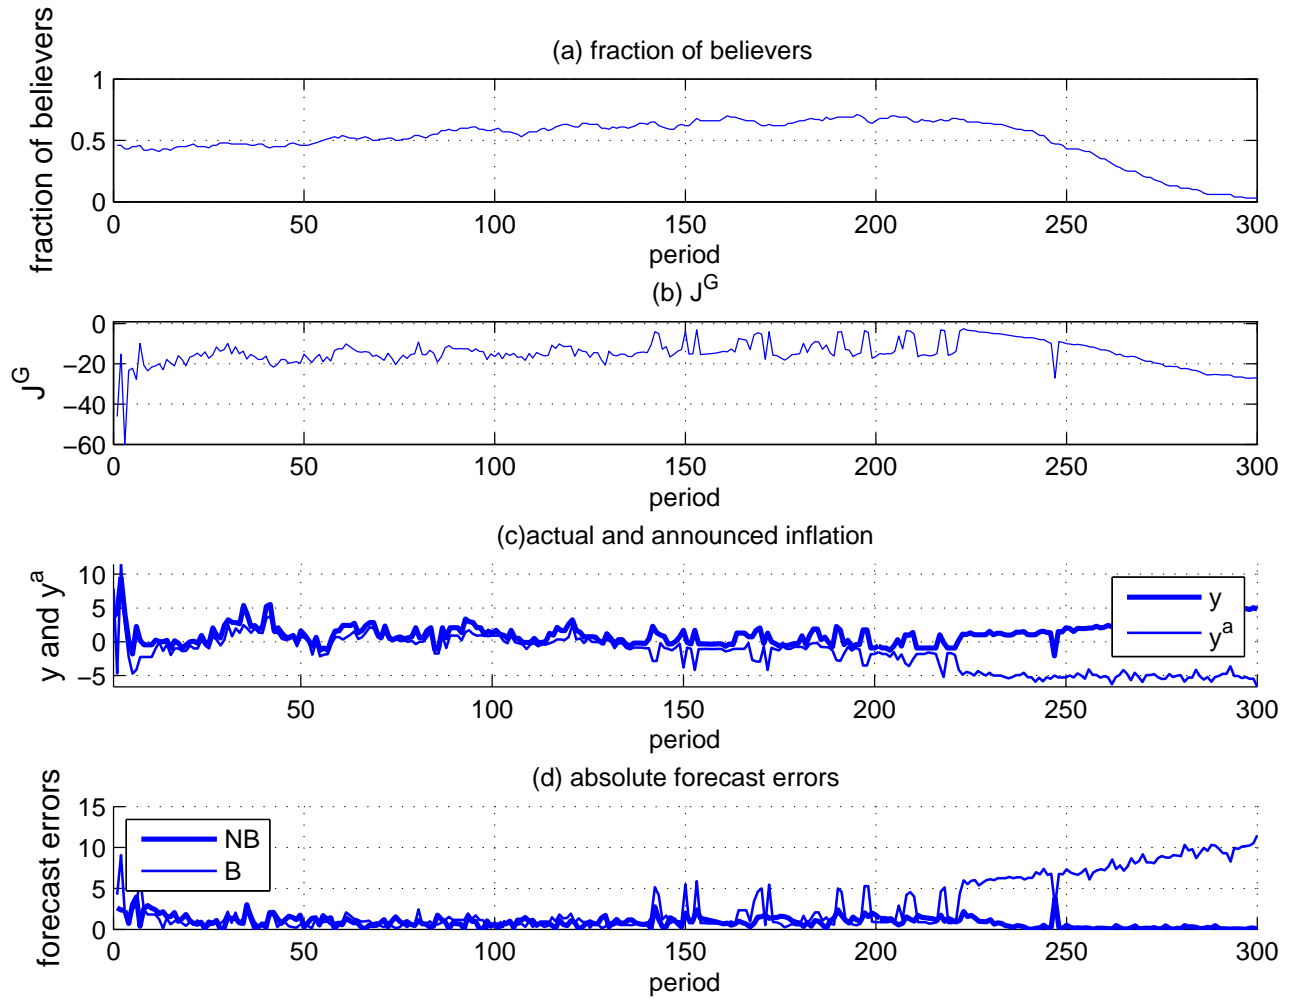


Figure 5: (a) Evolution of the fraction of believers , (b) payoff of the policy maker, (c) announced and realized inflation rate, (d) absolute forecast errors for $c^{NB} = 0.1, \gamma = 0.1$ under private forecasts.




Figure 6: (a) Payoffs of believers and nonbelievers and average private payoffs; (b) cumulated discounted and average payoff of the policy maker, for different costs c^{NB} under private forecasts.




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

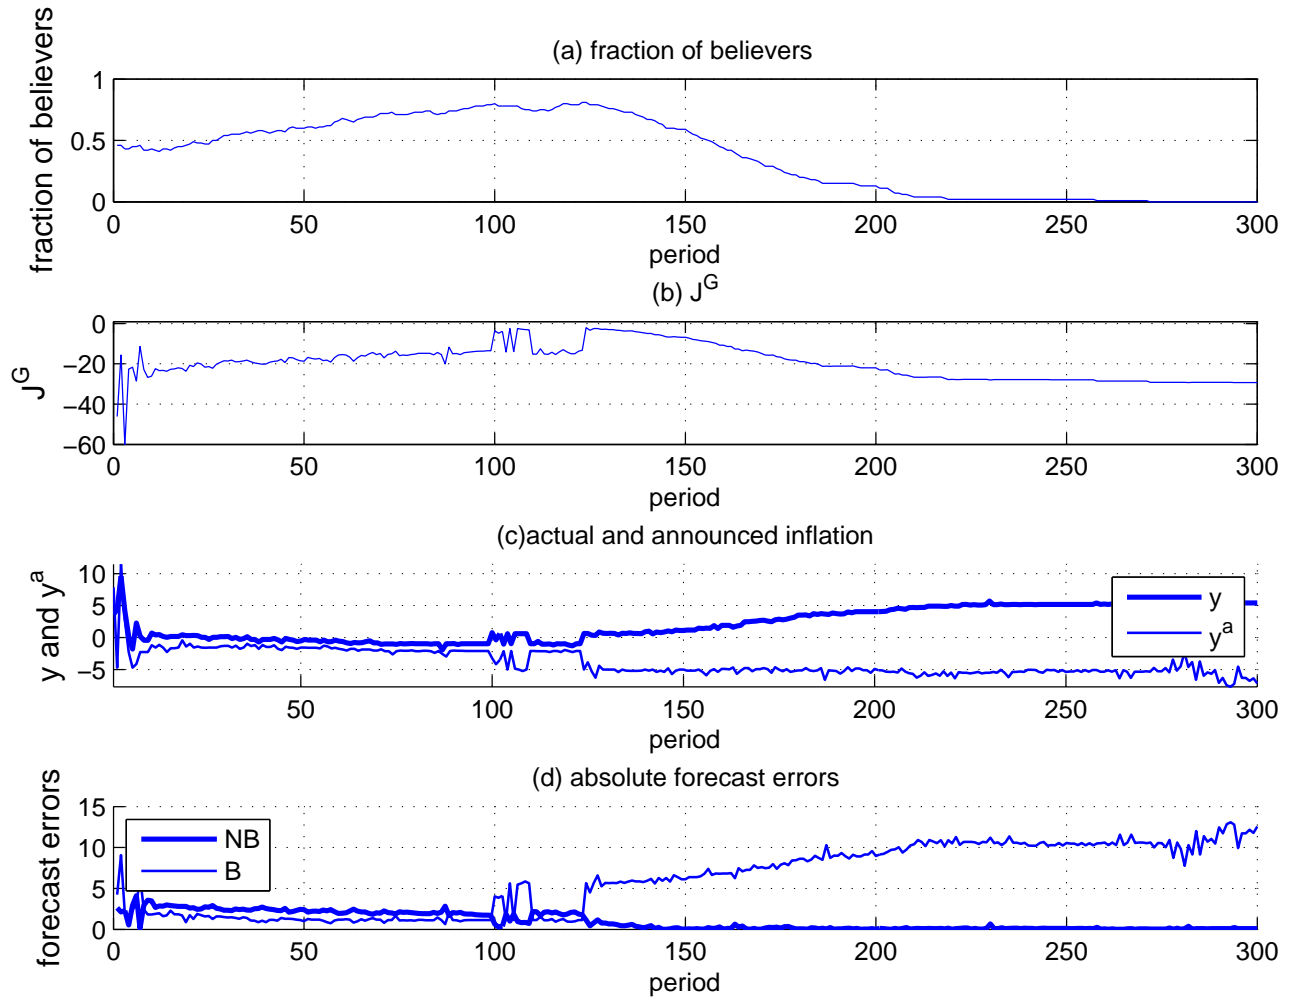


Figure 8: (a) Evolution of the fraction of believers , (b) payoff of the policy maker, (c) announced and realized inflation rate, (d) absolute forecast errors for $c^{NB} = 0.1, \gamma = 0.01$ under private forecasts.