CHEATING FOR THE COMMON GOOD

IN A MACROECONOMIC POLICY GAME

Christophe Deissenberg*[1]

Francisco Alvarez Gonzalez**

*CEFI, UMR CNRS 6126

Department of Economics

Université de la Méditerranée (Aix-Marseille II)

Château La Farge

Route des Milles

13290 Les Milles

France

**Universidad Complutense, Madrid, Spain

Abstract: This paper presents a simple repeated-game model of interaction between an optimizing government and the private sector. Two polar cases are considered: (a) the private sector is represented by a single agent; and (b) there is a continuum of heterogenous atomistic private agents. In both cases, the government starts each repetition by making a non-binding announcement about its future actions. The players have complete and perfect information, with one exception: the private agents do not know whether or not the government will act as announced. Thus, each private agent i either behaves with probability π_i as if it trusted the announcement, or plays with probability $1-\pi_i$ as a Stackelberg leader. After observing the reaction of the private sector, the government

[1]Corresponding author. E-mail: deissenb@univ-aix.fr

implements the actual policy measures. Finally, the private agent(s) update π_i as a function of the payoffs received.

We show that, although the government's announcements are never respected, acting as if they were true leads to an outcome that simultaneously improves the situation of the government and of those players that trust it, compared to the standard equilibrium solutions. An overwhelming majority of private agents very rapidly learns to behave as if the announcements were true. The other agents also experience higher payoffs due to an herding effect. The fact that the announcements are always reneged is crucial for the solution to be pareto-improving. These results are in stark contrast to the conclusions usually presented in the related economic literature.

Keywords: Macroeconomic policy-making, Kydland-Prescott model, time inconsistency, reinforcement learning, reversed Stackelberg games, optimal cheating strategies, reputation, credibility.

JEL classification: C69, C79, E5

Acknowledgement 1 *The authors would like to thank many friends and colleagues for kind and constructive comments, and most particularly Olivier Musy, Willi Semmler, the participants of the January 2000 Bielefeld workshop in honor of Edmond Malinvaud, and the participants of the October 2000 Workshop on Dynamics and Control in Rio de Janeiro. This paper was written while Francisco Alvarez Gonzalez was visiting at CEFI, that provided kind and much appreciated hospitality.*

1 Introduction

One of the most ubiquitous and stable characteristics of real-life policy-making appears to be that decision-makers repeatedly make announcements and promises that they later do not respect. Based on previous experience, private agents are fully aware that the promises they hear are unlikely to be kept. Yet, governmental announcements are not neutral. They have an impact on the private agents' decisions.

One may wonder why announcements that are suspected from the onset not to be respected are not totally disregarded. In this paper, we analyze a situation where these announcements are taken into account because, although it is known that they will be violated, (a) they contain useful information that may allow the agents increase their welfare; and (b) not acting as if they were true may prove to be costly. Thus, the announcements may help as a device to coordinate on a superior outcome that would not otherwise result from the interplay between government and private sector. Clearly, there are many other mechanisms that might explain the real impact of deceitful announcements. These additional or alternative explanations will not be reviewed here.

The potential usefulness of deliberately using misleading announcements to pareto-improve upon standard equilibrium solutions was suggested for general linear-quadratic dynamic games in Vallée and Deissenberg (1998) and subsequent papers by the same authors. We demonstrate it here in the context of a repeated game version of the archetypal Kydland and Prescott (1977) model – a model that is famed in the economic literature for having popularized a diametrically opposed point of view, namely, that the inability of a government to commit itself to its announcements necessarily implies a very poor economic outcome.

The paper is organized as follows. We first present the version of the Kydland-Prescott model used in this paper and the main conclusions usually drawn thereof in the related literature. Based on perceived ambiguities in the standard analysis, we then reinterpret the underlying static game between the government and the private sector, and present the concepts of optimal and pareto-improving cheating strategies, that are central to the further analysis. In the next section, the static game is used to define a repeated game between a government that makes optimally false announcements, and a single private agent that learns whether or not to disregard these announcements. In a further section, the model is extended to the case of heterogenous private agents. In the single as in the heterogenous agent cases, simulations are used to gain insights on the game's properties and emerging outcomes. The last section concludes.

2 The Analytical Background

2.1 The Kydland-Prescott one-period model

Numerous variants of the basic Kydland-Prescott model have been proposed, that do not crucially modify its basic message. This paper uses the Stokey (1989), (1990) formulation, as presented in the recent work by Sargent (1999), to which we refer at later places.

In the Stokey-Sargent formulation, the Kydland-Prescott model is as follows. A government, L, is engaged in a one-shot game with a continuum of private agents. There is perfect information. The player's objective functions and constraints are common knowledge. The government attempts to maximize through its choice of y its payoff function:

$$J^L = -\frac{1}{2}\left(U^2 + y^2\right), \tag{1}$$

where U is the unemployment rate and y the inflation rate.

Each private agent F^i tries to maximize through its choice of x_i the payoff function:

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

where x_i is the agent's expectation about y. Assuming that all private agents make the same choice, the average expectation, x, is equal to x_i. The game between the government L and the private agents reduces then to a two-players game between L and the aggregate private sector F, whereby the private sector uses x as instrument and has the payoff function:

$$J^F = -\frac{1}{2}\left[(y - x)^2 + y^2\right]. \tag{3}$$

The unemployment is supposed to depend upon inflation through an expectations augmented Phillips curve:

$$U = U^* - \theta\left(y - x\right), \tag{4}$$

where $U^* > 0$ is the natural rate of unemployment, and $\theta > 0$ is a parameter. This last equation asserts that unemployment U deviates from its natural rate only when the private sector anticipates incorrectly the inflation rate, $x \neq y$. Taking into account (4), J^L can be rewritten as:

$$J^L = -\frac{1}{2}\left[\left(U^* - \theta\left(y - x\right)\right)^2 + y^2\right]. \tag{5}$$

The optimal reaction functions $T^L : x \to y$ and $T^F : y \to x$ of L and F are given by:

$$T^L \triangleq \arg\max_y J^L = \frac{\theta}{1+\theta^2}U^* + \frac{\theta^2}{1+\theta^2}x, \tag{6}$$

$$T^F \triangleq \arg\max_x J^F = y. \tag{7}$$

The reaction function T^F implies that the private sector can – and will – always make a perfect prediction of whatever inflation rate the government chooses. This very strong property is usually interpreted as reflecting the private sector's ability to make rational expectations about the governments actions. As noted at a later point, it is a direct consequence of the very specific definition of the private sector's payoff function (3).

Two solution concepts have been traditionally studied in the context of the

Kydland-Prescott model, namely, the Nash equilibrium (N) and the Stackelberg equilibrium with L as a leader (SL).

The Nash equilibrium N is given when each player chooses a best response to the action of his opponent and, thus, has the nice property of being self-enforcing. It is given by:

$$x^N = y^N = \theta U^*, \tag{8}$$
$$J^{L,N} \triangleq J^L\left(y^N, x^N\right) = -\left(1+\theta^2\right)U^{*2},\ J^{F,N} \triangleq J^F\left(y^N, x^N\right) = -\frac{1}{2}\theta^2 U^{*2}.$$

The Stackelberg equilibrium with L as a leader SL, or Ramsey equilibrium, describes the outcome of a hierarchical situation where L chooses its action y knowing that F will react to this choice with its best response $T^F(y)$. Using from now on the notation introduced in (8), this equilibrium is given by:

$$x^{SL} = y^{SL} = 0,\ J^{L,SL} = -\frac{1}{2}U^{*2},\ J^{F,SL} = 0. \tag{9}$$

Clearly, this equilibrium strictly pareto-dominates the Nash equilibrium: both players are better off under Ramsey than under Nash. Assume, however, that the government had initially played Ramsey, and that the private sectors accordingly expects $x = x^{SL} = 0$. The best answer of the government to $x = 0$ is $y = \theta U^*/\left(1+\theta^2\right)$. That is, the government has ex post an incentive to deviate from Ramsey[2]. If the government does deviate, the private sector will revise its

[2] This "time inconsistency" of the solution is generic for Stackelberg equilibria. It reflects the fact that, contrary to a Nash equilibrium, a Stackelberg equilibrium does not correspond to a fixed point in the space of strategies.

expectations, leading to a new best answer by the government, and so on. It is easy matter to show that this iterative process converges towards the Nash equilibrium.

2.2 A few motivating remarks

The existence of ex post incentives to deviate from the Ramsey equilibrium led to the well-known pessimistic conclusion that, in the absence of binding commitments that force it to play Ramsey, a government would tend to renege on its previous engagements and act in a way that leads to the inferior Nash solution. An important and influential trend of research used this finding as a justification for advocating strict restrictions on government discretion in economic policy-making.

The story, however, is less clear-cut than above presentation may have led to believe. In particular, note that in the course of the presentation we departed from the original static, one-shot game description of the problem to argument within a dynamic framework. However, in a dynamic game, the curse of being forced to coordinate on the unfavorable Nash equilibrium is by no means inevitable, see e.g. McCallum (1997). In particular, the Ramsey outcome can be supported by the need for the government to maintain a favorable reputation, see e.g. Backus and Driffill (1985). Similarly, it can be supported if the private sector has the ability to punish the government whenever the latter deviates, see among others Rogoff (1987), Stokey (1989). The time inconsistency problem can also be mitigated by incentive contracts or by delegation, Personn and

Tabelinni (1993), Rogoff (1985), for example. More fundamentally, in a dynamic framework, almost any outcome can be supported as an equilibrium. This includes many outcomes that pareto-dominate both the Nash and the Ramsey equilibria.

Furthermore, note that it is not conceptually straightforward to interpret y as an actual economic policy decision if one defines x as an expectation on this decision. Consider for example the game sequence supporting the Ramsey outcome: L plays first by choosing y; F, that knows y, plays second by making an expectation x on y. The expectation is made after the action is realized, clearly a rather contrived scenario. This last observation (and more elaborate versions of it) led diverse authors to argue very early that the original Kydland-Prescott problem is not well-defined, see among others Basar (1989), Hall and Henry (1989), Miller and Salmon (1982). To avoid possible logical contradictions, it is necessary to make a clear differentiation between the announcement of a policy measure, y^a, and the measure actually implemented, y.

As a final remark, notice that both with the Ramsey and the Nash solutions, the private sector makes perfect anticipations, $x = y$. This is not accidental, but captures the assumption that economic agents are sufficiently rational and knowledgeable to always correctly anticipate the future if there is no exogenous uncertainty. Moreover, the Ramsey solution coincides with the unconstrained optimum of F. The equality $x = y$, however, is not robust with regard to slight modifications of the model. Among others, it is no longer given in the Stackelberg game with F as a non-atomistic leader. That is, the assumption of perfect prediction is not compatible with the natural sequence of play where anticipations on a variable are made on the basis of the currently available information prior to the realization of this variable. Furthermore, the equality $x = y$ is not satisfied either at the Nash nor at the Stackelberg game with F as a leader if one modifies even slightly the players' objectives – as can be e.g. seen by assigning to F the (more) plausible payoff function $J^F = -\frac{1}{2}\left[(y-x)^2 + y^2 + U^2\right]$. Such a modification also destroys the coincidence between the Ramsey solution and the unconstrained optimum of F. On the other hand, removing the term y^2 from the original payoff-function J^F has no impact on the game outcome. Thus, one may argue that Kydland-Prescott's private sector does not care about the real economic outcome: its only preoccupation is with not being fooled upon, a strong and not innocuous assumption. By contrast, our private agents will be exclusively or at least primarily interested with the real consequences of the game.

One can wonder why, in a repeated game, fully rational and perfectly informed agents like L and F in the Kydland-Prescott model would coordinate on N rather than on some of those superior outcomes. In this paper, we will show that even less ideally rational agents can easily coordinate on superior solutions in a set-up that appears both natural and robust.

We now present the framework used in the remainder of the paper for analyzing the interplay between the government and the private sector. Unless otherwise specified, all assumptions underlying the Kydland-Prescott model presented in the previous section are valid for the modified model. Notice, however, that in most of the Kydland-Prescott related literature it is the private sector that attempts to control a government which on its own would generate a poor macroeconomic outcome. In this paper, it is the government that influences the private in order to improve the economic situation. We shall argue that this is indeed the *raison d'être* of economic governance.

3 The single-shot game

Assume the following sequence of play in the single-shot game between L and F :

1. The government announces that it will chose some inflation rate y^a. The announcement is cheap talk, in the sense that it is not binding and does not enter as argument the payoff functions J^L and J^F.
2. After hearing this announcement, the private sector forms an anticipation x on the inflation rate y that will be effectively implemented.
3. Given this anticipation, the government chooses y.

In this so-called *reversed Stackelberg game*, L has not one, but two instruments at his disposition, y^a and y. The outcome of the game will depend on the way in which y^a influences F's anticipation x. We single out two important benchmarks:

- F fully disregards the announcement, $dx/dy^a = 0$. In that case, the announcement has no impact at all, neither on the actions nor on the payoffs. Thus, the natural outcome of the single-shot game is the Stackelberg equilibrium with F as a leader (SF)[3] :

$$x^{SF} = \frac{\theta\left(1-\theta^2\right)U^*}{1+\theta^4},\ y^{SF} = \frac{\theta U^*}{1+\theta^4},\ J^{L,SF} = -\frac{\left(1+\theta^2\right)U^{*2}}{2\left(1+\theta^4\right)^2},\ J^{F,SF} = -\frac{\theta^2 U^{*2}}{2+2\theta^4}. \tag{10}$$

- F believes (or: acts as if it believed) that L will realize y as announced, that is, that $y = y^a$ holds. In that case, $x = T^F\left(y^a\right)$. The best L can do in the context of a single-shot game is to maximize, with respect to y^a and y, (5) with x replaced by $T^F\left(y^a\right)$. The resulting *optimal cheating solution* (OC) is given by:

$$y^{a,OC} = x^{OC} = -\frac{U^*}{\theta},\ y^{OC} = 0,\ J^{L,OC} = 0,\ J^{F,OC} = -\frac{U^{*2}}{2\theta}. \tag{11}$$

The OC solution was introduced by Hämäläinen (1981). It was derived and analyzed in the case of linear-quadratic dynamic games in Vallée, Deissenberg, and Basar (1999), Vallée and Deissenberg (1998), and other papers by the same authors. The term "cheating" refers to the fact that under this solution generically one has $y \neq y^a$, that is, the announcement will not be respected.

[3] This definition of SF assumes that the private agents acts as a single, non-atomistic agent and accordingly recognize the impact of their aggregate choice on the government's behavior. Without prejudice for the main message of this paper, one could make the more usual assumption of private agents that suppose that their choices do not affect the government's decisions. In that case, SF coincides with N. If anything, by eliminating an externality from the argumentation, the alternative retained insures more powerful conclusions.

The optimal cheating solution OC is extremely attractive for L: there is no feasible pair (x, y) under which L fares better than under OC. However, OC is not a reasonable candidate for the coordination in a repeated interaction between L and F, since playing x^{SF} guarantees F the payoff $J^{F,SF} > J^{F,OC}$. Natural choices are strategies where both L and F have higher payoffs than under SF[4]. We are going to show that there are other cheating strategies with this property. Before characterizing these pareto-improving cheating strategies, however, a look in the geometry of the one-shot game may be useful.

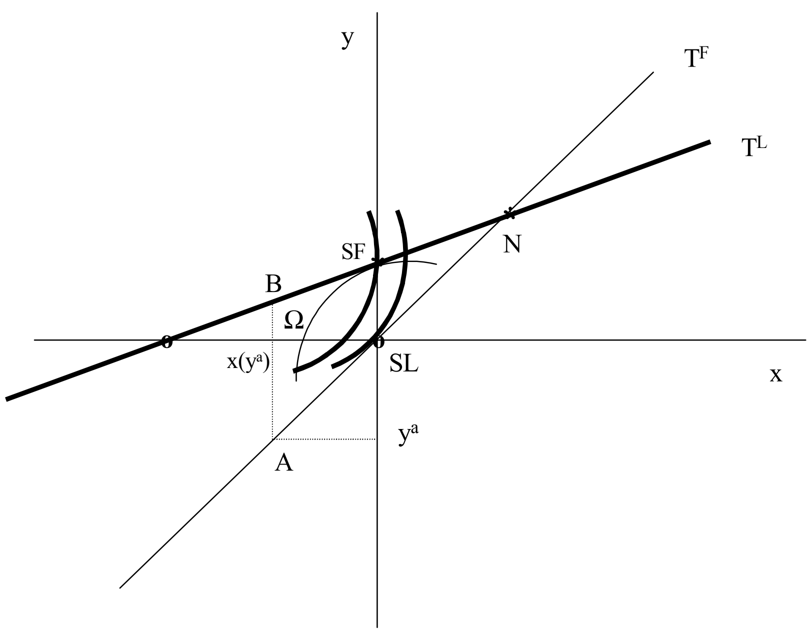

Figure 1: The geometry of the Kydland-Prescott model, $\theta = 1$.

In Figure 1, fat lines refer to L and thin ones to F. The straight lines represent

[4] Since y is now always realized after x is chosen, SL plays no longer any role in the argumentation.

the optimal reaction functions, the curved lines indifference curves (that is, level curves of their respective payoff function). The maximum possible payoffs of the two players are represented by large dots – moving away from these dots in any direction decreases payoffs. The Nash equilibrium N lies at the intersection of the reaction functions. The Stackelberg equilibrium SL (SF) is situated at the tangency point of T^F (T^L) with an indifference curve of F (L). That is, the Stackelberg solution allows the leader (L or F) to attain his best point on the reaction curve of his opponent. By contrast, the optimal cheating solution OC allows L to attain his best point on his own reaction function – that is for the class of problems we are considering here, his unconstrained optimum. The cheating mechanism allowing this outcome is apparent on the figure. An announcement y^a by L incites F to reply with $x = T^F(y^a)$ – see point A on F's reaction curve. Given x, L replies with $y = T^L(x)$ – see point B on L's own reaction curve. Since L's reaction curve goes by construction through L's unconstrained maximum, L can insure that the final outcome B coincides with this maximum trough a proper choice of y^a.

Now, neither one of the two equilibria N and SF is pareto-efficient. Indeed, there exists feasible points (x, y) that are both (a) efficient; and (b) *pareto-improving* in the sense that they pareto-dominate SF. Those are the points situated on that part Γ of the contract curve that lies within the convex lens Ω delimited by L's and F's indifference curves through SF. Any point within Γ is a candidate for the outcome of coordination over time. To characterize Γ, define:

$$J^{\alpha} \triangleq \alpha J^{L}(x,y) + (1-\alpha) J^{F}(x,y), \ x = T^{F}(y^{a}), \tag{12}$$

and consider the set O given by:

$$O \triangleq \left\{ (y^{a,\alpha}, x^{\alpha}, y^{\alpha}) : (y^{a,\alpha}, y^{\alpha}) \triangleq \arg\max_{y^{a},y} J^{\alpha}, \alpha \in [0,1], \ x^{\alpha} \triangleq T^{F}(y^{a,\alpha}) \right\}. \tag{13}$$

That is, O is the set of the optimal cheating solutions OC^{α} that would be generated if L based its decisions not on its original payoff function J^{L}, but on a convex combination J^{α} of J^{L} and of its opponent's payoff function J^{F}. As α decreases, this injects an increasing dose of altruism in L's actions. For $\alpha = 1$, OC^{α} coincides with the optimal cheating solution (11). For $\alpha = 0$, it coincides with the unconstrained optimum for F, that is here, with the origin. Since in the model considered $T^{F}(y^{a}) = y^{a}$, the set O can also be defined by:

$$O \triangleq \left\{ (x^{\alpha}, y^{\alpha}) : (x^{\alpha}, y^{\alpha}) \triangleq \arg\max_{x,y} J^{\alpha}, \alpha \in [0,1] \right\}. \tag{14}$$

In other words, O is the part of the contract curve located between the OC solution and the unconstrained maximum for F. A somewhat more involved argument can be used to show that this last result does not depend upon the fact that $T^{F}(y^{a}) = y^{a}$, but remains valid for arbitrary payoff functions J^{F} and J^{L} as long as they are strictly concave in the (scalar) actions.

It is straightforward to prove that:

There exists a non-empty interval $[\alpha_1, \alpha_2] \subset [0, 1]$ *such that* (x^α, y^α) *pareto-dominates* $\left(x^{SL}, y^{SL}\right)$ *iff* $\alpha \in [\alpha_1, \alpha_2]$, *with strict dominance whenever* $\alpha_1 < \alpha < \alpha_2$.

The set Γ of pareto-efficient and -improving cheating solutions is defined by $\Gamma \triangleq \{(y^{a,\alpha}, x^\alpha, y^\alpha), \alpha \in [\alpha_1, \alpha_2]\}$. As one would expect intuitively, α_1 is the unique value of α for which L is indifferent between playing SL or OC^α, that is for which $J^{L,\alpha} \triangleq J^L(x^\alpha, y^\alpha) = J^{L,SL}$. Similarly, α_2 is the unique value of α for which F is indifferent between SL and OC^α. The values of α_1 and α_2 do not depend upon U^* and are given by:

$$\alpha_1 = \max\left[0, \frac{1+\theta^4-\sqrt{1+\theta^2}}{1+\theta^4+(\theta^2-1)\sqrt{1+\theta^2}}\right], \quad \alpha_2 = \frac{1}{1-\theta^2+\sqrt{1+\theta^4}} \tag{15}$$

Notice that actual cheating is crucial in reaching a pareto-improving outcome. The outcome SF^α obtained by imposing that the announcement be respected (that is, by imposing $y^{a,\alpha} = y^a$) is the best point on T^F for the fictitious player $J^{L,\alpha}$. Clearly, for $\alpha < 1$, F has a higher payoff under SF^α than under SF. But SF^α does not pareto-dominate SF for $\alpha \neq 0$. Thus, L has no incentive to respect any benevolent ($\alpha < 1$) announcement, or to make such an announcement knowing that it will have to respect it. On the other hand, an incentive do exist if there is the possibility of subsequent cheating. Moreover, the private sector also have an incentive to act as if it believed the announcement, even if it knows that it will be cheated upon – as long as the cheating leads to a point in Γ.

4 The repeated game with a single agent

From the perspective of a single play of the game, the OC^{α} solutions suffer from not being equilibria. Once F has played x^{α} in response to an announcement $y^{a,\alpha}$, L can fare better by playing $T^{L}(x^{\alpha})$ than $T^{L,\alpha}(x^{\alpha})$, where $T^{L,\alpha}$ is the optimal reaction function associated with $J^{L,\alpha}$. Playing $T^{L}(x^{\alpha})$, however, leads to a very bad outcome for F. Being gullible is very costly. Moreover, in a repeated perspective, the OC^{α} solutions are associated with an apparent irrationality of the private sector: F repeatedly believes, or acts as if it believed, the announcement, although it is regularly cheated upon.

We shall argue that this does not necessarily disqualify the OC^{α} solutions as potential outcomes. For the ease of argumentation, we first consider the benchmark case of a single private agent F that knows exactly the objective functions J^{L} and J^{F} and the structure of the economy defined by the Phillips-curve (4). In particular, F knows x^{FS} and $J^{F,SF}$. However, F is not sure of the relationship between the announced and the realized inflation, y^{a} and y. Will L indeed chose $y^{a}=y$? Or, at least, will the realized y be such that it is rational for F to choose $x=T^{F}\left(y^{a}\right)$? In other words, the only incertitude lies in the fact that F does not know for sure the type of government it is facing.

4.1 The behavior of the private sector

Notice that, after observing an announcement y^{a}, F has fundamentally two alternatives, denoted OCA (for: Optimal Cheating Accommodate) and PAL (for: Play as A Leader):

1. OCA : It can believe (or act as if it believed) the announcement, and thus choose $x = x^{OCA} = T^F(y^a)$.

2. PAL : It can choose any other value of x. In particular, it can choose his best action in the one-shot game, x^{FS}. We assume in the following that this is the case.

The PAL alternative so defined essentially mimics dynamics that leads to the inferior Nash equilibrium in Sargent (1999). Other answers of F are clearly possible. At a later place, we will give some additional arguments to justify the fact that we are not considering them in this paper.

Consider now the following repeated game. At each repetition:

1. L chooses an $\alpha \in [0,1]$ and announces $y^{a,\alpha}$;

2. F plays OCA with probability $\pi \in (0,1)$, and PAL with probability $1-\pi$.

3. L observes the action x of F. (i) If $x = x^{\alpha}$, L plays y^{α}. (ii) If not, it plays its best answer $T^L(x)$.

4. F revises its probability of playing OCA or PAL at the next repetition according to whether or not its realized payoff is higher or lower than J^{SF}.

Notice that, under the very simplifying full information assumption made in this paper, the payoffs under 3. (i) will depend uniquely of the value of α chosen by L. The payoffs under 3. (ii) are always equal to $J^{L,SF}$ and $J^{F,SF}$.

Specifically, we assume that π is updated according to:

$$\pi_{t+1} = \phi(\pi_t, \rho_t, \delta_t), \quad t = 0, 1, 2, ..., \quad \pi_{-2} = \pi_{-1} \text{ given}, \tag{16}$$

$$\rho_t \triangleq \left(\max_{\tau<t} J_\tau^F\right) - J_t^F \tag{17}$$

whith J_t^F the payoff realized by F at the t-th repetition, and whith $\delta_t = 1$ if OCA is selected at repetition t, and $\delta_t = 0$ otherwise. The function ϕ is supposed to satisfy:

i. $\phi(\pi, 0, \cdot) = \pi$;

ii. $\phi(\pi, \rho, 1) \lesseqgtr \pi \Leftrightarrow \rho \gtreqless 0$ and $\phi(\pi, \rho, 0) \gtreqless \pi \Leftrightarrow \rho \gtreqless 0$;

iii. $\phi : (0, 1) \to (0, 1)$.

Equation (16) defines a so-called *reinforcement rule*, see e.g. Fudenberg and Levine (1998), Börgers and Sarin (1997), Brenner (1999) for a discussion of reinforcement learning in the context of repeated games. Roughly speaking, this rule implies that a positive experience with OCA (PAL) at the current repetition increases (and that a negative experience decreases) the probability with which F will play OCA (PAL) in later repetitions. Here, making a positive experience means: obtaining a higher payoff at the current repetition than ever in the past. Thus, the private agent's changes in behavior depend exclusively on the economic outcome, that is, on the values taken by U and y. In contrast to the original Kydland-Prescott model and to most of the related literature, cheating, that is the occurrence of a discrepancy between y^a and y, does not affect the payoffs – see our corresponding remarks towards the end of Section 2. Of course, we do not exclude that cheating causes per se real or psychological costs to the private agents. This can easily be taken into account by introducing a corresponding cost term in F's payoff function. As long that this cost term is not overwhelming compared to the other real payoffs, the results we are going to present remain qualitatively valid.

The assumptions i. and ii. on ϕ are self explanatory: favorable outcomes lead the private sector to increase its confidence in the government, wether or not the latter actually cheats (and similarly for unfavorable or neutral outcomes). The assumption iii. ensures that L will never get locked into playing exclusively one of the two alternatives OCA or PAL. The assumption is crucial, since it is necessary to insure that the private sector always (a) can be incited to move away from SF to a better equilibrium; and (b) cannot be lastingly lured into an unfavorable solution. In that sense, giving the government credit for favorable outcomes but never fully trusting it is a very strong element of rationality.

In the forthcoming simulations, we give the reinforcement rule the specific form:

$$\pi_{t+1} - \pi_t = \begin{cases} -\pi_t \frac{\rho_t}{1+|\rho_t|} & \text{if} \quad \rho_t \geq 0 \\ -(1-\pi_t) \frac{\rho_t}{1+|\rho_t|} & \text{if} \quad \rho_t < 0 \end{cases} \qquad t = 1, 2, ..., \tag{18}$$

if OCA is played, and, if PAL is chosen:

$$\pi_{t+1} - \pi_t = \begin{cases} (1-\pi_t) \frac{\rho_t}{1+|\rho_t|} & \text{if} \quad \rho_t \geq 0 \\ \pi_t \frac{\rho_t}{1+|\rho_t|} & \text{if} \quad \rho_t < 0 \end{cases} \qquad t = 1, 2, ... \tag{19}$$

This completes the description of F's behavior as an agent with (a) perfect information on all elements of the problem save on the actual choice of y, (b) perfect memory, but (c) simple, reinforcement-based adaptative behavior. The assumptions (a) and (b) can easily be weakened to imperfect information and limited memory.

4.2 The behavior of the government

In contrast to F, the government is supposed to be a maximizing agent. Specifically, we assume that L maximizes at any repetition t, through its choice of α_t, α_{t+1}, ..., α_{t+h}, its expected cumulated payoff over a revolving horizon $t+h$, $h \geq 0$ a constant:

$$\sum_{\tau=t}^{t+h} E\left[J_\tau^L\right] \longrightarrow \max_{\{\alpha_\tau\}} \tag{20}$$

subject to (16).

For $h = 0$, that is, in the case of a myopic government, the problem is trivial. The perfectly myopic government is exclusively interested in maximizing its payoff in the current repetition. Since the payoff under PAL is a constant, and since the probability that F plays OCA does not depend on L's choice of α, L's best choice is the value of α that maximizes the OCA payoff, that is, $\alpha = 1$. A myopic government is not benevolent. As a consequence, π decreases each time it is updated. On the long run, the private sector tends to systematically disregard the government's announcements.

For $h = 1$, L's problem is to find an optimal compromise between (a) obtaining a good (expected) payoff a the current repetition t, which implies a high value of α_t; and (b) obtaining a good payoff at the next repetition $t+1$, assuming that α_{t+1} will be set equal to one. To obtain a good payoff under (b), π_{t+1} must be high, which implies a low value of α_t. The optimal α_t is given by:

$$\alpha^* = \arg \max_{\alpha \in (0,1]} H(\alpha) - \frac{1}{4} \frac{\left(J_{t-1}^{F^*} - K(\alpha)\right) U^{*2}}{1 + \left|J_{t-1}^{F^*} - K(\alpha)\right|} \times \begin{array}{ll} \pi_t & \text{if} \quad J_{t-1}^{F^*} - K(\alpha) \geq 0 \\ 1 - \pi_t & \text{otherwise} \end{array} \tag{21}$$

where $J_{t-1}^{F^*} \triangleq \max_{\tau < t} J_\tau^F$, $H(\alpha) = -\frac{1}{2}(1-\alpha)^2 U^{*2}$, and $K(\alpha) = -\frac{1}{2}\alpha^2 U^{*2}$. The corresponding expression for $h = 2$ is very messy and will not be reproduced here. For $h > 2$, the expressions for α^* are no longer practically amendable, analytically and numerically.

The dependency of α^* upon U^*, π, and $J_{t-1}^{F^*}$ in the cases $h = 1$ and $h = 2$ is not monotonic, as one may recognize from the Figures 2 and 3. In these Figures, we plot α^* on the vertical axis against π (Figure 2) and J^{F^*} (Figure 3) on the horizontal axis when $h = 1$ (circles) and $h = 2$ (squares). In Figure 2, J^{F^*} is given the value -7.5625, in Figure 3 we set $\pi = 0.5$. Due to computational constraints, the indicated values of α^* are somewhat coarse approximations of the true values. They are restricted to multiples of 0.05, which explains the flat sections and the occasional exact coincidence of the value of α^* for $h = 1$ and $h = 2$.

Two features of the results are worth emphasizing. First, with increasing

J^{F^*}, α^* first decreases, and then increases until it takes the value $\alpha^* = 1$. As long as the aspiration level J^{F^*} is low (that is, strongly negative), L does not have to choose α significantly smaller than α_2 (defined in (15)) to make OCA attractive and therefore to insure that π will increase over time. On the opposite, when the private sector has very high expectations (when J^{F^*} is close to 0), it will be disappointed under OCA (that is, ρ will be a relatively large positive number) independently of the value of α chosen by the government. It turns out that under these conditions the government is better off by shortsightedly maximizing its payoff under OCA than by trying to build up a reputation, that is, to increase π by choosing a low α. Since F is assumed to have infinite memory and to use the best payoff it obtained in the past as benchmark J^{F*}, independently of how far in the past this payoff was realized, one cannot exclude a lock-in where L plays $\alpha = 1$ in any future repetition. This problem disappears if one more realistically assumes that F has a finite memory, and/or revise downwards the benchmark J^{F*} if it is not attained or surpassed over several repetitions. Second, contrary to what one might had expected, a government that has a longer planning horizon ($h = 2$ instead of $h = 1$) will not necessarily choose a lower value of α. The complexity of the trade-offs involved precludes giving a clear-cut explanation of this last result.

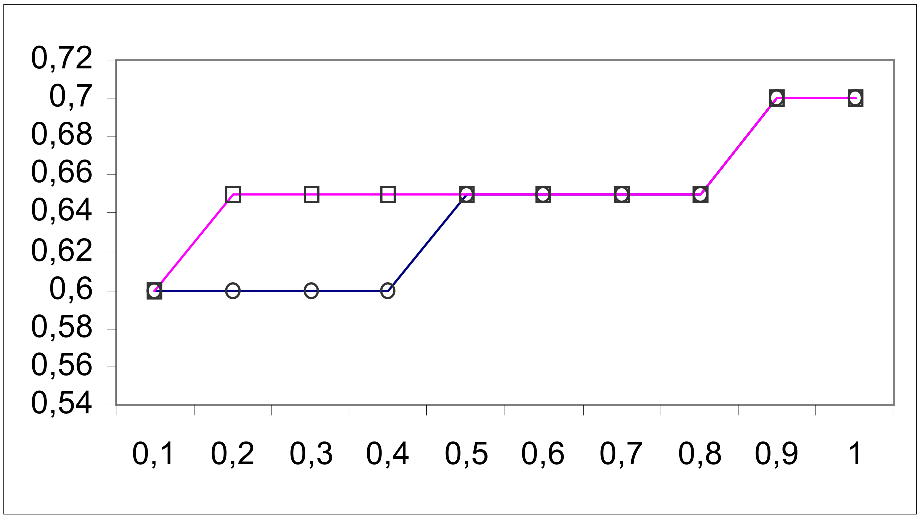

Figure 2: α^* as a function of π, $J^{F^*} = -7.5625$, $h = 1$ and $h = 2$

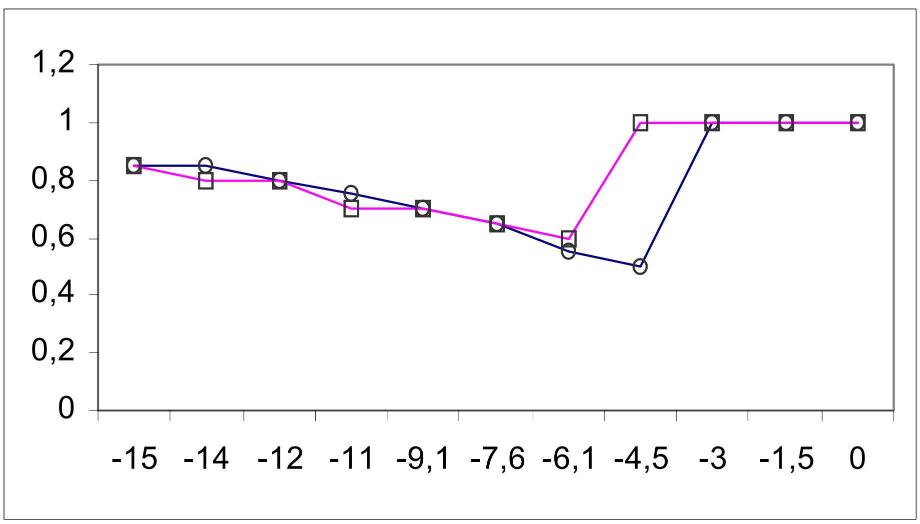

Figure 3: α^* as a function of J^{F^*}, $\pi = 0.5$, $h = 1$ and $h = 2$

4.3 Simulation results

In this Section, we present some illustrative simulation results when $h = 1$. The graphs pertain to the instantaneous averages for $t = 1, ..., 100$ over 50 Monte Carlo runs. The fat lines correspond to the averages themselves, the dotted ones to the averages +/- two standard deviations. The natural rate of unemployment U^* was set equal to 5.5, and θ equal to 1, as in Sargent (1999). Thus, $J^{L,N} = -30.25$ and $J^{F,N} = -15.125$; $J^{L,SF} = J^{F,SF} = -7.5625$; and $J_{\cdot}^{F*} \in [-15.125, 0]$.

The runs were initialized by conducting two repetitions assuming a constant value for π, $\pi_{-2} = \pi_{-1} = 0.5$. The best payoff for F obtained over these two repetitions was used as initial best past payoff, J_0^{F*}.

Figure 2 shows the evolution of the value of α chosen by L at each repetition. This value is slowly decreasing from about 0.64 to an almost constant value of about 0.587. It is worth noting that this value is considerably lower than the value $\alpha_2 = 0.88$ that makes F indifferent between playing OCA and PAL, see (15). Indeed, for α close to α_2, F will frequently play PAL, generating a very poor payoff for L. The latter is therefore motivated to significantly take into account the former's interest rather than making it indifferent between OCA and PAL.

Figure 3 shows the frequency with which F plays OCA. This frequency increases very rapidly, until stabilizing around a value only slightly lower than 1. That is, the private agents quickly learn to play almost always OCA. By construction, the corresponding outcome lies (approximately) on the contract curve.

Figures 4 and 5 show the evolution over time of J^L and J^F. Both payoffs increase over time, until practically stabilizing about some almost constant value. The payoff improvement over time is significantly more important for F than for L. From the onset, both players fare much better than under either N or SF. The payoff volatility decreases over time for both players. Statistical tests show that the volatility is similar for both players, contrary to what the graphs may suggest due to the use of different scales. The volatility of all the variables presented in the Figures is fairly small, so that the mean appears a good indicator of what might happen in a given historical situation.

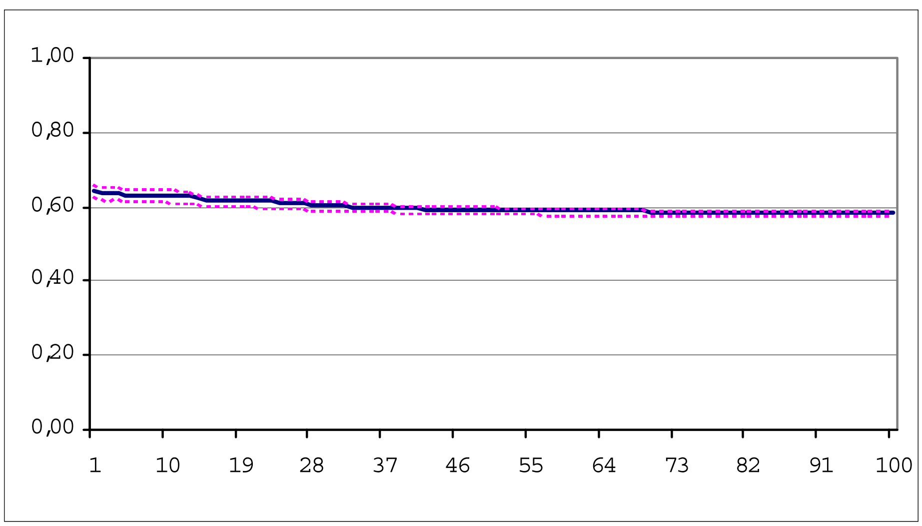

Figure 4: The time evolution of α

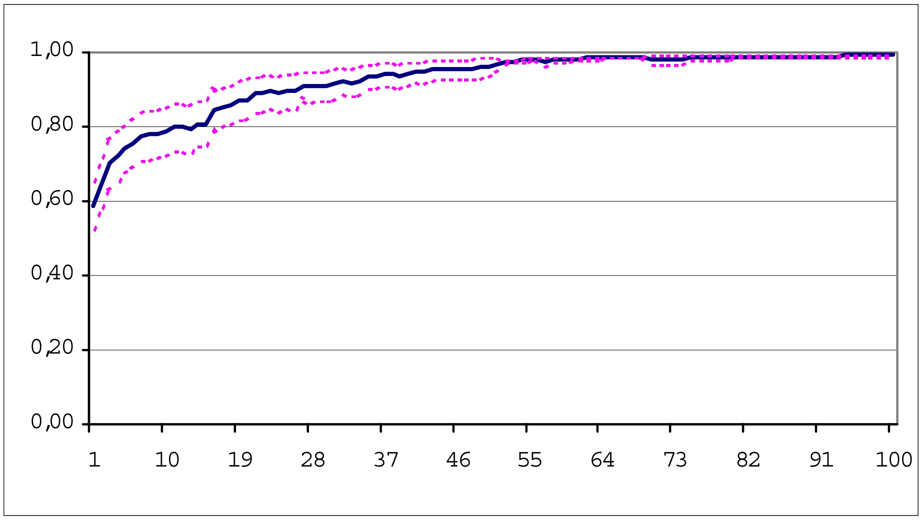

Figure 5: The time evolution of the probability of playing OCA

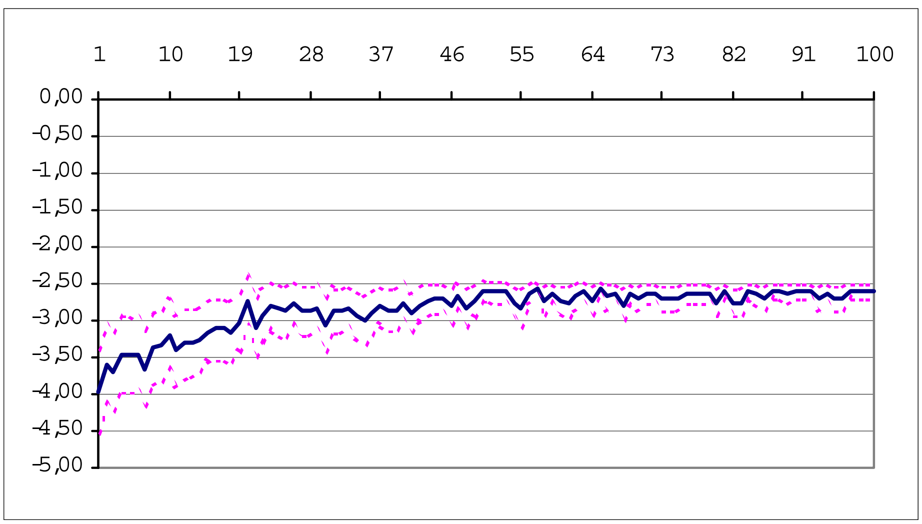

Figure 6: The time evolution of J_t^L

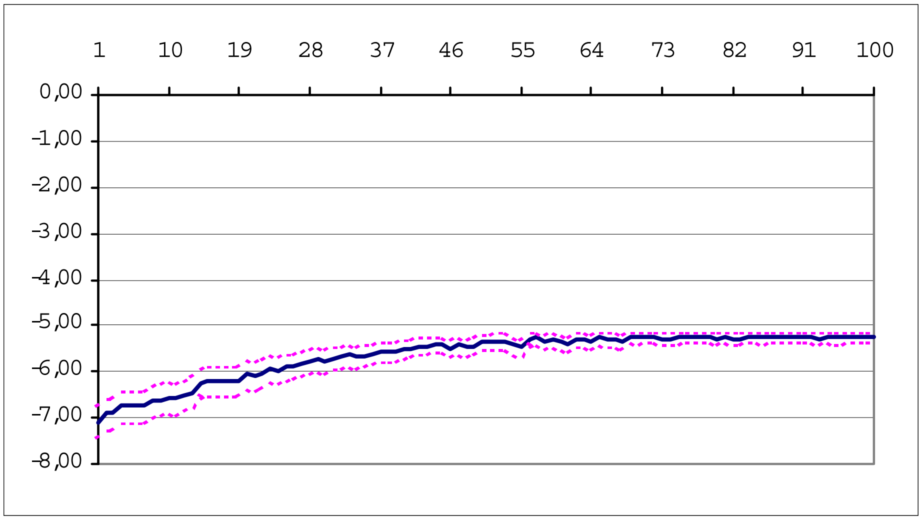

Figure 7: The time evolution of J_t^F

5 Heterogenous private agents

In this section, we extend the model to the case of a large number of private agents that independently decide to trust or to disregard the government's announcement, that is, to play OCA or PAL. This introduces heterogeneity among the private agents, since they build up different individual histories and different probabilities of playing OCA over time, even if they are initially identical in all respects. Moreover, and most importantly, when choosing y after having announced y^a, the government now simultaneously faces two types of agents, the $OCA-$ and the $PAL-$ players, namely: those private agents that believed the announcement and those that did not. This suggests looking for a cheating strategy that, at each repetition, pareto-improves the situation of the $OCA-$ players and of L, while penalizing the $PAL-$players. Such a strategy is derived below, in the spirit of the approach followed in the single agent case.

Taking into consideration multiple private agents also implies making specific assumptions about the information available on the π_is and ρ_is. To suppose that the different players know the individual probabilities π_i would clearly be too strong a conjecture. Making informal recourse to the law of large numbers, we make here the somewhat milder assumption that the government and the private agents share a common knowledge of the fraction π of players that will play OCA at the current repetition. This information might, for example, be provided by opinion polls and similar sources before the private agents actually take their decisions. Likewise, using the fact that we are going to consider only symmetric solutions where all $PAL-$ respectively all $OCA-$players make the same expectation, we assume that the payoffs previously obtained by both player categories are common knowledge. Thus, it is natural the re-define the variable ρ that drives the π-updating process, see (17), (18), (19), as:

$$\rho_t = J_t^{PAL} - J_t^{OCA}, \tag{22}$$

where J_t^{PAL} is the payoff obtained by each PAL-player, and J_t^{OCA} is the payoff obtained by each OCA-player in period t. That is, the updating of π is driven by the relative success of the PAL *vs.* the OCA strategy in the current period, rather than by the individual success histories. The definition of ρ can be changed to, say, $\rho_t = \max_{\tau,\theta \leq t}(J_\tau^{PAL} - J_\theta^{OCA})$, without major qualitative impact on the results.[5]

[5]Using the second definition opens the door to the lock-in effect previously mentioned for

The assumption of a known π, together with the other hypotheses previously made on the information structure, imply that there is no incertitude at the aggregate level within the single-shot game we are going to define. Accordingly, L will not switch between using T^L or $T^{L,\alpha}$ depending on the realized outcome. Nonetheless, the determinism at the aggregate level should not occult the fact that the updating behavior of the individual private agents is stochastic. There is full mixing in the sense that at each repetition any agent that previously played PAL may play OCA, and vice versa.

5.1 The structure of the game

With the exception of above modifications, we use as basic ingredients the very elements underlying the single agent case, or straightforward variants thereof. There is a government L and a continuum of mass 1 of private agents F^i. Agent's F^i payoff is given by (2), and L's payoff by (5). The sequence of play is:

1. The government makes a non-binding announcement $y^{a,\alpha}$.
2. After hearing $y^{a,\alpha}$, the agents F^i play, independently of each other, either OCA with probability π_i or PAL with probability $1-\pi_i$.
3. The government implements an action y^{α}.

The announcement $y^{a,\alpha}$ and the action y^{α} are determined simultaneously as the solution of an OC^{α} problem. As already mentioned, the main differ-

the single-agent case: convergence to the pareto-superior equilibrium will be impaired if a sizable fraction of the private agents have previously experienced very high payoffs. This lock-in effect cannot arise if one uses (22).

ence with the single agent case is that reaction of the private agents following the announcement $y^{a,\alpha}$ is (on the aggregate) deterministic, and that L always implement y^{α}.

Before defining the OCA and PAL options, and the OC^{α} problem, notice that, because: (a) the fraction π of private agents that play OCA (and thus, the fraction $1-\pi$ of agents that play PAL) is known at the onset of any repetition; and (b) all $OCA-$players choose the same value x^{a}, and all $PAL-$players the same value x^{SF}, the average inflation is:

$$x = \pi x^{a} + (1-\pi)\, x^{SF}, \quad \pi \text{ a known constant.} \tag{23}$$

As in the single agent case, we assume that the OCA choice x^{a} is given by the best response of an individual private agent to an announcement y^{a} believed to be true:

$$x^{a} \triangleq T^{F_i}\left(y^{a}\right) = y^{a}. \tag{24}$$

Since J^{F_i} depends only upon y and upon F^{i}'s own expectation x_i, any private agent can compute x^{a} without using any information on the actions of the others and in particular, without knowing either π or x.

On the other hand, the value of x_i under the PAL alternative is assumed to be given by:

$$x^{SF} \triangleq \arg\max_{x_i} J^{F}\left(x_i, T^{L}\left(x\left(x^{SF}\right)\right)\right). \tag{25}$$

with T^L given by (6).That is, x^{SF} is defined as the symmetric Stackelberg leadership solution for private agents that:

1. take for granted that the government will act rationally in the sense of replying to a given value of x with its best response $y = T^L(x)$;

2. rightly expect that π private agents will nonetheless believe the government's announcement and set $x^a = y^a$; and

3. recognize that their own individual choices have no influence on what is feasible or desirable for the government (that is, recognize that they are infinitesimally small). In other words, the $PAL-$players are perfectly rational and well-informed, except possibly with regard to point (a).

The thus defined behavior of the $PAL-$players has an interesting implication. When π is high, the value of x is almost entirely determined by the choice of the $OCA-$players, x^a. The $PAL-$players have no interest in choosing a x^{SF} much different from x^a. In other words, in a population where most agents trust the government, the optimal behavior for the agents that do not is nonetheless to act as if they did. Herding with the crowd is rational.

In a similar vein as in the single agent case, the government's OC^{α} problem is to maximize the artificial pay-off function (26) simultaneously with respect to $y^{a,\alpha}$ and y^{α} taking into account (23), (24), and (25):

$$J^{\alpha} \triangleq \alpha J^L(x, y^{\alpha}) + (1-\alpha) J^{F_i}(y^{a,\alpha}, y^{\alpha}) - (1-\pi)(y^{a,\alpha} - y^{\alpha})^2, \qquad (26)$$

where $\alpha \in [0,1]$. The number π of private agents playing OCA does not appear in this definition because of symmetry. The term $(1-\pi)(y^{a,\alpha}-y^{\alpha})^2$ is introduced to mitigate L's tendency, when π is small, to use the unconditional gullibility of the OCA-players (reflected in the fact that π is independent from y^a) to obtain a favorable value of x by making economically meaningless, extreme announcements $y^{a,\alpha}$.

The function J^{α} defined by (26) shares, with some restriction, the "contract curve generating" properties of its counterpart (12). Consider the set $\{(y^{a,\alpha}, x^a, y^{\alpha})\}$ of optimal cheating solutions OC^{α} generated by (26) for $\alpha \in [0,1]$. This set coincides with the set of all solutions that are pareto-efficient in terms of the payoffs received by a fictitious player with objective function $L = J^L(x, y^{\alpha}) - \delta(1-\pi)(y^{a,\alpha}-y^{\alpha})^2$ and by those F^is that played OCA, conditional on the PAL choice by the other players. However, for a given α, pareto-improvements are possible if the PAL players change of strategy.

Most importantly, numerical evidence shows that:

There exists, for any $\pi \in (0,1)$, a non-empty interval $[\alpha_1, \alpha_2] \subset [0,1]$ such that the solution where (a) π private agents play OCA; (b) $(1-\pi)$ private agents play PAL; and (c) L plays OC^{α}, has the following properties: (i) for L and the OCA-players, it (strictly) pareto-dominates the solution where all private agents play PAL and L replies with its best response $T^L(x)$; and (ii) the $PAL-$players fare no better (worst) than the $OCA-$players. Strict dominance (inequalities) are given whenever $\alpha_1 < \alpha < \alpha_2$.

In other words, believing misleading governmental announcements may pareto-improve the situation of the government and of the trustful agents, and insure that the later are better off than the $PAL-$players. Since the aggregate outcome is deterministic, L can always choose $\alpha \in [\alpha_1, \alpha_2]$ at the beginning of each repetition in order for (i) and (ii) to hold.

The white zone in Figure (8) gives the set of (α, π) combinations for which (i) and (ii) hold, for $\delta = 1$ and for different values of θ. This set does not depend upon U^*. The property (ii) is satisfied everywhere. The lower black zone corresponds to points at which L's payoff is lower than under the pure PAL solution. If nobody plays OCA, cheating cannot improve L's payoff. As the number of OCA players increases, more and more benevolent (low α) cheating suffices for such an improvement. Put somewhat crudely, for any value of α, it pays for the government to have many private agents that trust it. The lower black zone shifts to the left as θ increases. This reflects the fact that, when θ is large, even a slight discrepancy between realized (y) and expected (x) inflation suffices to insure an important increase in L's payoff.

In the upper black zone, the $OCA-$players realize a lower payoff than the $PAL-$players. One recognizes that hard cheating (high values of α) are less detrimental to the $OCA-$players when they are few. The upper black zone, too, shifts to the left as θ increases, for the reason already mentioned.

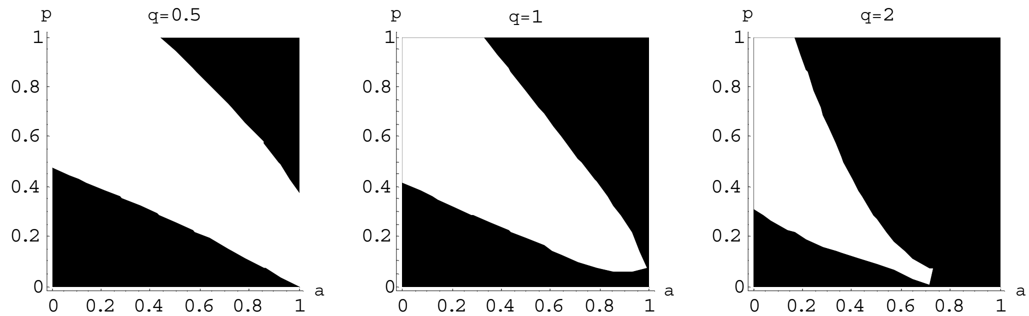

Figure 8: Sets of (α, π) combinations for which (i) and (ii) hold

Contrary to the single agent case, actual cheating is not necessary for a solution that is pareto-improving for the government and the private agents that trust it. However, reducing the amount of cheating (as expressed by some norm of $(y^a - y^{a,\alpha})$), decreases the payoff of the government and the $OCA-$players. In particular,the introduction of the term $(1-\pi)\left(y^{a,\alpha} - y^{\alpha}\right)^2$ in (26) leads to worst results than might otherwise be obtained.

The dynamics linking successive repetitions are very much as in the single agent case. The government maximizes at any repetition t, through its choice of α_t, α_{t+1}, ..., α_{t+h}, its expected cumulated payoff over a revolving horizon $t+h$, see (). However, L does not know the individual π_is and ρ_is. We therefore assume that it bases its optimization calculus on predictions of the future π values obtained with the help of an aggregate reinforcement rule. This rule, obtained by using the mean values π and ρ in (18), is defined by:

$$\pi_{t+1}^{OCA} = \pi_t - \frac{\rho_t}{1+|\rho_t|} \times \begin{cases} \pi_t & \text{if } \rho_t \geq 0 \\ 1-\pi_t & \text{if } \rho_t < 0 \end{cases}, \tag{27}$$

$$\pi_{t+1}^{PAL} = \pi_t + \frac{\rho_t^*}{1+|\rho_t^*|} \times \begin{cases} 1-\pi_t & \text{if } \rho_t^* \geq 0 \\ \pi_t & \text{if } \rho_t^* < 0 \end{cases}, \tag{28}$$

$$\pi_{t+1} = \pi_t \pi_{t+1}^{OCA} + (1-\pi_t)\,\pi_{t+1}^{PAL}, \tag{29}$$

where ρ_t is given by (22) and $\rho_t^* = -\rho_t$. Doing so, it makes a prevision error on π_{+1}. Stochastic simulations show that this error is small, and has almost no impact on the time behavior of the solution.

5.2 Simulation results

Again, the equations defining the optimal solution and the dynamics are sufficiently messy to justify the recourse to simulations. The value of the parameters are as before: $h = 1, \theta = 1$, and $U^* = 5.5$. Thus, the payoffs under the pure $PAL-$solution ($\pi = 0$) are $J^{L,PAL} = -30.25$ and $J^{F,N} = -15.125$. The results presented are averages over a population of 1,000 private agents and 50 Monte Carlo runs. The initial value of π is 0.1. That is, we consider a population that initially almost completely distrusts the government. The time horizon is of 20 periods. The simulation was conducted using a grid search algorithm for the needed optimizations, thus coarsifying somewhat the results. No standard deviations over the Monte Carlo runs are indicated, since their value is in the order of 10^{-16}, except in the case of J^L, where it takes the value of about 0.12 in the last 10 periods.

Figure 9 and 10 show that the optimal value of α very rapidly converges to a constant value of about 0.39, while the fraction of $OCA-$players rapidly stabilizes around 0.92%. Figure 11 illustrate the corresponding drastic improvement in payoff for L. Within a few period, this payoff increases from -32.73 to -7.54, clearly higher than the corresponding pure PAL or than the Ramsey value.

The evolution of the payoffs of the $OCA-$ and $PAL-$players, finally, is reproduced in Figure 12. Both converge, from above respectively from below, to the very good value of -1.82 for the $OCA-$ and -1.84 for the $PAL-$players.

These results prove remarkably robust with respect to variations in all parameters and can be considered generic, with only one exception. For very large values of θ, the percentage of OCA players becomes small: in that case, the government does not need that a large percentage of the population believes the announcement in order to obtain high payoffs.

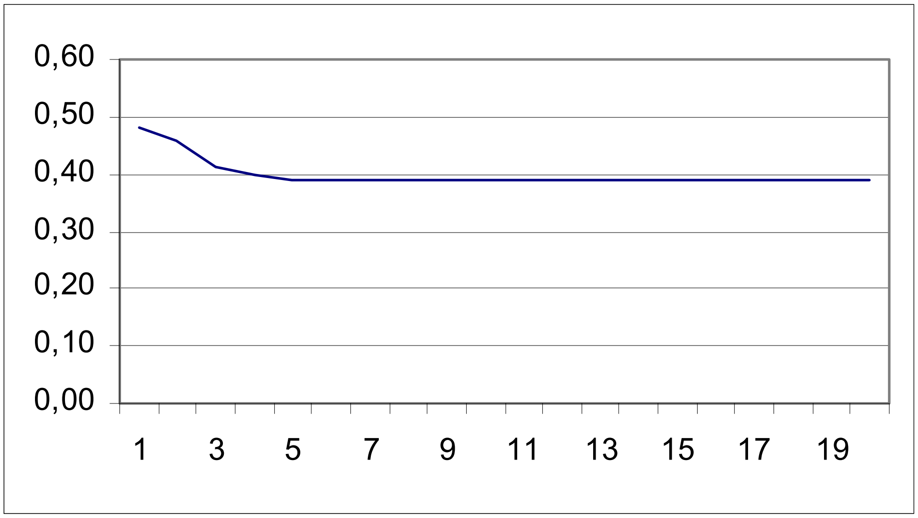

Figure 9: The time evolution of α

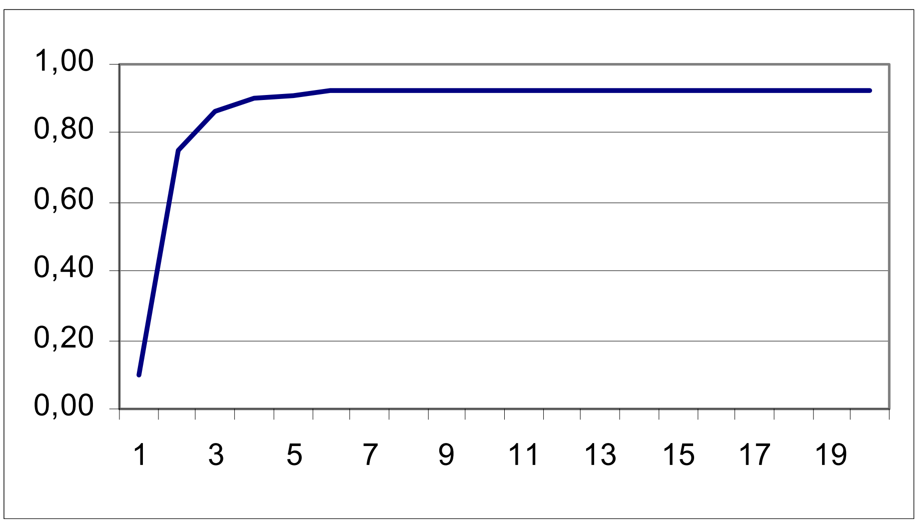

Figure 10: The time evolution of the fraction of private agents playing OCA

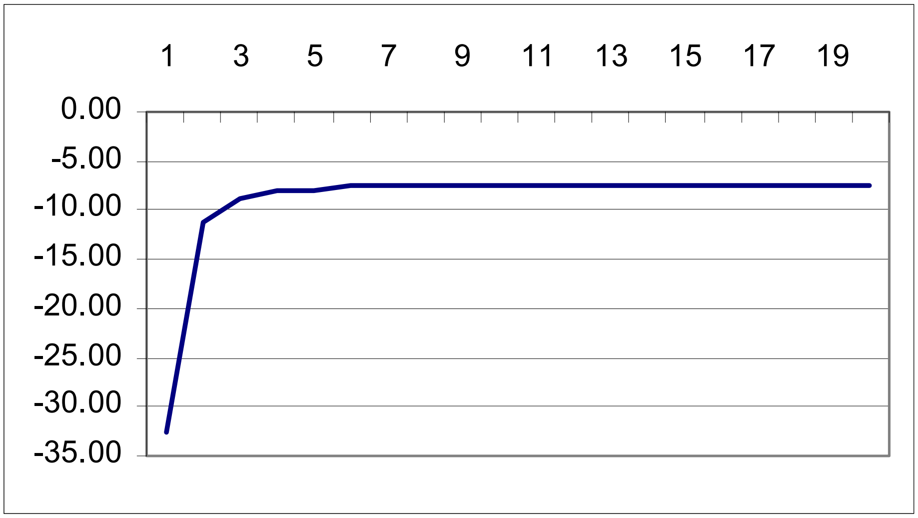

Figure 11: The time evolution of J_t^L

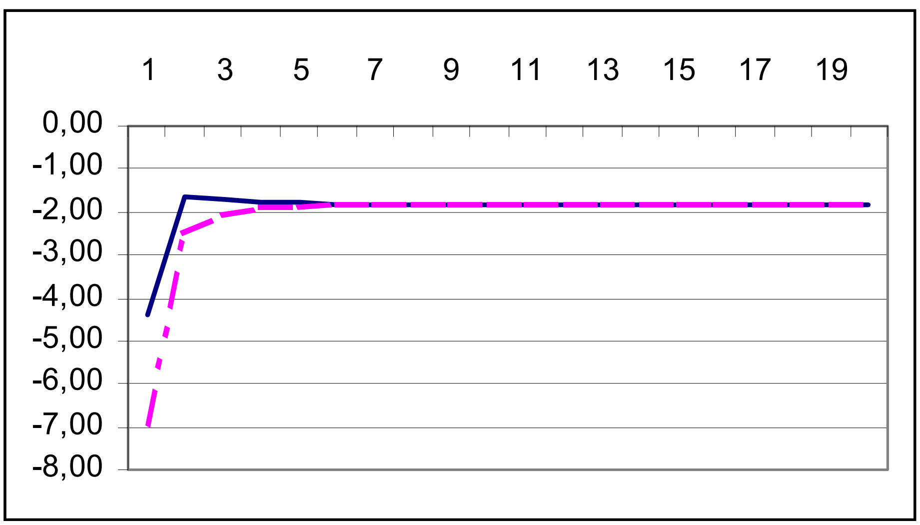

Figure 12: The time evolution of J_t^{OCA} (solid) and J_t^{PAL} (dashed).

6 Conclusions

The results presented here bring some strong messages. Compared to the Kydland-Prescott model, less sophisticated economic agents arrive to a better result, comparable to the one that could be attained by using for example complicated trigger strategies. Time inconsistency, understood as the discrepancy between a policy declaration and the subsequent action, is no longer a sufficient reason for rejecting governmental announcements. On the contrary, lies are necessary or at least useful for pareto-improvement. With their help, a government that is neither too myopic nor too stupid is able to better his and the private agents' payoffs – thus making it acceptable for the latter to be repeatedly cheated upon.

This outcome is paradoxical only in appearance. The rationality of the standard equilibria is based on the satisfaction of local first-order optimality conditions. Players that stick to these equilibria are fundamentally blind, in the sense that they are not considering the whole space of possibilities. Providing them for example with perfect foresight does not alleviate this fundamental limitation. By contrast, in this paper, we freed the government and the private agents from the tyranny of first-order conditions. Specifically, we gave the government both the ability to search a large subset of the state space and the wisdom to understand that (abstracting from the stable but unattractive Nash equilibrium) only a mutually favorable solution could be sustained over time. At the same time, we supplied the private sector with the freedom to assess solutions not in terms of their classical optimality properties, but with regard to the pay-offs they insure. Under these conditions, it comes as no surprise that the players could indeed coordinate on a superior outcome without recourse to complex information processing. Notice, incidentally, that our approach restores the government in its traditional role: to use a more global vision of the economy than the one of the micro-economic agents in order to help overall coordination and mitigate market failures.

As usual, numerous extensions of the model can be suggested. In particular, one might want to study the impact of uncertainty, both on the state of the economy following a choice of actions by the players, and on the assumed preferences and behavioral rules. Also, the consequences of other learning mechanisms than reinforcement should be thoroughly assessed. One extension, however, appears more fundamental, namely, to provide an explicit cognitive basis for the behavior of the $PAL-$players in the heterogenous agents case. Indeed, the results presented presuppose some kind of strong and rapid coordination of the $PAL-$players: no satisfactory macroscopic outcome is possible if these players act without very much taking into account what the others are doing. In this paper, we assumed one of the strongest possible types of coordination, namely, the timeless realization of a symmetric Stackelberg equilibrium. Most interesting would be to find constructive mechanisms for information exchange and coordination among private agents compatible with good macro-economic results. The field of investigation appears vast, and constitutes the reserved domain of future research.

7 References

Backus, D. and J. Driffill, 1985, Inflation and Reputation, American Economic Review, 75, 530-538.

Basar, T., 1989, Time Inconsistency and Robustness of Equilibria in Non-Cooperative Dynamic Games, in F. van der Ploeg and A.J. de Zeeuw, eds., Dynamic Policy Games in Economics, (North-Holland, Amsterdam).

Börgers, T. and R. Sarin, 1997, Learning Through Reinforcement and Replicator Dynamics, Journal of Economic Theory 77, 1-14.

Brenner, T., 1999, Modelling Learning in Economics (Edward Elgar, Cheltenham).

Fudenberg, D. and K. Levine, 1998, The Theory of Learning in Games, MIT Press (Cambridge, MA).

Hämäläinen, R., 1981, On the cheating problem in Stackelberg games, International Journal of Systems Science 12, 753-770.

Hall, S. and S. Henry, 1989, Macroeconomic Modelling (North Holland, Amsterdam).

Kydland, F. and E. Prescott, 1977, Rules Rather than Discretion: the Inconsistency of Optimal Plans, Journal of Political Economy. 85, 473-491.

McCallum, B., 1997, Crucial issues concerning central bank independence, Journal of Monetary Economics 39, 99-112.

Miller, M. and M. Salmon, 1982, Policy coordination and the time inconsistency of optimal policy in open economies, Economic Journal, 7, 124-137.

Personn, T. and G. Tabelinni, 1993, Designing Institutions for Monetary

Stability, Carnegie-Rochester Conference Series on Public Policy, 39, 53-84.

Rogoff, K., 1985, The Optimal Degree of Commitment to an Intermediate Monetary Target, Quarterly Journal of Economics, 100, 1169-1189.

Rogoff, K., 1987, Reputational Constraints on Monetary Policy, Carnegie-Rochester Conference Series on Public Policy, 26, 141-182.

Sargent, T.J., 1999, The Conquest of American Inflation, (Princeton University Press, Princeton, N.J.).

Stokey, N.L., 1989, Reputation and Time Inconsistency, American Economic Review, Papers and Proceedings, 79, 134-139.

Stokey, N.L., 1990, Credible Public Policy, Journal of Economic Dynamics and Control 15, 627-656.

Vallée, Th. and Ch. Deissenberg, 1998, Pareto efficient cheating in dynamic reversed Linear-Quadratic Stackelberg games: the open-loop linear-quadratic case, in S. Holly et S. Greenblatt, eds., Issues in Computational Economics, (Elsevier, New York).

Vallée, Th., Ch. Deissenberg, and T. Basar, 1999, Optimal open loop cheating in dynamic reversed Linear-Quadratic Stackelberg games, Annals of Operations Research 88, 247-266.