Eur. Phys. J. B **55**, 149–159 (2007)
DOI: 10.1140/epjb/e2007-00032-8

THE EUROPEAN PHYSICAL JOURNAL B

The emergence of coordination in public good games

W. Hichri[1] and A. Kirman[2,a]

[1] Gate & Université de Lyon 2, 93 Chemin des Mouilles, BP 167, 61131 Ecully Cedex, France
[2] IDEP-GREQAM, EHESS, IUF & Université d'Aix-Marseille III, 2 rue de la Charité, 13002 Marseille, France

Received 15 September 2006 / Received in final form 18 December 2006
Published online 2 February 2007 –

Abstract. In physical models it is well understood that the aggregate behaviour of a system is not in one to one correspondence with the behaviour of the average individual element of that system. Yet, in many economic models the behaviour of aggregates is thought of as corresponding to that of an individual. A typical example is that of public goods experiments. A systematic feature of such experiments is that, with repetition, people contribute less to public goods. A typical explanation is that people "learn to play Nash" or something approaching it. To justify such an explanation, an individual learning model is tested on average or aggregate data. In this paper we will examine this idea by analysing average and individual behaviour in a series of public goods experiments. We analyse data from a series of games of contributions to public goods and as is usual, we test a learning model on the average data. We then look at individual data, examine the changes that this produces and see if some general model such as the EWA (Expected Weighted Attraction) with varying parameters can account for individual behaviour. We find that once we disaggregate data such models have poor explanatory power. Groups do not learn as supposed, their behaviour differs markedly from one group to another, and the behaviour of the individuals who make up the groups also varies within groups. The decline in aggregate contributions cannot be explained by resorting to a uniform model of individual behaviour. However, the Nash equilibrium of such a game is a total payment for all the individuals and there is some convergence of the group in this respect. Yet the individual contributions do not converge. How the individuals "self-organsise" to coordinate, even in this limited way remains to be explained.

PACS. 05.65.+b Self-organized systems – 02.50.-r Probability theory, stochastic processes, and statistics

1 Introduction

Perhaps the most commonly cited feature of complex systems is that the behaviour of the aggregate is not the same as that of the component parts. In analysing the behaviour of a system of interacting particles in physics or a biological system the behaviour of the organism is not similar to that of the individuals, particles in the one case, or cells in the other. Yet, in economics we frequently build models based on hypotheses about individuals' behaviour and then test these on aggregate behaviour. The assumption is that the aggregate can reasonably be taken to behave like an individual. Thus, if we cannot reject our individual based model at the aggregate level then we conclude that it is a valid model of how people are behaving. Our purpose in this paper is to suggest that this is a misleading approach. To do this we examine data from a series of experiments on contributions to public goods and show that while the average behaviour seems to correspond to a reasonable individual model, neither the groups nor the individuals playing the game do so.

In such games individuals decide how much of their money to keep and how much to contribute to a public good which everybody will enjoy. If we look at a one-shot public good game there are two useful benchmarks. On the one hand there is that solution which maximises the total pay-off, the Collective Optimum (CO) and on the other there is the "Nash Equilibrium", (NE). In the latter, given what the other players contribute, each player chooses what is in his own best interest, i.e. his "best response" to the others actions. We should, of course, observe, that in all but the simplest games, both the CO and the NE of the one shot game are defined as a total contribution to be divided among the players. For this reason the literature has focused on the symmetric equilibria as a test since there are many asymmetric equilibria. Public goods games are useful since we have some

[a] *Present address*: GREQAM, Centre de la Charité, 2 rue de la Charité 13002, Marseille, France;
e-mail: `kirman@univmed.fr`

well-established stylised facts concerning them and these can be compared with the theoretical predictions. Indeed, there is a wealth of information from public goods experiments showing that, with respect to the Nash equilibrium of the one-shot game people initially over-contribute to public goods but that, with repetition, they contribute less. Although these experiments have too few rounds, in general, to make meaningful statements about the "convergence" of the total contributions, they do decline towards the Nash equilibrium. What the players in these experiments are faced with is a finite repeated game, which can be solved by backward induction. The solution to this problem in the standard linear model is one of a dominant strategy in which people contribute nothing from the outset. However, it is well established that this is not what actually happens (see e.g. Ref. [1]). Indeed since individuals do not play in this way, we cannot attribute the observed behaviour to that associated with an equilibrium of the finite repeated game. Furthermore players do not establish the cooperative, or socially optimal outcome (CO) and indeed, in general, move collectively away from such a solution. The way in which this is explained is that it takes some time for people to understand what is happening, and that they "learn to play Nash" or something approaching it. The purpose of this paper is to examine this idea by analysing average and individual behaviour in a series of public goods experiments.

There are two very different views of learning in games. Population learning suggests that the configuration of strategies in a population game will converge towards some limit, which may or may not be a solution of the one-shot game. This, it is argued, is because more successful strategies take over from those that perform less well. This simple evolutionary argument does not explain how the strategies are replaced. But this is the sort of idea invoked by Saijo and Yanaguchi [2], for example, who classified certain people, from their behaviour, as Nash and found that at the beginning of their public goods experiments 50% of players were "Nash" and, at the end 69% fell into this category. This makes it tempting to believe that the population was evolving towards Nash. But this raises an important question. Is it true that the people who switched had "learned to play Nash", and if so, how and why?

An alternative approach is to model the individual learning process and to see if observed behaviour, particularly that in experiments, corresponds to such a model. The usual approach is to assume that all individuals learn in the same way and then to test the learning model on the average observed data (see Ref. [3]). This is very common practice and often gives rise to rather convincing results. However, as Ho et al. [4], point out, the estimated parameters for the representative individual may not coincide with the average parameters of the population. Furthermore, this approach is fundamentally flawed. To assume that the average player behaves in a certain way is to give way to the same temptation as that offered by the "representative agent" in macroeconomics. It is not, in general logically consistent to attribute the characteristics of an individual to average behaviour, if for example, we reject the model, how do we know whether we are rejecting the model itself, or the hypothesis that all agents learn according to the same model?

One way out is to assume that individuals behave according to the same learning model but differ in their parameters. This is the approach adopted, for example, by Ho et al. [4]. Two basic classes of rules have been used. The first of these are the "reinforcement" models in which strategies are updated on the basis of their results in the past, (an approach based on the work of Bush and Mosteller, see e.g. Roth and Erev [3] and Mookerjhee and Sopher [5]). The second are the so-called "belief" models in which agents update their anticipation of their opponents' behaviour on the basis of their previous behaviour, fictitious play being a good example (see e.g. Fudenberg and Levine [6]) A more general model, (experience weighted attraction learning EWA), which incorporates both type of rule has been introduced by Camerer and Ho [7].

Another possibility is not to find a rule, which encompasses others as special cases, but to allow for different rules and simply to try, on the basis of observed behaviour, to assign agents to rules. This is the procedure followed by Cheung and Friedman [8], Stahl [9] and Broseta [10]. There are at least two problems with this sort of approach. Firstly, the rules specified are necessarily arbitrarily chosen, and secondly, the tests are not very powerful since, in such situations, the number of observations is, in general, not very large.

The last and most important point for this paper is that the Nash Equilibrium for the finitely repeated game is not unique. What is defined is the total sum that individuals should contribute at the equilibrium. However who should contribute what is not determined. If all agents use the same rule one might expect a symmetric result, but this is not what we observe. If each agent learns to contribute a certain amount and the total corresponds to the Nash Equilibrium (NE), then we have to explain how agents come to self organise in this way. More interestingly, if individuals contribute different amounts in different periods but the total still corresponds to the NE then this coordination has also to be explained. A cursory examination of our data reveals that individuals are not playing "mixed strategies" with fixed probabilities of contributing each of the possible sums. Thus the coordination mechanism is unexplained[1].

We proceed as follows in this paper. We analyse data from a series of games of contributions to public goods and firstly to see what happens, if we follow the standard ap-

[1] Efforts have been made to aid the coordination by introducing artificial labels which the players can choose between each round (see Page et al. [11]). In this situation the players sort themselves into groups by the amounts contributed. Sol et al. [12] allow players to sign binding agreements as to the contributing group they will form and find that the players do not coordinate on the Nash equilibrium of this game but paradoxically move in the direction of the collective optimum over time.

proach and first determine the trend of contributions and then test a learning model on the average data. We next look at group and then individual data, examine how this changes the results and then see if some general model such as the EWA with varying parameters can account for individual behaviour. Our case is rather favourable for this sort of test since, by telling agents how much was contributed to the public good in total, at each step, we allow them to know how much they would have obtained from foregone strategies. This avoids a fundamental problem raised by Vriend [13], as to how agents can update the weight they put on strategies which they have not played if they do not know how much these would have paid. We find that, nevertheless, behaviour differs across groups and individual behaviour is not easily categorized.

It is worth recalling that, in this type of experiment, individuals are divided into groups who play the game for a fixed number of periods. Thus the groups are unaffected by each other's behaviour. Yet the population, taken as a whole, seems to learn in a simple way. However, the separate groups do not learn as supposed, and their behaviour differs markedly from one group to another. Furthermore, behaviour of the individuals who make up the groups also varies within those groups.

The usual explanation for some of the discrepancies in strategies in the early rounds of public goods games, as Ledyard [1] explains, is that confusion and inexperience play a role. Indeed, this is one of the basic reasons why repetition has become standard practice in these experiments. Yet, as we will see, this would not be enough to explain some of the individual behaviour observed in our data.

We now specify the model used for the experiments and the rest of the paper will be structured as follows. Firstly we will explain the particular features of our model and their advantages. Then we will give the details of the experiments we ran. We will then describe the data from the experiments. In the following section we will perform some simple tests on the average data for the whole population, for the group averages and finally for the individuals. Having pointed out the differences between these, we proceed to an analysis of the performance of the EWA rule on individual data and compare this with our previous results. We then conclude.

2 The model

2.1 The basic public goods game

In the basic game of private contribution to a public good, each subject i, $(i = 1, ..., N)$ has to split an initial endowment E into two parts: the first part $(E - c_i)$ represents his private share and the other part c_i represents his contribution to the public good. The payoff of each share depends on and varies with the experimental design, but in most experiments is taken to be linear (Ref. [14]). The total payoff π_i of individual i, in that case is given by the following expression:

$$\pi_i = E - c_i + \theta \sum_{j=1}^{N} c_j.$$

This linear case gives rise to a corner solution. In fact, assuming that it is common knowledge that players are rational payoff maximisers, such a function gives a Nash equilibrium (NE) for the one shot game at zero and full contribution as social optimum. The dominant strategy for the finite repeated game, is to contribute zero at each step. Nevertheless, experimental studies show that there is generally over-contribution (30 to 70% of the initial endowments) in comparison to the NE.

Attempts to explain this difference between the theoretical and the experimental results are the main subject of the literature on private contribution to public goods. To do so, several pay-off functions with different parameters have been tested in various contexts to try to see the effect of their variation on subjects' contributions (for surveys, see Davis and Holt [15], Ledyard [1] and Keser [16]).

In the linear case, given that the NE is at zero, and giving that subjects could not contribute negative amounts to the public good, error can only be an over-contribution. To test the error hypothesis experimentally, Keser [17] performed a new experiment. She proposed a design in which the payoff function is quadratic and the equilibrium is a dominant strategy in the *interior of the strategy space*. With such a design, undercontribution becomes possible and error on average could be expected to be null. The results of Keser's experiment show that in each period, contributions are above the dominant solution, which leads to the rejection of this error hypothesis.

Another way to introduce an interior solution is to use a linear payoff for the private good and a concave function for the public one. Sefton and Steinberg [18] compare these two possible payoff structures. They call the first one "the Dominant strategy equilibrium treatment" and the second "the Nash equilibrium treatment". The results show that "average donations significantly exceed the predicted equilibrium under both treatments, falling roughly midway between the theoretical equilibrium and optimum... Donations are less variable under the dominant strategy equilibrium treatment than under the Nash equilibrium treatment".

2.2 Our model

The theoretical model and design used for the experiments we report in this paper concerns a public goods game with a "Nash equilibrium treatment". The individual payoff function is

$$\pi_i = E - c_i + \theta \left(\sum_{j=1}^{N} c_j \right)^{1/2}.$$

The Nash equilibrium and the social optimum corresponding to this payoff structure are not trivial solutions but in

Table 1. The NE and the CO values for the four treatments for one group.

Value of θ^*	Treatment	Endowment	Symmetric NE	CO
4	L	280	4	64
5.66	M	280	8	128
6.93	H	280	12	192
8.94	VH	280	20	280

* These are approximate values. The exact values are respectively: 4, 5.6568542, 6.9282032 and 8.9442719. We choose these values such that the CO corresponds respectively to 64, 128, 192 and 280.

Table 2. The symmetric NE and the CO values for the four treatments for one subject.

Value of θ	Treatment	Endowment	Symmetric NE	CO
4	L	70	1	16
5.66	M	70	2	32
6.93	H	70	3	48
8.94	VH*	70	5	70

the interior of the set of possible choices. The Nash equilibrium for individuals is not a dominant strategy for the finite repeated game. Indeed the solution for that game poses problems for a simple reason. There is a unique Nash equilibrium in the sense that for any Nash equilibrium the group contribution is the same. However, as we have said, that contribution can be obtained by several combinations of individual contributions. Since there are many Nash equilibria for the one-shot game, precisely what constitutes an equilibrium for the repeated game is unclear. For a group of N subjects, at the CO the total contribution is given by the following expression:

$$\overline{\overline{Y}} = \sum_{i=1}^{N} y_i = N^2 \cdot \left(\frac{\theta^2}{4}\right)$$

and at the NE is equal to:

$$Y^* = N \cdot y^* = \frac{\theta^2}{4}$$

where y^* is the symmetric individual Nash equilibrium.

With such a design, the Nash equilibrium and the social optimum vary with the value of θ. We shall see whether there is any difference between the evolution of individual and aggregate contributions under the different treatments.and how well learning models explain this evolution.

We gave θ four different values, which give four levels for the CO and the NE. The following tables summarize the four treatments (Low, Medium, High and Very High) the different levels of interior solutions for each group of four ($N = 4$) persons (Tab. 1) and for the individual subjects in each group (Tab. 2).

The set of possible group contributions in our model is very large. In fact, given that each one of the four individuals of a group is endowed with 70 tokens, each group can contribute an amount that varies between zero and 280.

We also examined situations in which individuals had more or less information about what the members of their group were doing or claimed they were going to do. It might be argued that knowing who had offered what or had given what would lead to more interactive behaviour and change the behaviour of the individuals. For this we added two other features and we again looked not only at the effect of variations in θ but also at the variability of individual behaviours. In the first variant, we replicated the same experiment as before with the introduction of promises before the decision period relative to the contribution to the public good. The four treatments with promises are called LP, MP, HP and VHP.

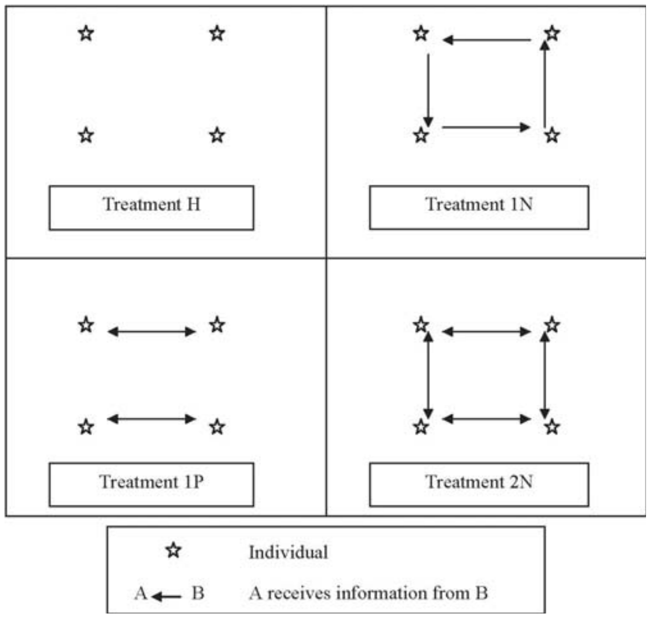

Fig. 1. Information about neighbours contributions with an H treatment Design.

In the last series of experiments, we introduced, as a parameter, the information about the contribution of some members of one group to the public good. Three treatments were compared to the treatment H. This treatment is used as a benchmark (no-information treatment) and is chosen because it allows us to keep the CO in the interior of the strategy space but, at a high level. The three new treatments with information are 1N, 1P and 2N and have the same theoretical benchmarks as treatment H (same CO and NE). In theory, the information which is provided after the end of the one-shot game should have no effect on the outcome.

While in treatment H individuals are informed only of the sum of contributions of their group, in treatment 1N each one of the four individuals in a group knows at the end of each period the contribution of his right neighbour, in addition to the sum of contributions of the group. In treatment 1P, each individual in a four persons group has a partner with whom he exchanges information about their own contributions. These Partners are the same for all the periods of the game. Finally, in the last treatment (2N), for each individual information concerns two Neighbours, the right one and the left one. The figure above (Fig. 1) depicts these four treatments.

In all treatments, and for every period, information given to players always concerns the same individuals.

The data from these three new treatments is also used to test the learning models at the population, group and individual level. This allows us to see whether the results in our reference framework are robust to different "institutional frameworks".

The following section presents the experimental results for the four treatments of both experiments with (treatments LP, MP, HP and VHP) and without (treatments L, M, H and VH) promises and those of the three new treatments with information (treatments 1N, 1P and 2N).

3 The experimental results

Our initial analysis will be at the aggregate level where we have for each treatment the average contribution of the six groups compared to the aggregate NE and to the aggregate CO. These results are reported in Figure 2 for the four treatments without promises, in Figure 3 for the four treatments with promises and in Figure 11 for treatments with information.

3.1 Basic results without promises

The first thing we observe when analysing the experimental results is the fact that the average group contribution (Y) decreases over time. In fact, the Very High treatment ($\theta = 8.94$), has 133.33 and 91.5 as values of the first and the last periods (see Fig. 2). In the case of the H and the M treatments, the average group contribution (Y) decreases during the 10 first periods and stays at a steady level during the rest of the periods of the game. In the last treatment L, this average group contribution starts at 67.67 and decreases steadily during the 25 periods of the game until finishing at 7.16. The decrease in contributions is however less evident in the VH treatment. In general, if we overlook the first five periods that could be assimilated to "*learning periods*", contributions are almost steady over the twenty last periods for the M, H and VH treatments.

Our results show that contributions vary with the CO level. There is overcontribution in comparison to the NE. As the CO level increases so does overcontribution. Nevertheless, average contributions as a proportion of the CO do not increase. Thus, computing these contributions in relative values by calculating an overcontribution index that takes into account the NE and the CO, shows that, except in the VH treatment, this ratio is constant.

In the L treatment, the average group contribution seems to decrease and to tend steadily to the NE value (which is 4). For this treatment, where the CO level is very low, subjects seem to tend to the theoretical predicted value for the one-shot game. The difference between the theoretical prediction and the experimental results is less evident in this framework as the game proceeds and indeed, such a close approximation to the NE is rarely observed in the experimental literature relative to public

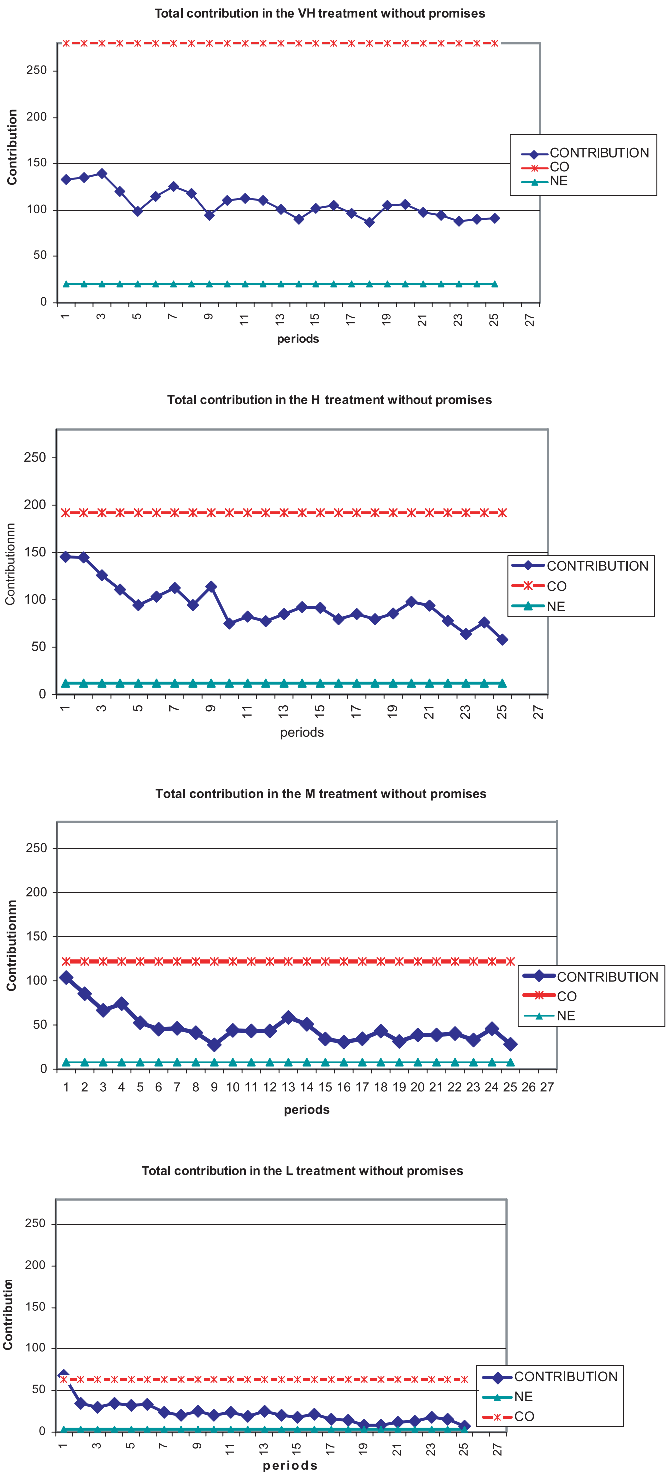

Fig. 2. Average total contribution in treatments VH, H, M and L without promises.

goods. It seems somewhat paradoxical that subjects contribute less and learn to play the NE value when the CO is low. For this is precisely the case in which the CO is easy to reach in the sense that it does not require a large contribution. For high levels of the CO, it is, in fact, risky for one subject to cooperate and to try to reach the social optimum by contributing a large amount to the public good. Taking such a risk can lead one subject to share his or her contribution with other subjects that choose not to contribute, and, in so doing, to lose, most of his or her private payoff. Thus the risk when faced with "free riding" behaviour is higher as the CO increases.

The other side of the coin is that the gains to be had from contributing more collectively are higher when the CO is higher. One natural idea is that individuals make generous contributions initially to induce others to do the same. This is, of course, not consistent with optimising non-cooperative behaviour but has been evoked in considerations of non-equilibrium behaviour.

3.2 Results with promises

Promises are introduced in the public goods game as a step preceding real contributions. In each group, and in each period, individuals are asked to announce their intentions as how much they will effectively contribute in the considered period. The sum of these intentions is revealed to the members of the group. Thus this information become common knowledge before the beginning of real game where individuals announce their effective contributions that will be considered when calculating their gains. Intentions or promises that are not binding are considered in game theory as "cheap talk", since the gains from different actions are not affected by their introduction. Also, the NE and the CO of the game are the same as in the game without promises.

As might be intuitively expected, and in concordance with findings in experimental literature, the introduction of promises does increase contributions to the public good, although this increase is not very marked. Consequently, the average group contributions are further from the Nash equilibrium than in treatments without promises.

What is interesting from our point of view is that the main difference between treatments with promises and those without promises is obviously the heterogeneity of groups and individuals behaviour when promises are allowed. In fact, the data permits to isolate, in treatments with promises, different strategies that do not exist in treatments without communication. Since the various levels of the treatment provide essentially similar results we show only one as an illustration.

3.3 Results with differing information

While the introduction of promises increases contributions to the public good, information about the contribution of the other members of one group seems to have no effect on the decision of contribution of individuals. In fact, as shown in the following figure (Fig. 4), where we represent average total contributions for the three treatments with information and for treatment H (already given in Fig. 2), the experimental results show that there is no difference between contributions with information and contributions without. The aggregate behaviour is very similar in the four treatment: overcontribution is evident during all the experiment and decreases over time, as contributions become closer to the NE.

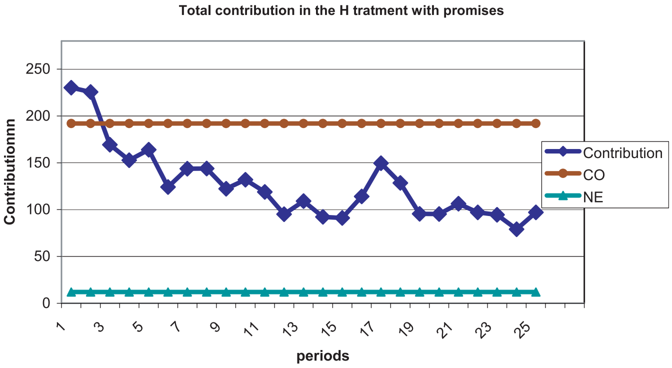

Fig. 3. Average total contribution in treatments HP with promises.

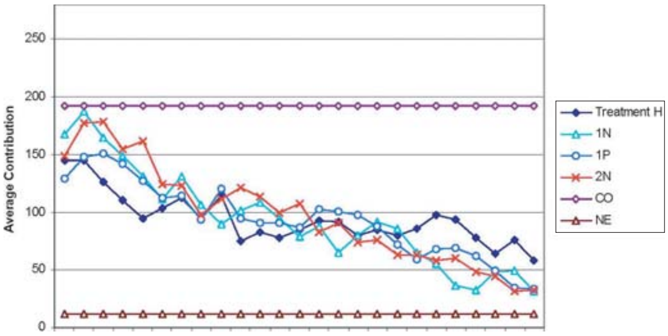

Fig. 4. Average total contribution in treatments H, 1N, 1P and 2N.

Table 3. Average total contribution in treatments H, 1N, 1P and 2N.

Treatment	First period	Average over 25 periods	Last period	Max.	Min.	St. Dev.
H	145.33	93.87	**58**	145.33	**58**	21.98
1N	**167.83**	94.12	31.83	**187.67**	31.83	42.97
1P	129.33	93.51	33.33	150.5	33.33	32.75
2N	149.17	**97.74**	32.83	178.17	31.83	**43.64**

Note that in all treatments, contributions start by increasing and are very close to the CO in the first periods of the game. Also, average total contribution is almost the same for three of the four treatment in the last period. Table 3 summarises the results for these four treatments.

These results are in concordance with other experiments where information is introduced in some way in a public goods game. See, for example Cason and Khan [19], and Ledyard [1].

In the following section, we will test a simple learning model both on the aggregate level to see whether this model fits the experimental data and whether this is still the case when we examine individual and group data.

The data that will be used for these tests are from treatment H without promises, treatment H with promises and treatments 1N, 1P and 2N with information. This choice is based on the fact that all these treatments have the same theoretical payoff function (high level of social

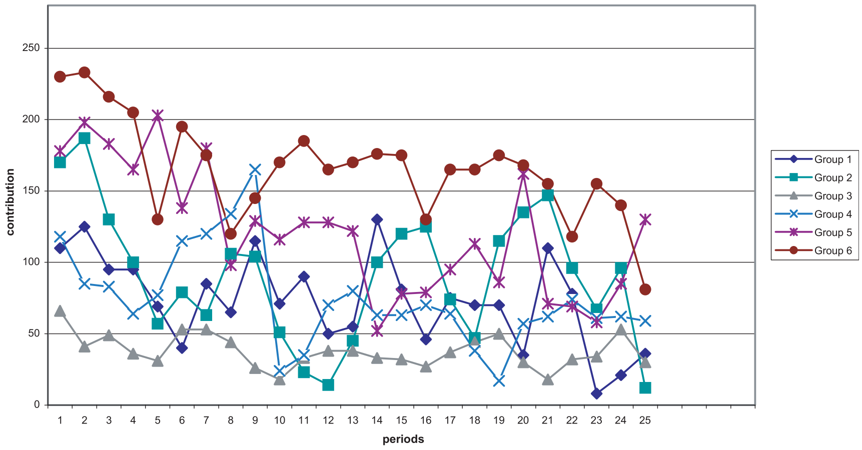

Fig. 5. Contributions of the six groups in treatment H (without promises).

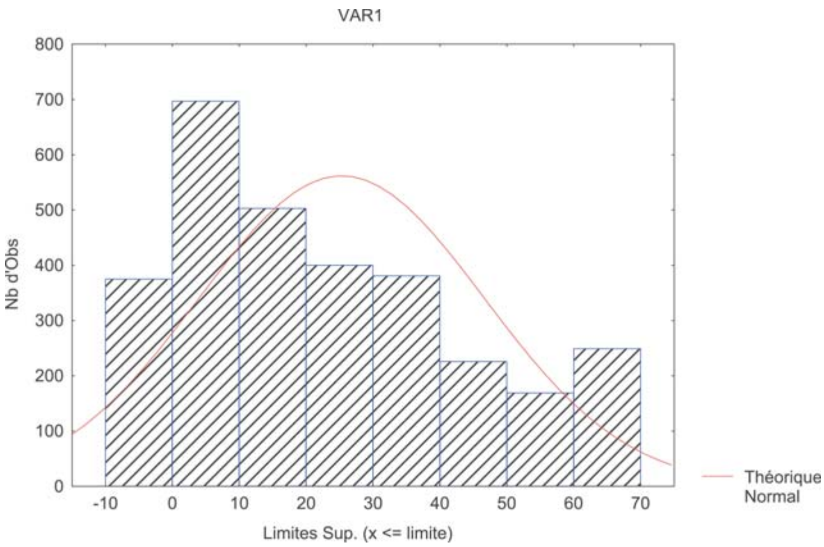

Fig. 6. Number of individual contributions for each interval for the 5 treatments (3000 decisions).

optimum). Since all these cases have the same theoretical solutions we have, in comparable situations, 24 persons playing 25 periods, and we have then 3000 observations or decision of contribution.

3.4 The group and the individual level

Although at the aggregate level behaviour seems rather consistent, in all treatments, there are at the group level different behaviours in the same group and between groups. The latter are not regular and we can clearly identify different attitudes to contribution. Figure 5 shows the contribution of the six groups in treatment H without promises and gives us an idea about this variation between groups.

Within groups, at the individual level, there is also a difference between the individual and the aggregate behaviour (see Ref. [20]). In fact, individuals behave differently. Moreover, the individual behaviour is more difficult to classify because of the great volatility of contributions of one subject during the 25 periods of the game.

To have an idea as to the different levels of contributions of individuals, we classify contributions in 8 intervals of ten each and we present in Figure 6 the number of times individuals make a contribution belonging to each interval. As in each of the 5 treatments there are 24 persons playing 25 periods, we have then 3000 observations or decision of contribution. As we can see in Figure 6, almost all the intervals are significant.

To isolate the different strategies that could explain the differences between the behaviour of different individuals, we also compare for the 5 treatments and for the 25 decisions of each individual her contribution in period (t) to her contribution in ($t-1$). This allows us to know whether individuals react in response to past contributions. We classify this comparison into three possibilities: contribution in t increases, decreases and remains unchanged in comparison to contributions in $t-1$. Figure 7 shows that all these strategies are significant.

When we test econometrically for treatment H the model:

$$C_t = a + \beta C_{t-1} + \varepsilon_t.$$

We find that β is significant for the aggregate and for the 6 groups while it is not for 13 subjects out of 24.

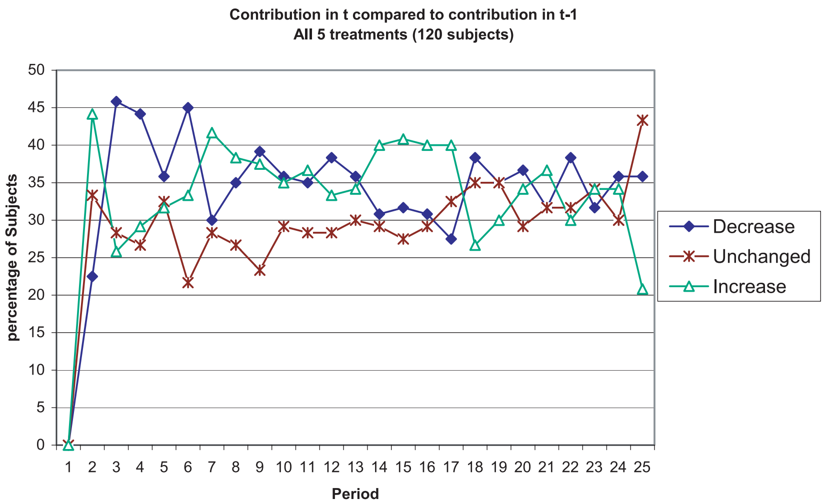

Fig. 7. Percentage of increasing, decreasing and unchanged individual contributions in t compared to $t-1$ for the 5 treatments (3000 observations).

The assumption $\beta = \beta_i$, $\forall i = 1, ..., 4$ is also rejected for all the individuals of the 6 groups. Thus we can conclude that individuals are heterogeneous.

4 Simple tests of the learning model

We will first present the "*Reinforcement learning model*", (R.L.) and we will then test this simple model at the aggregate level using the experimental data.

4.1 The reinforcement learning model

Two properties of human behaviour in the set of situations we analyse are mentioned in the psychology literature. The first one, known as the "Law of Effect" reflects the fact that choices that have led to good outcomes in the past are more likely to be repeated. The second one is called the "Power Law of Practice" and announces that learning curves tend to be steep initially, and then flatter. Another property is also observed, which is "recency", according to which recent experience may play a larger role than past experience in determining behaviour.

There exists a wide variety of learning models in the literature. We will use a simple learning model and apply it to our experiments. The model here is the basic reinforcement learning model used for example by Roth and Erev [3]. There are several variations of this model. In the one parameter reinforcement model, each player i, at time $t = 0$, before the beginning of the game, has an initial propensity to play his kth pure strategy. Let $A_k^i(0)$ be this initial propensity. When a player receives a payoff x after playing his kth pure strategy at time t, his propensity to play strategy k is updated. The rule for updating these propensities from a period to another is given by the following relation:

$$A_k^i(t+1) = A_k^i(t) + x.$$

The propensities to play the other pure strategies j are

$$A_j^i(t+1) = A_j^i(t).$$

These propensities allow player i to compute the probability that he plays his kth strategies at time t. Let this probability be

$$p_k^i(t) = \frac{\exp(\lambda \cdot A_k^i(t))}{\sum_{j=1}^{m_i} \exp(\lambda \cdot A_j^i(t))}$$

where the sum is over all of player i's pure strategies j[2].

4.2 The simple test of the reinforcement learning at the aggregate level

We will apply the reinforcement learning model to the aggregate level of the 5 treatments. We divide the set of possible contributions ([0; 280]) into ten equal intervals ([0; 28)]; [29; 56]; [57; 84]; ...; [252; 280]). At time $t = 0$, all the possible levels of contribution have the same attraction and the same probability to be chosen. In period 1, the strategy chosen in the data at the aggregate level receives the payoff x as explained above. At each of the 25 periods, the aggregate contribution of each treatment of one given period played in the real experiment is updated. This model is applied to each of the 5 treatments and the average of these results is calculated and presented in Figure 8. This figure shows that the strategy that is played the most according to the reinforcement learning model is to contribute an amount corresponding to the fourth interval, that is [85; 112]. This result obviously corresponds to the experimental data given that we are using this data to update the observed chosen strategies.

5 The performance of the EWA rule on aggregate and individual data

In this section, we will present the EWA learning model and we will next apply it to the aggregate and the individ-

[2] This well-known rule is also referred to as the "Quantal response" rule or the "logit" rule and can be justified as optimizing the trade-off between "exploration" and "exploitation".

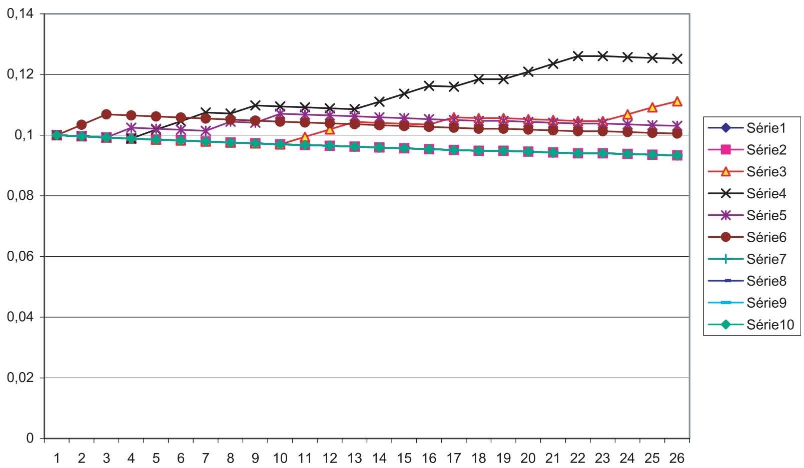

Fig. 8. The average probability of playing each of the 10 set of strategies when applying the Reinforcement learning model to the aggregate level of the six treatments.

ual level using the same data as those used in the previous section (the 5 treatments' data).

5.1 The EWA learning model

The details of the Expected Weighted Attraction (EWA) model can be found in Camerer and Ho [7,21] and Ho et al. [22]. Consider N players indexed by i, where $i = 1, 2, ..., N$. Let $A_i^j(t)$ denote the attraction of strategy j for player i once period t is played, where $j = 1, 2, ..., m_i$. The attraction of the different strategies in the EWA learning model are updated differently. The chosen strategies receive an attraction equal to the payoff π_i player i receives as a result of his choice, while the attractions of the unchosen strategies are updated by adding only a part δ of the foregone payoff. Thus, the parameter δ is the weight a player puts on the unchosen strategies. Let s_i^j denote the strategy j of player i, and $s_{-i}(t)$ be the strategies chosen by all the players, except for player I, in period t. Attractions of strategies are updated according to the payoffs these strategies provide, but also according to the payoff that unchosen strategies would have provided. Recall that it is not always obvious that this information will actually be available to the players. In our case the fact that the individuals are apprised of the total contributions allows them to work out what the foregone payoffs of unchosen strategies would have been. The rule for updating attractions in period t is:

$$A_i^j(t) = \frac{\phi \cdot N(t-1) \cdot A_i^j(t-1) + [\delta + (1-\delta) \cdot I(s_i^j, s_i^j(t))] \cdot \pi_i(s_i^j, s_{-i}(t))}{N(t)}$$

where $I(x, y)$ is an indicator function that is equal to 1 if $x = y$ and to 0 if not. $\pi_i(s_i^j, s_{-i}(t))$ is the payoff of player i when he plays strategy j while the other players play the combination of strategies $s_{-i}(t)$.

The parameter ϕ is a discount factor used to depreciate the previous attractions so that strategies become less attractive over time.

$N(t)$ is the second variable updated in the EWA learning model. It is the experience weight used to weight lagged attractions when they are updated. $N(t)$ is updated according to the rule:

$$N(t) = \phi(1-\kappa) \cdot N(t-1) + 1$$

where $t \geqslant 1$.

The parameter κ controls whether the experience weight depreciates more rapidly than the attractions.

At $t = 0$, before the game starts, the two variables $A_i^j(t)$ and $N(t)$ have initial values $A_i^j(0)$ and $N(0)$.

The probability of choosing a strategy j by player i in period $(t+1)$ is calculated by using a logit form:

$$p_i^j(t+1) = \frac{\exp(\lambda \cdot A_i^j(t))}{\sum_{k=1}^{m_i} \exp(\lambda \cdot A_i^k(t))}$$

where λ controls the reaction of players to the difference between strategies attractions. A low value of λ implies an equal probability for choosing strategies, while a high value supposes that players are more likely to chose strategies with higher attractions.

We note that the reinforcement learning model is a special case of the EWA model. In fact, when $\phi = 1$, $\kappa = 1$, $\lambda = 0$ and $N(0) = 1$, attractions are updated in the same way as in the cumulative reinforcement learning model. The EWA learning model includes the reinforcement learning model and other learning models such as belief learning as special cases (Ref. [7]).

5.2 The test of the EWA learning model at the aggregate level

We ran simulations using the same theoretical function as that used in our 5 treatments. We calculate the probability of playing each possible contribution by updating the initial propensities of each strategy using the EWA learning model. The parameters used are those which best fit the model.

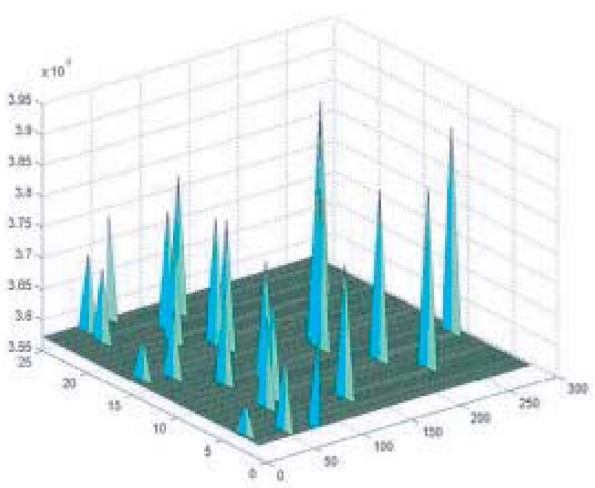

Fig. 9. The probability of each of the 25 chosen strategies using the EWA learning model.

First, we play the public goods game 25 periods where the simulation program chooses randomly in the first period one of the 281 possible levels of contribution ([0; 280]). This means that in the first period, all the strategies have the same probability of being chosen. The propensity of the strategy chosen in the first period is updated according to the EWA learning model. Next, the simulation program calculates for each strategy the probability of its being played in period 2. Obviously, the strategy chosen in period 1 has a greater probability of being chosen in period 2. In the second period, the simulation program chooses a new strategy using the new probabilities of being played for each strategy.

This program is run for 25 periods. Figure 9 is an example of 25 strategies belonging to the interval [0; 280] and chosen by the simulation program.

We repeat this simulation 1000 times and we calculate the average of the probabilities of all the chosen strategies. The result is presented in Figure 10.

5.3 The test of the EWA learning model at the individual level

We apply the same simulation program (run 1000 times) for a set of possible strategies that corresponds to the set available to one person playing the public goods game. This set of strategies is [0; 70]. The payoff of one player depends on his strategy but also on the choice of the three other players of the same group who are playing the game. The results are presented in Figure 11.

5.4 Comparison with the simple learning model

The simple learning model predicts that the strategies that have the highest probabilities of being played are

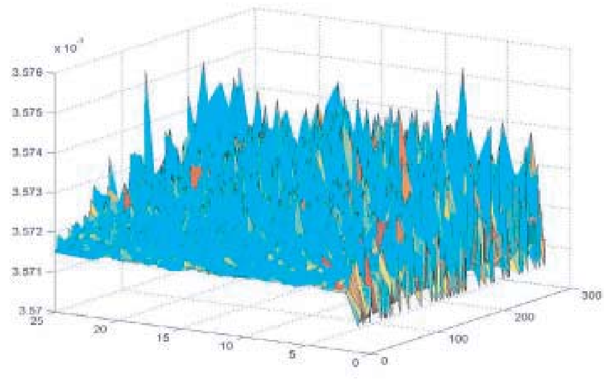

Fig. 10. The average of the probabilities of 1000 repeated simulation of the 25 periods public goods game using the EWA learning model.

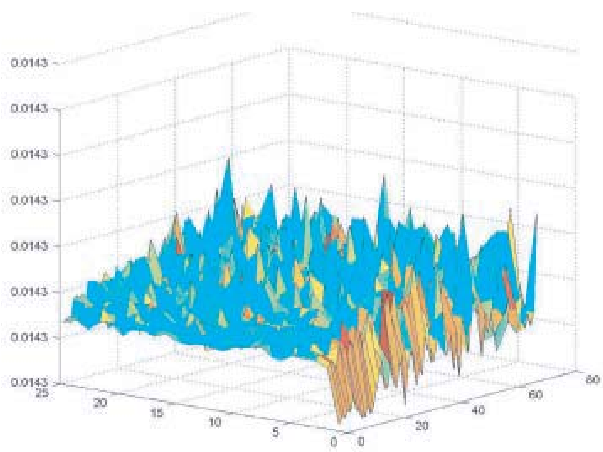

Fig. 11. The average of the probabilities of 1000 repeated simulation of the 25 periods public goods game for one player using the EWA learning model.

those which belong to the interval [85; 112]. The results of simulations applied at the aggregate level show that high contributions have the highest probabilities of being played. Obviously, this does not correspond to the experimental findings where low contributions are more likely to be played and where contributions decrease over time.

Simulating the EWA model for single players playing the public goods game shows that contributions belonging to the interval [40; 60] have the highest probabilities of being played (Fig. 11). This is not in accordance with the observed strategies. Furthermore, the EWA model would suggest an essentially monotone evolution. Whilst this is observed at the aggregate level for all treatments, this is far from being the case for the groups and the individuals. The behaviour of players seems to differ widely across individuals and what is more there is no convergence to any common behaviour. Players seem to play out of equilibrium strategies and this does not correspond to any simple learning model.

6 Conclusion

The simple point made in this paper is that what seems to be rather systematic behaviour at the aggregate level which persists in various treatments of the public goods game, does not reflect such systematic behaviour at the group or individual level. People within groups interact and react to the contributions of the other members of the group. This may lead to very different levels of total contributions across groups and over time. Furthermore, different groups do not necessarily exhibit the uniform decline in total contributions which is observed at the aggregate level across all treatments.

Rather different versions of our basic experiment yield similar results so the difference between the individual and aggregate level cannot be attributed to some specific institutional feature of our experiments.

What it is that causes the variation across indivduals remains an open question. Simplistic explanations such as "degrees of altruism" do not seem to be satisfactory. What does seem to happen is that some individuals try to signal to others by means of their contribution. They may hope, in so doing, to induce higher payments from their colleagues. If this is what they are doing then they violate some of the simple canons of game theory.

What is clear is that the players do not, in general manage to coordinate on cooperative behaviour. Thus, in the set-up here the optimistic conclusion that cooperation will emerge, found in the "Prisoner's Dilemma" literature does not seem to be justified. Free riding is something which on average increases over time but many individuals do not follow this pattern. To sum up, the behaviour of individuals varies considerably, but the complexity of the interaction is washed out in the average. Nevertheless, this should not lead us into the trap of attributing individual behaviour to the aggregate nor, and worse, of concluding from the apparent aggregate learning process, that individuals are learning in this way.

References

1. J. Ledyard, edited by J. Kagel, A. Roth, *The Handbook of Experimental Economics* (Princeton University Press, 1995)
2. T. Saijo, H. Yamaguchi, *The "spite" dilemma in voluntary contribution mechanism experiments*, Technical Work (University of Tsukuba, Japan 1992)
3. I. Erev, A. Roth, Am. Econ. Rev. **88**, 848 (1998)
4. T. Ho, C. Camerer, X. Wang, Individual differences in EWA learning with partial pay-off information, Forthcoming in The Economic Journal (2006)
5. D. Mookherjee, B. Sopher, Games and Economic Behavior **19**, 97 (1997)
6. D. Fudenberg, D. Levine, *The Theory of Learning in Games* (MIT Press, 1998)
7. C. Camerer, T. Ho, Econometrica **67**, 827 (1999)
8. Y. Cheung, D. Friedman, *Individual Learning in Normal Form Games: Some Laboratory Results*, Games and Economic Behavior, **19**, 46 (1997)
9. D. Stahl, International Journal of Game Theory **28**, 111 (1999)
10. B. Broseta, Games and Economic Behavior **32**, 25 (2000)
11. T. Page, L. Putterman, B. Unel, The Economic Journal **15**, 1032 (2005)
12. E. Sol, S. Thoron, M. Willinger, *Do binding agreements solve the social dilemma? An experiment*, Working Paper, GREQAM, Marseille, 2006
13. N. Vriend, Economics Working Papers **185**, Department of Economics and Business, Universitat Pompeu Fabra, 1996
14. J. Andreoni, Am. Econ. Rev. **85**, 891 (1995)
15. D. Davis, Ch. Holt, *Experimental Economics* (Princeton University Press, Princeton, N.J, 1993)
16. C. Keser, *Cooperation in Public Goods Experiments*, Scientific Series (CIRANO, 2000)
17. C. Keser, Economics Letters **50** (1996)
18. M. Sefton, R. Steinberg, Journal of Public Economics **61**, 263 (1996)
19. T. Cason, F. Khan, Journal of Development Economics **58**, 1999
20. W. Hichri, Applied Economics Letters, **13**, 969 (2006)
21. C. Camerer, T. Ho, Journal of Mathematical Psychology **42**, 305 (1998)
22. T. Ho, C. Camerer, J. Chong, *Functional EWA: A One-parameter Theory of Learning in Games* (unpublished)