Collective action, free-riding and evolution

Juan D. Montoro-Pons
Francisco Garcia-Sobrecases
Departamento de Economía Aplicada, Universitat de València, Spain

First Draft, May 2000

Abstract

Nash equilibrium is the point of departure of most of the standard literature on public goods. This stresses the sub-optimality of voluntary contributions towards the provision of a public good: in game theoretic terms, with unboundedly rational agents, individual best response is no cooperation in the provision of public goods. However, this is not a satisfactory conclusion, as empirical facts show that, to a certain degree, cooperative behavior as well as free riding emerge in collective action: privately provided public goods do in fact exist. This suggests rethinking the behavioral assumptions that support the conclusions of the conventional model and the extreme abilities and requirements that it imposes on economic agents. This is especially true in complex exchange situations such as the voluntary provision of a public good. Within this context, the aim of this paper is to sketch a behavior theory of non-market decision making in which agents choose a level of individual contribution for a public good. To this end, it departs from the concepts of bounded rationality and evolution, which help in explaining the outcomes of social interaction. We model agents as classifier systems which interact under two main scenarios, being the difference whether imitation is possible or not. This allows to reproduce individual and collective learning and to assess its impact in the outcomes of the game. The emergent properties of the social cooperation agree with most experimental literature findings: cooperation, although not optimal, is a fact. Moreover the institutional setting affects in a significant way the outcomes.

1 Introduction

Collective action is focused on the analysis of the institutional settings by which individual decisions are coordinated in order to get some defined collective outcomes. However spontaneous coordination is not always assured, and conflicts arise whenever individual rationality is opposed to collective rationality. This is the case for most collective action problems and has been shown as the justification for introducing an exogenous agent, the state, that will, using its coercive powers, reconcile individual and collective targets. Yet it must be noted that while accepting that exogenous solutions may, at least partially, solve this social dilemma (which we will call from now on free-riding without lost of generality), the efficiency of the outcomes relies heavily on extreme behavioral assumptions about the government[1]. If we loosen these constraints we may find that exogenous solutions may not improve the spontaneous coordination of the agents. This turns to be an interesting point of departure in order to analyze the evolutionary behavior of individuals in a collective action context. What we try to analyze is the rationale underlying cooperative behavior of a set of individuals in a generalized prisoner's dilemma (PD) framework. It is important to stress that the evolution of individual actions will be the point of departure for analyzing how this coordination is (or it is not) achieved. In order to simplify the expository aspects we will assume that we are in a public good provision setting, in which individuals decide their contributions, or level of cooperation, toward the provision of a pure public good. In this framework, individuals face a generalized PD setting in that the individually optimal amount of cooperation is zero, opposed to the efficient collective level of provision. However this assumes individuals exhibit unbounded rationality; they identify equilibrium states an peg to them; in a dynamic setting this will eliminate any transitional dynamics as agents, from the beginning, will know the optimum and will act accordingly. On the contrary, experiments have shown that real agents do not behave as normative game theory assumes them to, in a world in which adaptive behavior and bounded rationality seems to be the rule rather than the exception. Moreover, agents do not peg to strategies but adapt themselves to new information, use readily available data to update their forecasts, and rules of thumb to select new strategies, all this leading to an evolutionary pattern of behavior that discourages the indiscriminate application of first order maximization conditions and the modeling of omniscient agents.

The motivation of this paper is twofold. First is to analyze, within an evolutionary context, which are the emergent solutions of a generalized PD under different institutional settings. Second is to differentiate between the

[1] It is worth mention that the inefficient outcomes of the spontaneous order is also based on the axioms about the agents.

macro-behavior and micro-behavior of a population of boundedly rational agents and to test whether macro-behavioral programs do collide with our proposed micro-behavioral program. To this end we have structured the paper as follows: section 2 is a short introduction to the standard approach to the public goods provision[2] and its relation with the main findings in the experimental literature; section 3 shows the basic model of agent behavior as opposed and compared to population programs suggested sometimes in the experimental and evolutionary games literature; section 4 will analyze the experiments we run and its main results; finally section 5 will end with some conclusions.

2 Cooperation and collective action. Paradoxical outcomes?

It is a well-established fact that voluntary or private contribution to the provision of public goods does exist, and currently there is no question as to whether individual cooperation is carried out effectively in order to attain this provision. However this fact is not supported from the predictions of the conventional model. In order to abandon the Nash solution, the possible departures established from the conventional theory assumes that cooperative behavior to be introduced into the model should be done exogenously. However, those attempts have led toward an unsatisfactory theory to explain the voluntary provision of public goods on the traditional ground that individuals maximize their utility guided by selfish and outcome-oriented motivations based on rational cooperation. Thus, in an unbounded rationality scenario, only exogenously imposed solutions have meant cooperative behavior toward the public goods provision.

Some remarks must be cited. In the first place, these motivations exogenously introduce cooperative behavior of the individuals which, in fact, goes against the Nash assumption as a dominant strategy . In other words, what has been conventionally done is to setup a cooperative framework for noncooperative individuals. In the second place, the conventional theory has assumed the cooperative behaviors shown by individuals, when facing circumstances that require this type of collective action, as something of an extraordinary event. According to the construction of the theoretical model, some objections upraise about how individual rational behavior can be encouraged for the voluntary provision of public goods. Within the framework of public goods this rational behavior is characterized by (i) the nature of public goods, and (ii) the fact that rational behavior have been directed towards the acquisition of private goods given the rivalry for these goods . As Andreoni [3] remarks, this rational behavior is aimed at making the public

[2] That may be valuable for any other social dilemma situation.

goods only play the role of a mere opportunity cost for acquiring the private ones.

Furthermore, if the Nash or non-cooperative equilibrium is abandoned, and thus the individual takes a personal and independent decision consistent with his preferences, how can he continue to uphold the assumption of instrumental rationality within the framework of cooperation? The individual perceives the new framework correctly and is aware of the fact that other individuals are also playing the game, thereby altering the outcome with their own decisions. Therefore, the rational behavior, highlighted within the framework of cooperation, finds a strong obstacle in keeping Nash solutions as the only dominant ones. The assumptions to be made in order to abandon these solutions must necessarily be characterized as endogenous. Following Sugden [30],in a cooperative framework with utility maximizing individual behavior "public goods would never be supplied at all" .

Hence the solutions offered by the conventional theory are not valid for justifying approximations to cooperation from the Nash-Cournot or non-cooperative standpoint. Further, there is not a strong compatibility between the consistent equilibrium implied by the conventional theory and the bases that sustain it. Therefore, the departures proposed by the conventional theory in order to abandon the Nash-Cournot equilibrium can be characterized as more jeopardizing than beneficial.

However it is well known that the conventional model fails to predict real-world agents' behavior. Empirical facts[3] show that, to a certain degree, cooperative behavior among individuals emerges as free-riding does. In a similar way, experiments[4] show that participants tend to overcome the free-riding problem, with average contributions over the Nash level. Apart from this finding (which is by itself an important contribution toward the collective action literature), there are several other features of experiments that are worth mention. First, the design features of these experiments may affect the degree of cooperation that individuals show; this means that different institutional settings are likely to be related with different outcomes. Second, dynamic settings, those in which the interaction among players lasts for more than one round, allow for a decline in the contribution of individuals over time[5]. Third, communication between individuals may play a key role in promoting cooperation. And fourth, monitoring (or the information available to the players about the state of the system) seems to play a minor role in determining the individual level of contributions.

Whether cooperation in experiments is true altruism or error of agents is

[3]We may think of the existence of non-governmental organizations, or certain local services which are financed through private contributions.

[4]Interesting surveys are Dawes and Thaler [12], or more recently Ledyard [20].

[5]Which may also be linked with the experience of individuals playing the game. However it must be noted that most experiments find that cooperation is maintained over time for very large groups and with high per capita returns of the public good.

still under consideration but certainly efforts have been directed towards a better understanding of individual behavior. Palfrey and Prisbrey [24] find that most decisions are misled and attribute cooperation to mistakes by the agents; on the contrary Andreoni [2] finds some evidence of preferences for cooperation in the agents. Similarly Sefton and Steinberg [27] show that, in two public good experiments using different reward structures, the excessive giving can only partially be explained to players' confusion, whereas donations may be consistent with altruistic preferences. Institutions as the set of constraints that agents face are another element to take into consideration when analyzing the behavior of individuals in these experiments. Elliott et al. [14] generate two different framings: cooperative and non-cooperative or entrepreneur business strategy, which play the role of institutional factors influencing agents in their interaction. Their experiment show that different patterns of cooperation arise, being lower, as expected, the one related to the second, or enterpreneur, strategy. Willinger and Ziegelmeyer [35] identify a substantive aspect in public goods experiments that other kind of experiments with a similar framework, as is the case of oligopoly games, lack. In public goods games players identify actions of others as a positive externality[6] then they will be willing to cooperate. This positive externality may be somewhat linked to altruism.

It is worth noting that a change in the rules[7] implies a modification in the institutional setting, hence we must agree that each experiment addresses to a specific institutional environment. Modifying the specific conditions under which the game takes place is to modify the institutions. As an example, Cason and Khan [6] introduce imperfect monitoring and communication in the experiments; imperfect monitoring is to be understood as a limitation to the quantity and/or the quality of the information the agents face in order to make their decisions[8], whereas communication refers to the possibility of exchanging information between the individuals involved in experiments. With this modifications the authors find that communication is a powerful means to increase cooperation[9], while better monitoring[10] only increases cooperation when individuals can communicate.

[6] Which is not the case for oligopoly and common pool resources games. In these others' actions are sources of negative externalities.

[7] Rules in a broad sense, as the way a game is played.

[8] Which will complicate cooperation as it increases the incentives for deviating from the cooperative strategy.

[9] Which is a striking feature from the game theoretic standpoint, as cheap talk agreements cannot be enforced, and they do not modify the incentives structure of the game.

[10] Similar results, although in a different *institutional setting* were identified by Heijden et al. [31]. These authors develop a model of intergenerational transfers among individuals which resembles a public good game (or more specifically a social dilemma setting), as individual and collective rationality collide, being the latter pareto optimal. They observe that monitoring is of little importance for determining the outcomes, and a fairly efficient level of transfers emerges.

Experiments show that individuals do not behave strategically (or at least perfectly strategic) as long as Nash equilibria do not prevail. Yet some experiments show that contributions tend to decay over time, thus approaching asymptotically to the game theory concept of equilibrium. This poses the question whether learning is an important aspect of individual behavior: in this case agents could be using an adaptive process[11], and as the experiment goes on, errors are dampened through a learning process, which ultimately yield to free-riding as the dominant strategy. Several authors have analyzed this point as Andreoni [1] or more recently Sonnemans et al. [29]; while the former rejects the strategic and learning hypothesis, the latter found evidence on both strategic and adaptive behavior.

To summarize, over-contributions in experiments have been related to: (i) errors by the agents (agents follow rules that led them to systematic errors, as an adaptive pattern); (ii) altruism; or (iii) institutional factors, either formal or non-formal[12]. At this point there should be no doubt that empirically, cooperation can be justified on endogenous basis. Being this the case, a substantial revision of the basis of the non-cooperative behavior postulated by the conventional approach is needed. This suggests rethinking the behavioral assumptions that support the conclusions of the conventional model. Above all, the extreme abilities and requirements imposed by unbounded rationality to the agents poses serious doubts of its validity in a complex exchange situation such as the collective action problem that is reflected on the grounds of the voluntary provision of a public good. In what follows we propose a model of agents behavior in which cooperation may arise as the outcome of bounded rationality and evolution. Unlike early evolutionary game models which rely on some macro mechanisms[13] we will analyze the micro-behavior, and from this the emergent aggregated characteristics of the model.

3 An agent based model for collective action

Agents, as we know them in economics and political science, face decision problems. In a collective action framework as the one described they have to choose their individual levels of cooperation. This means that there has to be a mechanism by which information about the environment, the institutional setting, and the social system leads them to a unique decision about how

[11] Note that in this case we may conclude that agent's behavior is guided by a learning process, that leads them to commit systematic errors. This will lead us to conclude that cooperation is linked to errors rather than to altruism.

[12] Rules, social norms, education... in short, the environment in which individuals interact. However, it should be noted that most formal institutions may be a mean of inducing cooperation in an exogenous way.

[13] Like the replicator dynamics, as in Miller and Andreoni [23], more recently analyzed by Van Huyck et al. [32] in the context of strictly dominated strategies.

much contribute to the public good. In the conventional approach there is no doubt about individuals' actions as they are perfectly sketched by a maximization problem which leads to an optimal policy to be implemented, leading, in this setting, to the non-cooperative solution[14]. Nevertheless, it has been shown that unbounded rationality conditions and outcomes are systematically violated in real world and experiments[15]. Following these findings a substantial revision of our way of modeling individual behavior applies.

We propose such revision basing individual decision making on rules of thumb which can be simply stated as follows: assume a set of N agents in a collective action context; let E be the set of all possible states for the problem (as in our example, the collective provision of a public good), and let A be the set of all actions. Then, a decision problem is a mapping G that assigns an action to every state. By using rules of thumb or *classifier systems* individuals base their decisions on the strength of a set of rules, which evolve with the game. This process of evolution is guided with the knowledge and experience accumulation of the players, i.e. with learning. Whether this learning reflects the kind of learning of isolated agents or on the contrary is a social learning process will be discussed later.

3.1 Description of the basic setting

Definition 1 *Let E be a set of states, A a set of actions, and G a mapping from states to actions. A behavior theory is defined as a tuple $\Gamma = \langle E, A, G \rangle$.*

Departing from Definition 1, we can redefine the problem of voluntary provision of public goods in a bounded rationality context. Let $\langle N, \Gamma, U \rangle$ be an economy at time t, being N the total population, Γ a behavior theory and U the payoffs function or the individual value function. Hence, the model may be considered as a simultaneous N players game, with two main features: (i) at the beginning of each period, individuals are endowed with an identical income I, to be spent in a private good or in the provision of a public good[16]; (ii) individuals will follow a rule G such that, given the state of the system, an agent will perform the action with maximum strength[17].

[14] Although if we are in a dynamic context and the horizon is infinite and the discount factor is high enough, any equilibrium pareto-preferred to the Nash equilibrium will be a valid one. However this is only a possibility result and if any of the conditions stated is not fulfilled, specially the one referring to the temporal horizon of the interaction, we then return to the non-cooperative setting.

[15] A broad revision can be found in Conlisk [7].

[16] We will consider a 1:1 technical transformation relation between the public and private good.

[17] This is to say that agents will choose the action that performs best. This does not only include past payoffs but also future payoffs as individuals, through the updating mechanism discount the future. However, as it will be discussed later, agents will be subject to errors and imitation of other individuals which will modify this basic setting.

As in the conventional economic model, individuals face an allocative problem, where ex-post utility of agent i is given by $U_i(x_i, Z)$, being x_i individual allocation on private good, and Z the total public good available to a community. How individuals choose their level of cooperation is linked with how individuals update the fitness of the set of rules they are using. The basic features of the model can be described in an algorithmic way as follows:

1. At stage t individuals decide their cooperation level $\eta \in [0, 1]$. Depending on this and on the income endowment I, agent i contribution to the provision of the public good is given by $,\eta_i I$ while his consumption of private good is $x_i = (1 - \eta)I$. In this setting $Z = \sum I\eta_i$, and the instantaneous utility from the consumption of agent i is $U_i(x_i, Z)$. To choose their cooperation level individuals have to :
   (a) Identify the state of the system[18].
   (b) Given the state, the behavioral program that follow the agents implies choosing the action with the highest strength. The fitness or strength of an action $v(\cdot)$ defines a complete binary order relation[19]: $v(a)$ gives the relative position of action a with respect to any other action a', with $a \in A \quad , \forall a' \neq a$.
2. Update the strength of the chosen action, by using the following expression:

$$v(a)_{it+1} = v(a)_{it} + \phi\left(U_{it}(x_{it}, Z_t) + \beta v(b)_{it+1}\right) \tag{1}$$

[18] Agents do not posses global information about the system but indirectly through the average level of public good provision (or social cooperation). This level determines the state of the system E in an approximate way, as agents assign a finite set of states to the infinite potential states (as many as average levels of cooperation in the closed [0,1]) by means of a transformation $f : Z \to E$. This can be considered as a signal extraction mechanism from the system, or as an expectation formation rule.

[19] This value may be considered in two different ways: (i) in a deterministic setting, given a state, an individual will select the action with maximum value ; (ii) alternatively, in a probabilistic setting, it determines the relative probability that action a will be selected. It must be noted that many of the results obtained from evolutionary game theory and evolutionary economics exhibit what has been called path dependence. This implies that mechanisms based (as the one described above) on individual search (and even on imitation) may lead to an indeterminacy in the long-term equilibrium. In fact, for different initial conditions, the evolution of the system may be different. While this fact ("history matters") may be valuable in some settings, it is not the case of our model. We are seeking for the steady-state predictions of the model independent of the initial state. To this end, experimentation is introduced by means of a trembling hand effect: individuals base the latter on a mixed deterministic-probabilistic approach to the selection of actions. This means that there is a slight probability for the agents to select the wrong action, that is, not choosing the one with higher strength. The practical implementation is the following: an individual will choose based on a deterministic pattern with probability $(1 - \pi)$ and based on a probabilistic pattern with probability π .

$\forall a, b \in A$, where a is the action taken in t and b the action induced in $t+1$, β is the discount rate and ϕ a decreasing sequence that ensures that the agents corrections over time will finally dampen and set to a steady state. It is worth noting that expression (1) not only takes into account the utility derived from consumption but also the state induced at the next stage, that is at $t+1$: $v(b)_{t+1}$ is the fitness of the action that, given the actual state and the fitness of the action taken, will be taken at next stage which is discounted at a rate β , which depicts the temporal preferences of the agents[20].

4 Experimental results

Given the former model, we conducted some numerical simulations in order to give an experimental approach to the qualitative properties of the solutions that emerge in the interaction of a population of artificial agents facing a social dilemma. To this end, and unless otherwise stated, we modeled an economy composed of 1000 homogeneous agents interacting during 10000 time periods[21]. Agents start each sub-game with an initial endowment of $I = 10$ monetary units, which have to allocate either to private or public consumption. The net payoffs for a player at any stage will be given by:

$$P_i = I(1 - \eta_i) + \kappa \left(\frac{1}{N} \sum \eta_i I \right) = U_i(x_i, Z_i) \tag{2}$$

Any player receives what she kept for private consumption plus κ times the average contribution of the community. In this setting κ plays the role of the average returns of the public good. In a population of heterogenous donors, the higher κ the more individuals will benefit from their contribution toward the provision of the public good, and vice versa. In this game there is a corner solution for $\eta_i = 0$, that is the non-cooperative strategy, while the collectively optimum is given by $\eta_i = 1$, $\forall i$.

Given this setting we are interested in the evolution of the system and its outcomes. We will study the effects of a set of control parameters over the simulations. In order to do so, we will consider two alternative scenarios. First we will analyze a situation in which imitation is not available for the agents. Second we will allow the agents to imitate and will see how this affects the previous outcomes. Imitation will play the role of social learning.

[20] Recently, several papers, see Lettau and Uhlig [21], have shown the relation of expression (1) with the optimal policy solution to a dynamic programming problem expressed in terms of the Bellman equation.

[21] Other control parameters will be the discount rate β and the per capita returns κ. When no reference to these parameters is made we used $\beta = 0.6$ and $\kappa = 10$.

4.1 Cooperation with isolated agents

There are two aspects we want to test in this setting. First, whether per capita returns k are important when determining the level of cooperation of individuals. In fact we expect this to be important in a behavioral program based on individual learning. In this case if individual cooperation is heterogenous in the population, there will be agents cooperating over and below the average; for $\kappa \leq 1$ there will be at least a 50% of the population cooperating over its returns. This will cause cooperation to fall. On the other hand for $\kappa > 1$ there are more possibilities for individuals to benefit from collective contributions, although the public good nature of the returns will imply that individual rationality will drive cooperation levels to free-riding. To some extent κ controls the publicness of the setting. For low values (those below one) we find that the public characteristic of the good is low, either due to congestion effects or the private nature of the good. Higher values increase the publicness of the good; when $\kappa = N$ the situation is that of a pure public good: its quantity is available for consumption for every individual in the community.

Second we want to analyze whether discounting the future affects the agent's behavior. The *Folk Theorem* for repeated games is a classical in non-cooperative game theory which relates the impatience of the players with the possible outcomes of a game. One interesting feature about classifier systems is that they are driven not only for past information but for future events. In this way we would like to test if cooperation is enforced by a higher level of discounting, or if there is a type of *Folk Theorem* in the framework developed.

Result 1 *For individual learning algorithms, the system converges towards an equilibrium. The amount of contributions of the equilibrium is an increasing function of the per capita returns of the collective action process (κ).*

This first result can be summarized in Figure 1, which shows the convergence to a level of cooperation in both the private good case (with $\kappa = 0.2$) and in a public good setting ($\kappa = 3$). While in the first case the system converges towards non-cooperative behavior, which is the optimal strategy given the private aspect of the good, in the latter there is an equilibrium with a positive level of private contributions. More important, the equilibrium average contributions rise with the per capita returns, as it is shown in table 1. Hence partial cooperation holds in a social dilemma setting, which supports most experimental results as well as empirical facts. Finally it shall be noted that in some situations convergence will not be towards a fixed but a periodic equilibrium.

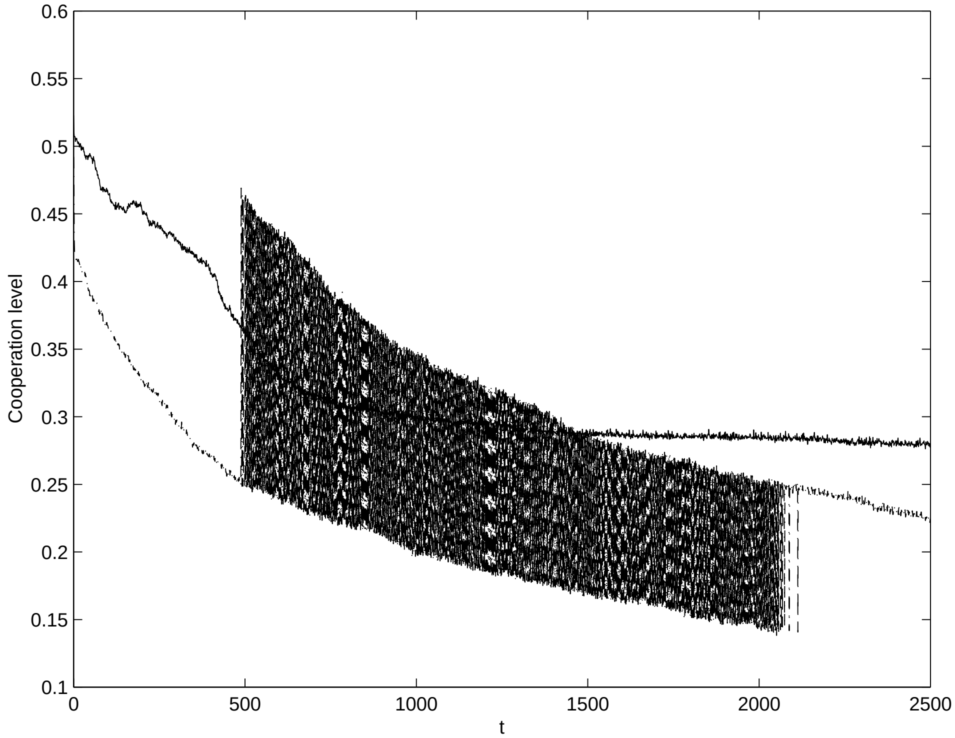

Figure 1: Convergence towards equilibrium: publicness of the good

k-value	Average cooperation level
0.2	0.1729
0.5	0.1750
1.0	0.1868
3.0	0.2406
10.0	0.2861
100.0	0.3502
1000.0	0.3640

Table 1: Average cooperation levels and per capita returns.

Result 2 *For individual learning algorithms, there is a discount level $\bar{\beta}$, such that:*

1. *For every $\beta \leq \bar{\beta}$ the system converges towards an equilibrium (either fixed point or periodic).*

2. *For every $\beta > \bar{\beta}$ there is no convergence an the model dynamics behaves erratically.*

Result 3 *The discount level $\bar{\beta}$ is a decreasing function of the per capita returns (k).*

Results 2 and 3 summarize our main findings with respect to the role of the impatience of the agents. This plays a major role in updating the strength of the classifier system, as can be seen by expression (1), in which the state induced by our present action is discounted at a rate β. Graphically, Figure 2 shows how, for low discount rates, the system settles to an equilibrium; however, for high enough rates, the average behavior wanders, in a bounded way, in the state space[22]. Additionally, it has been found that the threshold discount value over which the system dynamics changes from convergence to erratic is dependent on the returns of the game: the higher the returns, the lower the threshold value. This can be understood as an excess sensitivity of the system towards both variables, i.e. future payoffs and returns[23]. Figure 3 compares two simulations with the same parametric set ($\beta = 0.75$), being the only difference the κ-value; whereas in one setting (top) it is $\kappa = 50$, in the other (bottom) it is $\kappa = 500$. Here we can see that, for the same discount rate, higher instability are related to higher per capita returns.

4.2 Social exchange: the role of imitation

Up to this point we have had a population of isolated agents interacting in order to reach a level of cooperation. Now we will assume that individuals can communicate and exchange information about the fitness of their past strategies. In this way we will be assuming that a social learning process takes place within the population in order to reach to some level of cooperation. We have assumed two types of imitation, which will call *local* and *global*. In both a fixed percentage of the population is randomly chosen before the collective action stage; let n be the number of individuals which

[22] We tested against the null of unit root in both cases and results pointed towards this possibility.

[23] Although an empirical explanation is more difficult in the case of the discount rate, theoretically we are tempted to link it to the *Folk Theorem* for repeated games. There, for low discount rates Nash equilibrium is prevalent, while for high discount rates any equilibrium is possible; hence there is a determinacy and an indeterminacy of the equilibrium of a dynamic noncooperative game depending on the discount rate. Here we also find a determinacy and an indeterminacy of the equilibrium depending on the discount factor.

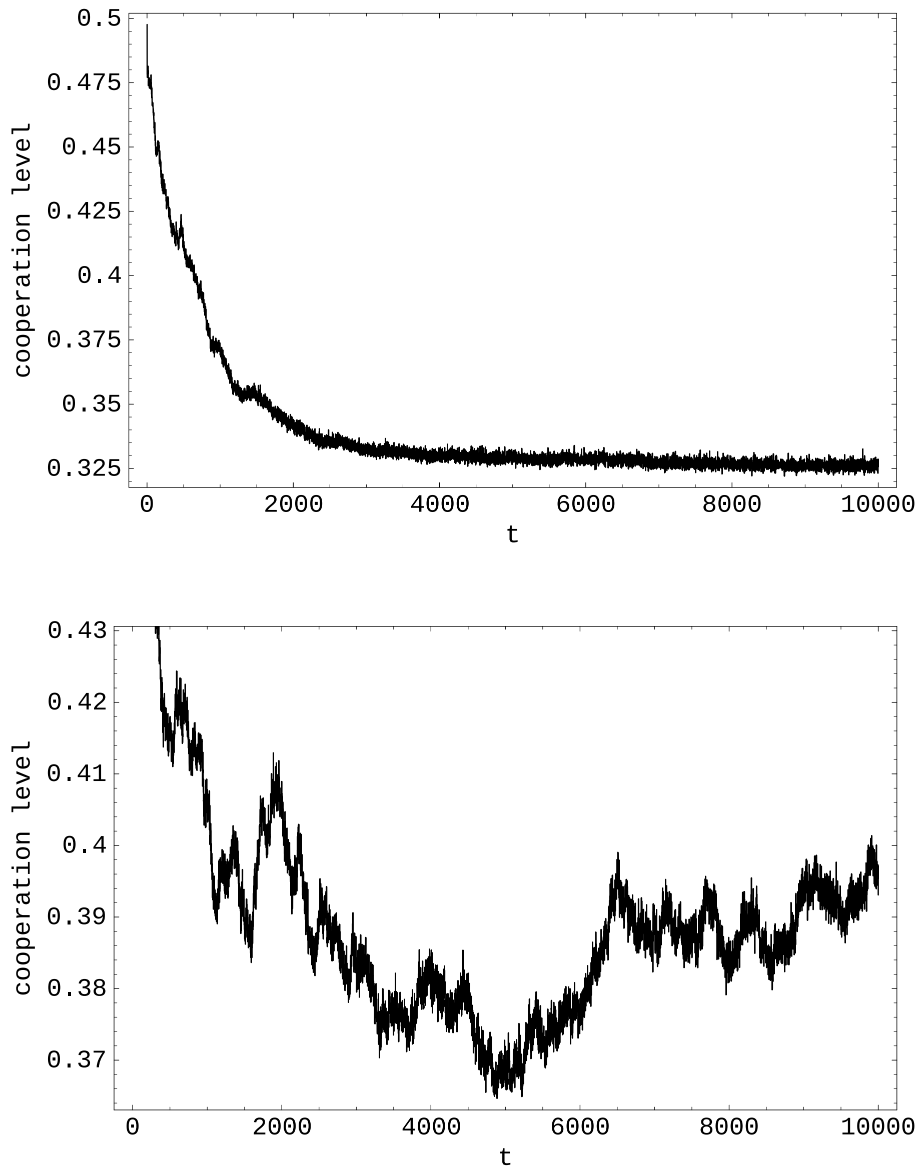

Figure 2: Population dynamics and discount rate: top ($\beta = 0.5$) convergence towards an steady state; bottom ($\beta = 0.95$) non-convergence.

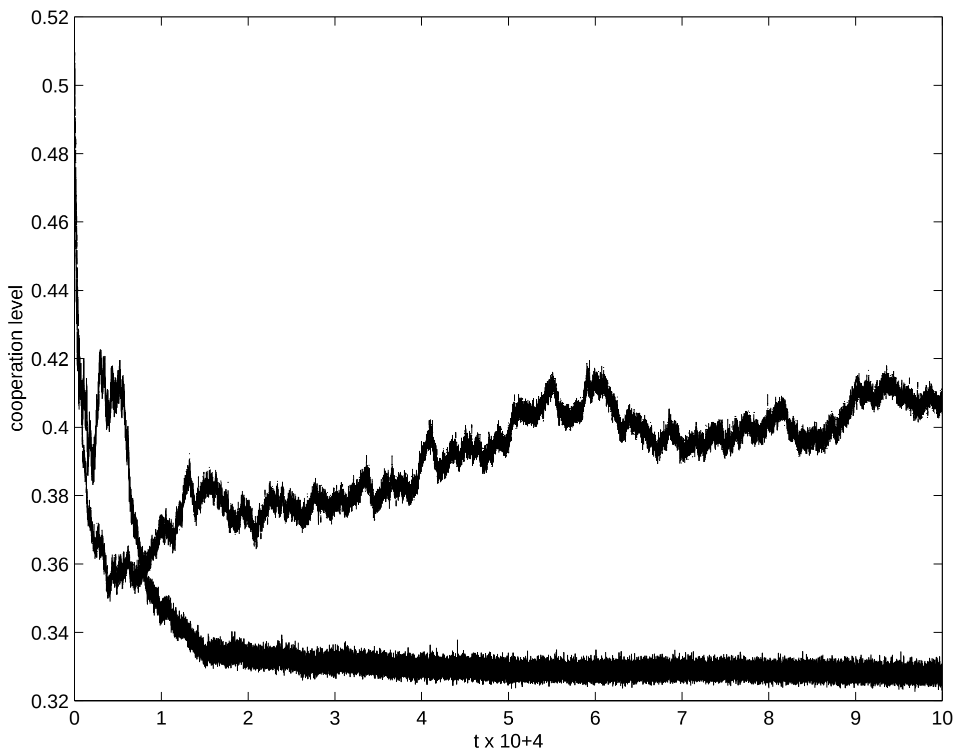

Figure 3: Population dynamics and per capita returns: top ($\kappa = 50$) convergence towards an steady state; bottom ($\kappa = 500$) non-convergence. In both cases $\beta = 0.75$.

will imitate. In the former type, individuals are paired; call them individual 1 and 2. Given the current state, both select the most fitted action which will be a_1 and a_2; obviously these can be the same or different actions. Next both agents compare the associated strengths of a_1 and a_2, $v(a_1)$ and $v(a_2)$. The one with higher strength will be chosen for both; for instance, if $v(a_1) > v(a_2)$ then both agents will implement a_1, with agent 2 imitating agent 1. In the latter type all the selected individuals will imitate the action of the agent with the highest strength of the population. To some extent this model implies a global imitation scheme, while the previous is a local imitation.

Result 4 *In the model under local imitation, agents :*

1. *move almost instantaneously towards non-cooperation for values of κ below a threshold $\bar{\kappa}$;*

2. *for $\kappa > \bar{\kappa}$ the system converges towards a (partially cooperative) equilibrium.*

Result 5 *In the model with local imitation, cooperative behavior is negatively related to the discount rate.*

Results 4 and 5 summarize the main conclusions with respect to the previously analyzed parameters. The former states that local imitation reinforces the previous results as convergence (either toward free-riding or positive contribution) takes place almost instantaneously, even for low imitation rates. The threshold value for κ was on most simulations close to 1, which supports our previous results. Most interesting, Table 2 shows the combined effects of the discount rate and the parameter κ on the average cooperation of the population. Now the discount parameter affects in a different way the outcomes of the simulations, as unambiguously can be shown that cooperation is inversely related to the discount rate of the population, and that higher per capita returns are related to higher average cooperative behavior.

Finally global interaction was analyzed and compared with local interaction. Main results appear in table 3. In it both types drive the population in opposite directions. While local learning tends to decrease the average level of cooperation, global learning significantly increases it. We will return to this in the conclusions.

5 Discussion and conclusions

The previous study has revisited the collective action problem (focusing in the voluntary provision of a public good) by revising its axiomatic foundations. In the conventional model, with unboundedly rational individuals,

	β				
	1	2	3	4	5
$\kappa = 5$	0.2499	0.1893	0.1816	0.1022	0.0409
$\kappa = 20$	0.3402	0.3190	0.2636	0.1861	0.1394
$\kappa = 100$	0.3592	0.3052	0.3265	0.2730	0.2590

Table 2: Discount rate and returns with local interaction.

Imitation rate	5%	15%	30%	50%	75%
Local imit.	0.2282	0.2234	0.2211	0.2025	0.1872
Global imit.	0.2252	0.2219	0.2607	0.5097	0.4915

Table 3: Local vs. global imitation and average cooperative behavior.

non-cooperation is the dominant strategy and only exogenous mechanisms (i.e. the state as a coercive agent) can solve the dilemma.

Given the apparent contradiction among theory and facts (both reality and experiments) , we tried to bring the setup closer to reality, by modeling agents as boundedly rational. Taking into account informational limitations, the limited computational abilities of agents, and the role of learning, we modeled agents as classifier systems. Within this setup results show that individuals partially cooperate. Although the optimum is not achieved, we can conclude that under bounded rationality, individuals do show a higher endogenous cooperation than the one in the conventional approach. Additionally we have shown that an agent-based model mimics some of the conclusions reached in laboratory experiments. One should be cautious as to extend the conclusion that all agents show the same willingness to cooperate, as aggregation shadows individual behavior. Although the average contribution follows the already mentioned patterns, there is a dispersion shown by maximum and minimum contributions which behave differently; in both cases minimum contributions keep to a low level, almost free-riding, while maximum stay well above the average [24]. Hence, even if there is a collective convergence towards an equilibrium, this process is not clear at the individual level.

Another interesting issue was to embody individual and social learning within a unique model. In a recent work, Vriend [34] compares the type of learning implied in a classifier system, which he calls *individual learning*, with the learning process of a genetic algorithm , which he calls *social learning*. While the main difference is whether there is an exchange of information

[24] This is not to say that there is an exploitation of some agents as no agent behaves always the same.

among the agents when updating an individual's rule[25], it is not generally true that a classifier system does not includes elements of social learning by allowing individuals to imitate[26]. While isolated agents were a point of departure, the exchange of information by means of the imitation of successful strategies is a means of social learning that may modify the original behavior in games. As communication or monitoring in experiments, imitation can be implemented from different perspectives and we used two models based on the type of information that agents gather, either global or local. To some extent, global imitation can be compared to a model of social learning, to use Vriend's terms. However it is not clear under which circumstances social learning may be a good approach for collective action processes as the one described. Whereas global interaction may describe the type of behavior of oligopoly-type games, in which the number of players is reduced and it is easy for participants of the game to identify successful strategies, collective action setups may imply a large number of players and individuals may have less incentives to gather global information. Take voting as an example; in the public choice literature simple models of voting behavior analyze it in cost-benefit terms. In this case political participation is seen as a public good: as the costs of voting exceed expected benefits rational individuals will not vote, and will free-ride on the actions of others. It is highly unlikely that global interaction plays a significant role in this case; on the contrary, local interaction may be better fitted to understand the role of information transmission, and may be an explanation of high turnout levels in elections[27].

To conclude, the emergent properties of an agent-based model resemble those outcomes of experiments: specifically when elements of learning and evolution, communication and imitation are included in the framework there is a strong parallelism between the model and facts. Classifier systems help in developing models of maximizing agents with limited computational abilities with a mixture of adaptive as well as look forward behavior which can help in understanding real world phenomena.

References

[1] James Andreoni. Why free ride? strategies and learning in public goods experiments. *Journal of Public Economics*, 37:291–304, 1988.

[25] Basically individuals interacting as GA use three basic operators: reproduction, mutation and crossover. The first creates new (and most fitted) rules as combination of previous ones, in a process that may be called as *social imitation*. The second and third operator allow for individual experimentation. Here the exchange of information is done by means of the reproduction operator: new rules are created as a result of previously successful ones.

[26] See for example Başçi [4].

[27] High for most public choice literature standards. Of course local interaction may be only one element to explain low turnout; additional factors can be bounded rational individuals, expected closeness of an election, etc.

[2] James Andreoni. Cooperation in public goods experiments: kindness or confusion? *American Economic Review*, 85:891–904, 1995.

[3] James Andreoni. Warm-glow versus cold prickle: the effects of positive and negative framing on cooperation in experiments. *The Quaterly Journal of Economics*, 110:1–21, 1995.

[4] Erdem Başçi. Learning by imitation. *Journal of Economic Dynamics and Control*, 23:1569–1585, 1999.

[5] Wolfgang Buchholz. A further perspective on neutrality in public goods economy with conjectural variations. *Public Finance Quaterly*, 21:115–118, 1993.

[6] Timothy N. Cason and Feisal U. Khan. A laboratory study of voluntary public goods provision with imperfect monitoring and communication. *Journal of Development Economics*, 58:533–552, 1999.

[7] John Conlisk. Why bounded rationality? *Journal of Economic Literature*, 34:669–700, 1996.

[8] Richard Cornes and Todd Sandler. The theory of public goods: Non-nash behavior. *Journal of Public Economics*, 23:367–379, 1984.

[9] Richard Cornes and Todd Sandler. On the consistency of conjectures with public goods. *Journal of Public Economics*, 27:125–129, 1985.

[10] Richard Cornes and Todd Sandler. The simple analeptics of pure public good provision. *Economica*, 52:103–116, 1985.

[11] Dipankar Dasgupta and Jun ichi Itaya. Comparative statics for the private provision of public goods in a conjectural variations model with heterogeneous agents. *Public Finance*, 47:17–31, 1992.

[12] Robin M. Dawes and Richard H. Thaler. Anomalies: cooperation. *Journal of Economic Perspectives*, 2:187–197, 1988.

[13] Herbert Dawid. On the convergence of the genetic learning in a double auction market. *Journal of Economic Dynamics and Control*, 23:27–53, 1999.

[14] Catherine S. Elliot, Donald M. Hayward, and Sebastian Canon. Institutional framing: some experimental evidence. *Journal of economic behavior and organization*, 35:455–464, 1998.

[15] Glenn Ellison and Drew Fudenberg. Rules of thumb for social learning. *Journal of Political Economy*, 101:612–643, 1993.

[16] Werber Güth and Shmuel Nitzan. The evolutionary stability of objections to free riding. *Economics and Politics*, 9:133–149, 1997.

[17] Elizabeth Hoffman, Kevin A. McCabe, and Vernon Smith. Behavioral foundations of reciprocity: experimental economics and evolutionary psychology. *Economic Inquiry*, 36:335–352, 1998.

[18] Claudia Keser. Voluntary contributions to a public good when partial contribution is a dominant strategy. *Economic Letters*, 50:359–366, 1996.

[19] Oliver Kirchkamp. Simultaneous evolution of learning rules and strategies. *Journal of Economic Behavior and Organization*, 40:295–312, 1999.

[20] John Ledyard. *Public goods: a survey of experimental research, in Roth, Kagel (Eds.),*The Handbook of Experimental Economics. Princeton University Press, Princeton, 1995.

[21] Martin Lettau and Harald Uhlig. Rules of thumb versus dinamic programming. *American Economic Review*, 89:148–174, 1999.

[22] George Mailath. Do people play nash equilibrium? lessons from evolutionary game theory. *Journal of Economic Literature*, 36:1347–1374, 1998.

[23] John Miller and James Andreoni. Can evolutionary dynamics explain free-riding in experiments? *Economic Letters*, 36:9–15, 1991.

[24] Thomas R. Palfrey and Jeffrey E. Prisbrey. Altruism, reputation and noise in linear public goods experiments. *Journal of Public Economics*, 61:409–427, 1996.

[25] Thomas R. Palfrey and Jeffrey E. Prisbrey. Anomalous behavior in public goods experiments: how much and why? *American Economic Review*, 87:829–846, 1997.

[26] Thomas Riechmann. Learning and behavioral stability: an economic interpretation of genetic algorithms. *Journal of Evolutionary Economics*, 9:225–242, 1999.

[27] Martin Sefton and Richard Steinberg. Reward structures in public good experiments. *Journal of Public Economics*, 61:263–287, 1996.

[28] Amartya Sen. Rational fools: a critique of the behavioral foundations of economic theory. *Philosophy and Public Affairs*, 6, 1977.

[29] Joep Sonnemans, Arthur Schram, and Theo Offerman. Strategic behavior in public good games: when partners drift apart. *Economic Letters*, 62:35–41, 1999.

[30] Robert Sugden. Consistent conjectures and voluntary contributions to public goods: why the conventional theory doesn't work? *Journal of Public Economics*, 27:117–124, 1985.

[31] E. C. M. van der Heijden, J .H. M. Nelissen, J. J. M. Potters, and H. A. A. Verbon. Transfers and the effect of monitoring in an overlapping-generations experiment. *European Economic Review*, 42:1363–1391, 1998.

[32] John van Huyck, Frederik Rankin, and Raymond Battalio. What does it take to eliminate the use of a strategy strictly dominated by a mixture? *Experimental economics*, pages 129–150, 1999.

[33] Fernando Vega-Redondo. Markets under bounded rationality: from theory to facts. *Investigaciones Económicas*, 23:3–26, 1999.

[34] Nicolas J. Vrien. An illustration of the essential difference between individual and social learning and its consequences for computational analysis. *Journal of Economic Dynamics and Control*, 24:1–19, 2000.

[35] Marc Willinger and Anthony Ziegelmeyer. Framing and cooperation in public good games: an experiment with an interior solution. *Economic Letters*, 65:323–328, 1999.