Emerging Cooperation in a Public Goods Game with Competition

Ana L. C. Bazzan* and Roberto da Silva
Inst. de Informática, UFRGS
Caixa Postal 15064
91501-970 P. Alegre, Brazil
{bazzan,rdasilva}@inf.ufrgs.br

Sílvio R. Dahmen
Inst. de Física, UFRGS
Caixa Postal 15051
90570-051 P. Alegre, Brazil
dahmen@if.ufrgs.br

Alexandre T. Baraviera
Inst. de Matemática, UFRGS
91509-900 P. Alegre, Brazil
baravi@mat.ufrgs.br

ABSTRACT

In public goods games individuals contribute to create a benefit for a group. Unfortunately they also attract free-riders who enjoy the benefits without contributing. Despite this, in reality, cooperation does not collapse. Several explanations exist for this phenomenon. In our work, individual behavioral rules play an important role in the emergence of a global, cooperative behavior. In our setting, the individual contribution depends on the motivational level of the individual, which depends on the wealth of neighbors. This in turn is associated with the wealth of the whole society. Previous works have used global persistence to analyze that wealth. Here we introduce new elements in the study of the public goods game, trying to bring the model closer to real world situations. We start by analysing the behavior of agents when they interact in a grid. Further, we introduce competition in the game, i.e. agents can select to contribute to two public goods, instead of just one. Results show an increase in wealth, hence in cooperation, when agents perform coordinated selections of goods, and when they imitate the neighbor with highest wealth.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence *Multiagent systems, Coherence and coordination*

General Terms

Economics

Keywords

Public Goods Game, Agent-based Simulation, Evolutionary Game Theory

1. INTRODUCTION

The issue of emerging macro-behavior from micro-rules has been studied in areas such as complex systems, Alife, microeconomics, cognitive science, AI, etc. In the latter, the agent-based approach has proven to be an effective and efficient one to analyze macro behavior arising from micro rules in classical scenarios of social sciences ranging from social simulation (e.g. of ancient civilizations) to artificial markets (economy). In the context of economy, there has been successful reports of modeling and simulation of economic processes as dynamic systems of interacting agents, the so-called agent-based computational economics [5], where cognitive science, evolutionary economics, and computer science play a role. Analytical methods are of little help here as the relationship among the micro rules are complex and so the forecast regarding the macro behavior. In order to analyze this behavior, we use agent-based simulation. This approach facilitates the modeling of what is sometimes referred to as "irrational" decision-making since they deal with the role emotions and motivational aspects play.

Specifically, we are interested on agent-based simulation of public goods games. In these games, individuals incur a cost to create a benefit for a group. Just think about blood donation, recycling, using solar energy, etc. They are problems because free-riders do enjoy the benefits created by the group without contributing themselves. Because free-riders are attracted by the benefits and proliferate, one may expect that eventually cooperation will collapse. However, human societies have somehow managed to solve this kind of problems. Therefore, there has been a great interest in public goods problems or dilemmas, and many researchers try to contribute to an understanding of the nature of these problems. The most popular explanations are based on signaling, reputation, and sanctions. See [2] for an overview.

[1] reports experiments with real subjects playing the public goods game (henceforth PGgame) when two subjects can select among two institutions to contribute. However they are primarily interested in studying which effects punishment has.

In a previous work [4], we have described the evolution of dynamics of the relationship among agents who are locally constrained, meaning that each agent has relationship with the two closest neighbors. The contribution by agents is modulated by a binary variable called "motivation" which is based on the actions of their nearest neighbors. In [3] we have studied the changes in persistence when agents are no longer locally constrained and interact in a small-world like scenario.

Here, we also use this idea of grounding the motivation on the state of close neighbors. However we extend this in three directions.

First, agents can inspect the profits collected by neighbors and imitate the one that has collected the most points. With this we wish to investigate whether there is an increase in the rate of cooperation in the population.

*Author partially supported by CNPq

Second, we perform the analysis of public goods games in grid-like structures.

Third, and more important, we investigate what happens when agents can select between *two* goods.

In the next subsection, we present the model for the PGgame. Section 3 discuss the scenarios and details of the simulation settings, as well as the results. Section 4 reports some preliminary conclusions.

2. MODEL FOR THE PUBLIC GOODS GAME

In its original formulation, this game deals with public spending on some work for the comunity: roads, libraries, etc. Players are offered the opportunity to contribute to a common pool; benefits (obtained from tolls, membership fees) are equally distributed among all participants irrespective of their contributions. Clearly it would be "fair" for people to pay the same quantity for those items. However individuals are different, as they have different social conditions and different stances which means that some contribute less than others. This being common-knowledge, if one assumes each player as rational s/he would default and contribute nothing. However this is not what occurs in reality.

In order to give this model a realistic taste, we let agents interact and contribute taking into account the actions of their immediate neighbors. Here, $L = 225$ individuals start the game with a quantity $w_0 = 5$ of money and can contribute a quantity $S_i \in [0, 2]$.

Actions of each individual are determined by a binary variable we call motivation, whose update depends on the return the agent gets. Motivation is modelled by a binary variable $\sigma_i \in \{0, 1\}$ where $\sigma_i = 1$ means an agent is motivated while $\sigma_i = 0$ means it is not. This abstraction aims at capturing issues such as return prospects as perceived by agents.

In our setting, agents are motivated when their return is positive, and unmotivated otherwise (except in the first round, when the motivation is determined randomly with equal probability, as in line 5 of Algorithm 1). Thus, motivation σ_i is updated at each time t in the following way: if $R_i(t) > 0$, then $\sigma_i(t+1) = 1$; otherwise $\sigma_i(t+1) = 0$.

Return, as in standard PGgames, depends on the quantity contributed by the whole society of agents, i.e. it is a function of the average contribution. Besides, it is modulated by a random variable, the interest rate $r \in \{0, 1\}$ (see line 14 in Algorithm 1).

We compute the average contribution of all L agents in the $t-$th iteration as:

$$Q(t) = \frac{1}{L} \sum_{i=1}^{L} S_i(t) \tag{1}$$

As said, to keep the model simple, we assume that the overall contribution is modulated by a random variable r uniformly distributed in $r \in [0,\ 1]$, and the return per agent is:

$$R_i(t) = Q(t) \left[\frac{1}{2} + r \right] - S_i(t) \tag{2}$$

According to this formula, profits ($1/2 < r < 1$) and losses ($0 < r < 1/2$) are allowed only within a range which depends on the mean contribution $Q(t)$. Practically individual agents can win or lose money. Besides, at each time, each agent has an accumulated wealth given by:

$$T_i(t+1) = T_i(t) + R_i(t) \tag{3}$$

Motivation Level	Contribution Level
$\sigma_i = 1$ and $\sigma_{i+1} = 1$	$S_i = 2$
$\sigma_i = 0$ and $\sigma_{i+1} = 1$	$S_i = 1$
$\sigma_i = 1$ and $\sigma_{i+1} = 0$	$S_i = 1$
$\sigma_i = 0$ and $\sigma_{i+1} = 0$	$S_i = 0$

Table 1: Contribution rules mapping motivation levels to contribution, for each pair of agents

Algorithm 1 Update of Motivation Level, Contribution Level, Return, and Wealth

```
1:  INPUT: global variable ring? // ring or grid configuration ?
2:  INPUT: global variable t_max // max. time steps
3:  INPUT: global variable w_0 // initial amount of money
4:  for each agent i do
5:     set motivation level randomly
6:  end for
7:  while not t_max do
8:     read global variable r // modulation factor (interest rate)
9:     for each agent i do
10:       for each neighbor j do
11:          get motivation of j
12:       end for
13:       compute contribution level according to the motivation
          rules in Table 1 and realize the contribution
14:       compute return (Eq. 2) and wealth (Eq. 3)
15:       update motivation according to the return
16:    end for
17:    for each agent i do
18:       find best neighbor
19:    end for
20:    for each agent i do
21:       copy motivation of best neighbor
22:    end for
23: end while
24: END
```

where $T_i(1) = w_0, i = 1, ..., L$.

In real life gains (or losses) in any kind of investment fluctuate, e.g. stocks/derivatives/bonds. So the idea of having a profit which varies is an attempt to give the model a realistic flavor.

In the *basic* version of the algorithm, each agent has relationship with just the two closest neighbors[1]. Thus interactions occur between agents labelled i and $i + 1$ ($i - 1$ also interact with i on its turn), and follow the rules presented in Table 1.

In each time step, agents (synchronously):

1. get motivation level of all their neighbors and compute a contribution level accordingly;
2. contribute;
3. compute return and wealth;
4. update their motivation level;
5. copy motivation from the best neighbor (est here means the one with highest return).

This is shown in Algorithm 1.

An important point is that if an agent is unmotivated, it influences its neighbors' motivational state and hence, their contribution level.

[1] Please notice that this situation will change when we discuss the grid configuration.

This is a first key difference to the present setting and the one in [4]. There, the contribution level of each agent i is regulated mainly by rules that account for both agent i's motivation and by each of its neighbors j. This means that the contribution level of i somehow reflects the *average* motivation in the neighborhood, while we simply copy the motivation of the *best* player. We have tested different ways to search for the best player.

The second distinction is that we have also tested grid scenarios besides ring ones. The third difference is that we allow agents to decide to select one among multiple goods to contribute an amount S_i. These differences are discussed in sections 3.2, 3.3, and 3.4.

3. SIMULATIONS AND RESULTS

In order to evaluate the wealth of the society, we compute the quantity $I(t) = \frac{1}{L}\sum_{i=1}^{L} T_i(t)$ (average accumulated wealth) over a certain number of runs or repetitions of the simulation, N_s. Regarding the results given in the next subsections we set $N_s = 100$.

3.1 Basic Scenario

In the basic setting [4], simulations were performed with the update of motivations according to Equation 1, with no further change.

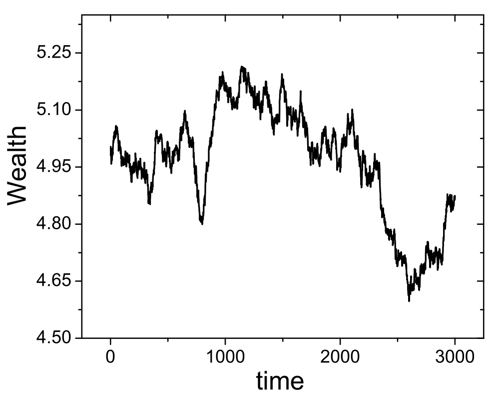

Figure 1: Average of the wealth as a function of time, $L = 250$.

The curves in Figure 1 depict a stochastic behavior similar to a random walk, achieved using the algorithm given in [4].

In the next subsection we compare this curve with our simulations.

3.2 No Competition

We have simulated the basic setting (described in [4]) with three modifications: first we use the idea of copying the motivation level from the best player in the neighborhood (see Algorithm 1). Further, we have two ways to search for the best player (line 18 of that algorithm). In the case described in this subsection, the copy is *probabilistic* based on either the *return* achieved by the neighbors, or on their *wealth*. This means that the player with the best return or wealth is selected with higher probability but not in a deterministic way. Finally, we have run configurations where agents interact in a grid, instead of in a ring.

Results of the simulations are shown in Figure 2. These show that, in general, the quantity $I(t)$ (average wealth, over all agents along time) is higher when agents interact in a ring. This might be explained by the fact that in the grid agents have more neighbors and hence, the chance that unmotivated agents are around is higher.

Another conclusion is that $I(t)$ fluctuates slightly more when agents copy the motivation of neighbors according to their wealth than when they copy motivation according to their returns. This result indeed goes counter to the intuition that wealth, being an accumulated quantity, would lead to less fluctuation.

Up to here we have tested cases in which there is only one good attracting agents's contributions. In the next section we discuss the situation in which agents may select between two goods to make their contributions. We only discuss the ring case there.

3.3 Competition, Best Neighbor Found Deterministically

In order to perform the simulations, we have changed the basic algorithm (Algorithm 1) in the following way: there are now two modulation quantities r_0 and r_1 attached to goods zero and one respectively. The basic algorithm must be changed accordingly (see Algorithm 2, where we omit lines that did not change).

Another difference to the previous situation is that we give a bonus to agents when two neighhors opt to contribute to the same good, as in real life when close people exchange information about their investment and contribution decisions. The reason for the inclusion of the bonus in the model is that we want to analyze whether a kind of coordinated selection would improve wealth. We call this bonus B and set B to different values, from 0 to 0.2 meaning that a bonus of up to 20% of the return can be paid when agents do coordinated selections. Of course this introduces a bias in Eq. 2 that changes the nature of the random walk.

Further, we have also run experiments on the deterministic copy of the motivation of the best neighbor, also based on return or wealth. This means that the best player is copied with 100% of probability as opposed to the case in the previous section.

Figure 3 shows the result of the simulation under these conditions for $B = 0$. We reproduce the curve "ring best return" from Figure 2 in order to facilitate the comparison. One can conclude that competition reduces wealth, as expected, because of the lack of coordination between agents.

Next we analyze what happens when we increase B. Figure 4 shows that, for $B = 0.2$ (20% of increase in the return if two neighbors make coordinated selections, i.e. same good), the wealth is

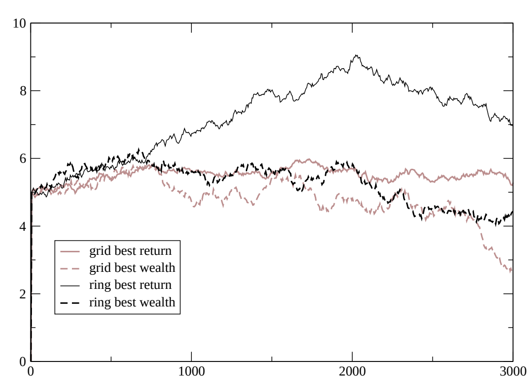

Figure 2: Average wealth $I(t)$ as a function of time; no competition; ring and grid.

Algorithm 2 Competition Between Two Goods: Update of Motivation Level, Contribution Level, Return, and Wealth

INPUT: global variable B // bonus for coordinated choice of good
...
for each agent i **do**
  set motivation level and good randomly
end for
while not t_{max} **do**
  read global variable r_0 and r_1 // modulation factor
  for each agent i **do**
    ...
    compute return (Eq. 2), add bonus, compute wealth (Eq. 3)
    update motivation according to the return
  end for
  for each agent i **do**
    find best neighbor
  end for
  for each agent i **do**
    copy motivation and choice of good from best neighbor
  end for
end while
END

tremendously higher. Returns are almost never negative and hence the wealth increases almost monotonically.

3.4 Competition, Best Neighbor Found Probabilistically

When the best neighbor is not necessarily the one with higher return or wealth, one may expect more fluctuation regarding the wealth. This is what happens and can be observed in Figures 5 and 6.

When the bonus is zero (Figure 5), results are similar to those depicted in Figure 3 (for the case of competition), except that in Figure 5 one observes both more fluctuation and a slightly higher wealth, especially as time passes. This increase is counter intuitive as one could have expected that selecting the best neighbor deterministically based on the higher return would lead the overall wealth to a higher value as well.

When the bonus is increase to $B = 0.01$ (also in Figure 5) which is only 1% of bonus for coordinated choices of goods between each two agents, one sees that the wealth increases as well. This means that very small bias in Eq. 2 are enough to make the wealth increase.

In Figure 6 we plot the increase in overall wealth when we increase the bonus to 5%, 10%, and 20%.

Also in Figure 6 we show results for finding the best neighbor based not on the return in the neighborhood but on the wealth achieved by neighbors. In this case we plot only for $B = 0.2$. As in Figure 2, copying players based on their wealth and not on their return brings more wealth to the society. This can be seen comparing both darker lines in Figure 6. Full line depicts copying best return while dashed line means copying best wealth.

4. CONCLUDING REMARKS

Networks of coupled individual elements are not only a paradigm for studying artificial systems, but also an artifact that appears often in Nature and social systems. Individual actions alone cannot usually alter significantly the dynamics of these systems. However, collective behavior, as for example the selling stampede caused by some expectation of political turmoil may drive markets to crash. In social sciences and economics, traditional methods of analysis

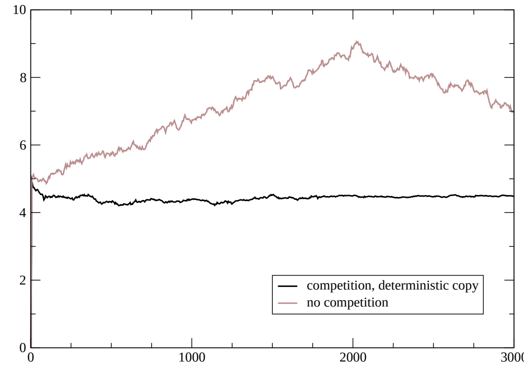

Figure 3: Average wealth as a function of time; competition; copy best return deterministically; $B = 0$.

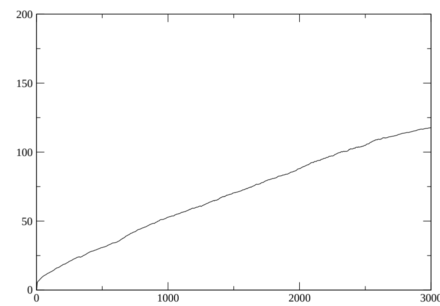

Figure 4: Average wealth as a function of time; competition; copy best return deterministically; $B = 0.2$.

in many-actor systems are being replaced by approaches able to explicitly deal with decision-making modulated by the interaction among individuals. This is important in many areas of AI such as multiagent systems and Alife. However the gap between individual rules and macro behavior is not very well studied as this problem has many facets and is domain dependent.

Here we explore this problem in a public goods game, a metaphor for many interactions among cooperative and non-cooperative agents, within well defined neighborhoods. The basic fact in all settings is that an agent might feel motivated by its peers to act as they do, based on feelings of belonging to a group (“to go with the pack”).

In this work we presented a simple model of a society of L economic agents, where each can invest a discrete quantity based on neighbors’ motivation level. The profit of a group fluctuates stochastically and influences the return and the motivation level of individual agents.

Our results on the analysis of the accumulated wealth seem to indicate that there is no big difference between interactions that occur on a grid or on a ring. Further, copying the motivation of the best neighbor is a way to ensure that cooperation arises, under certain conditions.

The idea of having agents selecting between two pools to which they can contribute money is an interesting one as it creates new sit-

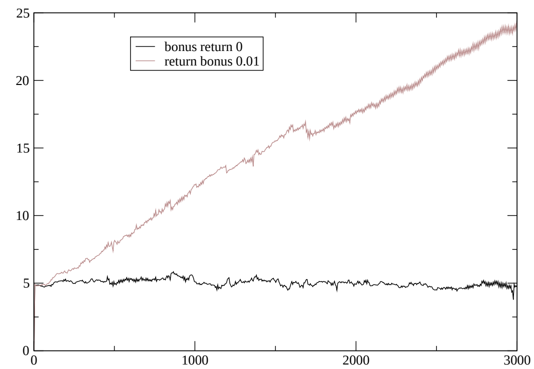

Figure 5: Average wealth as a function of time; competition; copy best return probabilistically; $B = 0$ **(darker) and** $B = 0.01$ **(lighter).**

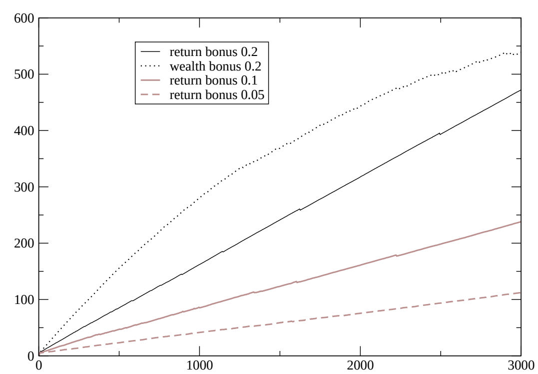

Figure 6: Average wealth as a function of time; competition; copy best return probabilistically; $B = 0.05$ **to** $B = 0.2$**.**

uations for the public goods game. When agents receive no bonus for coordinated selections of pools, the level of wealth tends to be lower than in the situation with only one pool. As soon as there is a bonus, no matter how little, cooperation tends to increase, and hence, wealth.

Our current effort is being directed to the reproduction of the results achieved in [1], where authors investigate the behavior of real subjects playing the PGgame with sanction. We are not running experiments in laboratory; rather we are doing this via a learning-based model which, if well calibrated, can be used for further investigations without the need of running the actual experiments.

Acknowledgments

The first author is partially supported by CNPq. We would like to thank the anonymous reviewers for their valuable suggestions.

5. REFERENCES

[1] O. Gürek, B. Irlenbusch, and B. Rockenbach. The competitive advantage of sanctioning institutions. *Science*, 312(5770):108–111, April 2006.

[2] J. Henrich. Cooperation, punishment, and the evolution of human institutions. *Science*, 312(5770):60–61, April 2006.

[3] R. Silva, A. Baraviera, S. R. Dahmen, and A. L. C. Bazzan. Dynamics of a public investment game: from nearest-neighbor lattices to small-world networks. In C. Bruun, editor, *Advances in Artificial Economics, The Economy as a Complex Dynamic System*, number 584 in Lecture Notes in Economics and Mathematical Systems, chapter 16, pages 211–233. Springer, 2006.

[4] R. Silva, A. L. C. Bazzan, A. Baraviera, and S. R. Dahmen. Emerging collective behavior in a simple artificial financial market. In F. Dignum, V. Dignum, S. Koenig, S. Kraus, M. P. Singh, and M. Wooldridge, editors, *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems*, volume 1, pages 313–319. ACM, July 2005.

[5] L. Tesfatsion. Agent-based computational economics: A constructive approach to economic theory. In K. L. Judd and L. Tesfatsion, editors, *Handbook of Computational Economics: Agent-Based Computational Economics*, volume 2 of *Handbooks in Economics Series*, chapter 16. North-Holland, 2006.