

# Punishment and Gossip: Sustaining Cooperation in a Public Goods Game

Francesca Giardini<sup>1</sup>, Mario Paolucci<sup>1</sup>, Daniel Villatoro<sup>2</sup>, and Rosaria Conte<sup>1</sup>

<sup>1</sup> ISTC-CNR, Roma, Italy

{francesca.giardini,mario.paolucci,rosaria.conte}@istc.cnr.it

<sup>2</sup> IIIA-CSIC, Barcelona, Spain

dvillatoro@iia.csic.es

**Abstract.** In an environment in which free-riders are better off than cooperators, social control is required to foster and maintain cooperation. There are two main paths through which social control can be applied: punishment and reputation. Our experiments explore the efficacy of punishment and reputation on cooperation rates, both in isolation and in combination. Using a Public Goods Game, we are interested in assessing how cooperation rates change when agents can play one of two different reactive strategies, i.e., they can pay a cost in order to reduce the payoff of free-riders, or they can know others' reputation and then either play defect with free-riders, or refuse to interact with them. Cooperation is maintained at a high level through punishment, but also reputation-based partner selection proves effective in maintaining cooperation. However, when agents are informed about free-riders' reputation and play Defect, cooperation decreases. Finally, a combination of punishment and reputation-based partner selection leads to higher cooperation rates.

**Keywords:** Reputation, Punishment, Agent-based simulation, Social Control, Public Goods Game.

## 1 Introduction

Social control is an emergent social phenomenon, which allows the costs of pro-social behavior to be redistributed over a population in which cooperators live side-by-side with non-cooperators. Some specific phenomena are usually subsumed under the large heading of social control, including ostracism [3,15] and altruistic punishment. The latter is defined as a costly aggression inflicted to cheaters by members of the group who did not necessarily undergo attacks from the punished, nor get direct benefits out of the sanction applied [9]. According to *strong reciprocity* theory, the presence of individuals who altruistically reward cooperative acts and punish norm violating behavior at a cost to themselves sustains cooperation and promotes social order [11]. However, the act of punishment results in an immediate reduction of welfare both for the punisher and for the punished individual, thus posing several problems, like efficiency [27,7], and the risk of counter-aggression [19].

An alternative solution can be found in weak reciprocity supported by reputation. Knowing about others' past behaviors is crucial to avoid cheaters and select good partners, especially when the group is large and it is not possible to directly witness all the interactions. Moreover, reputation allows the costs of social control to be reduced and distributed among individuals [12]. The importance of reputation in supporting cooperation has been proven in laboratory experiments [28], evolutionary models [20], and simulation studies [8]. Reputation is a signal that conveys socially relevant information about one's peers, and plays a fundamental role in identifying cheaters and isolating them.

Notwithstanding their importance in supporting cooperation, costly monetary punishment and reputation spreading have never been directly compared. Moving from the work of Carpenter [4], in which the provision of public goods is not negatively affected by the size of the group but by the ability of mutual monitoring among agents, we enriched the original model by designing agents who spread reputation about their previous partners. Extending previous research on cooperation and reputation [8] and on cooperation and punishment [27], we explore the performance of punishment and reputation as mechanisms for social control, and we test their effects on cooperation rates both in isolation and in combination.

Here we present a simulation platform to compare the effectiveness of costly punishment and reputation spreading in maintaining cooperation in a population in which defectors have a selective advantage because they exploit others' contributions, without paying the costs of cooperation. Using a Public Goods Game [4], we measure cooperation rates in mixed populations in which there are pure cooperators, pure cheaters and agents who play reactive strategies. Our contribution adds to existing literature in three ways:

1. we introduce a systematic exploration of two different social control mechanisms, and we test them also in combination with other parameters, like the group size and the costs of punishing;
2. we specify two different mechanisms for reputation: *Refuse* and *Defect* (they will be extensively explained in Section 3). *Refuse* is a partner choice mechanism which permits gossipers to avoid free-riders, whereas *Defect* is a social control mechanism that leads gossipers to defect against non-cooperators. Both these mechanisms are present in human societies, in which we use reputational information to avoid cheaters (when this is possible), or to treat them as we expect them to treat us. This difference between these reactions makes it important to compare them and to understand the conditions that make one mechanism more effective than the other.
3. we assess the extent to which reputation spreading and punishment are comparable mechanisms for social control, by comparing directly agents' average contributions when costly monetary punishment and reputation spreading are available.

In Section 2 we will introduce related work, in Section 3 we describe the simulation model, and Section 4 will present the simulation results. In Section 5 we will draw some conclusions and we will also sketch some ideas for future work.

## 2 Related Work

In their evolutionary history, humans have developed several mechanisms for the emergence and establishment of social norms [2]. Punishment and reputation are among the most widespread and effective mechanisms to sustain cooperation and they are specially interesting for virtual societies in which the efficacy of enforcing mechanisms is limited by a combination of factors (like their massive size, their spontaneity of creation and destruction, and dynamics). There is a large body of evidence showing that humans are willing to punish non-cooperators, even when this implies a reduction in their payoffs [10], and this is true also in simulation settings. Villatoro et al. [27] have analyzed the effect of sanctioning on the emergence of the norm of cooperation, showing that a monetary punishment accompanied with a norm elicitation, that is, sanctioning, allowed the system to reach higher cooperation levels at lower costs, when compared with other punishment strategies.

Reputation is, along with punishment, the other strategy used to support cooperation, even if it works in a completely different way. If punishing means paying a cost in order to make the other pay an even higher cost for his defection, reputation implies that the information about agents' past behavior becomes known, and this allows agents to avoid ill-reputed individuals. In Axelrod's words [1]: *"Knowing people's reputation allows you to know something about what strategy they use even before you have to make your first choice"* (p.151). The importance of reputation for promoting and sustaining social control is uncontroversial and it has been demonstrated both in lab experiments [23] and in simulation settings, in which reputation has proven to be a cheap and effective means to avoid cheaters and increase cooperators' payoffs [21].

When partner selection is available, reputation becomes essential for discriminating between good and bad partners, and then to be protected against exploitation. Giardini and Conte [13] presented ethnographic data from different traditional societies along with simulation data, showing how reputation spreading evolved as a solution to the problem of adaptation posed by social control, and highlighting the importance of gossip as a means to reduce the costs of cheaters' identification. The effect of partner selection has been studied also by Perrau and others [6], who have analyzed the effect of ostracism in virtual societies, obtaining high levels of tolerance against free-riders. However, in the work of Perrau, agents do not explicitly transmit information about other agents, they only reason about the interactions.

In the multi-agent field, several attempts have been made to model and use reputation, especially in two sub-fields of information technologies, i.e., computerized interaction (with a special reference to electronic marketplaces), and agent-mediated interaction (for a review, see [18]). Models of reputation for multi-agent systems applications [29,24,16] clearly show the positive effects of reputation, and there are also interesting cases in which trust is paired with reputation (for a couple of exhaustive reviews, see [22,26]).

More specifically, Sabater and colleagues [25,5] developed a computational system called REPAGE in which different kinds of reputational information were

taken into account and the role of information reliability in a market-like simulation scenario was addressed. Analogously, Giardini and colleagues [8] showed that reputation was a means to punish untruthful informers *without bearing the costs of further retaliation*, at the same time protecting the system from collapsing. In addition, for certain percentages of cheating rates, in both studies, the authors showed that reputation played a relevant role in enhancing the quality of production in an artificial cluster of interacting firms.

### 3 The Model

Moving from the simulation framework developed in [4], we designed a simulation platform in NetLogo in order to compare the performance of costly punishment and reputation spreading in mixed populations in which different types of agents play a Public Good Game (PGG), the classical experimental model used to investigate social dilemmas [17]. In this game, agents decide whether to free-ride or to contribute<sup>1</sup> a fixed amount (a contribution of 1 unit) to a public pool. The sum of all the contributions is multiplied by a benefit factor (set to 3 in the current model<sup>2</sup>) and the resulting quantity is divided amongst all the participants in the group, without considering their contributions. This is a classic public good where free-riding would be the utility maximizing strategy at the individual level; however, if all agents adopted that strategy, this would result in the overexploitation of resources and in a worse outcome at the group level (the so-called *Tragedy of the Commons* [14]).

---

#### Algorithm 1. Description of punisher's behaviors

---

```

for Number of Timesteps do
  Random group formation of the population;
  Agents take First Stage decision;
  Gather and Distribution of the Public good in each group;
  First Stage Decisions are made public within the group;
  Agents make Second Stage decision;
  Punishment Execution;
end for

```

---

Agents are either non-reactive or reactive types. In the former category we find *Cooperators* (C), who always contribute to the common pool, and *Free-riders* (FR), who never contribute to the common pool. Reactive agents change their

<sup>1</sup> Note that the decision in this framework is binary whether to cooperate or not. In other works, specially in those of experimental economics, this decision has to be taken in a continuum, deciding how much to contribute from a total amount of money common to all agents.

<sup>2</sup> According to the game design, in order for contribution to be irrational for a utility-maximizer individual, the tokens in the pot must be multiplied by an amount smaller than the number of players and greater than 1.

behaviors in response to the percentage of detected free-riders in their group: when the number of known defectors in a group is too high, i.e., it exceeds a certain threshold, agents become active and their strategy changes (as described in Algorithm 1). Each and every agent is endowed with an initial amount of 50 points that can be used to cooperate or to punish others; regardless of the strategy, agents are culled from the game when their payoff goes to zero and they are not replaced. The other strategies are:

- *Tit-for-Tat* (TFT): They start as cooperators but if active, they start free-riding until exiting from the active state.
- *Nice Punishers* (NP): They contribute in the passive state; once active, they punish free-riders at a cost to themselves and making *FRs* pay a cost, but they continue to cooperate in the PGG. This behavior will continue until the agent exits the active state (by being assigned in a group with a number of cheaters below the threshold).
- *Mean Punishers* (MP): They contribute in the passive state; once active, they punish free-riders and free ride themselves in the PGG until they exit the active state.

To the above types, developed by Carpenter [4], we add two more types, in order to compare punishment and reputation. Active Gossipers transmit and receive information and they integrate their personal experience and the reputational information received from other Gossipers in order to react against Free-Riders. Gossipers' behavior is described in Algorithm 2 (2). Gossipers can be:

- *Nice Gossipers* (NG): Agents contribute in the passive state; once active, they start spreading information about free-riders, and cooperate in the PGG.
- *Mean Gossipers* (MG): Agents contribute in the passive state; once active, they start spreading information about free-riders, but they always defect in the PGG when active.

---

**Algorithm 2.** Description of gossipers' behaviors

---

```

for Number of Timesteps do
  Random group formation of the population;
  if Group has bad reputation then
    Apply reputation strategy
  else
    Take First stage decision according to active/passive status.
  end if
  Gather and Distribution of the Public good in each group;
  First Stage Decisions are made public within the group;
  Agents make Second Stage decision;
  Reputation diffusion;
end for

```

---

In this work, we explore the efficacy of two different ways of using reputation. Reputation allows agents to be informed about other members of their group before interacting with them, hence this information can be used in two ways, either to retaliate against free-riders or to refrain from interacting with them. The former modality, called “Defect”, consists in Gossipers not contributing to the PGG when playing in a group containing too many cheaters. The latter modality, “Refuse” allows agents to refuse the interaction by skipping a turn, taking their contribution away from the pools, and paying the price of not receiving any dividend, if the number of known defectors in their current group assignment would make them active. Refusal cannot be performed twice in a row. In both modalities, Gossipers transmit information about cheaters, informing their peers of the identity of non-cooperators and using the information they receive from them.

## 4 Experiments and Results


We run three different sets of simulations, each one lasting for 100 time steps. Each experiment was repeated 20 times for each combination of the selected parameters. The cost of contributing to the Public Good was set to 1, and the sum of all the contributions was multiplied by a benefit factor set to 3. The public good, i.e., the resulting quantity, was divided among all the group members, without considering whether they contributed or not. Simulations started with equal proportions of each strategy. Variables of interest are summarized in Table 1.

**Table 1.** Parameters of the simulation

| PARAMETERS                          |     |    |    |    |
|-------------------------------------|-----|----|----|----|
| COST (cost of punishing)            | 0.2 | 1  |    |    |
| PUNISHMENT (cost of being punished) | 2   | 5  |    |    |
| GROUP SIZE                          | 5   | 10 | 25 | 50 |
| INFORMATION TRANSMITTED             | 0   | 1  | 10 |    |

In our first experiment we tested the effect of punishment as a partner control mechanism on cooperation rates measured as the total number of agents playing C divided by the total number of active agents per time step. Nice and Mean Punishers became active when they detected more than 20% of defections in their group. Punishing costed the punisher  $x$  and the punished agent  $y$ , with  $y \geq x$ . We identified 4 different combinations of Punishment and Cost: LpLc (low punishment, low cost), HpLc (high punishment, low cost), LpHc (low punishment, High cost) and HpHc (high punishment and high cost). The cooperation rates are affected by group size, and they change according to the different combinations of punishment and cost (Figure 1). When punishment is low and the cost is high, cooperation rates are the lowest for every group size.

Both HpHc and HpLc allows cooperation to reach 100%, no matter the group size and in a quite short amount of time. This is probably due to the fact that high punishment leads Free-Riders to use up all their resources, while protecting the other agents. Moreover, once the percentage of tolerated Free-riders goes below the fixed threshold, all the other populations contribute, thus maintaining cooperation stable and complete.




**Fig. 1.** Cooperation rates over time for different group sizes in mixed populations with C, FR, TFT, NP and MP. When punishment is costly (HpLc, HpHc), Free-riders are easily controlled and cooperation can be maintained. When the cost is high and the punishment is not (LpHc) the cooperation rate is always lower than in the other cases.

Our second experiment was designed to test the effects of information spreading in isolating Free-Riders and maintaining cooperation. We were also interested in comparing the efficacy of the two modalities, and to assess which one worked better. We tested the effectiveness of “Refuse” and “Defect” for different amounts of information ( $i$ ) available. Manipulating the number of gossips transmitted  $\{i = 0; 1; 10\}$  for each time-step, we wanted to test whether the amount of information had an effect on the ability of gossipers to react directly (playing “Defect”), or indirectly (playing “Refuse”) against Free-riders.

When agents used the “Defect” modality, cooperation rates declined for every group size. In the “Defect” modality (upper part of Figure 2), cooperation rates

were higher in the first 20 periods, but they showed a rapid decrease in the last 20 periods. Increasing the amount of information about other agents made this decrease steeper, even if this pattern remained stable for different group sizes. This is due to the fact that playing “Defect” triggered even more defections, thus reducing the overall cooperation rate and making the decrease even steeper when the amount of available information was higher.




**Fig. 2.** Cooperation rates over time for Defect (top) and Refuse (bottom) strategies for Gossipers. In larger groups (25 and 50 agents) the amount of information has a negative effect on cooperation rates when the reaction is based on defection, as compared to the Refuse strategy. In this latter case, cooperation rates are stable and quite high.


On the contrary, when gossipers played “Refuse”, cooperation levels were quite high for all group sizes. Small groups showed higher cooperation rates (Figure 2, bottom panel), because refusing dangerous interactions in small groups was effective in isolating cheaters and in making cooperators interact with other cooperators. Taking into account the first periods of the simulation experiments, we saw that there are interesting differences in the cooperation rates, not only between groups of different sizes, but also among situations with zero, one or ten items of information (3).

In a scenario in which all strategies are loaded, we compared the performance of populations in which all the strategies were present (C, FR, TFT, NP, MP, NG and MG) for different combinations of costs (HpHc, HpLc, LcHp, LcLp), for different group size (5 and 50). Figure 3 (4) shows that punishment and reputation can have a combined effect that boosts cooperation to 1 in large groups in every situation but in the LpHc. The negative effect of Defect is confirmed also when there are both Punishers and Gossipers, and it is also made more relevant by the higher amount of information.





**Fig. 3.** Cooperation rates in the first 40 ticks of the simulation experiment. When Gossipers can refuse the interaction, receiving more information (10 gossips) increases the cooperation rates in the initial phases of the simulation, especially for groups of 5 and 10 agents.



**Fig. 4.** Cooperation rates for all populations for different combinations of costs of punishment and number of gossips for small (5 agents) and large groups (25 agents). Both small and large groups achieve the highest percentages of cooperation when punishment is costly and being punished is costly as well (HpHc), even if this happens mostly in the larger group, where cooperation rate is 1 for all the populations except for the one including agents playing “Defect” and transmitting 10 gossips. This population has a good performance only when the group is big (25 agents) and the costs of punishment are low (LpLc). “Refuse” is the best strategy and its efficacy is increased by the amount of information exchanged.

## 5 Discussion and Future Work

In this work we proposed a direct comparison between reputation spreading and costly monetary punishment for controlling free-riding in a Public Goods Game. These strategies, punishment and reputation spreading, were tested in order to investigate their effects on cooperation rates in mixed populations, and to assess how effective they are in preventing cooperation from extinction in different settings. Our contribution added to existing literature by introducing a systematic exploration of different elements, in isolation and in combination. The effects of these strategies in isolation have been tested in a variety of settings and environments, but it is also interesting to compare them directly. Costly punishment is an effective strategy in promoting cooperation, and here it works quite well in selecting out the free-riders, especially when the cost of being punished is much higher than the cost of punishing, as expected. Even more interesting is our finding that in larger groups the effect of punishing, even when it is not costly, can support cooperation much better than in smaller groups.

We also introduced two different strategies based on reputation with the goal of evaluating two alternative ways of using information about cheaters. The *Refuse* strategy can be considered pro-active and it proved to be the best strategy in terms of cooperation rates, for every population composition. On the other hand, the *Defect* strategy was designed as a direct retaliation against cheaters, with negative consequences on the average cooperation rates in our populations. These two strategies worked in different ways and are not directly comparable, but their usage may shed new light on the conditions that make one mechanism more effective than the other, also considering that in humans societies the two options are usually available at the same time. It is worth noting that, on average, cooperation rates with Refuse were close to 70%, which is a very high percentage, especially because Gossipers were not able to affect Free-Riders' payoffs directly, as Punishers did, so their expected efficacy was much lower. Nonetheless, being able to use reputation to identify Free-Riders before interacting with them was really effecting in preventing exploitation and fostering cooperation.

This work represents a first step in a process of exploring different combinations of direct, i.e., punishment, and indirect, i.e., reputation, mechanisms for promoting cooperation in social dilemmas. Our platform will allow us to explore single parameters and different combinations of them, with the aim of understanding how they work, and in what way their combination may determine the success of a given strategy. We do not claim that reputation is a better mechanism, but our data show that it is worth exploring the possibility that, for given combinations of parameters, reputation would be more effective in sustaining social control.

Regarding future works, with the advent of social networking platforms like Facebook or Twitter, and motivated by the Living Labs philosophy, a possible extension would be to perform experiments on social networks sites with real subjects. Because of the public access to the messages shared by the users and the underlying social network that connects them, Twitter could be the best platform where to perform our experiment. We plan to identify active communities within

the social platform and made them play with bots developed by us. Our bots will allow us to generate controlled situations, and study the different treatments discussed in this paper. Measures like the virtual earnings of the players, the messages sent amongst users and the changes on the underlying social network will be observed and used as indicators to extract our results.

**Acknowledgements.** The authors wish to thank two anonymous referees for advice. This work has been partially supported by the PRISMA project (PON04a2 Aunder), within the Italian National Program for Research and Innovation (Programma Operativo Nazionale Ricerca e Competitività 2007-2013).

## References

1. Axelrod, R.: The evolution of cooperation. Basic Books (1984)
2. Axelrod, R.: An evolutionary approach to norms. *The American Political Science Review* 80(4), 1095–1111 (1986)
3. Bowles, S., Gintis, H.: The evolution of reciprocal preferences. Working papers. Santa Fe Institute (2000)
4. Carpenter, J.P.: Punishing free-riders: How group size affects mutual monitoring and the provision of public goods. *Games and Economic Behavior* 60(1), 31–51 (2007)
5. Conte, R., Paolucci, M., Sabater Mir, J.: Reputation for innovating social networks. *Advances in Complex Systems* 11(2), 303–320 (2008)
6. de Pinninck Bas, A.P., Sierra, C., Schorlemmer, M.: A multiagent network for peer norm enforcement. *Autonomous Agents and Multi Agent Systems* 21, 397–424 (2010)
7. Dreber, A., Rand, D., Fudenberg, D., Nowak, M.: Winners don't punish. *Nature* 452, 348–351 (2008)
8. Giardini, F., Di Tosto, G., Conte, R.: A model for simulating reputation dynamics in industrial districts. *Simulation Modelling Practice and Theory (SIMPAT)* 16(2), 231–241 (2008)
9. Fehr, E., Gächter, S.: Altruistic punishment in humans. *Nature* 415(6868), 137–140 (2002)
10. Fehr, E., Gächter, S.: Altruistic punishment in humans. *Nature* 415(6868), 137–140 (2002)
11. Fehr, E., Henrich, J.: Is strong reciprocity a maladaptation? on the evolutionary foundations of human altruism. IZA Discussion Papers 712. Institute for the Study of Labor, IZA (February 2003)
12. Giardini, F., Conte, R.: Gossip for social control in natural and artificial societies. *Simulation: Transactions of the Society for Modeling and Simulation International* (2011)
13. Giardini, F., Conte, R.: Gossip for social control in natural and artificial societies. *Simulation* 88(1), 18–32 (2012)
14. Hardin, G.: The Tragedy of the Commons. *Science* 162(3859) (1968)
15. Hirshleifer, D., Rasmusen, E.: Cooperation in a repeated prisoners' dilemma with ostracism (1989)

16. Huynh, T.D., Jennings, N.R., Shadbolt, N.R.: Certified reputation: how an agent can trust a stranger. In: Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS 2006, pp. 1217–1224. ACM, New York (2006)
17. Ledyard, J.O.: Public goods: A survey of experimental research. In: Handbook of Experimental Economics, pp. 111–194. Princeton University Press (1995)
18. Mui, L., Mohtashemi, M., Halberstadt, A.: Notions of reputation in multi-agents systems: a review. In: AAMAS, pp. 280–287 (2002)
19. Nikiforakis, N.: Punishment and counter-punishment in public good games: Can we really govern ourselves? *Journal of Public Economics* 92, 91–112 (2008)
20. Nowak, M.A., Sigmund, K.: Evolution of indirect reciprocity by image scoring. *Nature* 393(6685), 573–577 (1998)
21. Pinyol, I., Paolucci, M., Sabater-Mir, J., Conte, R.: Beyond accuracy. reputation for partner selection with lies and retaliation. In: Antunes, L., Paolucci, M., Norling, E. (eds.) MABS 2007. LNCS (LNAI), vol. 5003, pp. 128–140. Springer, Heidelberg (2008)
22. Ramchurn, Sarvapali, Sierra, C., Godo, L., Jennings, N.R.: Devising a trust model for multi-agent interactions using confidence and reputation. *Applied Artificial Intelligence* 18(9-10), 833–852 (2004)
23. Rockenbach, B., Milinski, M.: The efficient interaction of indirect reciprocity and costly punishment. *Nature* 444(7120), 718–723
24. Sabater, J., Sierra, C.: Reputation and social network analysis in multi-agent systems. In: Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2002), Bologna, Italy, pp. 475–482 (2002)
25. Sabater, J., Paolucci, M., Conte, R.: RePAGE: Reputation and image among limited autonomous partners. *Journal of Artificial Societies and Social Simulation* 9(2) (2006)
26. Sabater, J., Sierra, C.: Review on computational trust and reputation models. *Artif. Intell. Rev.* 24, 33–60 (2005)
27. Villatoro, D., Andrighetto, G., Conte, R., Sabater-Mir, J.: Dynamic sanctioning for robust and cost-efficient norm compliance. In: Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence, IJCAI (2011)
28. Wedekind, C., Milinski, M.: Cooperation Through Image Scoring in Humans. *Science* 288(5467), 850–852 (2000)
29. Yu, B., Singh, M.P.: Distributed reputation management for electronic commerce. *Computational Intelligence* 18(4), 535–549 (2002)